

Market Power in Coal Shipping and Implications for U.S. Climate Policy

Louis Preonas

Department of Agricultural and Resource Economics, University of Maryland, USA

First version received December 2019; Editorial decision February 2023; Accepted September 2023 (Eds.)

Economists have widely endorsed pricing CO₂ emissions to internalize climate change-related externalities. Doing so would significantly affect coal, the most carbon-intensive energy source. However, U.S. coal markets exhibit an additional distortion: the railroads that transport coal to power plants can exert market power. This article estimates how coal-by-rail markups respond to changes in coal demand. I identify markups in a major intermediate goods market using both reduced-form and structural methods. I find that rail carriers reduce coal markups when downstream power plant demand changes due to a drop in the price of natural gas (a competing fuel). My results imply that decreases in coal markups have increased recent U.S. climate damages by \$11.9 billion, compared to a counterfactual where markups did not change. Incomplete pass-through would likely erode the environmental benefits of an incremental carbon tax, shifting the tax burden towards upstream railroads. Still, a non-trivial tax would likely increase welfare.

Key words: Market power, Energy markets, Carbon tax pass-through

JEL codes: L41, Q41, Q54

1. INTRODUCTION

Under perfect competition, an optimal climate policy would tax carbon emissions at their marginal external cost (Pigou, 1932; Nordhaus, 1993). Market power reduces the efficiency of a Pigouvian tax (Buchanan, 1969; Barnett, 1980), and economists have long understood that firms with market power may adjust prices in response to a tax (Cournot, 1838). While most polluting industries are highly concentrated (Fowlie *et al.*, 2016), few studies have estimated how market power impacts the pass-through of an environmental tax.

This article investigates how market power in U.S. coal transportation would impact the efficacy and incidence of a carbon tax. Spatially concentrated mines supply coal to spatially dispersed power plants, and railroad intermediaries can exercise market power in coal shipping (Busse and Keohane, 2007; Hughes and Lange, 2020). If a carbon tax caused rail carriers to reduce coal markups, this would mute the carbon price signal received by power plants and erode the environmental benefits of the tax. It would also shift the carbon tax incidence upstream towards railroad shareholders, and away from electricity consumers.

The editor in charge of this paper was Francesca Molinari.

I start by estimating the size of the coal-by-rail market power distortion. I exploit predetermined cross-sectional variation in market structure: some coal plants are “captive” to a single monopolist rail carrier, while other plants may purchase coal from multiple rail carriers. Due to institutional factors—plant locations that predate railroad deregulation and consolidation; plants’ inability to resell coal and arbitrage spatial price differences—I can compare delivered coal prices at captive versus non-captive plants to identify differential markup levels. My estimates flexibly control for coal commodity value and rail freight costs, and I use nearest-neighbour matching to remove unobserved geographic confounders.

Next, I estimate how coal-by-rail markups respond to coal demand shocks: decreases in the natural gas price, which hurt coal’s competitiveness in electricity markets.¹ I use time-series gas price variation to estimate coal markup changes in a difference-in-differences (DD) design, for captive versus non-captive plants. Then, I refine my DD “treatment” definition using plant-specific coal demand estimates, in order to more accurately characterize the structural relationship between gas prices and coal markups. Since observed gas price shocks mimic the effect of a carbon tax on coal demand (Cullen and Mansur, 2017), these estimates predict how railroads might reoptimize coal markups under a carbon tax.

I find that coal plants facing the most market power paid \$2–5/ton higher average markups, compared to plants facing the least market power. This implies average markups of 13–34% above marginal shipping costs, and 4–13% above total marginal costs. For plants facing the most market power, I also find that a \$1/MMBTU decrease in gas prices caused coal markups to fall differentially by \$0.41–0.54/ton. These DD estimates understate the full extent of markup changes, since defining “treatment” using market structure alone ignores heterogeneity in plants’ coal demand shocks. When I incorporate such heterogeneity into DD “treatment”, I detect much larger effects: decreases of \$1.41–1.52/ton, for each \$1/MMBTU drop in gas prices. Since gas prices decreased by \$4/MMBTU from 2007 to 2011, this implies that coal markups fell by over \$5/ton (16%) for the average “treated” plant.

My results show that rail carriers reoptimize markups to effectively buffer coal plants against heterogeneous shocks to their competitiveness. As decreasing gas prices reduce the marginal cost of gas-fired electricity, rail carriers reduce markups *more* for coal plants that face a greater competitive threat from gas-fired rivals. By contrast, markups do not change for coal plants that are not threatened by low gas prices, or do not face market power in coal shipping. These findings are qualitatively consistent with the predictions of a static oligopoly model where rail carriers maximize profits separately for each coal plant.

Low gas prices disadvantage coal-fired plants in a manner similar to a carbon tax, which would penalize coal’s higher carbon intensity (compared to gas). Hence, I can convert my estimated markup changes into the pass-through rates of an implicit carbon tax, or the rates at which rail carriers would have passed a mine-mouth carbon tax on to delivered coal prices.² I estimate plant-specific pass-through rates as low as 0.74, and I find incomplete pass-through for 75% of coal shipments to plants that rely on rail. This suggests that market imperfections would likely mute the price signal of a U.S. carbon tax, for most (but not all) coal plants. It also implies that railroads would bear a substantial share of the tax burden, potentially reducing the regressivity of a carbon tax.

Finally, I predict how decreases to coal markups might erode the environmental benefits of a modest carbon tax. My counterfactual analysis implies that low gas prices could have yielded

1. I use “natural gas” and “gas” interchangeably. My analysis does not relate to gasoline.

2. The tax’s location in the supply chain should not affect the economic interpretation of pass-through, absent other distortions (Weyl and Fabinger, 2013). A mine-mouth tax follows the standard formulation of a cost shock passed forward to final goods prices. In practice, carbon taxes are typically levied downstream.

10% greater carbon abatement during my sample period, if coal markups had not changed. This translates to \$11.9 billion in realized climate damages that could have been avoided—via short-run coal-to-gas substitution alone.³ This suggests that the welfare consequences of incomplete carbon tax pass-through would be substantial: decreases to coal markups would similarly mute the carbon price signal for medium- and long-run abatement.

This article makes three main contributions. First, I provide the first empirical evidence that upstream market power in coal supply would likely erode the environmental benefits of U.S. climate policy. This contributes to a growing literature on the complex interactions between environmental policy and market power. Previous studies have found that emissions regulations may exacerbate market power distortions (Ryan, 2012; Fowlie *et al.*, 2016), and that market power can either increase or decrease the efficacy of environmental regulations (Mansur, 2007; Leslie, 2018). I find that an implied carbon tax reduces coal markups, which manifests as incomplete pass-through to coal prices. Contrary to previous findings of full carbon tax pass-through in electricity markets (Fabra and Reguant, 2014), my results provide the first evidence that incomplete pass-through would likely substantially reduce the efficacy of a carbon tax on U.S. electricity generation.

Second, my results show that upstream oligopoly rents would likely absorb a substantial share of the carbon tax burden. This adds to a nascent body of evidence that incomplete pass-through of energy cost shocks can render climate policies more progressive (Ganapati *et al.*, 2020; Stolper, 2021; Muehlegger and Sweeney, 2022). This has key implications for the equity-efficiency trade-off in climate policy design, especially regarding re-distribution of carbon tax revenues (Goulder *et al.*, 2019; Sallee, 2019; Goulder, 2020).

Third, I estimate transportation market power using a combination of reduced-form and structural methods. The literature on oligopolistic intermediaries typically relies on structural trade models to empirically separate markups from cost shocks (*e.g.* Atkin and Donaldson, 2015; Startz, 2021).⁴ In contrast, I leverage a unique feature of coal markets—limited spatial arbitrage between power plants—to credibly identify markups while making relatively few assumptions on coal demand. At the same time, I refine my reduced-form estimates using a novel coal demand estimation strategy. This advances the literature on rail market power (Busse and Keohane, 2007; Hughes, 2011; Hughes and Lange, 2020) by using economic theory to link observed market power to counterfactual climate policy outcomes.

2. U.S. COAL MARKETS

U.S. coal markets have three primary types of agents: mining firms, power plants, and transport intermediaries. Mines are concentrated near coal deposits, most notably in Appalachia (West Virginia) and the Powder River Basin (northeast Wyoming). By contrast, coal power plants are dispersed across regionally fragmented electricity markets. Coal is heavy relative to its commodity value, and plants located far from mines incur substantial shipping costs. Railroads are the dominant shipping mode; four carriers control nearly all coal-by-rail shipping (two firms each in the western and eastern U.S.). Figure 1 maps coal producing regions, coal power plants, and

3. I monetize damages using \$185/metric ton CO₂, following recent estimates of the social cost of carbon (Renert *et al.*, 2022). A previous version of this article used the now-outdated benchmark of \$50/metric ton CO₂, which would imply damages of \$3.2 billion.

4. A notable exception is Bergquist and Dinerstein (2020), who identify markups by experimentally manipulating intermediary market structure.

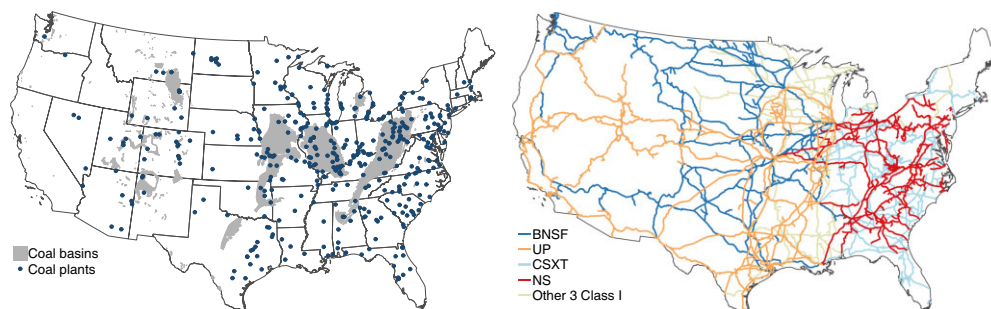


FIGURE 1
U.S. coal geography

Notes: The left panel maps productive deposits of bituminous and sub-bituminous coal. The vast majority of production comes from the Powder River Basin (northeast Wyoming), the Appalachian Basin (West Virginia and eastern Kentucky), and the Illinois Basin (southern Illinois and western Kentucky). Dots denote all 430 large coal-fired electric power plants that operated between 2002 and 2015. The right panel maps major rail lines owned and operated by the seven Class I rail carriers. Two rail carriers each dominate the West (BNSF, Union Pacific) and East (CSX Transportation, Norfolk Southern). I combine the smaller three carriers into a single shade (Canadian National, Canadian Pacific, Kansas City Southern).

major rail lines (by firm). Plants located on navigable waterways may also receive coal-by-barge shipments, which make up roughly 17% of coal deliveries.⁵

The 1980 Staggers Act substantially weakened rail price regulations, letting railroads set freight shipping rates with limited government oversight (MacDonald 1989, 2013). If one rail carrier exhibits “market dominance” along a given route, regulators may intervene to prevent rail revenues from exceeding 180% of total variable costs.⁶ This gives railroads significant leeway to exercise market power and negotiate complex long-term contracts with power plants (Joskow, 1988). By allowing carriers to extract oligopoly rents and exploit economies of scale, the Staggers Act also spurred a series of railroad mergers: the 33 “Class I” railroads of 1980 have consolidated into the 7 Class I railroads of today (Schmidt, 2001; Prater *et al.*, 2014).⁷

Three factors contribute to spatial dispersion in coal-by-rail markups. First, unlike most commodities, coal consumption must occur in precise geographic locations with potentially limited access to transportation networks. While some coal power plants have access to multiple rail carriers and/or barge shipments, many plants are captive to a single rail carrier for all coal deliveries. Second, many plants are constrained to buy a specific type of coal mined in a single region (Joskow, 1987). This further restricts their shipping options, as mines may also have limited access to rail and water networks.⁸ Third, the resale of coal is cost-prohibitive: infrastructure is

5. Trucks carry 3% of coal deliveries but cannot compete directly with rail and barges (Busse and Keohane, 2007). Power plants consume nearly all steam coal produced in the U.S. I ignore imports (1–2% of U.S. coal consumption), exports (3–5% of U.S. steam coal production), and other industrial end uses (*e.g.* steel).

6. In practice, regulators loosely interpret this threshold such that railroads may earn an adequate return on investment (Wilson, 1996). The Surface Transportation Board has reviewed just 34 rates challenges on coal shipping rates since 1996: <https://www.stb.gov/wp-content/uploads/Rate-Case-List-11-19-2019.pdf>

7. The Class I designation includes carriers with annual operating revenues exceeding \$453 million. These seven firms account for approximately 69% of rail mileage and 94% of rail freight revenues.

8. Coal’s attributes vary across (and within) regions. Plants typically value coal with high energy content (MMBTU/ton), and low sulfur and ash content (which create local air pollution). Plants self-calibrate to a pre-specified mix of coal attributes, and deviations can reduce the efficiency of their boilers (Kerkvliet and Shogren, 1992). Many plants comply with SO₂ regulations by burning low-sulfur coal from Wyoming’s Powder River Basin (Schmalensee and Stavins, 2013). If such a plant has access to two rail carriers, only one of which serves the Powder River Basin, then it has effectively one shipping option.

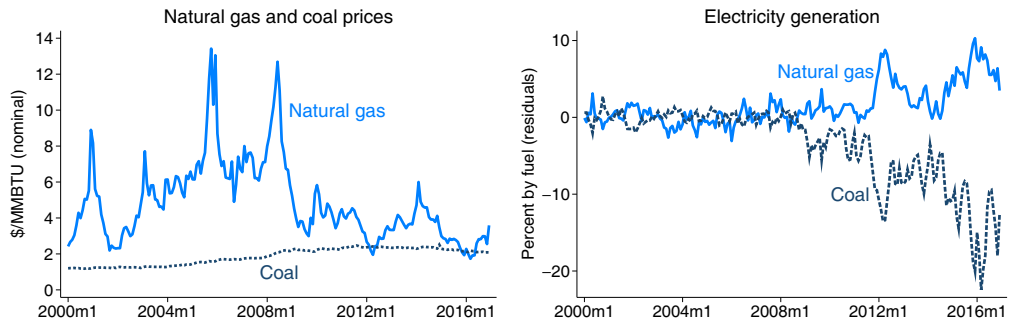


FIGURE 2
U.S. fuel prices and electricity generation

Notes: The left panel reports monthly average Henry Hub spot prices for natural gas, and monthly average costs of coal delivered to power plants (not accounting for heterogeneous coal attributes). The right panel plots monthly generation by fuel as a percent of total monthly generation, controlling for month-of-year fixed effects and a 2000–08 time trend for each fuel.

built to carry coal *to* (not *away from*) plants (Busse and Keohane, 2007; Jha, 2020). Hence, plants are unable to arbitrage spatial price differences, allowing railroads to charge higher markups to plants with fewer shipping options.⁹

U.S. coal consumption has declined over the past 15 years, largely due to decreases in the price of natural gas. Technological advances in hydraulic fracturing (“fracking”) have led to a boom in natural gas extraction, causing a historic drop in U.S. gas prices.¹⁰ Since coal plants compete directly with gas plants in electricity markets, low gas prices have crowded out coal-fired electricity generation. The left panel of Figure 2 shows how the fracking boom has depressed U.S. gas prices since 2008, and the right panel shows how the electricity sector has shifted towards gas and away from coal. The corresponding decrease in coal demand has likely caused rail oligopolists to reoptimize coal markups. Any observed changes in markups can predict what might occur under a carbon tax, which would similarly disadvantage coal relative to low-carbon natural gas (Cullen and Mansur, 2017). If coal markups decrease (increase), this would dampen (magnify) the carbon tax price signal as it passes along the coal supply chain. This effect would likely be heterogeneous across coal plants, due to variation in pre-existing markups *and* variation in plants’ exposure to gas-fired competition.¹¹

3. THEORETICAL FRAMEWORK

I develop a symmetric Cournot oligopoly model of railroad intermediaries who sell coal to power plants. This implies that markups should respond heterogeneously to gas price changes, given

9. Plants may buy coal directly from rail carriers; alternatively, they may separately purchase coal from mines and freight services from railroads. This distinction does not affect the economic interpretation of delivered coal markups. Without loss of generality, my model treats rail carriers as sellers of both the commodity and freight services.

