Appendices A–D for: Market Power in Coal Shipping and Implications for U.S. Climate Policy

Louis Preonas

February 6, 2023

Appendices E–G available at: https://www.louispreonas.com/s/preonas_coal_appxEFG.pdf

A Coal demand estimation

A.1 Demand estimation algorithm

I start by estimating a semi-parametric model of each coal generating unit's hourly operations, conditional on the ratio of marginal costs of coal vs. natural gas generation. Next, I construct a distribution of counterfactual coal prices at which each unit would have been marginal in electricity dispatch. Then, I transform and aggregate these distributions into quantity-price mappings, yielding plant-month-specific coal demand curves. Finally, I estimate how changes in natural gas prices affect both the level and the slope of each plant's inverse coal demand.

Step 1: I construct a coal-to-gas cost ratio by dividing each coal unit's marginal cost of generation by the generation-weighted average marginal cost of gas-fired units in its power control area (PCA). For both coal and gas units, I multiply unit-specific fuel prices (P for coal, Z for gas) by unit-specific heat rates (HR), and add the unit's marginal costs of environmental compliance (MC^{env}).[1] I assign each gas unit the daily spot price of its closest natural gas trading hub, which captures day-to-day price fluctuations.[2] I use the average delivered coal price at the plant-month level, which is the finest temporal resolution that EIA reports publicly.[3] I also assign unit-specific heat rates and environmental costs at the monthly level.

Indexing coal plants j, their constituent coal units u, gas units g, months m, and days d, the daily cost ratio is:

$$MC_{um}^{coal} \equiv HR_{um} \cdot (P_{jm} + MC_{um}^{env}) \tag{A1}$$

$$MC_{ud}^{gas} \equiv \sum_{g \in \mathrm{PCA}_u} \left(\frac{Q_{gm}^{elec} \cdot HR_{gm} \cdot (Z_{gd} + MC_{gm}^{env})}{\sum_{g \in \mathrm{PCA}_u} Q_{gm}^{elec}} \right) \tag{A2}$$

1. MC_{um}^{env} captures unit u's opportunity cost of SO_2, NO_x, and CO_2 emissions in month m, scaled by average monthly allowance prices (A_m^z) and unit-specific emissions rates per MMBTU (E_{um}^z), for each pollutant z. It also includes the non-energy operating costs of scrubbers (i.e. flue gas desulfurization to reduce SO_2 emissions), net of the marginal revenues from selling the gypsum byproduct of the desulfurization process: $MC_{um}^{env} = \sum_{z \in \{SO_2, NO_x, CO_2\}} A_m^z E_{um}^z \cdot \mathbf{1}[\text{in } z \text{ trading program}]_{um} + MC_{um}^{scrubber}$. I abstract from non-fuel, non-environmental variable operating costs, because these data are not reliable across all years and units.

2. Most gas plants have limited on-site storage capacity, meaning that short-run price changes can impact operating decisions. I use daily natural gas prices from SNL for 104 trading hubs (see Appendix G.5).

3. Coal's cheap storability let plants buffer sub-monthly price fluctuations. The relevant coal price for this application is plant j's opportunity cost of coal purchases. While other studies have characterized this opportunity cost using spot market purchases only (e.g. Cicala (2022); Chu, Holladay, and LaRiviere (2017)), I average P_{jm} across contract and spot transactions. Many plants buy exclusively on long-term contracts, and restricting P_{jm} to spot shipments would require populating many unobserved spot prices. Given that my analysis hinges on *plant-specific* delivered coal prices, I choose to pool all observed coal prices into an average P_{jm} for each plant-month. I report sensitivity analysis using alternative coal price variables in Figures A6 and E10.

$$\Rightarrow \quad CR_{ud} = \frac{MC^{coal}_{um}}{MC^{gas}_{ud}} \tag{A3}$$

Steps 1–3 of this estimation strategy treat the coal *unit* as the relevant unit of analysis, rather than the coal plant. This is because a single power plant may comprise multiple coal units, each with distinct heat rates, environmental costs, and operating protocols.

Step 2: For each coal unit u, I estimate Equation (6) (reproduced here) for all hours h, from 2002 to 2015:

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

Figure A1 shows that CF_{uh} is close to discrete: 77% of observations are under 5% or over 80%.

Figure A1: Histogram of hourly capacity factors for coal demand estimation

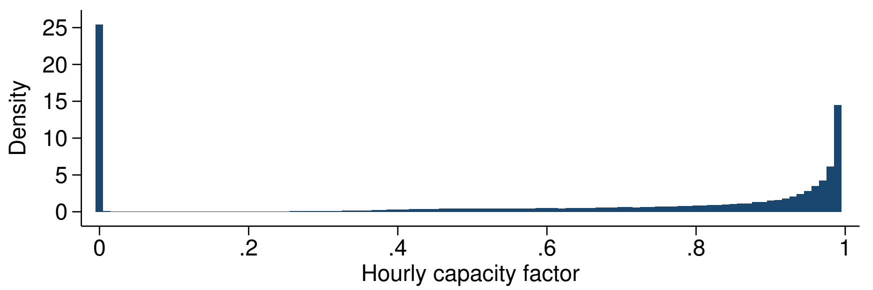

Notes: This is a histogram of coal unit capacity factors in hourly CEMS data (i.e. CF_{uh}), for all units and hours used in my demand estimation. $CF_{uh} \in [0, 1]$ by construction. 90% of non-zero capacity factors are greater the 0.5, which supports my use of $\widehat{CF}_{uh} = 0.5$ in Equations (A4)–(A5). In Figures A7 and E10, I report sensitivity analysis estimating Equation (6) as a probit and logit—using the outcome variable $\mathbf{1}[CF_{uh} > 0.5]$.

G_{uh} sums hourly net generation across all CEMS electric generating units in unit u's market region.[4] This is not equivalent to electricity "load", which includes non-CEMS generation such as nuclear, hydro, renewables. However, these other technologies are inframarginal; the marginal operating unit is almost always a CEMS unit.[5] Generation bins b allow me to flexibly estimate unit u's generation, both un-interacted ($\hat{\alpha}_{ub}$, following Davis and Hausman (2016)) and interacted with the cost ratio ($\hat{\gamma}_{ub}$). Because electricity demand is nearly perfectly inelastic, Equation (6) is unlikely to suffer from simultaneity bias between CF_{uh} and G_{uh}.

We might worry about endogeneity of P_{jm} as a component of CF_{uh}, if rail market power leads to idiosyncratic plant-specific coal price changes. I instrument for CR_{ud} with an analogous

4. I aggregate across "market regions" (i.e. ISOs or NERC regions) rather than PCAs because of inter-PCA trading: in a given hour, the marginal CEMS unit is less likely to reside u's PCA than in u's broader market region. (There is much less trading across market regions.) Appendix G.2.1 describes these market regions. I follow Linn and Muehlenbachs (2018) in defining average marginal costs at the PCA-level, because within-PCA comparisons more accurately exploit cross-sectional gas price differences driven by local pipeline constraints.

5. Nuclear, wind, and solar are (virtually) always inframarginal in hours when a CEMS unit is also operating (i.e. $G_{uh} > 0$). Hydro plants have complex dynamic operating constraints, and in certain regions hydro can be marginal in electricity dispatch. However, these regions have very little coal: from 2002–2015, six states (WA, OR, CA, NY, MT, ID) contributed 70% of U.S. hydro generation and only 2.5% of coal generation.

cost ratio CR^s_{ud} using state-level coal prices (i.e., replacing P_{jm} in Equation (A1) with the state-month average delivered coal price). This removes plant-specific price effects, since over 80% of plant-months average coal prices across at least 10 plants.[6] $\mathbf{G}_{uh}$ also includes year fixed effects, which absorbs any (potentially endogenous) trends in coal prices within unit u's time series.

$\mathbf{G}_{uh}$ includes several time-varying factors that affect unit u's probability of operating conditional on G_{uh} and CR_{ud}: daily G_{uh} controls, to account for dynamic operating constraints (coal boilers cannot instantaneously start or stop (Cullen and Mansur (2017))); daily maximum temperature, since outdoor temperature affects units' thermal efficiency; hour-of-day fixed effects, to control for diurnal operating patterns; quarter-of-year fixed effects, to control for seasonality in output demand, fuel prices, and maintenance; and year fixed effects, to capture long-run changes in unit u's operations. Finally, $\mathbf{G}_{uh}$ interacts year fixed effects with the daily sum of G_{uh}, to control for changes in generating capacity and unit u's position in the dispatch order.[7]

I do not conduct inference directly on the coefficient estimates or predictions from Equation (6). Instead, I use predictions from the fitted models to construct the dependent variables of OLS regressions that estimate coal demand parameters.

Step 3: Let $\tilde{P}_{uh}$ denote the coal price at which unit u would have had a 50% probability of operating at full capacity in hour h, or $\widehat{CF}_{uh} = 0.5$.[8] Rearranging Equation (6) post-estimation:

$$HR_{um} \cdot \left(\tilde{P}_{uh} + MC^{env}_{um}\right) = \frac{\left(0.5 - \hat{\alpha}_{ub,h} - \hat{\xi}_u \mathbf{G}_{uh}\right) \cdot MC^{gas}_{ud}}{\hat{\gamma}_{ub,h} + \hat{\delta}_u} \tag{A4}$$

$$\Rightarrow \quad \tilde{P}_{uh} = \left[\frac{\left(0.5 - \hat{\alpha}_{ub,h} - \hat{\xi}_u \mathbf{G}_{uh}\right) \cdot MC^{gas}_{ud}}{\hat{\gamma}_{ub,h} + \hat{\zeta}_u}\right] \Big/ HR_{um} - MC^{env}_{um} \tag{A5}$$

Here $\hat{\alpha}_{ub,h}$ and $\hat{\gamma}_{ub,h}$ denote the binned coefficients for each hourly realization of G_{uh}.

Step 4: Summing over all hours of month m, and across each of plant j's constituent coal units, I define plant j's monthly coal demand function $\widehat{Q}_{jm}$ as:

$$\widehat{Q}_{jm}(P) = \sum_{u \in j} \sum_{h \in m} \mathbf{1}\left[P < \tilde{P}_{uh}\right] \cdot \bar{Q}^{coal}_{um} \tag{A6}$$

where $\bar{Q}^{coal}_{um}$ is unit u's hourly coal consumption when operating at maximum capacity in month m. This simply assumes that for a given coal price P, plant j will demand the amount of coal required to operate each of its units at capacity, in inframarginal hours only. $\widehat{Q}_{jm}(\,\cdot\,)$ is invertible by construction; I define its inverse as $\widehat{P}_{jm}(\,\cdot\,)$, which I smooth using a kernel mean-smoothing algorithm. This gives me an estimate of plant j's inverse coal demand curve in month m. Figure A2 plots two $\widehat{P}_{jm}(\,\cdot\,)$ curves, for a representative plant j in two months m.

6. My results are also robust to a reduced-form version of Equation (A1), replacing CR_{ud} with either CR^s_{ud} or an version of CR^s_{ud} that averages delivered coal prices across ISOs/NERC regions (see Figures A5 and E10).

7. For example, the generation bin $G_{uh} \in [10, 20)$ GWh could imply dramatically different electricity prices in 2005 vs. 2015—and hence, dramatically different probabilities of whether unit u decides to operate.

8. Discretizing the counterfactual capacity factor at 0.5 is computationally much simpler than the alternative of solving for a schedule of counterfactual prices (e.g., a separate $\tilde{P}_{uh}$ for $\widehat{CF} \in \{0, 0.1, \ldots, 0.9, 1\}$). Using $\widehat{CF}_{uh} = 0.5$ is sufficient to characterize how higher coal prices make unit u less likely to operate in any hour h.

Figure A2: Example of estimated coal demand curves and the effect of gas prices

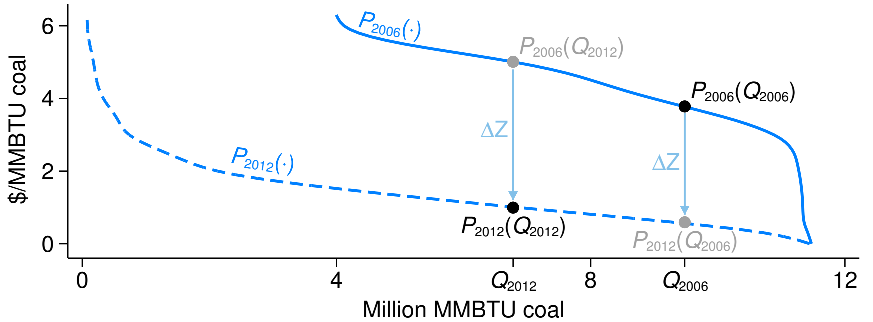

Notes: This figure plots two estimated inverse demand curves for a representative coal plant j. The solid curve is its estimated demand for May 2006, when the gas price was high; the dashed curve is its estimated demand for May 2012, when the gas price was low. Black points are estimated quantities of observed coal consumption. To be precise: they are not $\widehat{P}_{jm}(Q_{jm})$, but rather $\widehat{P}_{jm}(\widehat{Q}_{jm}(P_{jm}))$; this accounts for estimation error in $\widehat{P}_{jm}(\,\cdot\,)$ and forces each point to fall on its respective curve. Grey points plug in the opposite quantity (technically $\widehat{Q}_{j\ddot{m}}(P_{j\ddot{m}})$ not $Q_{j\ddot{m}}$) into each demand curve. Equations (7)–(8) populate all m-$\ddot{m}$ combinations (including $m = \ddot{m}$) and regress the level/slope of inverse demand at these points on the gas price in month m (embedded in each estimated demand curve) and fixed effects to control for the estimated quantity in month $\ddot{m}$.