10. Two technological advances have facilitated the “fracking boom” by enabling gas extraction from shale formations: horizontal drilling and hydraulic fracturing (Fitzgerald, 2013). The physical properties of natural gas make it expensive to export, which is why a domestic supply glut has depressed U.S. gas prices.

11. Low gas prices have not impacted all coal plants equally. For a coal plant located in an electricity market with many gas-fired competitors, a drop in gas prices should decrease its coal demand. For a coal plant in a market without any gas-fired competitors, the same gas price shock should not alter its coal demand. Low gas prices should also disproportionately hurt less inefficient (*i.e.* less competitive) coal plants.

the number of rail carriers and shocks to plants' coal demand. This framework makes several simplifying assumptions, which I relax in my empirical analysis below.

3.1. *Symmetric Cournot oligopoly*

Let power plant j be a price-taker for coal, which it may purchase from N_j identical rail carriers. Each rail carrier i chooses the best-response quantity of coal q_{ij} that maximizes its profits. In equilibrium, plant j consumes $N_j q_{ij} = Q_j$ units of coal at price P_j . Since plant j cannot resell its purchased coal, P_j is not restricted by a binding arbitrage constraint and rail carriers may effectively treat each plant as its own isolated coal market.

Rail carrier i 's profits from selling coal to plant j are:

$$\pi_{ij}(q_{ij}) = q_{ij} [P_j(Q_j; \mathbf{Z}_j) - C_j - S(\mathbf{T}_j)] - F_j \quad (1)$$

P_j is plant j 's inverse demand for coal, as a function of Q_j and a parameter vector \mathbf{Z}_j . Railroads purchase coal at a constant marginal cost C_j .¹² $S(\mathbf{T}_j)$ is the average cost of shipping coal to plant j , where \mathbf{T}_j is a vector of transportation cost parameters (e.g. rail mileage, diesel prices). F_j is a fixed cost of servicing plant j by rail.¹³ This simple model abstracts from railroad rate regulation, which may constrain firm i 's strategic behaviour.¹⁴

Rail carrier i 's first-order condition implies the following price-cost markup μ_j :

$$\mu_j \equiv P_j - C_j - S(\mathbf{T}_j) = - \left(\frac{1}{N_j} \right) \frac{\partial P_j}{\partial Q_j} Q_j. \quad (2)$$

If plant j is captive to a single rail carrier ($N_j = 1$), it should face weakly higher markups than if two carriers compete for its business. If plant j 's coal demand is relatively inelastic, it should face weakly higher markups, all else equal. For plants located on navigable waterways, barge competition should force railroads to set $\mu_j \approx 0$: barges have low barriers to entry, less restricted usage rights, and lower shipping costs (MacDonald, 1987; Wetzstein *et al.*, 2021).

3.2. *Comparative statics for coal markups*

Coal demand depends on the price of natural gas, because the two fuels compete in electricity dispatch. If the gas price decreases (increases), a coal plant may supply less (more) electricity at a given coal price. The gas price also influences the elasticity of coal demand, by determining the range of coal prices over which a coal plant is marginal in electricity supply. A marginal plant has (locally) elastic coal demand, because its coal consumption responds to small changes in coal price. At lower coal prices, a coal plant will be inframarginal and its strict capacity

12. C_j varies across plants, since coal is heterogeneous and coal markets are regionalized. In reality, coal supply may be upward-sloping, and need not be perfectly competitive (e.g. a few large firms dominate Wyoming's Powder River Basin). [Supplementary Material, Appendix B.2](#) considers a richer model with coal varieties and between-plant interactions. [Supplementary Material, Appendix C.4](#) discusses the welfare implications of alternate upstream market structures.

13. In reality, each firm's shipping routes are constrained by track ownership and trackage rights, implying non-identical costs $S(\mathbf{T}_j)$ and F_j . For simplicity, I assume that quantity q_{ij} does not enter into $S(\mathbf{T}_j)$, which precludes rail capacity constraints and increasing returns to scale in shipping. My empirical analysis relaxes the assumption of symmetric costs and also allows shipping costs to vary with shipment size.

14. [Supplementary Material, Appendix B.3](#) incorporates regulation into my oligopoly model. Unfortunately, I lack the data to empirically characterize the threat of regulation along a given coal-by-rail route.

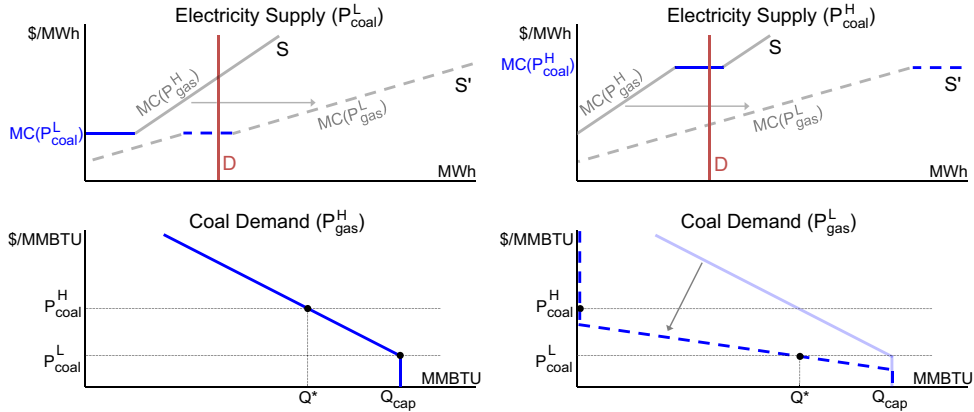


FIGURE 3
Coal demand and natural gas prices

Notes: This stylized electricity market has one coal generator with fixed capacity, and constant marginal cost at a given coal price ($MC(P_{coal})$). Natural gas generators have marginal costs that scale with the gas price ($MC(P_{gas})$, in light grey). Electricity demand (D) is perfectly inelastic, and deterministic (for simplicity). The top panels show electricity supply curves for four combinations of coal price (low in the left panel, high in the right panel) and gas price (high for solid lines, low for dashed lines). The bottom panels translate the coal plant's production into its coal input demand (MWh out as a function of MMBTU of coal in, given the plant's fixed production technology). Under a high gas price (P_{gas}^H), the coal plant consumes at full capacity (Q_{cap}) given a low coal price (P_{coal}^L) and at Q^* given a high coal price (P_{coal}^H). If the gas price decreases to P_{gas}^L , the coal plant becomes marginal given P_{coal}^L (where it had been inframarginal) and above the margin given P_{coal}^H (where it had been marginal). The decrease in gas price causes inverse coal demand at Q^* to shift down and become less steep.

constraint will bind; this translates to (locally) inelastic coal demand, as small changes in coal price will not change its coal consumption.

Figure 3 presents a stylized electricity market to illustrate how a negative gas price shock impacts both the level and slope of coal demand. There is a single coal plant with constant marginal cost, and an upward-sloping supply of gas-fired generation. Each technology's marginal costs scale with its fuel price. The top panels show four electricity supply curves, for four combinations of coal price (low, high) and gas price (high, low). In reality, electricity demand is stochastic and extremely inelastic; this stylized example assumes electricity demand is deterministic and perfectly inelastic.

At a given gas price, the plant's coal demand is the 1-to-1 mapping between coal price and coal consumption. For example, under a high coal price and high gas price (*i.e.* the solid supply curve in the top-right panel), the coal plant is marginal in the electricity market and generates at 70% capacity. Hence, it demands 70% of its throughput capacity for coal, or Q^* in the bottom-left panel. Comparing the bottom two panels, the gas price governs the range of coal prices for which the plant is marginal, and coal demand is not vertical. A negative gas price shock causes inverse coal demand to shift down *and* become less steep.¹⁵

Using my Cournot model, I derive how rail carriers should reoptimize coal markups in response to gas price changes. Let Z denote the Henry Hub spot price of natural gas, which enters plant j 's inverse coal demand as an element of Z_j . For plants without a coal-by-barge

15. In reality, electricity dispatch may not order plants from lowest-to-highest cost, and plants may not maximize short-run profits. Demand realizations come from a continuous probability distribution, and electricity is not storable. Coal storage enables plants to hedge against uncertainty in electricity markets, letting coal markets clear on a longer timescale. Hence, coal demand should not have sharp kinks.

option, the change in markup μ_j that results from a small change in gas price Z is:

$$\frac{d\mu_j}{dZ} = \frac{\frac{\partial P_j}{\partial Z} \left(2 + \frac{E_{D_j}}{N_j} - N_j\right) - \frac{\partial^2 P_j}{\partial Q_j \partial Z} Q_j}{2 + \frac{E_{D_j}}{N_j}}, \quad (3)$$

where $E_{D_j} \equiv \left(\frac{\partial^2 P_j}{\partial Q_j^2}\right) \left(\frac{\partial P_j}{\partial Q_j}\right)^{-1} Q_j$ is the elasticity of the slope of inverse demand.¹⁶ For plants where barge competition forces rail markups close to zero, $\frac{d\mu_j}{dZ} \approx 0$.

Equation (3) depends on the level, slope, and curvature of plant j 's inverse demand. $\frac{\partial P_j}{\partial Z}$ captures how gas price affects the *level* of inverse coal demand: if a negative gas price shock (*i.e.* $dZ < 0$) causes plant j 's inverse coal demand to shift down as in Figure 3, then $\frac{\partial P_j}{\partial Z} > 0$. $\frac{\partial^2 P_j}{\partial Q_j \partial Z}$ captures how gas price affects the *slope* of inverse coal demand: if lower gas prices make inverse coal demand less steep (*i.e.* if $dZ < 0$ causes $\frac{\partial P_j}{\partial Q_j}$ to become less negative), then $\frac{\partial^2 P_j}{\partial Q_j \partial Z} < 0$. Finally, the change in markup depends on the degree to which inverse demand is concave ($E_{D_j} > 0$) or convex ($E_{D_j} < 0$): more concave demand will tend to increase $\frac{d\mu_j}{dZ}$, while more convex demand will tend to decrease $\frac{d\mu_j}{dZ}$.¹⁷

These three features of coal demand interact with the rail market size (N_j) and structure (*i.e.* barge option) to determine how railroads should reoptimize plant j 's markups when the gas price changes. The sign of equation (3) is theoretically ambiguous, as $\frac{\partial P_j}{\partial Z}$, $\frac{\partial^2 P_j}{\partial Q_j \partial Z}$, and E_{D_j} may vary considerably across coal plants. Rail carrier behaviour may also depart from the predictions of this simple model—especially under the threat of regulation or if markups are not truly independent across plants.¹⁸ Below, I directly estimate plant-specific demand parameters $\frac{\partial P_j}{\partial Z}$, $\frac{\partial^2 P_j}{\partial Q_j \partial Z}$, and E_{D_j} , which I use to construct a prediction of $\frac{d\mu_j}{dZ}$ for each plant. Then, I take these predictions to the data to test whether cross-sectional variation in equation (3) causes rail carriers to reoptimize markups heterogeneously across plants.

4. DATA

I use publicly available data on coal deliveries to U.S. power plants, collected on the Energy Information Administration's (EIA) Form 923. These data are at the plant-supplier-month-purchase order level. For each observation, I observe the county of origin, coal attributes (*e.g.* sulfur content), total tons delivered, mode of transportation (*e.g.* rail, barge), and transaction type (long-term contract versus spot market). For deliveries to utility-owned plants, I also observe the average price inclusive of commodity costs, shipping costs, and markups. This serves as the

16. [Supplementary Material, Appendix B.1](#) provides a derivation of equation (3), which assumes Cournot competition among rail carriers on each route. This assumption on market conduct qualitatively matches observed markup changes (see Table 3). However, an alternate assumption of perfect collusion only slightly alters $\frac{d\mu_j}{dZ}$ for plants with multiple rail carriers and no barge option. These two assumptions are sufficiently similar to be empirically indistinguishable from each other (see [Supplementary Material, Appendix Table E6](#)).

17. This is a standard result in the pass-through literature on imperfectly competitive product markets, where the pass-through rate is closely related to the curvature of demand ([Weyl and Fabinger, 2013](#)).

18. Regulation may prevent rail carriers from extracting (unconstrained) oligopoly rents. This simple model ignores multiple-market negotiations between carriers, and dynamic interactions between carriers and plants.

outcome variable in my empirical analysis.¹⁹ I control for county-year average mine-mouth coal prices, from EIA's Annual Coal Report.

I merge coal shipment data with EIA data on power plant characteristics and operations (Forms 906, 923, 860, and 767). The EPA eGRID database reports each plant's power control area (PCA), or its region on the electricity transmission grid. To estimate plant-specific coal demand parameters, I use hourly generation data from the EPA's Continuous Emissions Monitoring System (CEMS). This allows me to estimate plants' coal consumption in each hour as a function of the relative prices of coal and natural gas.

The Bureau of Transportation Statistics (BTS) publishes detailed GIS shapefiles of the U.S. rail network, assigning an owning/operating rail carrier to each track segment. I use a graph algorithm to find the shortest route of segments that connects each coal plant to its coal-producing counties.²⁰ I proxy for congestion along each route using the average traffic density over its rail lines. I also control for time-series variation in shipping costs using the Association of American Railroads' (AAR) monthly fuel price index, which reflects changes in diesel prices paid by rail operators.²¹

I split plants into two time-invariant groups, "captive" and "non-captive", based on their locations on the rail network and the counties from which they purchase coal. "Captive" plants either (i) become unconnected from the network after removing any single Class I carrier, or (ii) become unconnected from *all* observed trading partners after removing the modal carrier along each origin–destination route. For example, suppose a plant only buys coal from two counties in Wyoming. I classify this plant as captive if a single Class I carrier controls all terminal nodes within 7 miles of the plant, or if after removing the modal carrier on its shortest route to *each* Wyoming county, the new shortest routes *both* increase by over 300 miles.²² I also calculate each plant's proximity to a navigable river, Great Lake, or coastline; this lets me determine the subset of plants with the option of barge deliveries.

5. EMPIRICAL STRATEGY

I match captive plants to nearby non-captive plants, which lets me estimate the effect of captiveness on coal-by-rail markup *levels*. Next, I use a difference-in-differences (DD) framework to estimate markup *changes* by interacting captiveness with the time series of natural gas prices. To better approximate the structural relationship between gas prices and coal markups, I estimate plant-specific coal demand curves and parameterize predictions of $\frac{d\mu_i}{dZ}$ from equation (3). Finally, I take those predictions to the data using my DD framework.

19. I focus on utility-owned plants, which received 77% of coal deliveries during my 2002–15 sample period. EIA redacts prices for non-utility plants that were divested during electricity market restructuring. Most coal plant divestments were in Pennsylvania, Illinois, Ohio, and New York; nearly all occurred prior to the fracking boom. Previous studies focusing on electricity deregulation have obtained non-disclosure agreements with EIA to unmask coal prices for non-utility plants (e.g. Cicala, 2015; Hughes and Lange, 2020).

20. Using a similar algorithm, Hughes (2011) finds that GIS-derived shortest distances closely approximate (yet slightly understate) actual rail shipping distances. [Supplementary Material, Appendix F](#) outlines my shortest-distance algorithm.

21. Diesel purchases represent roughly half of railroads' total variable transportation costs.

22. Seven miles is the 95th percentile of plants' distance to the closest rail node. A 300-mile increase in distance implies a 20% increase over the median delivered coal price. [Supplementary Material, Appendices F.2 and F.3](#) discuss these two distance thresholds; [Supplementary Material, Appendix Figures E2 and E8](#) report sensitivity analysis on each threshold.

5.1. *Matching captive versus non-captive plants*

While rail captiveness is not randomly assigned, three institutional factors facilitate causal comparisons between captive versus non-captive coal plants. First, nearly all plant locations were fixed prior to 1980, when the Staggers Act legalized rail price discrimination. Second, a wave of Class I mergers consolidated the rail network through 1999, increasing the likelihood that a given coal plant became captive to a single firm.²³ Hence, a plant could not have strategically influenced its own number of rail carriers—which depends on rail line ownership changes outside of plants' control. Finally, each plant is small relative to a Class I carrier's portfolio, which includes many commodities besides coal. Hence, it is unlikely that any rail merger decision hinged on strategic selection of an individual plants' captiveness.

However, geographic differences could confound captive versus non-captive comparisons. The rail network is relatively sparse in the western U.S., and western coal plants tend to have access to fewer rail service points. This sparseness may also have influenced rail mergers, if consolidating western track ownership created many newly captive customers. The first three columns of Table 1 show that captive plants are far more likely to be located west of the Mississippi River, compared to non-captive plants. Captive plants are also statistically more likely to be younger, have higher capacity factors, not sell into wholesale electricity markets, and consume low-sulfur sub-bituminous coal—all common characteristics of western plants.²⁴

I use nearest-neighbour matching to remove geographic confounders (Heckman *et al.*, 1997). I match each captive plant to its k nearest non-captive neighbours within 200 miles. I also force exact matches on plants' preferred coal type (bituminous versus sub-bituminous).²⁵ I omit plants that rely exclusively on non-rail shipping modes (*e.g.* barges), and non-utility plants (with no coal price data). Figure 4 reveals broad geographic overlap between captive versus non-captive plants and illustrates how matching corrects regional imbalances (*e.g.* only captive plants in Colorado). Table 1 reveals that nearest-neighbour weighting also yields balance on observable covariates prior to the fracking boom. This bolsters the credibility of my matching strategy, since geographically pruning the distribution of plants does not guarantee statistically indistinguishable captive and non-captive groups.²⁶

This strategy builds on Cicala (2015), who matches divested coal plants to non-divested plants using a 200-mile buffer. However, while the nature of divestment necessitates matching across state borders, 92% of captive/non-captive matches have a matched plant from the opposite group within their state. My matching strategy also yields 90% overlap by coal county and sample month—meaning that time-varying unobservables relating to coal production are unlikely to confound my estimates of differential coal markups.

23. The rail network was static during my sample period: 99.3% of Class I rail mileage has had constant ownership since 2006, the earliest year of BTS data. I exclude the few coal plants constructed after 1999.

24. [Supplementary Material, Appendix Table E1](#) reports summary statistics for plants west versus east of the Mississippi River. Within-region captive versus non-captive differences attenuate and lose statistical significance (except for plant vintage).

25. This ensures that matched plants do not receive most shipments from opposite sides of the country. Nearest-neighbour weights equal the inverse number of matches: if a non-captive plant is one of 3 matches for captive plant A and one of 2 matches for captive plant B, its weight is $\frac{1}{3} + \frac{1}{2} = \frac{5}{6}$. Matched captive plants have weight 1, unmatched plants have weight 0, and weights sum to twice the number of matched captive plants. I describe this matching strategy formally in [Supplementary Material, Appendix E.1](#).

26. [Supplementary Material, Appendix Table E2](#) reports summary statistics with $k = 1$ and $k = 5$ matches, which yield similar empirical results. [Supplementary Material, Appendix Figure E1](#) shows that the 200-mile buffer is not binding for most captive plants.

TABLE 1
Summary statistics—captive versus non-captive coal plants (2002–06)

	All coal plants			Matched sample ($k = 3$)		
	Captive	Non-captive	Difference	Captive	Non-captive	Difference
West of Mississippi (1/0)	0.47 (0.50)	0.24 (0.43)	0.23 [0.00]***	0.44 (0.05)	0.40 (0.06)	0.04 [0.67]
Coal-fired capacity (MW)	806.13 (738.72)	760.84 (703.91)	45.29 [0.52]	802.37 (67.07)	806.19 (87.29)	-3.82 [0.97]
Number of coal units	2.36 (1.32)	2.62 (1.64)	-0.26 [0.08]*	2.58 (0.15)	2.56 (0.14)	0.02 [0.94]
Vintage (year)	1968.85 (13.90)	1962.88 (13.34)	5.97 [0.00]***	1966.24 (1.39)	1962.52 (1.42)	3.72 [0.06]*
Annual capacity factor	0.63 (0.17)	0.60 (0.17)	0.03 [0.04]**	0.63 (0.01)	0.63 (0.01)	-0.00 [0.90]
Heat rate (MMBTU/MWh)	11.09 (1.40)	11.06 (1.52)	0.03 [0.86]	10.98 (0.14)	10.74 (0.13)	0.24 [0.20]
Scrubber installed (1/0)	0.36 (0.48)	0.29 (0.45)	0.07 [0.12]	0.26 (0.05)	0.27 (0.06)	-0.00 [0.95]
Elec. market participant (1/0)	0.49 (0.50)	0.71 (0.46)	-0.22 [0.00]***	0.45 (0.05)	0.49 (0.06)	-0.04 [0.62]
Coal bought (mil MMBTU/year)	48.82 (47.90)	44.00 (43.70)	4.82 [0.29]	46.63 (4.13)	43.85 (5.00)	2.77 [0.67]
Sulfur content (lbs/MMBTU)	0.87 (0.61)	1.02 (0.79)	-0.15 [0.03]**	0.79 (0.06)	0.85 (0.07)	-0.05 [0.55]
Ash content (lbs/MMBTU)	8.46 (4.21)	8.96 (8.24)	-0.50 [0.46]	8.00 (0.36)	7.75 (0.30)	0.25 [0.60]
Share spot market	0.19 (0.29)	0.19 (0.25)	-0.00 [0.87]	0.19 (0.03)	0.16 (0.02)	0.02 [0.57]
Share sub-bituminous	0.41 (0.47)	0.31 (0.42)	0.10 [0.03]**	0.42 (0.05)	0.39 (0.06)	0.02 [0.79]
Average rail distance (miles)	554.91 (385.90)	620.34 (417.90)	-65.43 [0.12]	573.73 (40.26)	588.77 (41.11)	-15.05 [0.79]
Non-rail plants	17	14	31	0	0	0
Utility plants	148	176	324	87	97	184
Total plants	190	240	430	87	97	184

Notes: This table compares coal plants captive versus non-captive to a single rail carrier. The left three columns include all CEMS coal-fired power plants from 2002 to 2015 that report coal deliveries in 2002–06 and 2007–15. The right three columns weight by nearest-neighbour matches: unmatched plants have weight 0; matched captive plants have weight 1; and matched non-captive plants have weights equal to the inverse number of matches. Matching criteria: up to $k = 3$ nearest neighbours within 200 miles; exact matches on preferred coal rank; and removing non-utility and non-rail plants. Standard deviations are in parentheses, and p -values [in brackets] are clustered by plant. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

5.2. Estimating markup levels

I begin by estimating differences in markups between captive and non-captive plants. The following OLS regression is analogous to the markup expression I derive in equation (2):

$$P_{ojms} = \tau D_j + \beta_C C_{ojms} + \beta_T S(\mathbf{T}_{ojms}) + \beta_X \mathbf{X}_{jm} + \eta_o + \delta_m + \varepsilon_{ojms} \quad (4)$$

P_{ojms} is the average delivered coal price, for purchase s by plant j in month m , from originating county o . D_j is an indicator for rail captiveness, and τ captures the average differential markup

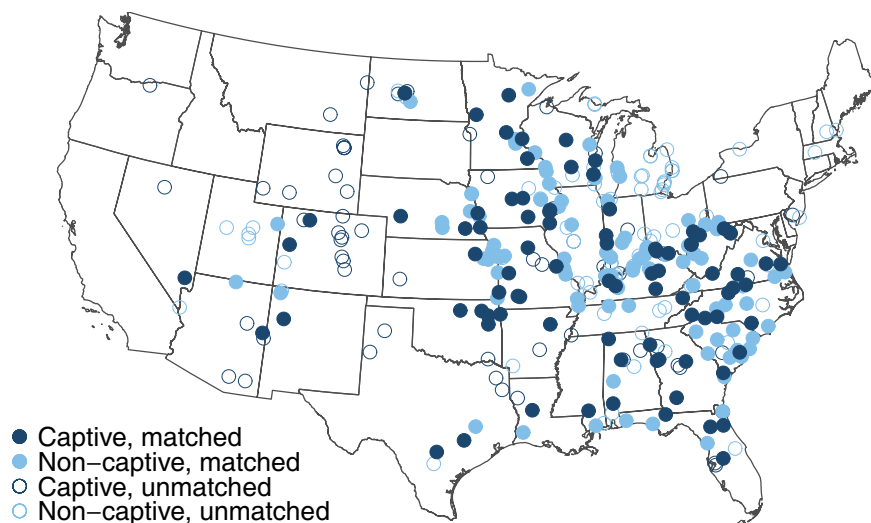


FIGURE 4

Nearest-neighbour matching captive versus non-captive coal plants

Notes: This map plots the locations all 324 utility-owned coal power plants in the full 2002–15 sample. Captive (non-captive) plants are dark (light); matched (unmatched) plants are filled (hollow). Matching criteria: up to $k = 3$ nearest neighbours, with a maximum distance of 200 miles; exact matches on coal rank; and removing non-utility and non-rail plants.

faced by captive plants, relative to non-captive plants. Since I do not directly observe markups, I use price as an outcome variable and control for commodity value and shipping costs (*i.e.* $C_j + S(\mathbf{T}_j)$ in equation (2)). I use nearest-neighbour weights and plant-specific controls (\mathbf{X}_{jm}) in the style of a doubly robust estimator (Wooldridge, 2007), plus coal county fixed effects (η_o) and month-of-sample fixed effects (δ_m). The remaining variation in P_{ojms} is close to the ideal experiment: comparing the price of two identical coal shipments to two otherwise identical coal plants, where only one plant is rail captive.

C_{ojms} controls for determinants of commodity value, including heat, sulfur, and ash content; coal rank; and the average annual mine-mouth price for coal produced in county o . C_{ojms} also includes dummies for spot market transactions and contracts expiring within 2 years, since plants pay higher prices for (longer) contracts that minimize the risk of supply disruptions.²⁷ \mathbf{T}_{ojms} includes four determinants of coal-by-rail shipping costs: rail distance between coal county o and plant j , average diesel price paid by rail carriers in month m , shipment size (since marginal freight costs may vary in tons shipped), and the share of route oj with high traffic density (to proxy for congestion). $S(\cdot)$ flexibly models shipping costs as the four-way interaction of the components of \mathbf{T}_{ojms} . \mathbf{X}_{jm} controls for predetermined and time-varying plant characteristics, including all covariates in Table 1. I cluster standard errors by plant, allowing for arbitrary within-plant serial correlation. I multiply nearest-neighbour weights by the quantity of coal transacted in each observation, in order to estimate differential markups for the average ton of coal.²⁸

27. Wolak (1996) finds that coal plants simultaneously purchase on long-term contracts and the spot market. Jha (2022) estimates that the average regulated coal plant is willing to trade a \$1.62 increase in expected delivered coal price for a \$1.00 decrease in the standard deviation of delivered coal price.

28. Coal quantity varies substantially across observations. I also control for $\log(\text{quantity})$ in \mathbf{T}_{ojms} .

If misspecification in C_{ojms} or $S(\mathbf{T}_{ojms})$ induced correlation between ε_{ojms} and D_j , I could not interpret $\hat{\tau}$ as causal. However, my results are robust to alternative cost parameterizations, implying that such a violation of mean independence is unlikely.²⁹ A causal interpretation also requires D_j to be independent of plant unobservables, after nearest-neighbour matching and conditioning on \mathbf{X}_{jm} . Since matching removes broad geographic confounders, and since the localized assignment of D_j is not strategic (either by railroads or the plants), any confounders would need to be local or non-geographic—for example, boilers at captive plants being less able to accommodate switching coal varieties. As such confounders seem unlikely, I interpret $\hat{\tau}$ as the causal effect of rail captiveness on markups.

5.3. Estimating markup changes

To estimate differential changes in coal markups, I interact the time series of natural gas prices with cross-sectional “treatment” indicators. Captiveness (D_j) is one such indicator, since rail markups in a monopoly regime are likely larger and more responsive to market conditions. I can refine this “treated” group by removing plants with a water delivery option W_j (i.e. $TREAT_j = D_j(1 - W_j)$), since markup changes are more likely for plants lacking this more competitive coal-by-barge option.

I modify equation (4) to estimate a lagged DD design:

$$P_{ojms} = \tau TREAT_j \cdot Z_{m-L} + \sum_{\ell=0}^{L-1} \tau_{\ell} TREAT_j \cdot \Delta Z_{m-\ell} \dots + \beta_C C_{ojms} + \beta_T S(\mathbf{T}_{ojms}) + \beta_X \mathbf{X}_{jm} + \eta_{oj} + \delta_m + \varepsilon_{ojms}. \quad (5)$$

Z_m is the average Henry Hub spot price in month m , and $\Delta Z_m = Z_m - Z_{m-1}$. The coefficient of interest τ captures the cumulative effect of a \$1/MMBTU change in gas price, over $L = 48$ months. Each lagged coefficient τ_{ℓ} captures the cumulative effect after ℓ months, for a plant with $TREAT_j = 1$ relative to a plant with $TREAT_j = 0$. I allow for delayed effects since most coal deliveries occur on long-term contracts, which may be slow to adjust to changing market conditions.³⁰ I also add route fixed effects η_{oj} , which control for the average markup of all rail shipments from county o to plant j ; this removes any compositional changes in coal purchases and isolates within- oj changes in markups.

I interpret $\hat{\tau}$ as the cumulative causal effect of gas price changes on coal-by-rail markups. The key identifying assumption is that gas price changes are independent of unobserved factors affecting the differential trajectory of coal markups. Technological advances of the fracking boom were unrelated to coal mining costs; the Henry Hub spot price is also uncorrelated with

29. Supplementary Material, Appendix Figure E2, Panels D–E report results using alternative versions of C_{ojms} and $S(\mathbf{T}_{ojms})$.