Step 5: I calculate local approximations of the first and second derivatives of $\widehat{P}_{jm}(\,\cdot\,)$, which I denote as $\Delta\widehat{P}_{jm}(\,\cdot\,)$ and $\Delta^2\widehat{P}_{jm}(\,\cdot\,)$, respectively. This lets me estimate empirical analogs of three components of Equation (3): $\frac{\partial P_j}{\partial Z}$, $\frac{\partial^2 P_j}{\partial Q_j \partial Z}Q_j$, E_{D_j}. Since these terms are *partial* derivatives, I estimate their empirical analogs based on *realized* coal quantities for each plant-month.

I define $\ddot{m}$ as the month of each realized coal quantity, and m as the month of each estimated inverse demand curve. This lets me plug each estimated coal quantity into each estimated demand curve, created an m-$\ddot{m}$ "panel" for each plant j.[9] This forms the dependent variable of two plant-specific OLS regressions (reproducing Equations (7)–(8)):[10]

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

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

The monthly Henry Hub price Z_m is closely correlated with Z_{gd} in Equation (A2), which enters into my $\tilde{P}_{uh}$ predictions.[11] $\phi_{j\ddot{m}}$ are month $\ddot{m}$ fixed effects, which control for endogenous factors affecting plant j's coal quantity and isolate variation in the m-dimension. $\hat{\lambda}_{0j}$ and $\hat{\lambda}_{1j}$ estimate how variation in the natural gas price affects the level and slope of plant j's estimated inverse coal demand, at each observed quantity. Figure A2 illustrates how time-series variation in the height and slope of $\widehat{P}_{jm}(\,\cdot\,)$—driven by changing gas prices—identifies λ_{0j} and λ_{1j}.

Other factors also contributed to time series variation in plant j's coal demand, such as more stringent regulations on criteria pollutants or increased renewables generation. However, the non-monotonic time profile of gas price changes is not correlated with more gradual trends in pollution regulations and renewables penetration. Moreover, the dependent variables in Equations (7)–(8) come from a model that holds the electricity market and environmental

9. For a plant that consumed coal in all 168 sample months, this would yield 168^2 m-$\ddot{m}$ "observations".

10. To be precise: I plug in $\widehat{Q}_{j\ddot{m}}(P_{j\ddot{m}})$, the estimated quantity implied by month $\ddot{m}$'s observed coal price. I plug actual price instead of actual quantity since $\widehat{Q}_{jm}(P_{jm})$ predicts Q_{jm} better than $\widehat{P}_{jm}(Q_{jm})$ predicts P_{jm}.

11. I construct CR_{ud} using more geographically explicit natural gas prices, in order to maximize the predictive power of Equation (6). However, here I use the Henry Hub price time series which are (more) plausibly exogenous (absent time fixed effects), while also aligning with my DD specification in Equation (5).

Figure A3: Mapping coal plants by $\hat{M}_j$

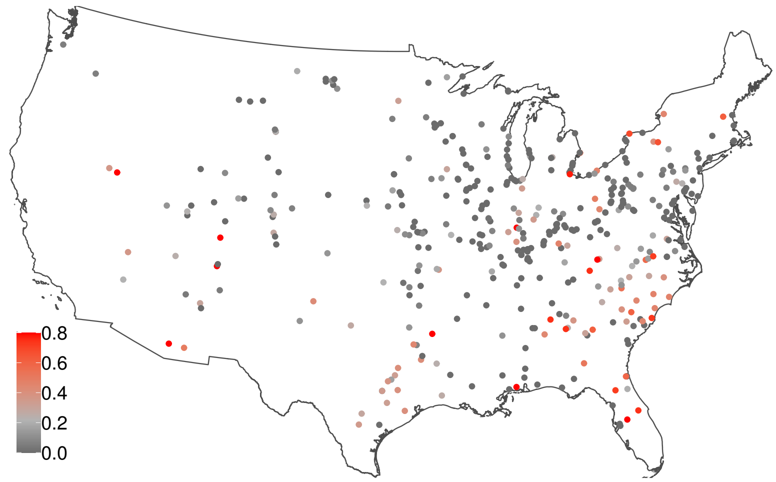

Notes: This map plots 428 coal plants, color-coded by their value of $\hat{M}_j$ (as defined in Equation (10)). I top-code (bottom-code) the color scale at 0.8 (0.0) for ease of presentation.

regulations fixed (i.e., Equation (6)). This means that my predicted $\widehat{P}_{jm}(\,\cdot\,)$ should not vary systematically with these other time-series factors—the way that true $P_{jm}(\,\cdot\,)$ would.[12]

Finally, I estimate the time series Equation (9) (reproduced here) for each plant j:

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

Here, the partial derivative of interest, E_{D_j}, does not relate to gas prices, obviating the need for a second dimension $\ddot{m}$. $\hat{\lambda}_{2j}$ captures the average elasticity of the slope of plant j's estimated inverse coal demand over all months, at the realized (estimated) quantities. Identification of $\hat{\lambda}_{2j}$ follows from exogeneity of the cost ratio in Equation (6).

The standard errors from Equations (7)–(9) let me simulate draws of $\langle\hat{\lambda}_{0j}, \hat{\lambda}_{1j}, \hat{\lambda}_{2j}\rangle$, to account for generated regressors on the right-hand side of Equation (5) and incorporate uncertainty from this demand estimation algorithm. Appendix E.5 outlines this bootstrapping procedure. Since estimation error only enters Equations (7)–(9) on the left-hand side, the standard errors on $\hat{\lambda}_{0j}$, $\hat{\lambda}_{1j}$, and $\hat{\lambda}_{2j}$ are likely unbiased if this measurement error is classical.

This procedure recovers counterfactual coal prices ($\tilde{P}_{uh}$) that hold the electricity market constant, including prices faced by other coal plants. In reality, plant-specific markups make up only a small portion of delivered coal prices, and large changes to plant j's coal price (e.g. due to a global coal price shock, or a regional diesel shortage) likely affect many plants simultaneously. This means that my demand estimates are likely only informative for small idiosyncratic changes in plant-specific coal prices. If rail carriers *jointly* reoptimize markups across multiple plants (i.e. plant j's markups move in the same direction as the markups of rival coal plants), then my estimated demand functions $\widehat{Q}_{jm}(\cdot)$ may be too large (small) at low (high) coal prices.

Figure A3 maps coal plants by $\hat{M}_j$: while high $\hat{M}_j$ plants are more concentrated in the southeastern U.S., they are not unique to any particular region of the country. Figure A4 summarizes uncertainty in my plant-specific $\hat{M}_j$ estimates. I plot $\hat{M}_j$, or the mean of the $\hat{M}_j^{(S)}$

12. Even if Equations (7)–(8) do not perfectly identify the effect of changes in gas price, I ultimately feed $\hat{\lambda}_{0j}$ and $\hat{\lambda}_{1j}$ into Equation (5), where DD identification purges these other confounding trends.

Figure A4: Summarizing the precision of my $\hat{M}_j$ estimates

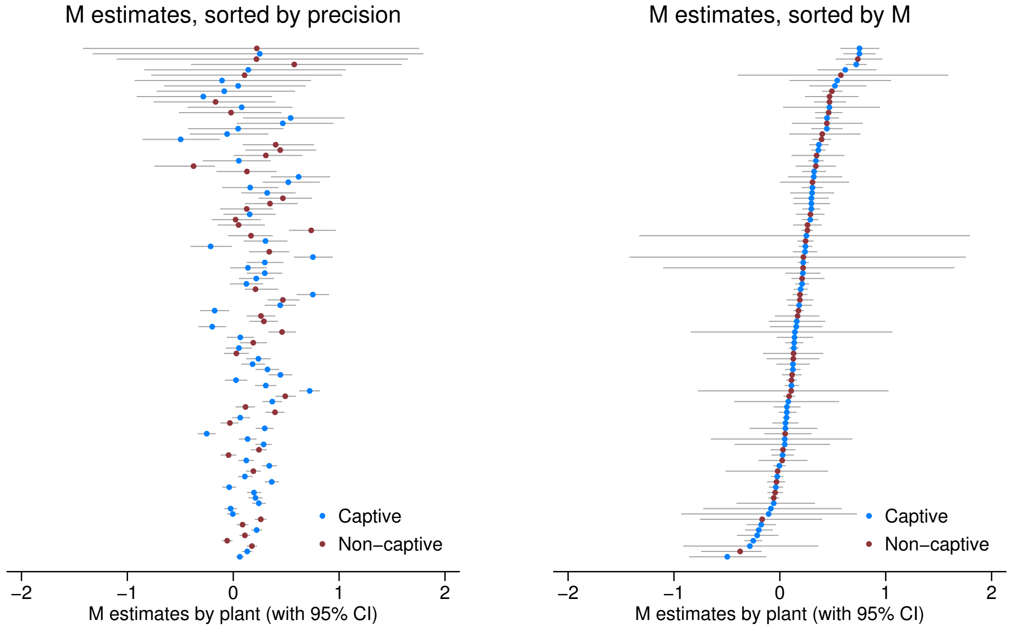

Notes: Each panel plots $\hat{M}_j$ for a separate coal plant in my main estimation sample (with $k = 3$ nearest neighbors). Here I present the same $\hat{M}_j$ estimates as in the bottom-right panel of Figure 5, except whiskers now plot the 5th and 95th percentiles of each $\hat{M}_j$'s sampling distribution. The left panel sorts plants by precision: for 42% of plants, the interval spanning the 5th-to-95th percentiles is less than 0.25; for 64% of plants, this interval does not include zero. The right panel reports identical information, but sorts plants by $\hat{M}_j$. Both panels omit plants with a coal-by-barge option, for which $\hat{M}_j = 0$.

sampling distribution (see Equation (E8)); whiskers plot the 5th and 95th percentiles of each $\hat{M}_j^{(S)}$ sampling distribution. The left panel sorts plants from least to most precise $\hat{M}_j$; the right panel presents identical information, but sorts by $\hat{M}_j$. Both panels omit plants with $W_j = 1$, for which I set $\hat{M}_j = 0$. This figure highlights three key patterns. First, captive and non-captive plants appear equally likely to have precise or imprecise $\hat{M}_j$, and to have low or high $\hat{M}_j$. Second, $\hat{M}_j$ estimates tend to be reasonably precise: for 42% of these plants, the interval spanning the 5th and 95th percentiles of its sampling distribution is narrower than 0.25. Third, $\hat{M}_j$'s sign tends to be unambiguous: for 64% of these plants, the interval spanning the 5th and 95th percentiles of its sampling distribution does not include zero.[13]

A.2 Sensitivity analysis

Here, I compare my preferred $\hat{M}_j$ estimates with those derived from alternative versions of Equation (6). Figure A5 shows robustness to my decision to instrument CR_{ud} (using plant-specific coal prices) with CR^s_{ud} (using state-average coal prices). The top-left panel estimates the OLS version of Equation (6), without purging any plant-specific coal price endogeneity (e.g., plant-specific markups); the resulting $\hat{M}_j$ estimates are not systematically higher or lower. The top-right panel estimates the reduced-form of Equation (6), replacing CR_{ud} with CR^s_{ud}; this yields $\hat{M}_j$ estimates that are similar to my preferred IV, signalling that most of the state-level variation in coal prices predicts plant-specific coal prices. The bottom-left panel averages coal

13. Each of these patterns would appear even more pronounced if I included $W_j = 1$ plants in the denominator: these plants are not systematically captive or non-captive, and their $\hat{M}_j = 0$ is not uncertain (by construction).

Figure A5: Alternative coal price aggregations, without instrumenting for coal price

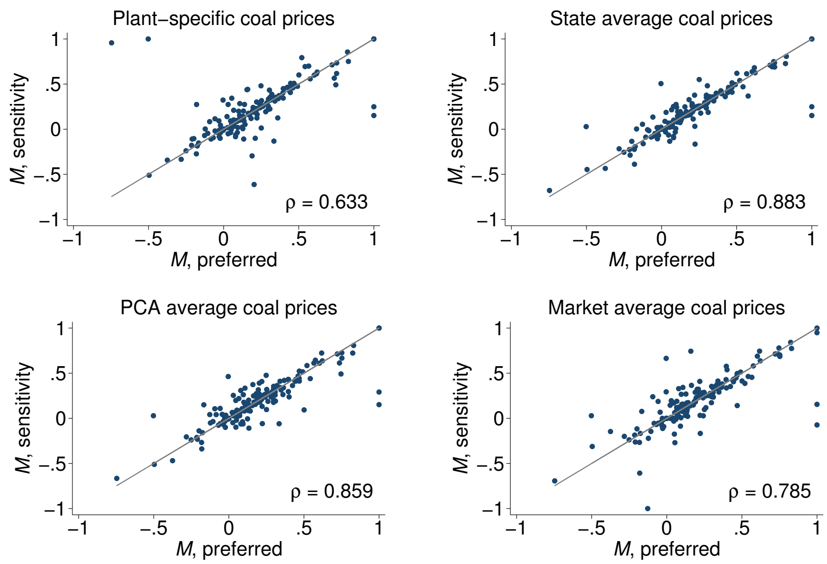

Notes: Each scatter plot compares an alternative $\hat{M}_j$ (vertical axes) to my preferred $\hat{M}_j$ (horizontal axes), which uses predictions from Equation (6) that instrument for CR_{ud} using CR^s_{ud}. The top-left panel constructs $\hat{M}_j$ without instrumenting for CR_{ud}, where plant-specific coal prices may cause mis-specification in Equation (6). The top-right panel uses a reduced-form Equation (6), replacing CR_{ud} with CR^s_{ud} (using state average coal prices). The bottom two panels present alternate versions of this reduced-form, averaging coal prices by PCA (bottom-left; similar level of aggregation as state) and market region (bottom-right; aggregated up to ISOs and NERC regions). Each panel plots the 45-degree line, winsorizes $\hat{M}_j \in [-1, 1]$ for ease of presentation, and reports the pairwise correlation coefficient for plants with $W_j = 0$.