30. Delayed pass-through is common in settings where price changes are not instantaneous (e.g. Borenstein et al., 1997). Equation (5) is algebraically equivalent to a standard (non-differenced) distributed lag model, $\sum_{\ell=0}^L \beta_{\ell} D_j \cdot Z_{m-\ell}$, where $\sum_{\ell=0}^L \beta_{\ell} = \tau$. Many coal contracts include flexible price-adjustment provisions that enable rail carriers to adjust markups before contract expiration (Joskow, 1988; Kosnik and Lange, 2011). I estimate equation (5) separately for contract and spot market shipments.

U.S. diesel prices, which drive coal shipping costs.³¹ A violation of the parallel counterfactual trends assumption would occur if some unobservable factor correlated with coal prices (*e.g.* how electricity regulators monitor plants' coal purchase costs at low gas prices) changed differentially for plants with $TREAT_j = 1$ versus $TREAT_j = 0$.³²

While I can estimate equation (5) defining "treatment" only based on market structure (*i.e.* D_j and W_j), my oligopoly model shows that markup changes should also reflect plant-specific demand shocks. If plant j 's coal demand is not sensitive to gas price changes, we should not expect rail carriers to reoptimize plant j 's markups absent any demand shock. In this sense, failing to incorporate plant-specific demand shocks into $TREAT_j$ mischaracterizes the structural relationship between gas prices and coal markups. The next section describes how I estimate plant-specific coal demand parameters, in order to construct versions of $TREAT_j$ that more accurately represent this structural relationship.

5.4. Estimating coal demand

I estimate plant-specific coal demand curves using a semi-parametric policy function approach, following Davis and Hausman (2016).³³ I predict electricity generation conditional on fuel prices, allowing me to infer plant-specific coal demand curves (as in Figure 3). For each coal generating unit, I estimate the following time series regression, where CF_{uh} is unit u 's capacity factor (*i.e.* generation divided by capacity) in hour h :

$$CF_{uh} = \sum_b \alpha_{ub} \mathbf{1}[G_{uh} \in b] + \sum_b \gamma_{ub} \mathbf{1}[G_{uh} \in b] \cdot CR_{ud} + \zeta_u CR_{ud} + \xi_u \mathbf{G}_{uh} + \omega_{uh} \quad (6)$$

Each unit-specific regression predicts \widehat{CF}_{uh} conditional on the daily ratio of u 's marginal costs relative to the marginal costs of gas generators (CR_{ud}), aggregate fossil generation in u 's electricity market region (G_{uh} , in discrete bins b), and controls (\mathbf{G}_{uh}).³⁴ This accounts for two first-order determinants of unit u 's operating decisions: its relative position in the supply curve

31. The fracking boom may have impacted coal labour markets; C_{ojms} includes average county-year mine-mouth prices, subsuming any differential wage pass-through to commodity costs. Fracked oil increased rail congestion in western states (Covert and Kellogg, 2018); my results are robust to dropping western coal shipments. During 2002–15, the correlation between Henry Hub and U.S. average monthly diesel prices was -0.01 ; this suggests that multicollinearity between Z_m and diesel prices in \mathbf{T}_{ojms} is unlikely.

32. Christian and Barrett (forthcoming) show that even spurious time trends can induce bias for a DD treatment variable that interacts a cross-sectional characteristic with an exogenous time series. Supplementary Material, Appendix Figure E9 shows that P_{ojms} exhibits parallel pre-trends across three definitions of $TREAT_j$. I also assume that SUTVA holds, and that gas price changes do not cause markup spillover effects across plants with different values of $TREAT_j$. Such spillovers are unlikely for $TREAT_j = D_j(1 - W_j)$, since the "untreated" group is composed mostly of plants with a coal-by-barge option (who likely face markups close to zero).

33. This approach treats coal demand estimation as a prediction problem, rather than an optimization problem. An optimization approach would be extremely challenging since (i) regulated plants do not necessarily minimize short-run costs (Cicala, 2015); (ii) non-market plants do not respond to wholesale electricity prices (Cicala, 2022); (iii) complex transmission constraints impact plant output (Borenstein *et al.*, 2000); and (iv) plants have state-dependent, non-convex operating costs (Mansur, 2008).

34. Plants comprise units (boilers) with distinct variable costs and operating decisions. $CF_{uh} \in [0, 1]$ by construction; $CF_{uh} = 0$ if unit u does not operate in hour h . G_{uh} sums hourly CEMS generation across all units in u 's market region. CR_{ud} divides unit u 's marginal cost (including coal price and environmental costs) by the average marginal cost of gas units in the same PCA. \mathbf{G}_{uh} includes the daily maximum, minimum, and standard deviation of G_{uh} ; daily maximum temperature; hour-of-day, quarter-of-year, and year fixed effects; and year dummies interacted with the daily sum of G_{uh} . Supplementary Material, Appendix A.1 provides more detail on my demand estimation procedure, including equations (6)–(9). Supplementary Material, Appendix A.2 conducts sensitivity analysis.

(captured by CR_{ud}), and whether hour h has high versus low electricity demand (captured by G_{uh}).³⁵ G_{uh} is plausibly exogenous, since electricity demand is perfectly inelastic in the short run and a fossil plant is nearly always at the margin. To account for coal price endogeneity in CR_{ud} , I instrument using the average coal price in each state.³⁶

After estimating equation (6) for each unit, I solve each fitted model for the range of counterfactual coal prices at which unit u would likely have produced at capacity in hour h (i.e. any price P where $\widehat{CF}_{uh}(P) > 0.5$).³⁷ Then, I integrate discretized $\widehat{CF}_{uh}(P)$ over all hours in each month and across each of plant j 's units, scaling up to coal quantities based on each unit's size and efficiency. This creates a quantity-price mapping for each plant-month, which I invert to arrive at an estimated inverse demand curve $\widehat{P}_{jm}(\cdot)$.

I use $\widehat{P}_{jm}(\cdot)$ to estimate empirical analogues of the three partial derivatives in equation (3). For each coal plant j , I separately estimate the following OLS regressions:

$$\{\widehat{P}_{jm}(\widehat{Q}_{j\ddot{m}})\}_{j\ddot{m}} = \lambda_{0j}Z_m + \phi_{j\ddot{m}} + \epsilon_{j\ddot{m}} \rightarrow \hat{\lambda}_{0j} \sim \frac{\partial P_j}{\partial Z} \quad (7)$$

$$\{\Delta \widehat{P}_{jm}(\widehat{Q}_{j\ddot{m}}) \cdot \widehat{Q}_{j\ddot{m}}\}_{j\ddot{m}} = \lambda_{1j}Z_m + \kappa_{j\ddot{m}} + v_{j\ddot{m}} \rightarrow \hat{\lambda}_{1j} \sim \frac{\partial^2 P_j}{\partial Q_j \partial Z} Q_j \quad (8)$$

$$\{\Delta^2 \widehat{P}_{jm}(\widehat{Q}_{jm}) / \Delta \widehat{P}_{jm}(\widehat{Q}_{jm}) \cdot \widehat{Q}_{jm}\}_{jm} = \lambda_{2j} + \iota_{jm} \rightarrow \hat{\lambda}_{2j} \sim E_{D_j}. \quad (9)$$

In equation (7), the dependent variable plugs plant j 's predicted coal consumption in month \ddot{m} into its estimated inverse demand curve for every month m (creating an m - \ddot{m} "panel"). Equation (8) is analogous, using discrete approximations of the slope of inverse demand in month m (denoted $\Delta \widehat{P}_{jm}(\cdot)$) and normalizing by $\widehat{Q}_{j\ddot{m}}$. Z_m is the average Henry Hub gas price for month m ; $\phi_{j\ddot{m}}$ and $\kappa_{j\ddot{m}}$ are month \ddot{m} fixed effects. $\hat{\lambda}_{0j}$ and $\hat{\lambda}_{1j}$ estimate how variation in Z_m alters the level and slope of plant j 's inverse coal demand in month m (at each quantity $\widehat{Q}_{j\ddot{m}}$).³⁸ Finally, equation (9) is a time series regression on a constant; $\hat{\lambda}_{2j}$ captures the average (discrete approximation of the) elasticity of the slope of plant j 's inverse coal demand.³⁹

Since the relevant observation for modelling power plant utilization is the unit-hour, while the relevant observation for modelling coal demand is the plant-month, I cannot simply derive $\langle \hat{\lambda}_{0j}, \hat{\lambda}_{1j}, \hat{\lambda}_{2j} \rangle$ analytically from equation (6).⁴⁰ However, by aggregating unit-hour predicted capacity factors up to plant-month coal demand curves, I am able to parsimoniously account for non-linearities in hourly electricity production. Using monthly gas price variation to identify $\hat{\lambda}_{0j}$ and $\hat{\lambda}_{1j}$ yields time-invariant parameter estimates for each plant j .⁴¹

35. Interacting G_{uh} with CR_{ud} accounts for joint changes in electricity supply and demand: an increase in CR_{ud} may reduce CF_{uh} when demand is low but have no effect on CF_{uh} when demand is high.

36. This is a strong instrument: the correlation between plant-specific and state-average coal prices is 0.86.

37. The coal price P_{jm} (for plant j in month m) enters equation (6) as a component of CR_{ud} . Discretizing at $\widehat{CF}_{uh} > 0.5$ greatly reduces computation time without sacrificing realism: within-unit operating decisions are close to binary (i.e. $CF_{uh} \approx 0$ or $CF_{uh} \approx 1$; see [Supplementary Material, Appendix Figure A1](#)). Aggregating across units and hours smooths out kinks created by this discretization.

38. [Supplementary Material, Appendix Figure A2](#) shows how gas price variation identifies $\hat{\lambda}_{0j}$ and $\hat{\lambda}_{1j}$.

39. Δ^2 denotes a discrete approximation of the second derivative. I two-way cluster the error terms $\epsilon_{j\ddot{m}}$ and $v_{j\ddot{m}}$ by m and \ddot{m} ; for ι_{jm} , I use Newey–West standard errors with 12 lags. I use these standard error assumptions when bootstrapping to correct for generated regressors (discussed in [Supplementary Material, Appendix E.5](#)).

40. Plants can operate units separately, electricity markets clear hourly, and power is not feasibly storable. Coal is cheaply storable, purchased periodically, and not easily attributed to specific units within the plant.

41. I lack the statistical power to estimate time-varying versions of $\langle \hat{\lambda}_{0j}, \hat{\lambda}_{1j}, \hat{\lambda}_{2j} \rangle$. Moreover, rail carriers likely cannot adjust long-term contract markups at a high temporal frequency.

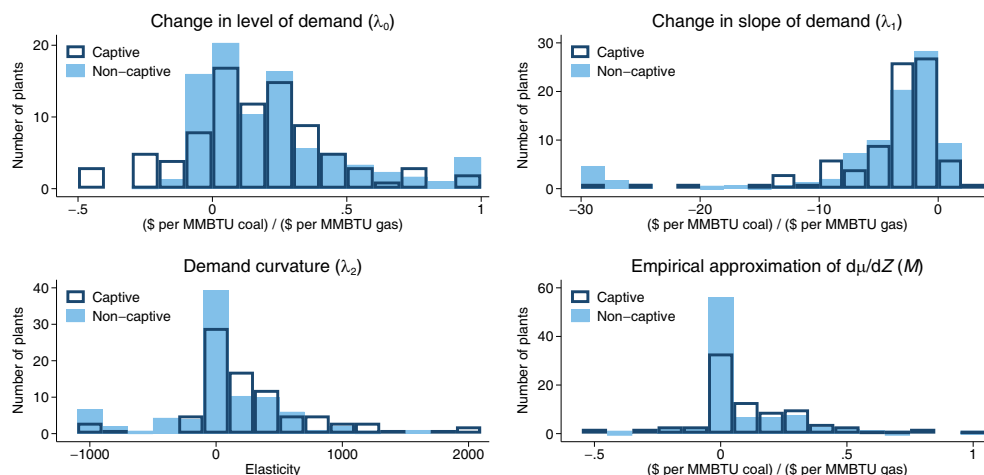


FIGURE 5

Coal demand estimation results

Notes: These histograms report the distributions of estimated demand parameters ($\hat{\lambda}_{0j}$, $\hat{\lambda}_{1j}$, $\hat{\lambda}_{2j}$ from equations (7)–(9)), and the empirical approximation of the comparative static $\frac{d\mu}{dZ}$ (M_j from equation (10)). Each histogram includes one observation per plant and applies nearest-neighbour weights ($k = 3$). Matching criteria: up to k nearest neighbours within 200 miles; exact matches on coal rank; and removing non-utility and non-rail plants. I winsorize the outer bins of each histogram for ease of presentation. All plots include 87 captive plants and 97 matched non-captive plants.

Figure 5 plots histograms of these estimated parameters, for captive versus non-captive plants. The distribution of $\hat{\lambda}_{0j}$ has a median of 0.18, implying that for a \$1/MMBTU decrease in gas price, coal price would need to fall by \$0.16/MMBTU to maintain the median plant's baseline coal consumption.⁴² $\hat{\lambda}_{0j} \in [0, 0.5]$ for 66% of plants, consistent with moderate substitution between coal and gas. $\hat{\lambda}_{1j} < 0$ for 91% of plants, implying that nearly all coal plants have become more marginal in electricity markets at lower gas prices, with more elastic coal demand (*i.e.* less steep inverse demand). Finally, $\hat{\lambda}_{2j} > 0$ for 69% of plants, suggesting that coal demand tends to be (locally) concave.⁴³

Importantly, $\hat{\lambda}_{0j}$, $\hat{\lambda}_{1j}$, and $\hat{\lambda}_{2j}$ come from linear predictions that impose no assumptions on plant j 's objective function, or on the shape of coal demand. My estimated demand curves do hold the rest of the market fixed, including coal prices faced by other plants.⁴⁴ This means that my demand estimates could not predict the effects of a common coal price shock. However, they *can* predict variation in plant j 's idiosyncratic opportunity cost of coal—the very type of price change that occurs when a rail carrier reoptimizes plant j 's markup.

42. Whereas my coal shipment regressions use \$/ton (weight matters most for rail freight), my demand parameters use \$/MMBTU (energy content matters most for coal combustion). The mean (standard deviation) BTU content in my estimation sample is 19.7 (3.4) MMBTU/ton.

43. [Supplementary Material, Appendix B.4](#) shows that $\hat{\lambda}_{2j}$ tends to be consistent with estimates of the second-order condition.

44. In reality, rail carriers may jointly reoptimize markups across multiple plants in the same electricity market. If markups move in the same direction for multiple plants, then my coal demand estimates may be too large (small) at low (high) coal prices. I also assume plants operate their units independently and impose discreteness in counterfactual coal consumption (*i.e.* either zero or at unit u 's capacity).

5.5. Incorporating coal demand parameters into my DD framework

Figure 5 reveals that fracking-induced shocks to coal demand were not correlated with rail captiveness. Hence, defining DD “treatment” using only market structure indicators (*i.e.* captiveness D_j , water option W_j) ignores both non-captive plants with negative demand shocks (*i.e.* $\hat{\lambda}_{0j} > 0$) and captive plants without demand shocks (*i.e.* $\hat{\lambda}_{0j} \approx 0$). I can refine my “treated” group by interacting plant j ’s market structure with an indicator for an above-median demand shock: $TREAT_j = D_j \mathbf{1}[\hat{\lambda}_{0j} \geq 0.16]$, and $TREAT_j = D_j(1 - W_j) \mathbf{1}[\hat{\lambda}_{0j} \geq 0.16]$. For these definitions of “treatment”, equation (5) more accurately targets the set of plants likely to have experienced markup changes.

Going one step further, I can use my estimated demand parameters $\langle \hat{\lambda}_{0j}, \hat{\lambda}_{1j}, \hat{\lambda}_{2j} \rangle$ to construct an empirical approximation of my comparative static in equation (3):

$$\hat{M}_j \equiv \begin{cases} \frac{\hat{\lambda}_{0j} [D_j + \hat{\lambda}_{2j}(2 - D_j)^{-1}] - \hat{\lambda}_{1j}}{2 + \hat{\lambda}_{2j}(2 - D_j)^{-1}} & \text{if } W_j = 0 \\ 0 & \text{if } W_j = 1 \end{cases} \quad (10)$$

I use the captiveness indicator D_j to parameterize N_j (plant j ’s number of rail carriers), setting $N_j = 2 - D_j = 1$ for captive plants and $N_j = 2 - D_j = 2$ for non-captive plants.⁴⁵ I set $\hat{M}_j = 0$ for plants with a more competitive coal-by-barge option ($W_j = 1$), which likely prevents rail carriers from exercising market power (or reoptimizing markups).

The bottom-right panel of Figure 5 plots a histogram of \hat{M}_j —the predicted change in coal markups for a \$1/MMBTU increase in gas price, under the assumptions of my oligopoly model. The mass at $\hat{M}_j = 0$ includes: plants with coal-by-barge options; most non-captive, rail-only plants; and many captive, rail-only plants with negligible demand shocks. However, $\hat{M}_j > 0$ for 54% of captive plants, and for 35% of non-captive plants that have thus far been classified as $TREAT_j = 0$. For this subset of plants with $\hat{M}_j > 0$, the interquartile range is $\hat{M}_j \in [0.12, 0.35]$, implying predicted markup changes that vary by a factor of 3.

I take these predictions to the data by estimating equation (5) using two additional “treatment” variables. First, I use the indicator $TREAT_j = \mathbf{1}[\hat{M}_j \geq 0.29]$, where 0.29 is the mean of \hat{M}_j ’s positive support. Second, I use the continuous $TREAT_j = \hat{M}_j$, which captures variation across the range of non-zero \hat{M}_j predictions. Importantly, \hat{M}_j does not provide accurate *quantitative* predictions of markup changes: my stylized oligopoly model abstracts from rail regulation, heterogeneous coal attributes, and other factors controlled for in equation (5). However, \hat{M}_j does provide substantial signal for identifying differential markup changes. For definitions of $TREAT_j$ that include generated regressors ($\hat{\lambda}_{0j}$ or \hat{M}_j), I block bootstrap by plant and simulate random draws from the $\langle \hat{\lambda}_{0j}, \hat{\lambda}_{1j}, \hat{\lambda}_{2j} \rangle$ sampling distributions to construct confidence intervals for $\hat{\tau}$.⁴⁶

45. I set $N_j = 2$ for non-captive plants since both the western and eastern U.S. have near duopolies in coal-by-rail shipping (see Figure 1). For my analysis, I average \hat{M}_j across simulated draws of $\langle \hat{\lambda}_{0j}, \hat{\lambda}_{1j}, \hat{\lambda}_{2j} \rangle$, since equation (10) is non-linear in estimated parameters; I winsorize to reduce the influence of extreme draws with a near-zero denominator (see [Supplementary Material, Appendix E.5](#) for details).

46. I construct bias-corrected accelerated bootstrap confidence intervals, following [Efron \(1987\)](#) and [Cameron et al. \(2006\)](#). [Supplementary Material, Appendix E.5](#) describes my bootstrapping procedure in detail.

TABLE 2
Markup levels—captive versus non-captive coal plants

	Outcome: delivered coal price (\$/ton)					
	(1)	(2)	(3)	(4)	(5)	(6)
I[Captive] _j	2.37*** (0.71)	1.87*** (0.61)	1.71*** (0.58)	2.50** (0.96)	2.75*** (0.71)	2.30*** (0.76)
<i>k</i> nearest neighbours	1	3	5	1	3	5
Removing plants with water option				Yes	Yes	Yes
Avg price in omitted group (\$/ton)	40.80	38.06	39.70	45.30	40.70	43.30
Plants	142	184	195	85	99	106
Observations	66,336	88,529	93,968	46,814	54,841	59,793

Notes: Regressions estimate equation (4) at the county-plant-month-transaction level, for rail shipments only. I control for shipping costs using the 4-way interaction of rail distance, diesel price, tons shipped, and rail traffic density. I also control for plant- and delivery-specific controls (listed in Table 1), average coal price in the originating county, BTU content of each shipment, distance to plants' closest rail terminal, baseload natural gas capacity in each plant's PCA, origin county fixed effects, and month-of-sample fixed effects. Matching criteria: up to *k* nearest neighbours within 200 miles; exact matches on coal rank; and removing non-utility and non-rail plants. Regressions use nearest-neighbour weights, and also weight by the quantity of coal transacted. Standard errors are clustered by plant. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

6. RESULTS

6.1. Markup levels

Table 2 reports results from estimating equation (4), showing that captive plants indeed face higher markups than non-captive plants. In Columns (1)–(3), point estimates of \$2 translate to average differential markups of 4–6%, on an average delivered price of \$38–40/ton for non-captive plants. This implies that markups for captive plants contribute 13–19% of the spatial gap between mine-mouth prices (averaging \$26–28/ton) and delivered prices. Columns (4)–(6) remove plants with the outside option to receive barge shipments. While this reduces the sample size by 30%, my markup estimates retain statistical precision and increase in magnitude. These results are robust to the number of nearest-neighbour matches, and to a wide range of sensitivity checks reported in [Supplementary Material, Appendix Figure E2](#).

Given measurement error in how I assign captiveness, and since non-captive plants likely face non-zero markups, these estimated *differentials* likely understate the average markup levels faced by captive plants. Modifying equation (4) to interact captiveness with an indicator for a coal-by-barge option, I estimate differential markups as large as \$5/ton for captive non-water plants—relative to the most-competitive comparison group of non-captive plants with a barge option (see [Supplementary Material, Appendix Table E3](#)). This implies markup levels as high as 13% of the delivered price, and up to 34% of rail carriers' shipping cost.⁴⁷

How does this market power distortion compare to coal's climate damages? I can reject differential coal-by-rail markups greater than \$7/ton, which corresponds to \$2–5/metric ton CO₂. This is far below recent social cost of carbon estimates of \$185/metric ton CO₂ ([Rennert et al., 2022](#)).⁴⁸ Hence, the welfare gains from a Pigouvian carbon tax would likely dwarf any welfare

47. These magnitudes align with \$4/ton coal-by-rail markup estimates from [Busse and Keohane \(2007\)](#).

48. These estimates (centred on \$185/metric ton) form the basis of the Environmental Protection Agency's regulatory impact analyses for new greenhouse gas regulations (U.S. EPA, 2022).

TABLE 3
Markup DD results—linear \hat{M}_j interacted with gas price

DD estimates with 48 lags	Outcome: delivered coal price (\$/ton)		
	(1)	(2)	(3)
$\hat{M}_j \times (\text{Gas price})_m$	3.66 [1.69, 9.73]	2.94 [0.51, 9.22]	3.01 [0.43, 9.44]
DD effect for $\hat{M}_j = 0.29$	1.06	0.85	0.87
DD effect for $\hat{M}_j = 0.44$	1.61	1.29	1.32
k nearest neighbours	1	3	5
Mean of dep var ($\hat{M}_j = 0$)	31.34	32.21	32.53
Plants	140	182	192
Observations	65,856	88,038	93,476

Notes: Regressions estimate equation (5) with $TREAT_j = \hat{M}_j$, for all shipments (contract and spot-market). I report the cumulative change in markups (over 48 months) caused by a \$1/MMBTU increase in gas price, for a plant with $\hat{M}_j = 1$ compared to a plant with $\hat{M}_j = 0$. Regressions are identical to Figure 6, except that $TREAT_j$ is continuous not binary. $\hat{M}_j Z_m$ is in \$/MMBTU of coal; BTU content ranges from 14 to 30 MMBTU/ton. I report effects for plants with $\hat{M}_j = \{0.29, 0.44\}$, the minimum and mean of the “treated” group in regressions (9)–(10) of Figure 6. Brackets report bias-corrected accelerated bootstrapped 95% confidence intervals. Regressions include month-of-sample and plant-by-origin-county fixed effects; commodity, shipping cost, and plant controls; and $k \in \{1, 3, 5\}$ nearest neighbour weights. Regressions drop two plants with extreme outlier values of \hat{M}_j , which I address in [Supplementary Material, Appendix E.6](#). See notes under Figure 6 for further details.

loss from exacerbating market power (echoing [Oates and Strassmann, 1984](#)).⁴⁹ In practice, carbon taxes tend to be less stringent than this Pigouvian ideal of \$185/metric ton CO₂ ([Metcalf, 2021](#)). Coal markups would complement a suboptimally low U.S. carbon tax: increasing welfare by internalizing an additional share of marginal climate damages.⁵⁰

6.2. Markup changes

Next, I estimate heterogeneous *changes* in coal markups due to changes in the gas price. I estimate equation (5) using five binary definitions of $TREAT_j$, and plot 48-month cumulative effects in Figure 6. While regressions (1)–(2) do not detect differential markup changes using captiveness alone, regressions (3)–(4) find statistically significant DD effects for captive, rail-only plants. Refining $TREAT_j$ by interacting rail market structure with an indicator for an above-median demand shock (*i.e.* $\hat{\lambda}_{0j} \geq 0.16$), my DD estimates increase in magnitude (regressions (5)–(8)). This reflects how both supply- and demand-side factors govern the structural relationship between gas prices and coal markups. Regressions (9)–(10) parameterize this structural relationship by defining $TREAT_j = \mathbf{1}[\hat{M}_j \geq 0.29]$: plants with large predicted markup changes. This yields DD estimates of $\hat{\tau} = 1.41$ for all shipments and $\hat{\tau} = 1.52$ for contract

49. By contrast, the distortion above marginal cost pricing is large relative to pollution externalities in U.S. retail natural gas ([Davis and Muehlegger, 2010](#)), and cement markets ([Fowle et al., 2016](#)).

50. As I discuss below in Section 7.3, fracking-induced decreases in the natural gas price (*i.e.* my “natural experiment”) mimicked an implicit carbon tax of roughly \$24/metric ton CO₂. Even under a larger tax equal to marginal climate damages, markups would still increase welfare by internalizing local air pollution externalities ([Muller et al., 2011](#)). [Goodkind et al. \(2019\)](#) find that average marginal damages from criteria pollutants (SO₂ and NO_x) may exceed \$100/ton of coal. Since this is much larger than the markup distortion, a tax of \$185/metric ton CO₂ would likely still be smaller than first-best.

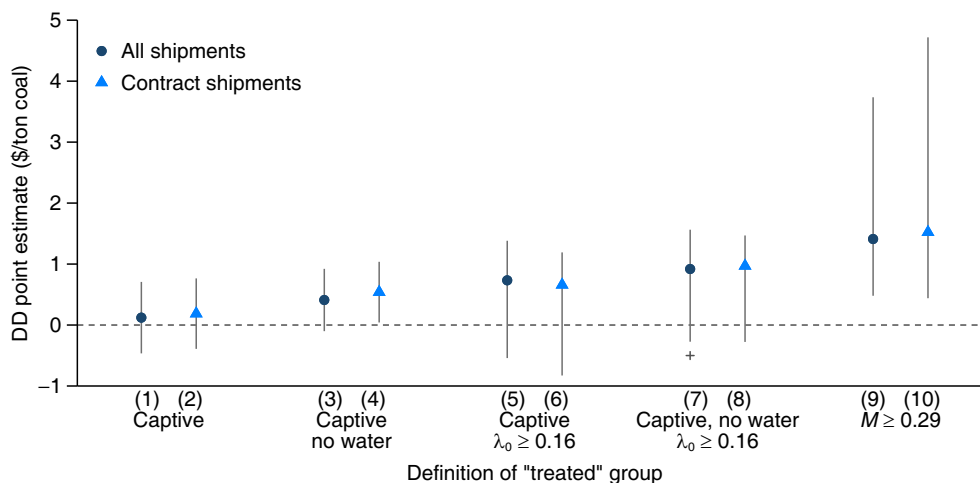


FIGURE 6

Markup DD results—“treatment” interacted with gas price

Notes: This figure plots DD point estimates and 95% confidence intervals from 10 regressions of equation (5). I define $TREAT_j$ 5 ways (reading left to right): D_j ; $D_j(1 - W_j)$; $D_j\mathbf{1}[\hat{\lambda}_{0j} \geq 0.16]$; $D_j(1 - W_j)\mathbf{1}[\hat{\lambda}_{0j} \geq 0.16]$; $\mathbf{1}[\hat{M}_j \geq 0.29]$. The outcome variable is delivered coal price, averaging \$33–38/ton for the omitted group. Triangles use only the 78% of shipments on long-term contracts. I interact $TREAT_j$ with the Henry Hub natural gas spot price in a lag-differenced model; I report $\hat{\tau}$, or cumulative effects over $L = 48$ months. Regressions include month-of-sample and plant-by-county fixed effects, and the same commodity, shipping cost, and plant controls as Table 2. All regressions use $k = 3$ matches, applying nearest-neighbour weights and also weighting by the quantity of coal transacted. Regressions (1)–(4) cluster standard errors by plant. To account for generated regressors in regressions (5)–(10), I block bootstrap by plant using draws from simulated $(\hat{\lambda}_{0j}, \hat{\lambda}_{1j}, \hat{\lambda}_{2j})$ sampling distributions, then construct bias-corrected accelerated bootstrap confidence intervals (Efron, 1987). Pluses denote estimates where the 95% confidence interval includes zero, but the 90% confidence interval does not. [Supplementary Material, Appendix Table E4](#) reports these results numerically. Regressions (5)–(10) drop two plants with extreme outlier values of \hat{M}_j , which I address in [Supplementary Material, Appendix E.6](#).

shipments. Since gas prices fell by \$4/MMBTU during the fracking boom, this implies that rail carriers reduced coal markups by \$6/ton for plants with high \hat{M}_j (compared to plants with low \hat{M}_j).

Table 3 reports results using $TREAT_j = \hat{M}_j$, the predicted change in coal markups for a \$1/MMBTU increase in gas price, for all shipments and $k \in \{1, 3, 5\}$ nearest neighbours. I find positive, statistically significant point estimates, indicating that plants with greater \hat{M}_j experienced larger markup reductions.⁵¹ If my \hat{M}_j predictions were quantitatively accurate, I would recover $\hat{\tau} \approx 20$: \hat{M}_j is in (\$/MMBTU coal)/(\$/MMBTU gas), P_{ojms} is in \$/ton coal, and coal averages 19.7 MMBTU/ton. However, I estimate DD effects of $\hat{\tau} \approx 3$. This mismatch in magnitudes underscores the limitations of my simple oligopoly model, which ignores rail regulation (among other factors).⁵² At the same time, Table 3 shows that my model’s predictions are *qualitatively* consistent with observed markup changes: when “treatment” better captures variation in

51. These results are robust to a wide range of sensitivity checks (see [Supplementary Material, Appendix Figure E8](#)). [Supplementary Material, Appendix Table E7](#) reports similar results using a two-step DD estimator that more closely matches equations (2)–(3).