prices within PCAs (roughly as aggregate as states), and the bottom-right panel aggregates up to much larger market regions (i.e., ISOs and NERC regions). All three aggregations yield similar estimates, which supports the exogeneity of state average prices as an instrument.[14]

Figure A6 tests for robustness to alternate definitions of marginal costs in Equations (A1)–(A2). The top-left panel adds technology-specific defaults for non-fuel variable costs (see Appendix G.2.3); this yields similar $\hat{M}_j$ estimates, confirming that such costs (e.g., labor, maintenance) are second order relative to fuel costs (Cicala (2022)). The top-right panel removes marginal environmental compliance costs (MC^{env}_{um} and MC^{env}_{gm}); this only slightly alters my $\hat{M}_j$ estimates, assuaging concerns about measurement error in permit prices (see Appendix G.6). The bottom-left panel defines P_{jm} as plant j's *minimum* (rather than average) coal price paid in month m, which may better characterize the plant's opportunity cost of coal;[15] the resulting $\hat{M}_j$ estimates are quite similar. The bottom-right panel lags coal prices by one month; similar $\hat{M}_j$ estimates suggest that coal storage is not creating systematic misspecification in CR_{ud}.[16]

Finally, Figure A7 uses alternate specifications of Equation (6). The top-left panel adds month-of-year fixed effects; the top-right panel removes year-specific controls for daily generation, which allow for changes in the fossil electricity supply curve. Both sets of $\hat{M}_j$ estimates change only slightly. The bottom panels discretize CF_{uh} (i.e. $\mathbf{1}[CF_{uh} > 0.5]$) and estimate the reduced form of Equation (6) as a probit and logit; this yields the $\hat{M}_j$ estimates that most

14. Equation (6) uses year fixed effects to control for (potentially endogenous) trends in coal prices.

15. If a plant purchases 75% of its coal on a (relatively expensive) long-term contract, and 25% of its coal on the (cheaper) spot market, the spot price may be more relevant for marginal operating decisions.

16. Plants with coal stockpiles my have short-run opportunity costs of zero (Jha (2022)). I lack data on coal inventories, which would be necessary to test this hypothesis more directly.

diverge from my preferred (linear-IV) model. Figure E10 shows that my DD results using $TREAT_j = \hat{M}_j$ are broadly robust to the sensitivities in Figures A5–A7.

Figure A6: Alternative assumptions on marginal costs

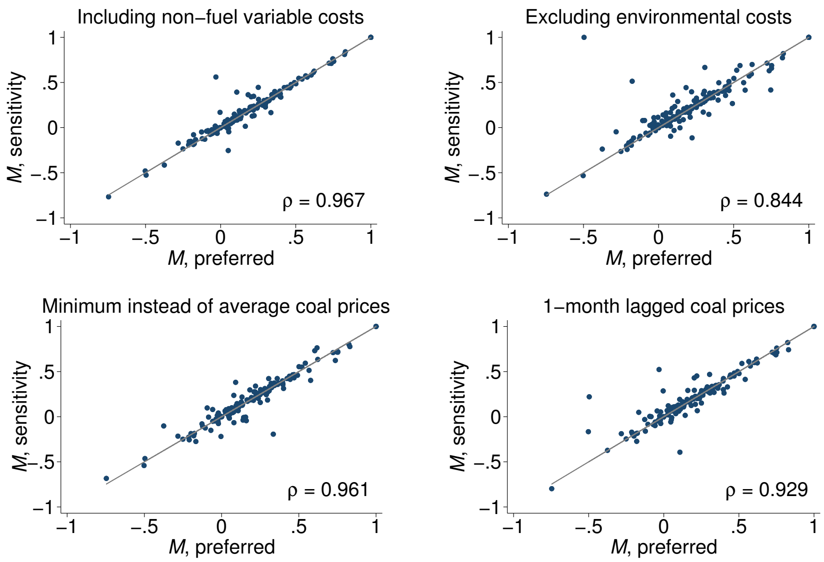

Notes: Each scatter plot compares an alternative $\hat{M}_j$ (vertical axes) to my preferred $\hat{M}_j$ (horizontal axes), which constructs marginal costs using Equations (A1)–(A2). The top-left panel adds an estimate of non-fuel variable costs common to each generating technology. The top-right panel removes marginal environmental costs (MC^{env}_{um} and MC^{env}_{gm}). The bottom-left panel defines P_{jm} as the minimum (rather than the average) coal price paid in month m. The bottom-right panel lags coal price by one month (i.e., $P_{j(m-1)}$) to account for potential storage. Each panel plots the 45-degree line, winsorizes $\hat{M}_j \in [-1, 1]$ for ease of presentation, and reports the pairwise correlation coefficient for plants with $W_j = 0$.

Figure A7: Alternative specifications for Equation (6)

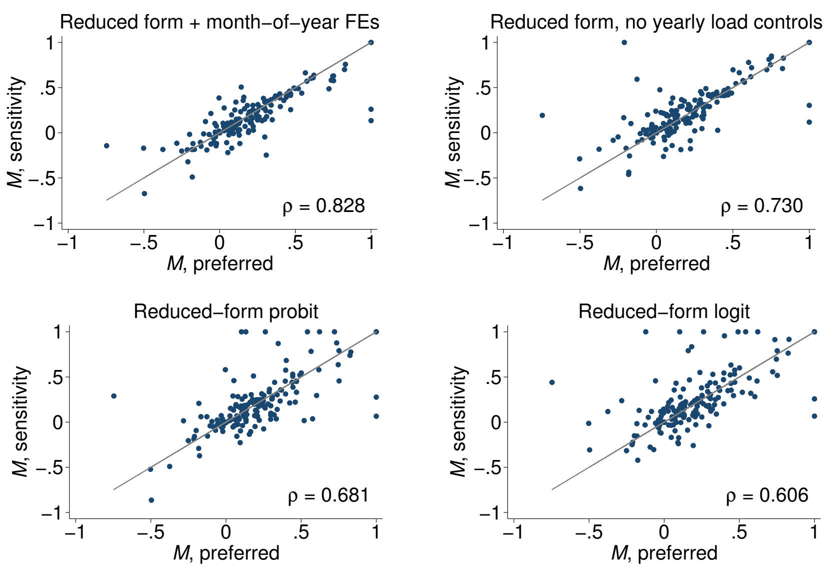

Notes: Each scatter plot compares an alternative $\hat{M}_j$ (vertical axes) to my preferred $\hat{M}_j$ (horizontal axes). The top panels use a reduced-form version of Equation (6) that replaces CR_{ud} with CR^s_{ud}, either adding month-of-year fixed effects (top-left) or removing year fixed effects interacted with the sum of daily CEMS generation in each market region (top-right). The bottom panels discretize CF_{uh} and estimate Equation (6) as either a probit (bottom-left) or a logit (bottom-right). Each panel plots the 45-degree line, winsorizes $\hat{M}_j \in [-1, 1]$ for ease of presentation, and reports the pairwise correlation coefficient for plants with $W_j = 0$.

B Rail oligopoly model

B.1 Derivation of comparative static in Equation (3)

Here I provide a full derivation of the comparative static $\frac{d\mu_j}{dZ}$ from Equation (3) of the main text. I start with rail carrier i's profit function from selling to plant j without a coal-by-barge option (reproduced from Equation (1)):

$$\pi_{ij}(q_{ij}) \,=\, q_{ij}\Big[P_j(Q_j;\mathbf{Z}_j) \,-\, C_j \,-\, S(\mathbf{T}_j)\Big] \,-\, F_j \tag{B1}$$

Firm i earns revenue $q_{ij}P_j$ from selling coal to plant j, while incurring commodity costs $q_{ij}C_j$, shipping costs $q_{ij}S(\mathbf{T}_j)$, and a fixed cost F_j. Plant j's inverse demand is a function of $Q_j = N_j q_{ij}$, the total quantity of coal purchased across all N_j symmetric oligopolists. It also depends on the parameter vector $\mathbf{Z}_j$, which includes Z, the price of natural gas price.

Firm i's first-order condition is:

$$\frac{\partial \pi_{ij}}{\partial q_{ij}} \,=\, P_j(Q_j;\mathbf{Z}_j) \,+\, q_{ij}\frac{\partial P_j}{\partial Q_j}\frac{\partial Q_j}{\partial q_{ij}} \,-\, C_j \,-\, S(\mathbf{T}_j) \tag{B2}$$

For simplicity, I assume that $S(\mathbf{T}_j)$ does not depend on q_{ij}, which abstracts from rail capacity constraints and increasing returns to scale in shipping.[17] Totally differentiating Equation (B2) by q_{ij} and Z, and rearranging:[18]

$$\frac{dq_{ij}}{dZ} \,=\, \frac{\dfrac{\partial P_j}{\partial Z} \,+\, \dfrac{\partial^2 P_j}{\partial Q_j \partial Z}\dfrac{\partial Q_j}{\partial q_{ij}} q_{ij}}{-\left(2\dfrac{\partial P_j}{\partial Q_j}\dfrac{\partial Q_j}{\partial q_{ij}} \,+\, \dfrac{\partial^2 P_j}{\partial Q_j^2}\left(\dfrac{\partial Q_j}{\partial q_{ij}}\right)^2 q_{ij} \,+\, \dfrac{\partial P_j}{\partial Q_j}\dfrac{\partial^2 Q_j}{\partial q_{ij}^2} q_{ij}\right)} \tag{B3}$$

Invoking symmetry, I substitute $q_{ij} = \frac{Q_j}{N_j}$. I also define the "conduct parameter" $\theta_j \equiv \frac{\partial Q_j}{\partial q_{ij}}$.[19] Rewriting (B3):

$$\frac{dq_{ij}}{dZ} \,=\, \frac{\dfrac{\partial P_j}{\partial Z} \,+\, \dfrac{\partial^2 P_j}{\partial Q_j \partial Z}\dfrac{Q_j \theta_j}{N_j}}{-\left(2\dfrac{\partial P_j}{\partial Q_j}\theta_j \,+\, \dfrac{\partial^2 P_j}{\partial Q_j^2}\dfrac{Q_j \theta_j^2}{N_j} \,+\, \dfrac{\partial P_j}{\partial Q_j}\underbrace{\dfrac{\partial \theta_j}{\partial q_{ij}}}_{=0}\dfrac{Q_j}{N_j}\right)} \tag{B4}$$

I assume $\frac{\partial \theta_j}{\partial q_{ij}} = 0$, because small changes in q_{ij} are unlikely to change the relationship between carrier i's quantity q_{ij} and total demand Q_j.[20]

Next, I derive the comparative static for $Q_j = \sum_i q_{ij}$. Totally differentiating Q_j by Z, and invoking symmetry across all N_j rail carriers:

$$\frac{dQ_j}{dZ} \,=\, \sum_i \frac{\partial Q_j}{\partial q_{ij}}\frac{dq_{ij}}{dZ} \,=\, N_j \theta_j \frac{dq_{ij}}{dZ} \tag{B5}$$

17. My empirical analysis relaxes this assumption by allowing rail transport costs to vary with shipment size. Since plants are small relative to the coal mining sector, I also assume that C_j is independent of q_{ij}.

18. This applies the Implicit Function Theorem: $\frac{\partial \pi_{ij}}{\partial q_{ij}}$ is a function of q_{ij} and Z, for level set $\frac{\partial \pi_{ij}}{\partial q_{ij}} = 0$.

19. I use this "conduct parameter" formulation only for notational convenience (following Atkin and Donaldson (2015)). Below, I replace $\theta_j = 1$, which is consistent with Cournot competition. Calibrating θ_j as a structural parameter can be problematic, as it only has a well-defined interpretation at a few values (Corts (1999)).

20. A small change in q_{ij} should not change whether firm i behaves as Cournot competitive ($\theta_j = 1$).

$$\Rightarrow \quad \frac{dQ_j}{dZ} = \frac{\frac{\partial P_j}{\partial Z} N_j + \frac{\partial^2 P_j}{\partial Q_j \partial Z} Q_j \theta_j}{-\left(2 \frac{\partial P_j}{\partial Q_j} + \frac{\partial^2 P_j}{\partial Q_j^2} \frac{Q_j \theta_j}{N_j}\right)} \tag{B6}$$

The first term in the numerator, $\frac{\partial P_j}{\partial Z} N_j$, captures the level-effect of the demand shift, which should be weakly positive: an inward demand shift (i.e. $dZ < 0$) should reduce plant j's coal consumption Q_j. The second term in the numerator, $\frac{\partial^2 P_j}{\partial Q_j \partial Z} Q_j \theta_j$, captures the extent to which the demand shift dZ changes the slope of inverse demand: if demand becomes more elastic as gas prices fall (i.e. $\frac{\partial^2 P_j}{\partial Q_j \partial Z} < 0$), rail carriers should increase their best-response quantities. The denominator is positive by the second-order condition (see Appendix B.4 below).