52. [Supplementary Material, Appendix B.3](#) shows that accounting for the threat of regulation could increase or decrease my \hat{M}_j predictions. My model also assumes that rail carriers optimize markups independently for each plant; to the extent that this is not true, \hat{M}_j would likely over-predict markup changes. Even if my oligopoly model perfectly characterized $\frac{d\mu_j}{dZ}$, I would likely estimate $\hat{\tau} < 20$ for two reasons: (i) attenuation bias due to measurement error in \hat{M}_j ; and (ii) markup decreases that stop at (or close to) the lower bound $\mu_j = 0$.

the structural relationship between gas prices and coal markups, I can detect much larger, more statistically precise markup changes. The magnitudes in Table 3 are consistent with Figure 6: for a plant with $\hat{M}_j = 0.44$ (the mean for the discrete “treated” group with $\hat{M}_j \geq 0.29$), Column (2)’s DD effect is $0.44 \times 2.94 = 1.29$. This implies that when gas prices fell by \$4/MMBTU, coal markups fell by $1.29 \times 4 = \$5.16/\text{ton}$.⁵³

My results imply that rail carriers heterogeneously reoptimized markups to eliminate most of the market power distortion for a subset of plants. Unlike previous estimates of coal-by-rail markups (Busse and Keohane, 2007; Hughes and Lange, 2020), my DD design explicitly incorporates the theoretical predictions from an oligopoly model. This enables me to (i) detect a larger markup distortion, (ii) attribute markup changes to strategic reoptimization consistent with profit maximization, and (iii) decompose supply-side versus demand-side mechanisms.⁵⁴ Each factor improves the credibility of my results, especially in using my DD estimates to infer outcomes under a counterfactual climate policy.

In addition, my “structural” DD estimates impose relatively few restrictions on demand, compared to the broader literature estimating transportation markups and cost pass-through. While such studies tend to identify demand using complex structural models,⁵⁵ I estimate coal demand for each plant without even specifying its objective function. Similar techniques could be used to estimate demand for other commodities that face transportation market power due to geographic specificity and high freight costs—including crude oil (Covert and Kellogg, 2018); ethanol (Hughes, 2011); cement (Hortaçsu and Syverson, 2007; Miller and Osborne, 2014); and metals (Ellis and Halvorsen, 2002).

7. IMPLICATIONS FOR CLIMATE POLICY

7.1. *Heterogeneous pass-through of an implicit carbon tax*

A negative gas price shock makes coal plants less competitive in electricity supply. A tax on CO₂ emissions similarly disadvantages coal, the more carbon-intensive fuel. Cullen and Mansur (2017) show that under reasonable assumptions, the coal-to-gas price ratio is a sufficient statistic for CO₂ emissions from the electricity sector under a counterfactual carbon tax. If electricity demand is perfectly inelastic, and if only coal or gas generators can be marginal in electricity supply, then a decrease in gas price ($\Delta Z < 0$) should yield the same short-run emissions outcomes as the equivalent carbon tax t :

$$\text{CR} = \frac{MC_{\text{coal}}}{MC_{\text{gas}}} = \frac{P}{Z + \Delta Z} = \frac{P + tE_{\text{coal}}}{Z + tE_{\text{gas}}}. \quad (11)$$

53. Supplementary Material, Appendix Figure E3 reports similar results using discrete bins of \hat{M}_j (i.e. relaxing linearity in \hat{M}_j). Supplementary Material, Appendix Figure E4 shows that these DD effects accumulated gradually after a 1-year lag, for both contract and spot-market shipments. I find similar results using a DD event-study specification (Supplementary Material, Appendix Figures E5–E7).

54. To illustrate the advantage of my more structural approach, Hughes and Lange (2020) attribute differential markup changes to information gains from market-based (as opposed to traditional) electricity dispatch. However, I show that heterogeneous shocks to coal demand impact markups for plants in both groups—by sharpening the signal in $TREAT_j$ using variation in market structure.

55. Examples include Nakamura and Zerom (2010), Goldberg and Hellerstein (2013), and Brancaccio *et al.* (2020).

Using Cullen and Mansur's framework, CR is the ratio of marginal costs (MC_{fuel}), P is the coal price paid by power plants, and E_{fuel} are fuel-specific CO₂ emissions factors.⁵⁶

My DD results demonstrate that P is not fixed, and I can rewrite equation (11) to allow plant-specific coal markups (μ_j) to endogenously respond to ΔZ :

$$\frac{P_j + \Delta\mu_j}{Z + \Delta Z} = \frac{P_j + \rho_j t E_{\text{coal}}}{Z + t E_{\text{gas}}} \quad (12)$$

ρ_j is the pass-through rate of the implicit tax t . If plant j 's markups do not change ($\Delta\mu_j = 0$), then its pass-through rate is $\rho_j = 1$ (i.e. full pass-through, as in equation (11)). If plant j 's markups decrease in response to a negative gas price shock ($\Delta\mu_j < 0$, for $\hat{M}_j > 0$), then $\rho_j < 1$ and pass-through of t is incomplete.⁵⁷

I translate my DD coefficients from Table 3 into estimates of implied pass-through rates $\hat{\rho}_j$ by rearranging equation (12) and substituting $\Delta\mu_j = \hat{\tau} \hat{M}_j \left(\frac{\text{MMBTU}}{\text{ton}}\right)_j \Delta Z$.⁵⁸ Figure 7 plots $\hat{\rho}_j$ across the range of predicted markup changes \hat{M}_j , weighting by the quantity of coal shipped and $k = 3$ nearest neighbours. This reveals substantial heterogeneity in predicted pass-through, driven both by variation in predicted markup changes and variation in P_j : lower coal prices (e.g. due to shorter rail distances) will magnify markup changes, pulling $\hat{\rho}_j$ away from 1. For 50% of coal-by-rail shipments, I estimate $\hat{\rho}_j < 1$: implying that rail carriers reoptimized markups during the fracking boom to effectively insulate most plants from the negative shock to relative costs. For 17% of shipments, $\hat{M}_j < 0$ implies $\hat{\rho}_j > 1$ (i.e. overshifting); for the 33% of rail shipments to plants with a coal-by-barge option, $\hat{M}_j = 0$ implies $\hat{\rho}_j = 1$.⁵⁹ To my knowledge, this is the first empirical evidence that predicts heterogeneous and incomplete pass-through of a carbon tax—as low as 74 cents on the dollar—in either U.S. coal markets or the U.S. electricity sector.⁶⁰

My results contribute to a growing body of research finding heterogeneous pass-through of price-based climate policies. Previous work has shown that variation in market structure either across industries (Ganapati *et al.*, 2020), or across space within an industry

56. Here, all costs are in \$/MMBTU and emissions rates are in tons CO₂/MMBTU from combustion. Rail shipping contributes additional CO₂ emissions from diesel locomotives, which are two orders of magnitude smaller than combustion emissions per MMBTU of coal (U.S. EPA, 2008).

57. I assume full pass-through of t for gas plants. Supplementary Material, Appendix C.3 provides evidence to support this assumption and shows that relaxing it would not meaningfully change estimated pass-through for coal plants.

58. For $TREAT_j = \hat{M}_j$, $\hat{\tau} \hat{M}_j \left(\frac{\text{MMBTU}}{\text{ton}}\right)_j$ estimates $\Delta\mu_j$ in (\$/MMBTU coal)/(\$/MMBTU gas), converting from \$/ton coal using plant j 's BTU intensity by weight. Hence, $\Delta\mu \sim \hat{\tau} \hat{M}_j \left(\frac{\text{MMBTU}}{\text{ton}}\right)_j \Delta Z$ for any ΔZ (by linearity). This assumes $\hat{\rho}_j = 1$ for $\hat{M}_j = 0$. Supplementary Material, Appendix C.1 outlines these assumptions in further detail; describes how I parameterize P_j , Z , E_{coal} , and E_{gas} ; and provides derivations of t and $\hat{\rho}$.

59. Supplementary Material, Appendix C.2 reports sensitivity analysis on $\hat{\rho}_j$, for different $\hat{\tau}$ estimates, adding environmental compliance costs to equations (11)–(12), and including gas pipeline costs in the parameterization of Z .

60. Kim *et al.* (2010) conceptually illustrate how variation in power plants' costs may lead to incomplete carbon tax pass-through. Chu *et al.* (2017) estimate incomplete pass-through from coal spot prices to delivered coal prices; the authors caution that their analysis is not predictive of long-term price changes that would occur under a carbon tax. Hughes and Lange (2020) apply my framework for inferring pass-through for regulated versus deregulated coal plants; they recover pass-through estimates close to 1 for regulated plants, and close to 0.85 for plants in deregulated electricity market regions.

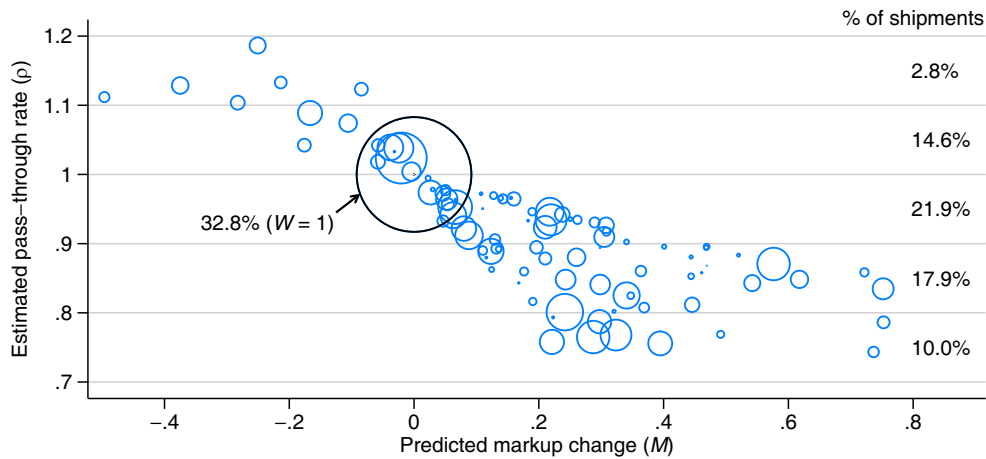


FIGURE 7
Estimated pass-through of an implied carbon tax

Notes: This figure converts my DD estimate from Column (2) of Table 3 into estimated pass-through rates of an implied carbon tax. I rearrange equation (12) to solve for ρ_j ; substitute $\Delta\mu_j = \hat{M}_j \Delta Z$; set P_j and Z to their 2007–08 averages; and set $E_{\text{coal}} = 0.095$ and $E_{\text{gas}} = 0.053$, each fuel's average metric tons CO_2 /MMBTU. Markers apply $k = 3$ nearest-neighbour weights, and scale by the total 2007–15 coal deliveries. On the right, I report the percentage of coal-by-rail shipments for which $\hat{\rho}_j \in (0.7, 0.8)$, $(0.8, 0.9)$, etc. Each marker denotes a separate plant; the big circle aggregates all plants with a barge option (*i.e.* $W_j = 1$ implies $\hat{\rho}_j = 1$). [Supplementary Material, Appendix C.1](#) provides a derivation for $\hat{\rho}_j$, while [Supplementary Material, Appendix C.2](#) presents sensitivity analysis.

(Lade and Bushnell, 2019), can generate substantial heterogeneity in pass-through rates.⁶¹ Similarly, I find that heterogeneous pass-through of a carbon tax in U.S. coal markets would arise largely from spatial variation in the competitiveness of coal shipping. However, coal markups also adjust heterogeneously to plant-specific demand shocks. By incorporating this second dimension of heterogeneity into \hat{M}_j , I am able to detect both lower average pass-through rates and greater dispersion in $\hat{\rho}_j$ across plants.⁶²

My findings contrast with the existing literature on carbon tax pass-through in electricity markets. Fabra and Reguant (2014) estimate full pass-through of carbon prices in the Spanish wholesale electricity market, which they attribute to highly correlated cost shocks among marginal plants.⁶³ In my setting, incomplete pass-through arises due to a weaker correlation of cost shocks (*i.e.* markup changes) across coal plants. This echoes Muehlegger and Sweeney (2022), who estimate incomplete pass-through of firm-specific cost shocks in petroleum refining, but full pass-through of cost shocks that are common across firms. My results underscore the value of using economic theory to predict cost shock heterogeneity: ignoring or mischaracterizing such heterogeneity will likely yield overestimates of the average pass-through rate, while understating dispersion in pass-through across firms.

61. Ganapati *et al.* (2020) find heterogeneous energy cost pass-through in manufacturing due to imperfect competition. Lade and Bushnell (2019) find that pass-through of ethanol subsidies varies with spatial market structure (see also Poulriot *et al.*, 2020; Knittel *et al.*, 2017; Li and Stock, 2019). Heterogeneous pass-through may also arise from spatial variation in production capacity (*e.g.* petroleum refining: Marion and Muehlegger, 2011; Borenstein and Kellogg, 2014).

62. Linn and Muehlenbachs (2018) and Kim (2022) also find heterogeneous pass-through in U.S. electricity markets due to heterogeneous cost shocks. Neither paper considers the pass-through of a carbon tax.

63. Subsequent studies have estimated full carbon price pass-through in German (Hintermann, 2016), Australian (Nazifi, 2016), and Greek (Dagoumas and Polemis, 2020) electricity markets.

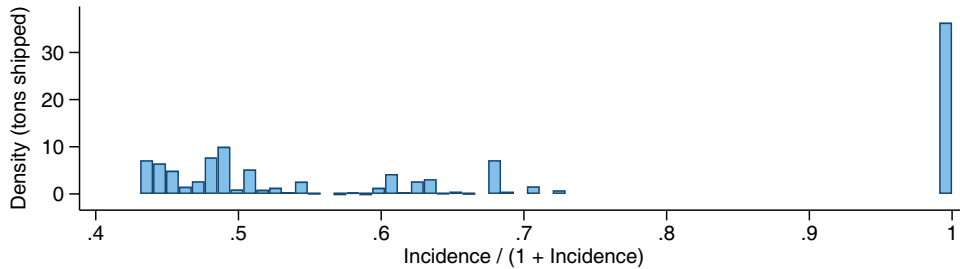


FIGURE 8

Estimated share of implied tax burden borne by coal plants

Notes: This figure plots a histogram of $\hat{I}_j/(1 + \hat{I}_j)$, parameterizing equation (13) using the $\hat{\rho}_j$ estimates from Figure 7. I weight observations (*i.e.* plants) using the product of $k = 3$ nearest-neighbour weights and total 2007–15 coal deliveries.

7.2. Heterogeneous carbon tax incidence

Weyl and Fabinger (2013) show how pass-through under imperfect competition is closely linked to economic incidence. In my setting, the pass-through rate (ρ_j) and market structure are sufficient to characterize the share of implied tax incidence (I_j) borne by coal plants:

$$\frac{I_j}{1 + I_j} = \frac{dCS_j/dt}{dCS_j/dt + dPS_j/dt} = \begin{cases} \frac{\rho_j}{1 + \rho_j/N_j} & \text{if } W_j = 0 \\ 1 & \text{if } W_j = 1 \end{cases}, \quad (13)$$

where CS_j and PS_j are consumer and producer surplus. Lower pass-through rates imply that consumers (*i.e.* coal plants) bear relatively less of the tax burden than producers (*i.e.* rail carriers). For a given pass-through rate ρ_j , a less competitive market structure (*i.e.* smaller N_j) shifts more tax burden away from plants and towards rail oligopolists. Plants who pay competitive prices due to a coal-by-barge option (*i.e.* $W_j = 1$) bear the full burden of an implied carbon tax—since railroads' producer surplus is zero.

Figure 8 plots a histogram of implied tax incidence, parameterizing equation (13) using my $\hat{\rho}_j$ estimates from Figure 7. This reveals substantial variation in implied tax burden: for 33% of coal-by-rail shipments, competitive freight pricing caused plants to pay 100% of lost surplus; for 36% of shipments, over 50% of lost surplus came via foregone rail oligopoly rents. This understates the full costs to coal plants, 91% of which saw their electricity sales decrease during the fracking boom. At the same time, Figure 8 highlights railroads as major energy stakeholders, who have lost substantial economic surplus as coal has declined.⁶⁴

My results add to a nascent body of evidence that the assumption of homogeneous incidence can obscure the true distributional impacts of energy taxes. Stolper (2021) uncovers heterogeneous tax incidence for Spanish transportation fuels, which renders a seemingly regressive tax unambiguously progressive. Similarly, Ganapati *et al.* (2020) show that a carbon tax appears less regressive after accounting for variation in the competitiveness of intermediate product markets. In my setting, heterogeneous incidence suggests that under a carbon tax, certain coal plants

64. Supplementary Material, Appendix Figure C9 plots changes in capacity factors against predicted incidence. Supplementary Material, Appendix C.4 provides a more detailed discussion of implied carbon tax incidence as it pertains to my theoretical framework.

would stand to lose relatively less than others.⁶⁵ By shifting a share of the tax burden further upstream from electricity consumers, market imperfections in coal shipping may also reduce the regressivity of a carbon tax. This finding contributes to an extensive literature on the distributional impacts of climate policy (*e.g.* [Goulder et al., 2019](#)) by identifying a new channel through which a carbon tax would harm producers: lost oligopoly rents in the fossil fuel supply chain.

7.3. Erosion of CO₂ emissions reductions

Figure 2 shows how U.S. electricity generation shifted away from coal as gas prices fell. Several studies have estimated the environmental benefits of fracking-induced switching from coal to gas.⁶⁶ My analysis is the first to show that coal markups have adjusted to partially offset this change in relative fuel prices. This suggests that had coal markups not changed, the fracking boom could have yielded even greater reductions in CO₂ emissions from electricity.

To quantify how decreases in coal markups moderated CO₂ abatement, I consider two counterfactual scenarios: (1) if the fracking boom never happened and gas prices remained high; (2) if the fracking boom did happen but coal markups were fixed. I use pre-2009 Henry Hub futures prices for scenario (1), and plug these counterfactual gas prices into my fitted DD model (with $TREAT_j = \hat{M}_j$) to predict counterfactual coal prices for scenarios (1) and (2). Then, I estimate unit-specific time series regressions similar to equation (6), conditioning on the factual coal-to-gas cost ratio. Plugging counterfactual cost ratios into each fitted model and summing predicted CO₂ emissions across units, I can quantify short-run CO₂ abatement from the fracking boom—both with and without changes to coal markups.⁶⁷

I estimate that low gas prices caused CO₂ emissions to fall by 4.7%, via short-run coal-to-gas substitution alone. However, if coal markups had not changed, this would have been a 5.2% emissions reduction. This implies that decreases in coal markups eroded roughly 10% of the fracking boom's short-run abatement potential—equal to \$11.9 billion in climate damages that could have been avoided.⁶⁸ Importantly, these projections only capture short-run changes on the intensive margin of fossil generation. Since several other margins have contributed to the 20–25% fracking-induced drop in CO₂ emissions from electricity, \$11.9 billion is likely a substantial underestimate of the damages caused by falling coal markups.⁶⁹

65. All coal plants would likely see profits decrease under a carbon tax, yet some plants would likely bear *relatively* less burden in the short run. [Muehlegger and Sweeney \(2022\)](#) find that a carbon tax on petroleum refiners would imply heterogeneous firm-specific cost shocks, also creating relative winners and losers.

66. [Knittel et al. \(2019\)](#) find that a \$1/MMBTU decrease in gas price caused CO₂ emissions from coal to fall by 5–12%. [Holladay and LaRiviere \(2017\)](#) estimate short-run changes in the marginal CO₂ emissions rates that vary substantially across regions. [Fell and Kaffine \(2018\)](#) attribute the decline in coal generation to a combination of low natural gas prices and increased wind generating capacity.

67. This short-run exercise holds generating capacity fixed. I assume perfectly inelastic electricity demand with gas crowding out coal 1-for-1; I sum CO₂ emissions from both fuels. Since actual retail electricity demand is extremely inelastic, any pass-through of higher gas prices in the no-fracking counterfactual would be unlikely to cause misspecification by reducing electricity demand. Following [Cullen and Mansur \(2017\)](#), I include a cubic spline in the average cost ratio across unit u 's PCA; unlike in equation (6), I use the *average* cost ratio to capture *common* shocks to fuel prices. [Supplementary Material, Appendix D](#) discusses my counterfactuals in further detail.

68. I monetize 5.2% – 4.7% = 0.5% unrealized abatement at \$185/metric ton CO₂—the central estimate from the most up-to-date analysis of the social cost of carbon ([Rennert et al., 2022](#)). The EPA has incorporated the [Rennert et al. \(2022\)](#) estimates into its framework for regulatory impact analysis of greenhouse gas regulations (U.S. EPA 2022). Prior to 2022, the EPA had assumed a value closer to \$50/metric ton CO₂, which aligned with pre-2016 integrated assessment modelling (Interagency Working Group 2016).

69. Relevant medium- and long-run margins where falling coal markups likely eroded CO₂ abatement include efficiency improvements at coal plants ([Linn et al., 2014](#)); investment in new gas plants ([Brehm, 2019](#)), and accelerated

This illustrates the extent to which coal price endogeneity might undermine the efficacy of a carbon tax. At the same time, continued competitive pressures might “harvest” markups at the remaining inframarginal coal plants. This suggests that previous retrospective studies—which have treated coal prices as exogenous (*e.g.* Cullen and Mansur, 2017; Fell and Kaffine, 2018)—have likely underestimated CO₂ abatement from a future carbon tax: a sufficiently stringent climate policy would “harvest” nearly all remaining coal markups, largely eliminating the countervailing effect of endogenous coal prices.

My counterfactual analysis also highlights the welfare consequences of market power in coal shipping. Falling coal markups have not only transferred economic surplus from railroads to power plants. They have also meaningfully impacted the allocation of electricity generation by attenuating the displacement of coal, resulting in greater global (and local) pollution damages. Going forward, the welfare effects of this market power distortion will depend on the stringency of U.S. climate policy. As a natural experiment, the \$4/MMBTU drop in gas prices from 2007 to 2011 corresponds to a tax of \$24/metric ton CO₂.⁷⁰ While this implicit tax is far weaker than the Pigouvian ideal (*i.e.* closer to \$185/metric ton CO₂), it aligns with real-world carbon prices that tend to be suboptimally low (Carl and Fedor, 2016; Metcalf, 2021). Until U.S. climate policy becomes stringent enough to make all coal plants marginal (or extramarginal), falling markups will likely continue to slow coal’s decline.

8. CONCLUSION

I provide the first empirical evidence that market power in coal shipping may substantially undermine the efficacy of U.S. climate policy. My findings contribute to the literature that estimates how pre-existing market distortions impact welfare under environmental policies (*e.g.* Ryan, 2012). More broadly, my results underscore how the behaviour of a few upstream firms (in this case, rail carriers) can meaningfully alter market outcomes in the presence of downstream regulation (Kellogg and Reguant, 2021). My pass-through analysis reveals that upstream rent dissipation would likely mute the price signal of a carbon tax in the U.S. electricity sector—attenuating CO₂ abatement while shifting the tax burden towards rail oligopolists. This adds to the literature on market power and climate policy incidence (Ganapati *et al.*, 2020) and could influence how policymakers choose to redistribute carbon tax revenues (Goulder *et al.*, 2019).

While previous studies have documented coal-by-rail market power, my estimates are the first to directly attribute heterogeneous changes in coal markups to variation in both supplier market power and coal demand shocks. These findings have key implications for climate policy instrument choice. For example, heterogeneous pass-through implies that a uniform carbon price may not incentivize an efficient allocation of CO₂ abatement (Montgomery, 1972), and that the optimal second-best climate policy may feature a non-uniform carbon price. Moreover, the relative desirability of second-best climate policies hinges on the magnitudes of existing price distortions (Borenstein and Kellogg, 2023). My results highlight the importance of anticipating changes to coal’s competitiveness, which still have the potential to enhance or undermine policy effectiveness.⁷¹

coal plant retirements (Davis *et al.*, 2022). While low gas prices have increased both climate-damaging methane emissions and non-electricity CO₂ emissions (*e.g.* from space heating; Hausman and Kellogg, 2015), the opposite would occur under a carbon tax.

70. Equation (11) implies $t = \$24/\text{ton}$, using average prices at the start of the fracking boom and $\Delta Z = -4$.

71. In August 2022, the Henry Hub spot price rose to \$8.81/MMBTU, its highest level since July 2008. My results suggest that rail carriers may have increased markups in response to this positive coal demand shock.

Future research should address the long-run implications of coal-by-rail markups, which could delay power plant retirements (Davis *et al.*, 2022) or impact the fiscal health of coal mining communities (Morris *et al.*, 2021). Similar transportation market imperfections may also exist in developing countries that still heavily rely on coal. Finally, future research should seek to incorporate market power distortions into climate policy projects for other carbon-intensive industries (*e.g.* petroleum refining, aluminum).

Acknowledgments. I thank Meredith Fowlie, Severin Borenstein, Francesca Molinari, Maximilian Auffhammer, James Archsmith, Lint Barrage, Erich Battistin, Peter Berck, Susanna Berkouwer, Joshua Blonz, Fiona Burlig, Steve Cicala, Lucas Davis, Karl Dunkle Werner, Harrison Fell, Don Fullerton, Stephen Jarvis, Akshaya Jha, Jeremy Magruder, Dave McLaughlin, Elisabeth Sadoulet, James Sallee, Leo Simon, James Stock, Sofia Villas-Boas, Reed Walker, Matt Woerman, Catherine Wolfram, and Derek Wolfson. I also thank seminar participants at UC Berkeley, University of Illinois Urbana-Champaign, University of Nevada Reno, University of Maryland, Cornell, UC Santa Barbara, University of Utah, Georgia State University, University of Pennsylvania, Toulouse School of Economics, NBER EEE Summer Institute, the ASSA Meetings, the TREE Seminar, the Congressional Budget Office, and the Federal Energy Regulatory Commission. Derek Wietelman provided exceptional research assistance. Resources for the Future's Joseph L. Fisher Dissertation Fellowship provided generous support of this research. All remaining errors are mine. Online appendices for this article are available at http://www.louispreonas.com/s/preonas_coal_appxABCD.pdf and http://www.louispreonas.com/s/preonas_coal_appxEFG.pdf.

Supplementary Data

Supplementary data are available at *Review of Economic Studies* online.

Data Availability Statement

The data used throughout this analysis are all publicly available, with one exception: commodity prices and electricity supply curves maintained by S&P Global Market Intelligence. Researchers interested in accessing these proprietary data may purchase a subscription from S&P Global Market Intelligence. The replication archive (<https://doi.org/10.5281/zenodo.7920690>) describes how to download these data, along with other publicly available datasets that were too large to include in the archive.

REFERENCES

- ATKIN, D. and DONALDSON, D. (2015), "Who's Getting Globalized? The Size and Implications of Intra-national Trade Costs" (NBER Working Paper 21439).
- BARNETT, A. H. (1980), "The Pigouvian Tax Rule under Monopoly", *American Economic Review*, **70**, 1037–1041.
- BERGQUIST, L. F. and DINERSTEIN, M. (2020), "Competition and Entry in Agricultural Markets: Experimental Evidence from Kenya", *American Economic Review*, **110**, 3705–3747.
- BORENSTEIN, S., BUSHNELL, J. and STOFT, S. (2000), "The Competitiveness Effects of Transmission Capacity in a Deregulated Electricity Industry", *RAND Journal of Economics*, **31**, 294–325.
- BORENSTEIN, S., CAMERON, A. C. and GILBERT, R. (1997), "Do Gasoline Prices Respond Asymmetrically to Crude Oil Price Changes?", *Quarterly Journal of Economics*, **112**, 305–339.
- BORENSTEIN, S. and KELLOGG, R. (2014), "The Incidence of an Oil Glut: Who Benefits from Cheap Crude Oil in the Midwest?", *Energy Journal*, **35**, 15–33.
- BORENSTEIN, S. and KELLOGG, R. (2023), "Carbon Pricing, Clean Electricity Standards, and Clean Electricity Subsidies on the Path to Zero Emissions", in Deryugina, T., Kotchen, M. and Stock, J. H. (eds) *Environmental and Energy Policy and the Economy* (Vol. 4) (Chicago: University of Chicago Press) 125–176.
- BRANCACCIO, G., KALOUPSIDI, M. and PAPAGEORGIU, T. (2020), "Geography, Transportation, and Endogenous Trade Costs", *Econometrica*, **88**, 657–691.
- BREHM, P. (2019), "Natural Gas Prices, Electric Generation Investment, and Greenhouse Gas Emissions", *Resource and Energy Economics*, **58**, 101106.
- BUCHANAN, J. M. (1969), "External Diseconomies, Corrective Taxes, and Market Structure", *American Economic Review*, **59**, 174–177.
- BUSSE, M. R. and KEOHANE, N. O. (2007), "Market Effects of Environmental Regulation: Coal, Railroads, and the 1990 Clean Air Act", *RAND Journal of Economics*, **38**, 1159–1179.
- CAMERON, A. C., GELBACH, J. B. and MILLER, D. L. (2006), "Bootstrap-based Improvements for Inference with Clustered Errors" (Working Paper).
- CARL, J. and FEDOR, D. (2016), "Tracking Global Carbon Revenues: A Survey of Carbon Taxes versus Cap-And-Trade in the Real World", *Energy Policy*, **96**, 50–77.
- CHRISTIAN, P. and BARRETT, C. B. (2024), "Spurious Regressions and Panel IV Estimation: Revisiting the Causes of Conflict", *Economic Journal*, **134**, 1069–1099.