The final step converts Equation (B6) into the total derivative of P_j with respect to Z:

$$\frac{dP_j}{dZ} = \frac{\partial P_j}{\partial Q_j} \frac{dQ_j}{dZ} + \frac{\partial P_j}{\partial Z} \tag{B7}$$

$$\frac{dP_j}{dZ} = \frac{\partial P_j}{\partial Q_j} \left[\frac{\frac{\partial P_j}{\partial Z} N_j + \frac{\partial^2 P_j}{\partial Q_j \partial Z} Q_j \theta_j}{-\left(2 \frac{\partial P_j}{\partial Q_j} + \frac{\partial^2 P_j}{\partial Q_j^2} \frac{Q_j \theta_j}{N_j}\right)} \right] + \frac{\partial P_j}{\partial Z} \tag{B8}$$

$$\frac{dP_j}{dZ} = \left[\frac{\frac{\partial P_j}{\partial Z} N_j + \frac{\partial^2 P_j}{\partial Q_j \partial Z} Q_j \theta_j}{-\left(2 + \frac{\partial^2 P_j}{\partial Q_j^2} \left(\frac{\partial P_j}{\partial Q_j}\right)^{-1} \frac{Q_j \theta_j}{N_j}\right)} \right] + \frac{\partial P_j}{\partial Z} \tag{B9}$$

Let $E_{D_j} \equiv \left(\frac{\partial^2 P_j}{\partial Q_j^2}\right) \left(\frac{\partial P_j}{\partial Q_j}\right)^{-1} Q_j$, or the elasticity of the slope of inverse demand:

$$\frac{dP_j}{dZ} = \frac{\frac{\partial P_j}{\partial Z} N_j + \frac{\partial^2 P_j}{\partial Q_j \partial Z} Q_j \theta_j}{-\left(2 + E_{D_j} \frac{\theta_j}{N_j}\right)} + \frac{\partial P_j}{\partial Z} \left(\frac{2 + E_{D_j} \frac{\theta_j}{N_j}}{2 + E_{D_j} \frac{\theta_j}{N_j}} \right) \tag{B10}$$

$$\Rightarrow \quad \frac{dP_j}{dZ} = \frac{\frac{\partial P_j}{\partial Z} \left(2 + E_{D_j} \frac{\theta_j}{N_j} - N_j\right) - \frac{\partial^2 P_j}{\partial Q_j \partial Z} Q_j \theta_j}{2 + E_{D_j} \frac{\theta_j}{N_j}} \tag{B11}$$

I define markups as $\mu_j \equiv P_j - C_j - S(\mathbf{T}_j)$, assuming C_j and $S(\mathbf{T}_j)$ are independent of q_{ij} and Z.[21] Hence, it follows that:

$$\Rightarrow \quad \frac{d\mu_j}{dZ} = \frac{\frac{\partial P_j}{\partial Z} \left(2 + E_{D_j} \frac{\theta_j}{N_j} - N_j\right) - \frac{\partial^2 P_j}{\partial Q_j \partial Z} Q_j \theta_j}{2 + E_{D_j} \frac{\theta_j}{N_j}} \tag{B12}$$

21. During my sample period, U.S. natural gas prices were uncorrelated with diesel prices, the time series component of $\mathbf{T}_j$. Figure 2 shows that delivered coal prices do not respond to medium-run gas price shocks.

Equation (3) replaces $\theta_j = 1$, which is consistent with Cournot competition:

$$\textbf{Cournot competition:} \qquad \frac{d\mu_j}{dZ} = \frac{\dfrac{\partial P_j}{\partial Z}\left(2 + \dfrac{E_{D_j}}{N_j} - N_j\right) - \dfrac{\partial^2 P_j}{\partial Q_j \partial Z} Q_j}{2 + \dfrac{E_{D_j}}{N_j}}$$

An alternate market structure of perfect collusion would replace $\theta_j = N_j$:

$$\textbf{Perfect collusion:} \qquad \frac{d\mu_j}{dZ} = \frac{\dfrac{\partial P_j}{\partial Z}\left(2 + E_{D_j} - N_j\right) - \dfrac{\partial^2 P_j}{\partial Q_j \partial Z} Q_j N_j}{2 + E_{D_j}}$$

Under either assumption on the conduct between rail carriers, the outside option of barge transportation should eliminate rail carriers' ability to set positive markups:

$$\textbf{Barge option:} \qquad \frac{d\mu_j}{dZ} \approx 0$$

The literature treats barge shipping as (close to) competitive (e.g., MacDonald (1987); Busse and Keohane (2007); Wetzstein et al. (2021)). EIA data report lower average transportation costs for coal-by-barge (\$5–8/ton) than for coal-by-rail (\$17–23/ton).[22] I observe prices for both coal-by-barge and coal-by-rail deliveries to the same plant for 94 plant-years in my coal delivery dataset; for 65 plant-years (69%), coal-by-barge has lower average shipping costs than coal-by-rail.[23] Together, this evidence suggests that even if barge shipping were not (close to) competitive, its cost advantage over rail would make rail carriers' residual demand quite elastic—leaving them little ability to exercise market power by setting rail markups.

B.2 Arbitrage, coal attributes, and cross-plant dependencies

Here, I address three simplifying assumptions in my Cournot model. First, I assume that plant j cannot resell purchased coal to other plants, which effectively lets rail carrier i optimize each plant independently. This conforms with reality, where it is cost-prohibitive for plants to circumvent the rail carriers by reselling coal using trucks.[24] A more explicit formulation of the rail carrier's problem would include arbitrage constraints that bind when price wedges are just large enough to make coal resale cost-effective. This would follow a standard representation of 3rd-degree price discrimination (e.g., Schmalensee (1981)).

Second, my model abstracts from coal's heterogeneous attributes and geography. In reality, plant j has preferences over coal varieties, given its boiler specifications, pollution control devices, and environmental compliance costs. If plant j is serviced by two rail carriers and prefers coal from Wyoming, but one carrier's track does not extend to Wyoming, then the plant faces an effective rail monopoly. A richer model could specify a separate profit function for each rail carrier i, plant j, and coal-producing county o. This would account for heterogeneous demand across coal varieties $P_{oj}(Q_{oj}; \mathbf{Z}_{oj})$, commodity costs C_o that reflect coal attributes (e.g.,

22. See Table 1 here: https://www.eia.gov/coal/transportationrates/.

23. Here, I approximate shipping costs by subtracting county-year average mine-mouth prices from delivered prices. This captures roughly 80% of the variation in physical coal attributes (i.e., BTU, sulfur, and ash content).

24. Using barges to resell coal would be more feasible and cost-effective. However plants with this waterborne option likely face markups close to zero, obviating the need to arbitrage.

sulfur content), shipping costs $S(\mathbf{T}_{oj})$ that reflect origin-to-destination-specific rail mileage, and route-specific entry costs F_{oj}.

Third, I ignore dependencies in coal demand across plants and coal types. In reality, plant j's demand for type-o coal depends on the prices it faces for coal from other counties, and likely also on the prices faced by other plants (Varian (1985); Katz (1987)). I can incorporate this vector of non-oj prices into the parameter vector $\mathbf{Z}_{oj}$, or explicitly include this vector of prices as an argument entering inverse demand: $P_{oj}(Q_{oj}; \mathbf{P}_{-(oj)}, \mathbf{Z}_{oj})$.

I can modify Equation (B1) to explicitly account for each of these assumptions, in order to illustrate what the above comparative statics assume away:

$$\Pi_i(\mathbf{q}_i) \;=\; \mathbf{q}_i \cdot \Big[\mathbf{P}(\mathbf{Q};\mathbf{Z}) \;-\; \mathbf{C} \;-\; S(\mathbf{T}_i)\Big] \;-\; \mathbf{F}_i \cdot \mathbf{1}\Big[\mathbf{q}_i > \mathbf{0}\Big] \tag{B13}$$
$$\text{s.t.} \quad \Big|P_{oj}(Q_{oj}; \mathbf{P}_{-(oj)}, \mathbf{Z}) \;-\; \mathbf{P}(\mathbf{Q};\mathbf{Z})\Big| \;\leq\; \mathbf{A}_j \quad \forall\, o,\ \forall\, j$$

Here, rail carrier i jointly optimizes across all county-plant pairs, and $\mathbf{q}_i$ is an $[OJ \times 1]$ vector of coal quantities (indexed by oj). $\mathbf{P}(\mathbf{Q};\mathbf{Z})$ is the OJ-dimensional inverse demand function, which depends on the $[OJ \times 1]$ vector of coal quantities $\mathbf{Q}$ and a matrix of demand parameters $\mathbf{Z}$ with OJ rows. $\mathbf{C}$ is an $[OJ \times 1]$ vector of mine-mouth coal costs, repeating each C_o J times. $S(\mathbf{T}_i)$ is an $[OJ \times 1]$ vector of carrier i's coal shipping costs from origin o to destination j. $\mathbf{F}_i$ is an $[OJ \times 1]$ vector of carrier i's fixed costs of maintaining each oj route, which is multiplied by an $[OJ \times 1]$ vector of indicators for carrier i's entry decision along each route. $\mathbf{A}_j$ is an $[OJ \times 1]$ vector of the price wedges at which arbitrage becomes feasible (i.e. reflecting plant j's costs of shipping coal to/from all other plants).[25]

If the OJ arbitrage constraints never bind, and if the first-derivative matrix of $\mathbf{P}(\mathbf{Q};\mathbf{Z})$ is diagonal,[26] then firm i faces a separate unconstrained maximization problem for each plant:[27]

$$\max_{\mathbf{q}_{ij}} \pi_{ij}(\mathbf{q}_{ij}) \;=\; \sum_o q_{ioj}\Big[P_{oj}(Q_{oj};\mathbf{Z}_{oj}) \;-\; C_o \;-\; S(\mathbf{T}_{oj})\Big] \;-\; \mathbf{F}_{ij} \cdot \mathbf{1}\Big[\mathbf{q}_{ij} > \mathbf{0}\Big] \;\; \forall\, j \tag{B14}$$

Setting profits to zero for oj routes that carrier i does not service, Equation (B14) becomes:

$$\max_{\{q_{ioj}\}\,|\,q_{ioj}>0} \pi_{ij} \;=\; \sum_{o\,|\,q_{ioj}>0} \Big\{q_{ioj}\Big[P_{oj}(Q_{oj};\mathbf{Z}_{oj}) \;-\; C_o \;-\; S(\mathbf{T}_{oj})\Big] \;-\; F_{oj}\Big\} \;\; \forall\, j \tag{B15}$$

Suppose that plant j only consumes coal from the set of counties $\mathcal{O}_j$, and treats all coal from these counties as perfectly substitutable: $P_{oj}(Q_{oj};\mathbf{Z}_{oj}) = P_j(Q_j;\mathbf{Z}_j)\ \forall\, o \in \mathcal{O}_j$. Further suppose that it has zero demand for coal from other counties: $Q_{oj}(P_{oj};\mathbf{Z}_{oj}) = 0\ \forall\, o \notin \mathcal{O}_j$. Then, the rail carrier can choose the cheapest quantity q^*_{ij} from across counties $o \in \mathcal{O}_j$:

$$\max_{q^*_{ij}} \pi_{ij} \;=\; q^*_{ij}\Big[P_j(Q^*_j;\mathbf{Z}_j) \;-\; C^*_{ij} \;-\; S(\mathbf{T}^*_{ij})\Big] \;-\; F_j \;\; \forall\, j \tag{B16}$$

25. $\mathbf{A}_j$ is $[OJ \times 1]$ to conform with the dimension of $\mathbf{P}(\mathbf{Q};\mathbf{Z})$, meaning that there are O null j-to-j no-arbitrage constraints within plant j's set of arbitrage constraints. Vertical lines denote the absolute value operator, applied element-wise to the difference between the scalar $P_{oj}(Q_{oj};\mathbf{P}_{-(oj)},\mathbf{Z})$ and the vector $\mathbf{P}(\mathbf{Q};\mathbf{Z})$. Busse and Keohane (2007) note that in order to arbitrage around the railroads, plants would need to transfer coal from their on-site storage piles onto trucks, a more costly mode of transportation. While coal resale is quite rare in practice, the *threat* of arbitrage may limit railroads' willingness to price discriminate.

26. That is, $\frac{\partial P_{oj}}{\partial Q_{nk}} = 0$ for $k \neq j$ or $n \neq o$; and $\frac{\partial P_{oj}}{\partial Z_{nk}} = 0$ for $k \neq j$ or $n \neq o$, for any element Z of $\mathbf{Z}$.

27. Here, $\mathbf{F}_{ij}$ and $\mathbf{q}_{ij}$ are $[O \times 1]$ vectors.

Here, $Q_j^* = \sum_i q_{ij}^*$; and C_{ij}^* and $\mathbf{T}_{ij}^*$ are carrier i's minimized costs of supplying q_{ij}^*.[28] Equation (B16) collapses to Equation (B1) if $C_j = C_{ij}^*$ and $\mathbf{T}_j = \mathbf{T}_{ij}^*$ for all carriers that service plant j.

B.3 Accounting for the threat of regulation

My model ignores the potential for binding rail price regulation, which is unrealistic. Since 1996, the Surface Transportation Board has reviewed only 34 cases disputing the "reasonableness" of coal-by-rail shipping rates; only 10 of these 34 cases ended with a decision that the rail carrier's rate was "unreasonable". This likely represents only a small fraction of rail rates constrained by the threat (or even the *perceived* threat) of a regulatory challenge. I can modify Equation (B1) to account for the threat of regulation:

$$\pi_{ij}^R(q_{ij}) \;=\; q_{ij}\Big[P_j(Q_j;\mathbf{Z}_j) \;-\; C_j \;-\; S(\mathbf{T}_j)\Big] \;-\; F_j \;-\; R_j(\mu_j) \tag{B17}$$

$R_j(\mu_j)$ is the expected penalty (in dollars) from regulatory oversight (e.g., a rate case brought by the Surface Transportation Board). $R_j(\cdot)$ is an increasing function of the markup μ_j, since higher markups are more likely to lead to action from the regulator. This reflects the Surface Transportation Board's practice of loosely interpreting its statutory standard (180% above total variable costs), while exercising discretion in allowing railroads to earn an adequate return on investment (Wilson (1996); MacDonald (2013); Mayo and Sappington (2016); InterVISTAS (2016)).[29] π_{ij}^R represents expected profits, net of $R_j(\mu_j)$.

Firm i's first-order condition is now:

$$\frac{\partial \pi_{ij}^R}{\partial q_{ij}} \;=\; P_j(Q_j;\mathbf{Z}_j) \;+\; q_{ij}\frac{\partial P_j}{\partial Q_j}\frac{\partial Q_j}{\partial q_{ij}} \;-\; C_j \;-\; S(\mathbf{T}_j) \;-\; \frac{\partial R_j}{\partial P_j}\frac{\partial P_j}{\partial Q_j}\frac{\partial Q_j}{\partial q_{ij}} \tag{B18}$$

Totally differentiating by q_{ij} and Z, rearranging, substituting $q_{ij} = \frac{Q_j}{N_j}$ and $\theta_j \equiv \frac{\partial Q_j}{\partial q_{ij}}$, and assuming $\frac{\partial \theta_j}{\partial q_{ij}} = 0$:

$$\frac{dq_{ij}}{dZ} \;=\; \frac{\dfrac{\partial P_j}{\partial Z}\left(1 - \dfrac{\partial P_j}{\partial Q_j}\theta_j\dfrac{\partial^2 R_j}{\partial P_j^2}\right) \;+\; \dfrac{\partial^2 P_j}{\partial Q_j \partial Z}\theta_j\left(\dfrac{Q_j}{N_j} - \dfrac{\partial R_j}{\partial P_j}\right)}{-\left[2\dfrac{\partial P_j}{\partial Q_j}\theta_j \;+\; \dfrac{\partial^2 P_j}{\partial Q_j^2}\theta_j^2\left(\dfrac{Q_j}{N_j} - \dfrac{\partial R_j}{\partial P_j}\right) \;-\; \theta_j^2\left(\dfrac{\partial P_j}{\partial Q_j}\right)^2\dfrac{\partial^2 R_j}{\partial P_j^2}\right]} \tag{B19}$$

Substituting using Equation (B5), and multiplying through by $N_j\theta_j$:

$$\frac{dQ_{ij}}{dZ} \;=\; \frac{N_j\dfrac{\partial P_j}{\partial Z}\left(1 - \dfrac{\partial P_j}{\partial Q_j}\theta_j\dfrac{\partial^2 R_j}{\partial P_j^2}\right) \;+\; \dfrac{\partial^2 P_j}{\partial Q_j \partial Z}\theta_j\left(Q_j - N_j\dfrac{\partial R_j}{\partial P_j}\right)}{-\left[2\dfrac{\partial P_j}{\partial Q_j} \;+\; \dfrac{\partial^2 P_j}{\partial Q_j^2}\theta_j\left(\dfrac{Q_j}{N_j} - \dfrac{\partial R_j}{\partial P_j}\right) \;-\; \theta_j\left(\dfrac{\partial P_j}{\partial Q_j}\right)^2\dfrac{\partial^2 R_j}{\partial P_j^2}\right]} \tag{B20}$$

Substituting using Equation (B7), rearranging, substituting $E_{D_j} \equiv \left(\frac{\partial^2 P_j}{\partial Q_j^2}\right)\left(\frac{\partial P_j}{\partial Q_j}\right)^{-1} Q_j$, and invoking the assumption of exogenous C_j and $S(\mathbf{T}_j)$:

28. Let $G_{oj} \equiv C_o + S(\mathbf{T}_{oj})$. Then, $q_{ij}^* \equiv \arg\min_{q_{ioj}\,|\,q_{ioj}>0\,,\,o\in\mathcal{O}_j} G_{oj}$, $G_{ij}^* \equiv \min_{q_{ioj}\,|\,q_{ioj}>0\,,\,o\in\mathcal{O}_j} G_{oj}$, and $C_{ij}^* + S(\mathbf{T}_{ij}^*) = G_{ij}^*$.

29. If the regulator strictly enforced this standard, a more realistic formulation of Equation (B17) would replace $R_j(\mu_j)$ with a constraint that $q_{ij}P_j(Q_j;\mathbf{Z}_j) \le 1.8q_{ij}\Big[C_j \;+\; S(\mathbf{T}_j)\Big]$.

$$\frac{d\mu_j}{dZ} = \frac{\frac{\partial P_j}{\partial Z}\left[2 + E_{D_j}\theta_j\left(\frac{1}{N_j} - \frac{\partial R_j}{\partial P_j}\frac{1}{Q_j}\right) - N_j + (N_j - 1)\theta_j\frac{\partial P_j}{\partial Q_j}\frac{\partial^2 R_j}{\partial P_j^2}\right] - \frac{\partial^2 P_j}{\partial Q_j \partial Z}\theta_j\left(Q_j - N_j\frac{\partial R_j}{\partial P_j}\right)}{2 \,+\, E_{D_j}\theta_j\left(\frac{1}{N_j} - \frac{\partial R_j}{\partial P_j}\frac{1}{Q_j}\right) \,-\, \theta_j\frac{\partial P_j}{\partial Q_j}\frac{\partial^2 R_j}{\partial P_j^2}}$$

I assume that $\frac{\partial R_j}{\partial P_j} > 0$ and $\frac{\partial^2 R_j}{\partial P_j^2} \geq 0$, since the probability of regulatory action should increase in μ_j, and likely convexly. If I restrict $\frac{\partial^2 R_j}{\partial P_j^2} = 0$, this expression becomes:

$$\frac{d\mu_j}{dZ} = \frac{\frac{\partial P_j}{\partial Z}\left[2 + E_{D_j}\frac{\theta_j}{N_j}\left(1 - \frac{\partial R_j}{\partial P_j}\frac{N_j}{Q_j}\right) - N_j\right] - \frac{\partial^2 P_j}{\partial Q_j \partial Z}Q_j\theta_j\left(1 - \frac{\partial R_j}{\partial P_j}\frac{N_j}{Q_j}\right)}{2 \,+\, E_{D_j}\frac{\theta_j}{N_j}\left(1 - \frac{\partial R_j}{\partial P_j}\frac{N_j}{Q_j}\right)} \tag{B21}$$

I can rearrange Equation (B21) in order to sign terms in the numerator and denominator:

$$\frac{d\mu_j}{dZ} = \frac{\overbrace{\overbrace{\frac{\partial P_j}{\partial Z}}^{>0}\left(2 + \overbrace{E_{D_j}}^{>0}\frac{\theta_j}{N_j} - N_j\right) - \overbrace{\frac{\partial^2 P_j}{\partial Q_j \partial Z}}^{<0}Q_j\theta_j}^{\text{Equation (B12) numerator: } >0} + \overbrace{\overbrace{\frac{\partial R_j}{\partial P_j}}^{>0}\left(\overbrace{\frac{\partial^2 P_j}{\partial Q_j \partial Z}}^{<0}N_j\theta_j - \overbrace{\frac{\partial P_j}{\partial Z}}^{>0}\overbrace{E_{D_j}}^{>0}\frac{\theta_j}{Q_j}\right)}^{\text{Added threat of regulation: } <0}}{\underbrace{2 \,+\, \underbrace{E_{D_j}}_{>0}\frac{\theta_j}{N_j}}_{\text{Equation (B12) denominator: } >0} \quad + \quad \underbrace{-\underbrace{\frac{\partial R_j}{\partial P_j}}_{>0}\underbrace{E_{D_j}}_{>0}\frac{\theta_j}{Q_j}}_{\text{Added threat of regulation: } <0}}$$

I sign $\frac{\partial P_j}{\partial Z} > 0$, $\frac{\partial^2 P_j}{\partial Q_j \partial Z} < 0$, and $E_{D_j} > 0$, which is consistent with both the intuition of Figure 3 and my $\langle\hat{\lambda}_{0j}, \hat{\lambda}_{1j}, \hat{\lambda}_{2j}\rangle$ estimates from Figure 5.[30] The threat of regulation decreases the numerator *and* the denominator, creating an ambiguous net effect on markup changes.

It might seem counter-intuitive that the threat of regulation could simultaneously dampen markup levels and induce *larger* changes in markups. This effect is driven by the curvature of demand. Regulation unambiguously reduces markup changes if I set $E_{D_j} = 0$:

$$\left.\frac{d\mu_j}{dZ}\right|_{E_{D_j}=0} = \frac{\overbrace{\overbrace{\frac{\partial P_j}{\partial Z}}^{>0}(2 - N_j) - \overbrace{\frac{\partial^2 P_j}{\partial Q_j \partial Z}}^{<0}Q_j\theta_j}^{\text{Modified Equation (B12) numerator: } >0} \quad + \quad \overbrace{\overbrace{\frac{\partial R_j}{\partial P_j}}^{>0}\overbrace{\frac{\partial^2 P_j}{\partial Q_j \partial Z}}^{<0}N_j\theta_j}^{\text{Added threat of regulation: } <0}}{2}$$

Comparing these two expressions illustrates how concave demand ($E_{D_j} > 0$) interacts with the threat of regulation in two ways. First, it moderates the demand shock, since regulation induces lower markups prior to the change in Z (i.e., the numerator effect). Second, it magnifies the penalty risk associated with maintaining markups (i.e., the denominator effect).

For suggestive evidence of how the threat of rail regulation impacts markup changes, I can compare my $\hat{M}_j$ predictions—the empirical analog of $\frac{d\mu_j}{dZ}$ constructed *without* accounting for regulation—with my DD estimates using $TREAT_j = \hat{M}_j$. Regression coefficients of $\hat{\tau} \approx 20$

30. I also assume $N_j \in \{1, 2\}$, $\theta_j > 0$, and $Q_j > 0$.

would indicate that my predictions (ignoring regulation) match the average observed markup change (inclusive of regulation): $\hat{M}_j$ is in (\$/MMBTU coal)/(\$/MMBTU gas), τ is in (\$/ton coal)/(\$/MMBTU gas), and coal's average BTU content is 19.7 MMBTU/ton. My much smaller DD estimates of $\hat{\tau} \approx 3$, suggest that regulation tends to dampen coal-by-rail markup changes.

B.4 Empirically validating the second-order condition

The second-order condition of the rail carrier's problem (Equation (B1)) is:

$$\frac{\partial^2 \pi_{ij}}{\partial q_{ij}^2} = 2\frac{\partial P_j}{\partial Q_j}\frac{\partial Q_j}{\partial q_{ij}} + q_{ij}\frac{\partial^2 P_j}{\partial Q_j^2}\Big(\underbrace{\frac{\partial Q_j}{\partial q_{ij}}}_{=\theta_j}\Big)^2 + q_{ij}\frac{\partial P_j}{\partial Q_j}\underbrace{\frac{\partial^2 Q_j}{\partial q_{ij}^2}}_{=\frac{\partial \theta_j}{\partial q_{ij}}=0} < 0 \tag{B22}$$

$$= 2\frac{\partial P_j}{\partial Q_j}\theta_j + q_{ij}\frac{\partial^2 P_j}{\partial Q_j^2}\theta_j^2 < 0 \tag{B23}$$

I can empirically approximate this expression using results from my coal demand estimation—namely, the first and second derivatives of inverse coal demand at the observed quantity of consumption, for each plant-month. If I restrict my sample to plants without a water option (where $\theta_j = 1$), I observe 3,184 plant-months with a rail shipment—72% of which satisfy the empirical analog of this second-order condition. For the 1,841 plant-months where the plant is captive (and $q_{ij} = Q_j$), 73% satisfy the empirical analog of Equation (B23).

Figure B1 plots the distribution of this restricted subset of plant-months, separately for plants with above- vs. below-mean $\hat{M}_j$. While the second-order condition holds for most plant-months, it holds more often for plant-months with $\hat{M}_j \geq 0.29$ (77%, compared to 69% of plant-months with $\hat{M}_j < 0.29$). This suggests that if violations of the second-order condition cause misspecification of $\hat{M}_j$, such misspecification would bias my DD estimates towards zero (since plants with larger $\hat{M}_j$ would be relatively less misspecified). This aligns with my demand parameter estimates: $\hat{\lambda}_{2j} > 0$ for 69% of plants, indicating that the relevant portion of the coal demand curve tends to be concave. Since Equation (B23) can only fail if coal demand is convex, these empirical patterns are reassuring in the context of my theoretical framework.

Figure B1: Empirical approximation of the second-order condition

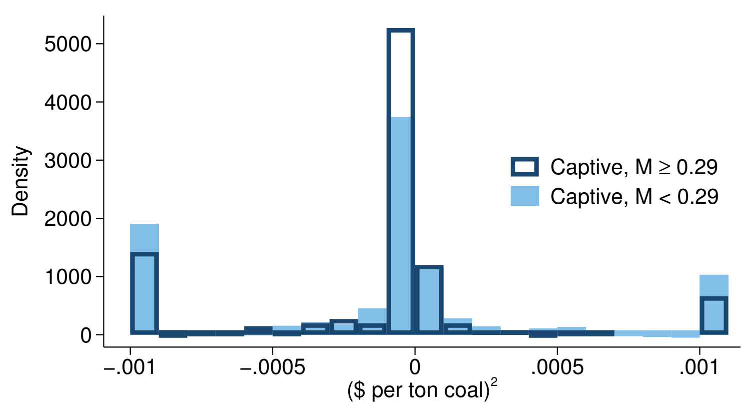

Notes: These are histograms of the empirical approximation of the Equation (B23), at the plant-month level. I populate the first and second derivatives using my preferred demand curve estimates (see Appendix A.1), and restrict sample to plants with $W_j = 0$ (so that $\theta_j = 1$) and $D_j = 1$ (so that $q_{ij} = Q_j$). I weight observations using $k = 3$ nearest-neighbor weights, and winsorize at $[-0.001, 0.001]$. The second-order condition holds for 77% (69%) plant-months with $\hat{M}_j$ above (below) the mean of 0.29.