- CHU, Y., HOLLADAY, J. S. and LARIVIERE, J. (2017), "Opportunity Cost Pass-Through from Fossil Fuel Market Prices to Procurement Costs of the U.S. Power Producers", *Journal of Industrial Economics*, **65**, 842–871.
- CICALA, S. (2015), "When Does Regulation Distort Costs? Lessons from Fuel Procurement in U.S. Electricity Generation", *American Economic Review*, **105**, 411–444.
- CICALA, S. (2022), "Imperfect Markets versus Imperfect Regulation in U.S. Electricity Generation", *American Economic Review*, **112**, 409–441.
- COURNOT, A. (1838), "Of the Influence of Taxation on Commodities Produced under a Monopoly", in *Researches into the Mathematical Principles of the Theory of Wealth*, translated by N. T. Bacon (London: Macmillan, 1897) 67–78.
- COVERT, T. R. and KELLOGG, R. (2018), "Crude by Rail, Option Value, and Pipeline Investment" (NBER Working Paper 23855).
- CULLEN, J. A. and MANSUR, E. T. (2017), "Inferring Carbon Abatement Costs in Electricity Markets: A Revealed Preference Approach Using the Shale Revolution", *American Economic Journal: Economic Policy*, **9**, 106–133.
- DAGOUMAS, A. S. and POLEMIS, M. L. (2020), "Carbon Pass-Through in the Electricity Sector: An Econometric Analysis", *Energy Economics*, **86**, 104621.
- DAVIS, L. and HAUSMAN, C. (2016), "Market Impacts of a Nuclear Power Plant Closure", *American Economic Journal: Applied Economics*, **8**, 92–122.
- DAVIS, R. J., HOLLADAY, J. S. and SIMS, C. (2022), "Coal-Fired Power Plant Retirements in the United States", in Deryugina, T., Kotchen, M. and Stock, J. H. (eds) *Environmental and Energy Policy and the Economy* (Vol. 3) (Chicago: University of Chicago Press) 4–36.
- DAVIS, L. W. and MUEHLEGGGER, E. (2010), "Do Americans Consume Too Little Natural Gas? An Empirical Test of Marginal Cost Pricing", *RAND Journal of Economics*, **41**, 791–810.
- EFRON, B. (1987), "Better Bootstrap Confidence Intervals", *Journal of the American Statistical Association*, **82**, 171–185.
- ELLIS, G. M. and HALVORSEN, R. (2002), "Estimation of Market Power in a Nonrenewable Resource Industry", *Journal of Political Economy*, **110**, 883–899.
- FABRA, N. and REGUANT, M. (2014), "Pass-Through of Emissions Costs in Electricity Markets", *American Economic Review*, **104**, 2872–2899.
- FELL, H. and KAFFINE, D. T. (2018), "The Fall of Coal: Joint Effects of Fuel Prices and Renewables on Generation and Emissions", *American Economic Journal: Economic Policy*, **10**, 90–116.
- FITZGERALD, T. (2013), "Frackonomics: Some Economics of Hydraulic Fracturing", *Case Western Reserve Law Review*, **63**, 1337–1362.
- FOWLIE, M., REGUANT, M. and RYAN, S. P. (2016), "Market-Based Emissions Regulation and Industry Dynamics", *Journal of Political Economy*, **124**, 249–302.
- GANAPATI, S., SHAPIRO, J. S. and WALKER, R. (2020), "Energy Cost Pass-Through in U.S. Manufacturing: Estimates and Implications for Carbon Taxes", *American Economic Journal: Applied Economics*, **12**, 303–42.
- GOLDBERG, P. K. and HELLERSTEIN, R. (2013), "A Structural Approach to Identifying the Sources of Local Currency Price Stability", *Review of Economic Studies*, **80**, 175–210.
- GOODKIND, A. L., TESSUM, C. W., COGGINS, J. S., *et al.* (2019), "Fine-Scale Damage Estimates of Particulate Matter Air Pollution Reveal Opportunities for Location-Specific Mitigation of Emissions", *Proceedings of the National Academy of Sciences*, **116**, 8775–8780.
- GOULDER, L. H. (2020), "Timing Is Everything: How Economists Can Better Address the Urgency of Stronger Climate Policy", *Review of Environmental Economics and Policy*, **14**, 143–156.
- GOULDER, L. H., HAFSTEAD, M. A. C., KIM, G., *et al.* (2019), "Impacts of a Carbon Tax across U.S. Household Income Groups: What Are the Equity-Efficiency Trade-Offs?", *Journal of Public Economics*, **175**, 44–64.
- HAUSMAN, C. and KELLOGG, R. (2015), "Welfare and Distributional Implications of Shale Gas", *Brookings Papers on Economic Activity*, **Spring**, 71–125.
- HECKMAN, J. J., ICHIMURA, H. and TODD, P. E. (1997), "Matching as an Econometric Evaluation Estimator: Evidence from Evaluating a Job Training Programme", *Review of Economic Studies*, **64**, 605–654.
- HINTERMANN, B. (2016), "Pass-Through of CO₂ Emission Costs to Hourly Electricity Prices in Germany", *Journal of the Association of Environmental and Resource Economists*, **3**, 857–891.
- HOLLADAY, J. S. and LARIVIERE, J. (2017), "The Impact of Cheap Natural Gas on Marginal Emissions from Electricity Generation and Implications for Energy Policy", *Journal of Environmental Economics and Management*, **85**, 205–227.
- HORTAÇSU, A. and SYVERSON, C. (2007), "Cementing Relationships: Vertical Integration, Foreclosure, Productivity, and Prices", *Journal of Political Economy*, **115**, 250–301.
- HUGHES, J. E. (2011), "The Higher Price of Cleaner Fuels: Market Power in the Rail Transport of Fuel Ethanol", *Journal of Environmental Economics and Management*, **62**, 123–139.
- HUGHES, J. E. and LANGE, I. (2020), "Who (Else) Benefits from Electricity Deregulation? Coal Prices, Natural Gas, and Price Discrimination", *Economic Inquiry*, **58**, 1053–1075.
- Interagency Working Group on Social Cost of Greenhouse Gases (2016), "Technical Update of the Social Cost of Carbon for Regulatory Impact Analysis", United States Government, Technical Support Document.
- JHA, A. (2020), "Dynamic Regulatory Distortions: Coal Procurement at U.S. Power Plants" (Working Paper).
- JHA, A. (2022), "Regulatory Induced Risk Aversion in Coal Contracting at U.S. Power Plants: Implications for Environmental Policy", *Journal of the Association of Environmental and Resource Economists*, **9**, 51–78.

- JOSKOW, P. L. (1987), "Contract Duration and Relationship-Specific Investments: Empirical Evidence from Coal Markets", *American Economic Review*, **77**, 168–185.
- JOSKOW, P. L. (1988), "Price Adjustment in Long-Term Contracts: The Case of Coal", *Journal of Law and Economics*, **31**, 47–83.
- KELLOGG, R. and REGUANT, M. (2021), "Energy and Environmental Markets, Industrial Organization, and Regulation", in Ho, K., Hortaçsu, A. and Lizzeri, A. (eds) *Handbook of Industrial Organization* (Vol. 5) (Amsterdam: Elsevier B.V.) 615–742.
- KERKVLIIET, J. and SHOGREN, J. F. (1992), "The Impacts of Environmental Regulation on Coal Procurement Strategies: Design Coal and Multi-Attributed Quality", *Journal of Environmental Management*, **35**, 83–91.
- KIM, H. (2022), "Heterogeneous Impacts of Cost Shocks, Strategic Bidding, and Pass-Through: Evidence from the New England Electricity Market", *American Economic Journal: Microeconomics*, **14**, 370–407.
- KIM, W., CHATTOPADHYAY, D. and PARK, J.-B. (2010), "Impact of Carbon Cost on Wholesale Electricity Price: A Note on Price Pass-Through Issues", *Energy*, **35**, 3441–3448.
- KNITTEL, C. R., MEISELMAN, B. S. and STOCK, J. H. (2017), "The Pass-Through of RIN Prices to Wholesale and Retail Fuels under the Renewable Fuel Standard", *Journal of the Association of Environmental and Resource Economists*, **4**, 1081–1119.
- KNITTEL, C. R., METAXOGLU, K. and TRINDADE, A. (2019), "Environmental Implications of Market Structure: Shale Gas and Electricity Markets", *International Journal of Industrial Organization*, **63**, 511–550.
- KOSNIK, L. and LANGE, I. (2011), "Contract Renegotiation and Rent Re-Distribution: Who Gets Raked over the Coals?", *Journal of Environmental Economics and Management*, **62**, 155–165.
- LADE, G. E. and BUSHNELL, J. (2019), "Fuel Subsidy Pass-Through and Market Structure: Evidence from the Renewable Fuel Standard", *Journal of the Association of Environmental and Resource Economists*, **6**, 563–592.
- LESLIE, G. (2018), "Tax Induced Emissions? Estimating Short-Run Emissions Impacts from Carbon Taxation under Different Market Structures", *Journal of Public Economics*, **167**, 220–239.
- LI, J. and STOCK, J. H. (2019), "Cost Pass-Through to Higher Ethanol Blends at the Pump: Evidence from Minnesota Gas Station Data", *Journal of Environmental Economics and Management*, **93**, 1–19.
- LINN, J., MASTRANGELO, E. and BURTRAW, D. (2014), "Regulating Greenhouse Gases from Coal Power Plants under the Clean Air Act", *Journal of the Association of Environmental and Resource Economists*, **1**, 97–134.
- LINN, J. and MUEHLENBACHS, L. (2018), "The Heterogeneous Impacts of Low Natural Gas Prices on Consumers and the Environment", *Journal of Environmental Economics and Management*, **89**, 1–28.
- MACDONALD, J. M. (1987), "Competition and Rail Rates for the Shipment of Corn, Soybeans, and Wheat", *RAND Journal of Economics*, **18**, 151–163.
- MACDONALD, J. M. (1989), "Railroad Deregulation, Innovation, and Competition: Effects of the Staggers Act on Grain Transportation", *Journal of Law and Economics*, **32**, 63–95.
- MACDONALD, J. M. (2013), "Railroads and Price Discrimination: The Roles of Competition, Information, and Regulation", *Review of Industrial Organization*, **43**, 85–101.
- MANSUR, E. T. (2007), "Do Oligopolists Pollute Less? Evidence from a Restructured Electricity Market", *Journal of Industrial Economics*, **55**, 661–689.
- MANSUR, E. T. (2008), "Measuring Welfare in Restructured Electricity Markets", *Review of Economics and Statistics*, **90**, 369–386.
- MARION, J. and MUEHLEGGGER, E. (2011), "Fuel Tax Incidence and Supply Conditions", *Journal of Public Economics*, **95**, 1202–1212.
- METCALF, G. E. (2021), "Carbon Taxes in Theory and Practice", *Annual Review of Resource Economics*, **13**, 245–265.
- MILLER, N. H. and OSBORNE, M. (2014), "Spatial Differentiation and Price Discrimination in the Cement Industry: Evidence from a Structural Model", *RAND Journal of Economics*, **45**, 221–247.
- MONTGOMERY, W. D. (1972), "Markets in Licenses and Efficient Pollution Control Programs", *Journal of Economic Theory*, **5**, 395–418.
- MORRIS, A. C., KAUFMAN, N. and DOSHI, S. (2021), "Revenue at Risk in Coal-Reliant Counties", in Kotchen, M., Stock, J. H. and Wolfram, C. D. (eds) *Environmental and Energy Policy and the Economy* (Vol. 2) (Chicago: University of Chicago Press) 83–116.
- MUEHLEGGGER, E. and SWEENEY, R. L. (2022), "Pass-Through of Own and Rival Cost Shocks: Evidence from the U.S. Fracking Boom", *Review of Economics and Statistics*, **104**, 1361–1369.
- MULLER, N. Z., MENDELSON, R. and NORDHAUS, W. (2011), "Environmental Accounting for Pollution in the United States Economy", *American Economic Review*, **101**, 1649–1675.
- NAKAMURA, E. and ZEROM, D. (2010), "Accounting for Incomplete Pass-Through", *Review of Economic Studies*, **77**, 1192–1230.
- NAZIFI, F. (2016), "The Pass-Through Rates of Carbon Costs on to Electricity Prices Within the Australian National Electricity Market", *Environmental Economics and Policy Studies*, **18**, 41–62.
- NORDHAUS, W. D. (1993), "Reflections on the Economics of Climate Change", *Journal of Economic Perspectives*, **7**, 11–25.
- OATES, W. E. and STRASSMANN, D. L. (1984), "Effluent Fees and Market Structure", *Journal of Public Economics*, **24**, 29–46.
- PIGOU, A. C. (1932), *The Economics of Welfare, 1920* (London: Macmillan).
- POULIOT, S., SMITH, A. and STOCK, J. H. (2020), "RIN Pass-Through at Gasoline Terminals" (Working Paper).

- PRATER, M., SPARGER, A. and O'NEIL JR., D. (2014), "Railroad Concentration, Market Shares, and Rates" (Technical Report), U.S. Department of Agriculture, Agricultural Marketing Service.
- RENNERT, K., ERRICKSON, F., PREST, B. C., *et al.* (2022), "Comprehensive Evidence Implies a Higher Social Cost of CO₂", *Nature*, **610**, 687–692.
- RYAN, S. P. (2012), "The Costs of Environmental Regulation in a Concentrated Industry", *Econometrica*, **80**, 1019–1061.
- SALLEE, J. M. (2019), "Pigou Creates Losers: On the Implausibility of Achieving Pareto Improvements from Efficiency-Enhancing Policies" (NBER Working Paper 25831).
- SCHMALENSEE, R. and STAVINS, R. N. (2013), "The SO₂ Allowance Trading System: The Ironic History of a Grand Policy Experiment", *Journal of Economic Perspectives*, **27**, 103–122.
- SCHMIDT, S. (2001), "Market Structure and Market Outcomes in Deregulated Rail Freight Markets", *International Journal of Industrial Organization*, **19**, 99–131.
- STARTZ, M. (2021), "The Value of Face-to-Face: Search and Contracting Problems in Nigerian Trade" (Working Paper).
- STOLPER, S. (2021), "Local Pass-Through and the Regressivity of Taxes: Evidence from Automotive Fuel Markets" (Working Paper).
- U.S. Environmental Protection Agency (2008). "Optimal Emissions from Commuting, Business Travel and Product Transport" (Technical Report), PA430-R-08-006.
- U.S. Environmental Protection Agency (2022). "Report on the Social Cost of Greenhouse Gases: Estimates Incorporating Recent Scientific Advances" EPA External Review Draft of Report on the Social Cost of Greenhouse Gases: Estimates Incorporating Recent Scientific Advances Docket ID No. EPA-HQ-OAR-2021-0317.
- WETZSTEIN, B., FLORAX, R., FOSTER, K., *et al.* (2021), "Transportation Costs: Mississippi River Barge Rates", *Journal of Commodity Markets*, **21**, 100123.
- WEYL, E. G. and FABINGER, M. (2013), "Pass-Through as an Economic Tool: Principles of Incidence under Imperfect Competition", *Journal of Political Economy*, **121**, 528–583.
- WILSON, W. W. (1996), "Legislated Market Dominance in Railroad Markets", *Research in Transportation Economics*, **4**, 49–67.
- WOLAK, F. A. (1996), "Why Do Firms Simultaneously Purchase in Spot and Contract Markets? Evidence from the United States Steam Coal Market" in Martimort, D. (ed) *Agricultural Markets: Mechanisms, Failures, and Regulations* (Amsterdam: Elsevier Science) 109–168.
- WOOLDRIDGE, J. M. (2007), "Inverse Probability Weighted Estimation for General Missing Data Problems", *Journal of Econometrics*, **141**, 1281–1301.