C Pass-through of a carbon tax

C.1 Derivation of pass-through rate ρ_j

Here, I derive expressions for the carbon tax pass-through (ρ) implied by a change in natural gas prices (ΔZ) and a reoptimization of coal markups ($\Delta\mu$). Consider a coal plant j in a market with many natural gas plants. The ratio of coal-to-gas marginal costs governs the relative ordering of plants on the electricity supply curve, which is the primary factor influencing plant j's operating decisions (see Figure 3). Plant j's cost ratio is (suppressing subscripts):

$$CR = \frac{MC^{coal}}{MC^{gas}} = \frac{HR_{coal} \cdot P}{HR_{gas} \cdot Z} \tag{C1}$$

HR_{fuel} is the heat rate, or the rate at which each plant converts fuel into electricity (in MMBTU/MWh).[31] P is the coal price, and Z is the gas price, both in \$/MMBTU.

Let ΔZ denote a change in the gas price Z, which implies a new cost ratio:

$$CR' = \frac{HR_{coal} \cdot P}{HR_{gas} \cdot (Z + \Delta Z)} \tag{C2}$$

A hypothetical carbon tax t would yield an identical CR', holding gas prices constant:

$$CR' = \frac{HR_{coal} \cdot P}{HR_{gas} \cdot (Z + \Delta Z)} = \frac{HR_{coal} \cdot (P + tE_{coal})}{HR_{gas} \cdot (Z + tE_{gas})} \tag{C3}$$

The tax t is in \$ per metric ton CO_2, and E_{fuel} is the CO_2 emissions rate for each fuel, in metric tons CO_2/MMBTU. Since $E_{coal} > E_{gas}$, t exists for any feasible change in gas prices ΔZ.[32]

Solving Equation (C3) for t:

$$\begin{aligned}
\frac{P}{Z + \Delta Z} &= \frac{P + tE_{coal}}{Z + tE_{gas}} \\
P(Z + tE_{gas}) &= (Z + \Delta Z)(P + tE_{coal}) \\
P \cdot tE_{gas} &= (Z + \Delta Z)tE_{coal} + \Delta Z \cdot P \\
t[P \cdot E_{gas} - (Z + \Delta Z)E_{coal}] &= \Delta Z \cdot P \\
\Rightarrow \quad t(\Delta Z) &= \frac{\Delta Z \cdot P}{P \cdot E_{gas} - (Z + \Delta Z)E_{coal}} \qquad \text{(C4)}
\end{aligned}$$

$t(\Delta Z)$ represents the equivalent carbon tax implied by the gas price change ΔZ.

However, rail carriers may reoptimize coal markups in response to ΔZ. If markups change by $\Delta\mu$, I can rewrite Equation (C2):

$$CR' = \frac{HR_{coal} \cdot (P + \Delta\mu)}{HR_{gas} \cdot (Z + \Delta Z)} \tag{C5}$$

31. HR_{gas} is weight-averaged across all gas plants that compete with plant j in electricity dispatch.

32. I adapt this mapping from fuel prices to carbon tax from Cullen and Mansur (2017). Natural gas has homogeneous emissions of 0.053 metric tons CO_2/MMBTU (https://www.eia.gov/tools/faqs/faq.php?id=73&t=1). Coal's average emissions are 0.095 metric tons CO_2/MMBTU, which are fairly homogeneous. EPA assumes 0.00025 metric tons CO_2/ton-mile of rail freight (U.S. EPA (2008)); this implies the CO_2 emitted from diesel combustion in coal-by-rail is two orders of magnitude smaller than the CO_2 emitted from burning coal.

Pass-through of the implicit tax $t(\Delta Z)$ is a function of $\Delta\mu$. If $\Delta\mu = 0$, then the pass-through rate is $\rho = 1$: the coal plant faces the full implicit tax $t(\Delta Z)$, without any changes in markups that weaken or strengthen this price signal. If $\text{sign}(\Delta\mu) = \text{sign}(\Delta Z)$, then $\Delta\mu$ weakens the effect of ΔZ on CR', translating to incomplete pass-through (i.e. $\rho < 1$). If $\text{sign}(\Delta\mu) = -\text{sign}(\Delta Z)$, then $\Delta\mu$ strengthens the effect of ΔZ on CR', translating to overshifting (i.e. $\rho > 1$).

Modifying Equation (C3) to allow for changes in markups ($\Delta\mu$) and incomplete pass-through (i.e. $\rho \neq 1$):

$$\frac{HR_{coal} \cdot (P + \Delta\mu)}{HR_{gas} \cdot (Z + \Delta Z)} = \frac{HR_{coal} \cdot (P + \rho\, t(\Delta Z) E_{coal})}{HR_{gas} \cdot (Z + t(\Delta Z) E_{gas})} \tag{C6}$$

Here, $t(\Delta Z)$ is the carbon tax implied by ΔZ under full pass-through (i.e. $\Delta\mu = 0$, $\rho = 1$). If markups adjust (i.e. $\Delta\mu \neq 0$), this causes coal plants to face a different proportion (i.e. $\rho \neq 1$) of this implicit tax. Note that ρ appears only in the numerator, as I assume full tax pass-through for natural gas (I relax this assumption in Appendix C.3 below). Solving Equation (C6) for ρ:

$$\begin{aligned}
\frac{P + \Delta\mu}{Z + \Delta Z} &= \frac{P + \rho\, t(\Delta Z) E_{coal}}{Z + t(\Delta Z) E_{gas}} \\
[P + \rho\, t(\Delta Z) E_{coal}][Z + \Delta Z] &= [P + \Delta\mu][Z + t(\Delta Z) E_{gas}] \\
\rho &= \frac{[P + \Delta\mu][Z + t(\Delta Z) E_{gas}] - P[Z + \Delta Z]}{t(\Delta Z)[Z + \Delta Z] E_{coal}} \\
\rho &= \frac{\Delta\mu \cdot Z + t(\Delta Z)[P + \Delta\mu] E_{gas} - \Delta Z \cdot P}{t(\Delta Z)[Z + \Delta Z] E_{coal}} \\
\rho &= \frac{(P + \Delta\mu) E_{gas}}{(Z + \Delta Z) E_{coal}} + \frac{\Delta\mu \cdot Z - \Delta Z \cdot P}{(Z + \Delta Z) E_{coal}} \left[\frac{1}{t(\Delta Z)}\right]
\end{aligned}$$

Substituting $t(\Delta Z)$ from Equation (C4), and rearranging:

$$\begin{aligned}
\rho &= \frac{(P + \Delta\mu) E_{gas}}{(Z + \Delta Z) E_{coal}} + \frac{\Delta\mu \cdot Z - \Delta Z \cdot P}{(Z + \Delta Z) E_{coal}} \left[\frac{P \cdot E_{gas} - (Z + \Delta Z) E_{coal}}{\Delta Z \cdot P}\right] \\
\rho &= \frac{(P + \Delta\mu) E_{gas}}{(Z + \Delta Z) E_{coal}} + \frac{\Delta\mu \cdot Z \cdot E_{gas}}{\Delta Z (Z + \Delta Z) E_{coal}} - \frac{\Delta\mu \cdot Z}{\Delta Z \cdot P} + 1 - \frac{P \cdot E_{gas}}{(Z + \Delta Z) E_{coal}} \\
\rho &= 1 + \frac{\Delta\mu \cdot E_{gas}}{(Z + \Delta Z) E_{coal}} + \frac{\Delta\mu \cdot Z \cdot E_{gas}}{\Delta Z (Z + \Delta Z) E_{coal}} - \frac{\Delta\mu \cdot Z}{\Delta Z \cdot P}
\end{aligned}$$

Restoring the coal plant subscript j:

$$\Rightarrow \quad \rho_j(\Delta\mu_j, \Delta Z) = 1 + \frac{\Delta\mu_j}{\Delta Z} \left(\frac{E_{gas}}{E_{coal}} - \frac{Z}{P_j}\right) \tag{C7}$$

This expression shows that $\Delta\mu_j$ leads to incomplete pass-through via two channels. The first term adjusts for the wedge in emissions factors, while the second term rescales for the baseline difference in the fuel costs. This second term is j-specific, depending on plant j's coal price.

For expositional clarity, I simplify this derivation in the main text by removing heat rates (which cancel out in Equation (C6)). I parameterize Equation (C7) by substituting $\Delta\mu = \hat{\tau} \hat{M}_j \left(\frac{\text{MMBTU}}{\text{ton}}\right)_j \Delta Z$; set P_j equal to plant j's average delivered coal price in 2007–08 (the start of the fracking boom); and set Z equal to the 2007–08 average Henry Hub spot price.

Figure C1: Electricity supply with gas price decrease or carbon tax

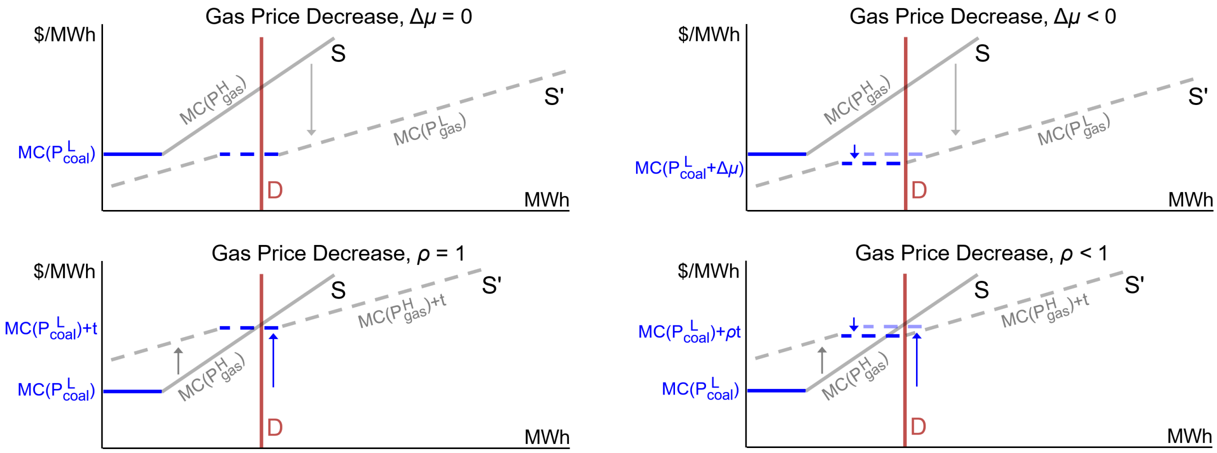

Notes: This figure shows how a gas price decrease mimics a carbon tax on the electricity sector in the short-run, using the same stylized electricity market as Figure 3. The top-left panel reproduces the top-left panel of Figure 3. My empirical results find that coal markups decrease due to a decrease in gas prices; I illustrate this decrease in markups ($\Delta\mu < 0$) in the top-right panel. Absent a gas price decrease, there exists a carbon tax (t) that yields the same supply curve as the top-left panel, vertically shifted upwards (i.e. bottom-left panel). Incomplete pass-through of that carbon tax (i.e. $\rho < 1$ in the bottom-right panel) can result in the same generation allocation under a (counterfactual) carbon tax as decreasing coal markups after a (factual) gas price drop.

I can rewrite Equation (C3) to incorporate marginal environmental compliance costs:

$$\frac{HR_{coal} \cdot (P + \Delta\mu + MC^{env}_{coal})}{HR_{gas} \cdot (Z + \Delta Z + MC^{env}_{gas})} = \frac{HR_{coal} \cdot (P + MC^{env}_{coal} + \rho\, t(\Delta Z) E_{coal})}{HR_{gas} \cdot (Z + MC^{env}_{gas} + t(\Delta Z) E_{gas})} \tag{C8}$$

MC^{env}_{fuel} represents the marginal costs of environmental compliance per MMBTU of fuel, as defined in Step 1 of Appendix A.1. The pass-through expression in Equation (C8) becomes:

$$\Rightarrow \quad \rho_j(\Delta\mu_j, \Delta Z) \;=\; 1 \;+\; \frac{\Delta\mu_j}{\Delta Z}\left(\frac{E_{gas}}{E_{coal}} - \frac{Z + MC^{env}_{gas,\,j}}{P_j + MC^{env}_{coal,\,j}}\right) \tag{C9}$$

Figure C4 estimates $\hat{\rho}_j$ inclusive of environmental costs, following Equation (C9). I parameterize $MC^{env}_{coal,\,j}$ as the 2007–08 plant-specific average of MC^{env}_{um} (from Equation (A1)); and $MC^{env}_{gas,\,j}$ as the 2007–08 average of $\sum_{g \in \text{PCA}_j} MC^{env}_{gm}$ (from Equation (A2)).[33]

The mapping from cost ratio (CR) to implicit carbon tax (t) relies on two key assumptions. First, electricity demand must be perfectly inelastic. In reality, U.S. electricity demand is close to perfectly inelastic, as there is extremely limited demand response capacity. Second, the marginal fuel in electricity markets must be either coal or gas. In reality, other technologies (e.g. diesel or hydro) are rarely marginal in U.S. regions with non-trivial coal generation. Under these two assumptions, electricity generation depends only on the *ordering* of plants along the supply curve. The "equivalent" carbon tax $t(\Delta Z)$ would produce the same ordering, yielding the same generation outcomes at higher electricity prices (since the tax would raising marginal costs). Figure C1 illustrates how a gas price decrease can yield the same *allocation* of generation as a carbon tax, comparing the top-left vs. bottom-left panels. Comparing the top-right vs. bottom-right panels, a decrease in coal markups ($\Delta\mu < 0$) can increase the allocation of coal generation in a way that mimics incomplete pass-through of a carbon tax to coal plants ($\rho < 1$).

33. I ignore non-fuel variable costs (e.g., labor) since (i) electricity production is Leontief in fuel (Fabrizio, Rose, and Wolfram (2007)), and non-fuel inputs are of second-order importance to marginal operations of fossil plants (Cicala (2022)); and (ii) reliable data on non-fuel variable costs are unavailable (see Appendix G.2.3).

The CR-to-t mapping also assumes that higher electricity prices would not alter plants' bidding strategies, due to either dynamic operating constraints (which Cullen (2015) finds to be second-order), exercise of market power (Mansur (2013)), or differential pass-through of shocks to fuel vs. carbon prices (Fabra and Reguant (2014)). Finally, this mapping only holds in the short-run, where both the stock of generators and their CO_2 emissions rates are fixed.

C.2 Sensitivity of pass-through results

Figure 7 plots $\hat{\rho}_j$ using my DD estimate for $TREAT_j = \hat{M}_j$ with $k = 3$ nearest neighbors (i.e., $\hat{\tau} = 2.94$ from Column (2) of Table 3) to parameterize Equation (C7). I weight each plant's marker in the scatter plot using the $k = 3$ nearest-neighbor weights (consistent with the sample of plants in the regression) times the total quantity of 2007–15 coal deliveries (scaling plants by size/importance). Figure C2 constructs the analogous scatter plot of $\hat{\rho}_j$ estimates with $k = 1$ nearest neighbors, using $\hat{\tau} = 3.66$ from Column (1) of Table 3. Figure C3 does the same for $k = 5$ nearest neighbors, using $\hat{\tau} = 3.01$ from Column (3) of Table 3. All three $\hat{\tau}$ estimates imply qualitatively similar distributions of $\hat{\rho}_j$, with greater dispersion for larger $\hat{\tau}$. Figure C4 plots my preferred $\hat{\rho}_j$ estimates against $\hat{\rho}_j$ estimates using Equation (C9), which includes marginal environmental compliance costs; this slightly compresses the distribution of $\hat{\rho}_j$ towards 1.

Figure C2: Estimated pass-through, for $k = 1$ nearest neighbors

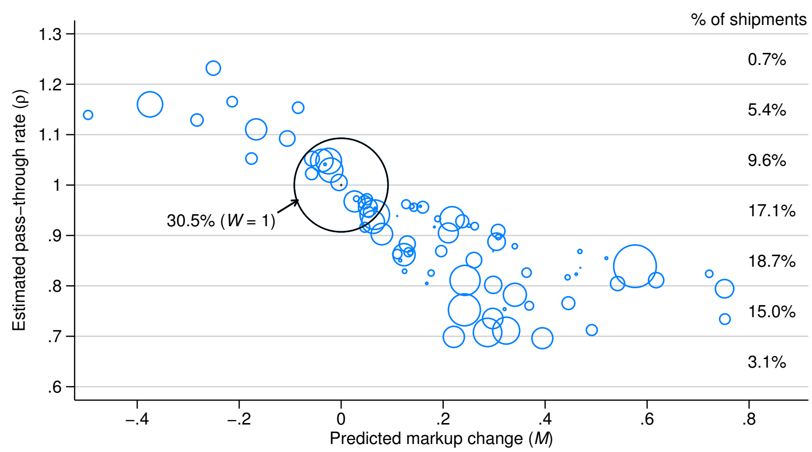

Notes: This figure is identical to Figure 7, except that I use $\hat{\tau}$ from Column (1) of Table 3, with $k = 1$ nearest-neighbor weights.

Figure C3: Estimated pass-through, for $k = 5$ nearest neighbors

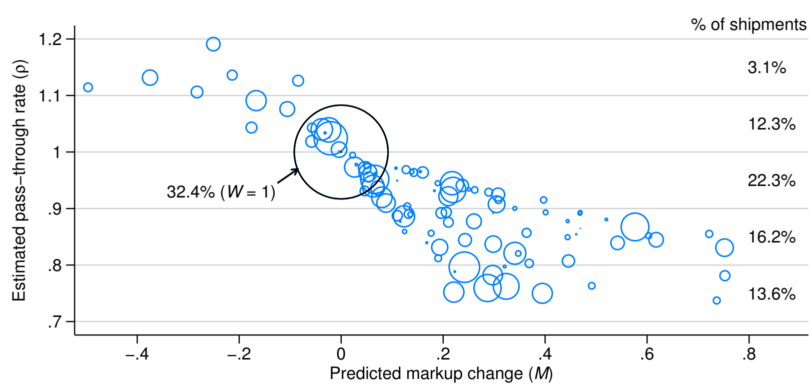

Notes: This figure is identical to Figure 7, except that I use $\hat{\tau}$ from Column (3) of Table 3, with $k = 5$ nearest-neighbor weights.

Figure C4: Accounting for environmental costs in $\hat{\rho}_j$

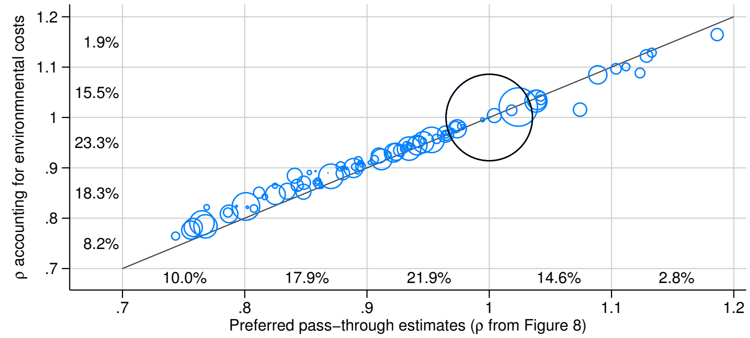

Notes: This figure plots my preferred $\hat{\rho}_j$ estimates from Figure 7 (horizontal axis) against $\hat{\rho}_j$ estimates that account for marginal environmental compliance costs (using Equation (C9), vertical axis). I plot the 45-degree line, and report percentages of 2007–15 coal-by-rail shipments in each $\hat{\rho}_j$ grid cell along each axis (suppressing the 32.8% label for shipments to plants with a barge option).

Figure C5: Estimating $\hat{\rho}_j$ using $\hat{\tau}$ estimates in \$/MMBTU

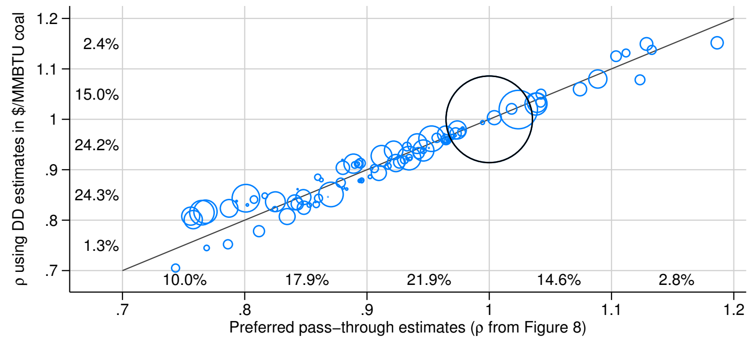

Notes: This figure plots my preferred $\hat{\rho}_j$ estimates from Figure 7 (horizontal axis) against $\hat{\rho}_j$ estimates based on a DD model with coal prices in \$/MMBTU (vertical axis). I plot the 45-degree line, and report percentages of 2007–15 coal-by-rail shipments in each $\hat{\rho}_j$ grid cell along each axis (suppressing the 32.8% label for shipments to plants with a barge option).

Figure C6: Estimating $\hat{\rho}_j$ using state-specific gas prices

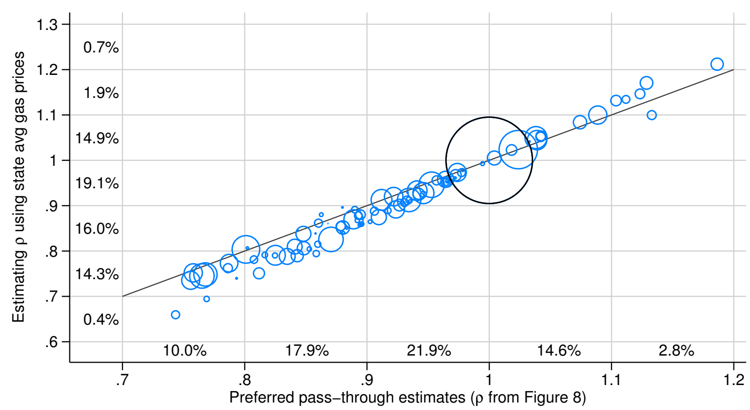

Notes: This figure plots my preferred $\hat{\rho}_j$ estimates from Figure 7 (horizontal axis) against $\hat{\rho}_j$ estimates using state-specific gas prices (vertical axis). The latter use $\hat{\tau}$ estimates that replace Z_m in Equation (5) with state-month average gas prices paid by power plants; they likewise parameterize Equation (C7) using 2007–08 average delivered prices to gas plants in the same state as plant j. I plot the 45-degree line, and report percentages of 2007–15 coal-by-rail shipments in each $\hat{\rho}_j$ grid cell along each axis (suppressing the 32.8% label for shipments to plants with a barge option).

For Figure C5, I estimate Equation (5) after converting the dependent variable (P_{ojms}) from \$/ton to \$/MMBTU of coal. The resulting $\hat{\tau}$ estimate is already in the right units, meaning that I can substitute $\Delta\mu = \hat{\tau}\hat{M}_j\Delta Z$ into Equation (C7). This yields slightly different $\hat{\rho}_j$ estimates, due to heterogeneity in coal's MMBTU/ton conversion rate. I prefer $\hat{\rho}_j$ estimates that convert to \$/MMBTU *after* estimating Equation (5), since per-ton prices more accurately reflect variation in coal shipping costs (and since rail carriers price by ton).

Finally, Figure C6 replaces Henry Hub spot prices in both Equations (5) and (C7) with average prices paid by gas power plants in each state. This better captures cross-sectional variation in gas prices paid by coal plant j's competitors, but makes the econometric identification of $\hat{\tau}$ less clean (e.g., due to potential regional gas price endogeneity). The resulting $\hat{\tau}$ estimate increases in magnitude (see Panel G of Figure E8), and adding pipeline costs to Z increases the last term in Equation (C7). Both factors pull the resulting $\hat{\rho}_j$ estimates further away from 1.

C.3 Accounting for natural gas markups

In formulating my estimates of pass-through of an implicit carbon tax for coal plants, I assume full pass-through of the tax for gas plants. This does not necessarily require competitive pricing in natural gas transportation; rather, it assumes that any non-marginal-cost pricing in the gas pipeline network does not respond to shocks to natural gas demand.

I can modify Equation (C6) to allow for incomplete pass-through of the implicit carbon tax for gas plants (denoted as ψ):

$$\frac{HR_{coal} \cdot (P + \Delta\mu)}{HR_{gas} \cdot (Z + \Delta Z)} = \frac{HR_{coal} \cdot (P + \rho\, t(\Delta Z) E_{coal})}{HR_{gas} \cdot (Z + \psi\, t(\Delta Z) E_{gas})} \tag{C10}$$

Solving for ρ_j:

$$\Rightarrow \quad \rho_j(\Delta\mu_j, \Delta Z, \psi) = 1 + \frac{\Delta\mu_j}{\Delta Z}\left[\frac{E_{gas}}{E_{coal}}\left(\frac{Z + \psi\Delta Z}{Z + \Delta Z}\right) - \frac{Z}{P_j}\right] + \frac{E_{gas}(\psi - 1)P_j}{E_{coal}(Z + \Delta Z)} \tag{C11}$$

This is equivalent to Equation (C7) for $\psi = 1$ (i.e. full pass-through for gas plants). Differentiating Equation (C11) by ψ:

$$\frac{\partial\rho_j}{\partial\psi} = \frac{E_{gas}}{E_{coal}}\left(\frac{P_j + \Delta\mu_j}{Z + \Delta Z}\right) \tag{C12}$$

This shows that (i) $\hat{\rho}_j$ is increasing in the pass-through rate for gas plants; and (ii) a larger negative ΔZ magnifies this effect (shrinking the denominator faster than the numerator).

Either $\psi < 1$ or $\psi > 1$ is plausible. Incomplete pass-through would be consistent with average-cost pricing, where pipeline operators recover a portion of their fixed costs volumetrically. Pass-through greater than 1 would be consistent with pipeline operators raising markups as lower gas prices made gas plants more inframarginal in electricity supply (i.e., the reverse of the bottom-right panel of Figure 3). However, Figures C7–C8 demonstrate that relaxing the assumption of $\psi = 1$ has only a small effect on my $\hat{\rho}_j$ estimates. If I set $\psi = 0.9$ (as in Figure C7), $\hat{\rho}_j$ decreases for all coal plants, but only slightly; if I set $\psi = 1.1$ (as in Figure C8), $\hat{\rho}_j$ increases for all coal plants, but also only slightly.

My DD estimates also suggest that $\psi \neq 1$ would be unlikely to qualitatively impact my results: when I estimate Equation (5) replacing the Henry Hub price with average prices paid by gas power plants *that are inclusive of endogenous changes to gas markups*, my $\hat{\tau}$ estimate

Figure C7: Estimating $\hat{\rho}_j$ assuming $\psi = 0.9$ for gas plants

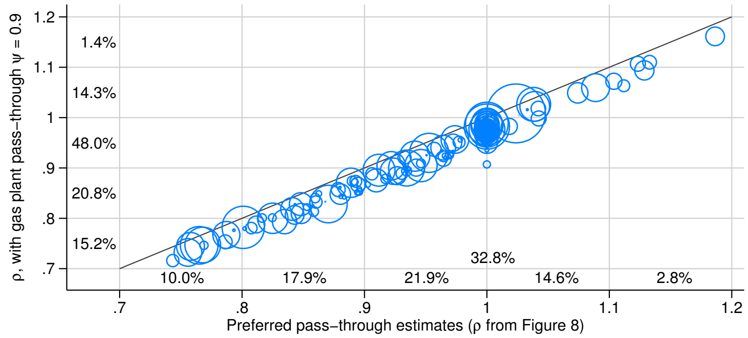

Notes: This figure plots my preferred $\hat{\rho}_j$ estimates from Figure 7 (horizontal axis) against $\hat{\rho}_j$ estimates assuming $\psi = 0.9$ for gas plants (vertical axis). The latter parameterize Equation (C11) assuming $\psi = 0.9$ and $\Delta Z = -4$, but are otherwise identical. I plot the 45-degree line, and report percentages of 2007–15 coal-by-rail shipments in each $\hat{\rho}_j$ grid cell along each axis. Here, I plot a separate marker for each plant with $W = 1$, since $\hat{\rho}_j < 1$ for these plants when $\psi < 1$.

Figure C8: Estimating $\hat{\rho}_j$ assuming $\psi = 1.1$ for gas plants

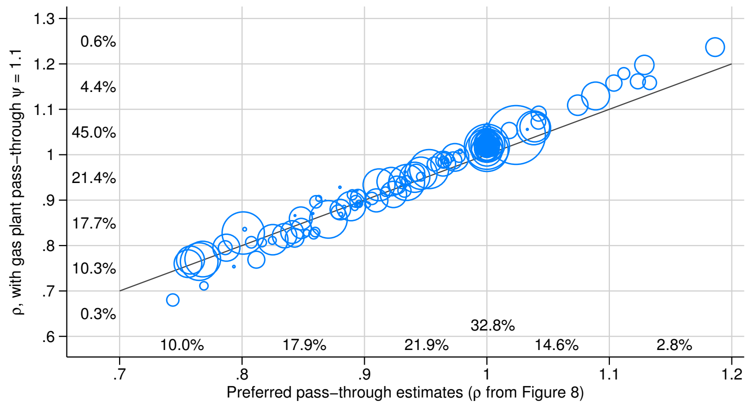

Notes: This figure plots my preferred $\hat{\rho}_j$ estimates from Figure 7 (horizontal axis) against $\hat{\rho}_j$ estimates assuming $\psi = 1.1$ for gas plants (vertical axis). The latter parameterize Equation (C11) assuming $\psi = 1.1$ and $\Delta Z = -4$, but are otherwise identical. I plot the 45-degree line, and report percentages of 2007–15 coal-by-rail shipments in each $\hat{\rho}_j$ grid cell along each axis. Here, I plot a separate marker for each plant with $W = 1$, since $\hat{\rho}_j > 1$ for these plants when $\psi > 1$.

slightly increases in magnitude (see Figure E8, Panel G). Figure C6 shows that using these markup-inclusive gas prices pulls $\hat{\rho}_j$ even further below 1 for the majority of coal plants.

Finally, I can directly estimate whether the prices of gas deliveries to power plants respond to changes in these plants' competitiveness. The following DD specification is analogous to Equation (5), for monthly deliveries to gas plants (indexed by g):

$$P_{gm} = \tau^C GW^{\text{coal}}_{g(m-L)} + \sum_{\ell=0}^{L-1} \tau^C_\ell GW^{\text{coal}}_{g(m-\ell)} + \tau^H P^{\text{hub}}_{g(m-L)} + \sum_{\ell=0}^{L-1} \tau^H_\ell P^{\text{hub}}_{g(m-\ell)} + \eta_g + \delta_m + \varepsilon_{gm} \quad \text{(C13)}$$

P_{gm} is the delivered gas price to plant g in month m. GW^{coal}_{gm} is the operating coal-fired capacity in plant g's PCA: decreases in GW^{coal}_{gm} should weakly improve plant g's competitiveness, making it more inframarginal in electricity supply. P^{hub}_{gm} is the natural gas price at the trading hub nearest to plant g: this serves the same function as commodity and shipping controls in Equation (5), by removing variation in natural gas costs that is common across (nearby) plants. η_g are plant fixed effects, and δ_m are month-of-sample fixed effects. τ^C captures the cumulative DD effect of coal capacity changes on delivered gas markups, for plants in different PCAs.[34]

Table C1 reports cumulative DD effects over $L \in \{24, 36, 48\}$ months. Differential changes in hub-specific gas trading prices have a strong effect on the prices paid by power plants, and I cannot reject full pass-through (i.e. $\hat{\tau}^H = 1$). I find no evidence that differential coal retirements across PCAs (which improve gas plants' competitiveness) lead to differential changes in markups, and I can reject $|\hat{\tau}^C| > 0.02$ (i.e. 0.5% of the average gas price of \$5.32–5.53/MMBTU). These results provide empirical support for my assumption of $\psi = 1$.

Table C1: Markup DD results for natural gas plants

	Outcome: delivered gas price (\$/MMBTU)		
	(1)	(2)	(3)
GW coal capacity in g's PCA ($\hat{\tau}^C$)	−0.001	0.000	−0.001
	(0.009)	(0.008)	(0.011)
Price at g's nearest gas hub ($\hat{\tau}^H$)	0.994***	1.100***	1.309***
	(0.134)	(0.148)	(0.184)
Monthly lags (L)	24	36	48
Mean of dep var (\$/MMBTU)	5.533	5.485	5.322
Plants	485	461	438
Observations	35,134	30,262	25,797

Notes: Each regression estimates Equation (C13) using a monthly panel of all plants that reported natural gas deliveries between 2002–2015. Regressions control for plant and month-of-sample fixed effects; results are similar if I replace plant fixed effects with state fixed effects and control for each plant's distance to nearest gas hub. I weight observations by the quantity of gas transacted. Standard errors are clustered by plant. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

C.4 Incidence and alternate market structures

Weyl and Fabinger (2013, p. 547) derive the following expression for the incidence (I) of a tax (t)—or the ratio of changes in consumer surplus (CS) vs. producer surplus (PS)—in a symmetric oligopoly:

$$I = \frac{dCS/dt}{dPS/dt} = \frac{\rho}{1 - (1 - \theta/N)\,\rho} \tag{C14}$$

Here, ρ is the pass-through rate of the tax, θ is the conduct parameter, and N is the number of symmetric firms in the market.[35] As pass-through becomes more incomplete (i.e., as ρ decreases from 1), incidence decreases and consumers pay proportionately less of the tax burden. For a given pass-through rate ρ, a less competitive market structure (i.e., greater θ/N) also implies lower incidence, because producers stand to lose more under a tax if they are already extracting more oligopoly rents.[36]

I assume all plants faces either an effective rail monopoly (i.e. captive, $N_j = 1$) or an effective symmetric rail duopoly (i.e. non-captive, $N_j = 2$). I also assume that market conduct

34. The correlation between GW^{coal}_{gm} and P^{hub}_{gm} is 0.04. This suggests that differential changes in gas competitiveness are not influencing prices at gas trading hubs. Most gas plants do not purchase directly from tradings hubs, however I use hub-specific prices to control for cross-sectional variation in pipeline costs and congestion.

35. Using my notation, θ/N corresponds to θ in the notation of Weyl and Fabinger.

36. That is, $\partial I/\partial\rho > 0$; and $\partial I/\partial(\theta/N) < 0$, unless $I < 0$, which can only occur if $\rho > 1$.

is Cournot (i.e. $\theta_j = 1$) for plants without a coal-by-barge option, and competitive (i.e. $\theta_j = 0$) for plants with a coal-by-barge option (i.e. $\theta_j = 0$, $\rho_j = 1$). These assumptions reduce Equation (C14) to three possible mappings between pass-through and incidence:

$$I_j = \begin{cases} \rho_j & \text{if } N_j = 1\,,\; W_j = 0 \\ \frac{\rho_j}{1-\rho_j/2} & \text{if } N_j = 2\,,\; W_j = 0 \\ \infty & \text{if } W_j = 1 \end{cases} \tag{C15}$$

If pass-through is incomplete (i.e. $\rho_j < 1$), then for a given pass-through rate, a less competitive market (i.e. lower N_j) reduces the share of the tax borne by coal plants (i.e. consumers) relative to rail carriers (i.e. producers). For example, $\rho_j = 0.8$ would imply $I_j = 0.8$ for a plant with $N_j = 1$ and $I_j = 1.3$ for a plant with $N_j = 2$. Infinite incidence occurs under perfect competition, since perfectly elastic supply implies $\rho_j = 1$ and $dPS/dt = 0$.

I can reformulate incidence to summarize the proportion of the tax burden in the coal-by-rail market that is borne by coal plants (i.e. consumers):

$$\frac{I_j}{1+I_j} = \frac{\left(\frac{dCS}{dt}\right)_j}{\left(\frac{dCS}{dt}\right)_j + \left(\frac{dPS}{dt}\right)_j} = \frac{\rho_j}{1+(\theta_j/N_j)\,\rho_j} = \begin{cases} \frac{\rho_j}{1+\rho_j} & \text{if } N_j = 1\,,\; W_j = 0 \\ \frac{\rho_j}{1+\rho_j/2} & \text{if } N_j = 2\,,\; W_j = 0 \\ 1 & \text{if } W_j = 1 \end{cases} \tag{C16}$$

For a pass-through rate of $\rho_j = 0.8$, a plant could bear 44% (if $N_j = 1$, $W_j = 0$), 57% (if $N_j = 2$, $W_j = 0$), or 80% (if $W_j = 1$) of the lost surplus in coal markets. Importantly, the full tax burden would depend on the extent to which coal plants could pass on marginal emissions costs via higher wholesale electricity prices. If emissions tax pass-through in wholesale electricity markets is 1 (consistent with Fabra and Reguant (2014)), then a carbon tax *could* increase profits for coal plants that are relatively clean/efficient and have low pass-through rates in coal markets.

While the fracking boom simulated the effect of a carbon tax on the *relative* costs of coal vs. gas plants, it had the opposite effect on electricity prices, as low gas prices caused electricity prices to fall (Linn and Muehlenbachs (2018)). Lower electricity prices meant that coal plants were very unlikely to be "winners" in the fracking boom, though incomplete pass-through in coal shipping likely caused certain coal plants to be "smaller losers". Figure C9 shows that across the distribution of predicted incidence, capacity factors fell by an average of 21 percentage points; 91% of coal plants sold less electricity during 2011–15 than prior to the fracking boom.

My theory model assumes that rail carriers buy coal from a perfectly competitive mining sector. While this greatly simplifies the derivations in Appendix B.1, this assumption is not crucial for evaluating markups for coal deliveries.[37] However, the upstream market structure does impact how mines and railroads share the tax burden. Given that coal is both spatially and physically heterogeneous, the assumptions of perfect competition likely do not hold.

I assume that C_j in Equations (1) and (B1) is exogenous, implying perfectly elastic coal supply (a standard assumption for *homogeneous* commodity markets). This assumes that coal mines earn zero economic rents, and would incur no lost profits due to a downstream carbon tax. However, mines would face non-zero tax burden under any of three alternate market structures.

37. Equations (4)–(5) control for the average *equilibrium* coal price by county-year, taking the market structure of the mining sector as given.

Figure C9: Predicted incidence and changes in coal plants' capacity factors

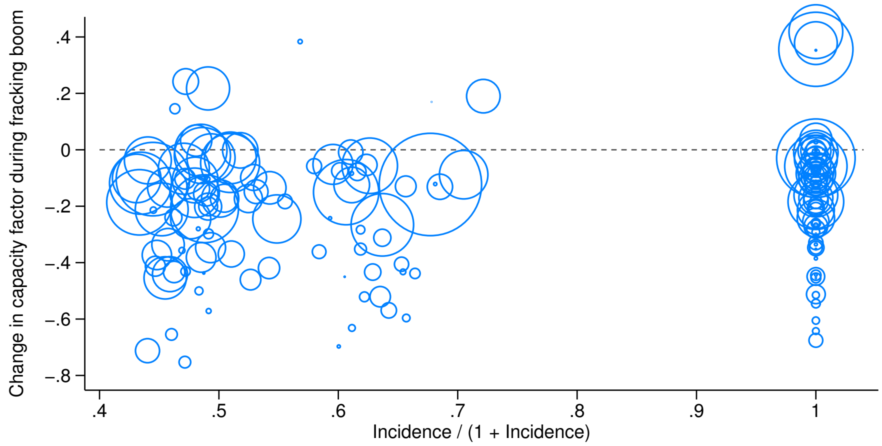

Notes: This figure plots predicted share of implied tax burden (from Figure 8) against plants' observed change in average capacity factor from 2002–06 (pre-fracking boom) to 2011–15. A change of $\Delta\overline{CF}_j = -0.4$ would be consistent with a plant's capacity factor falling from 0.8 to 0.4; $\Delta\overline{CF}_j < 0$ for 91% of plants. I calculate $\Delta\overline{CF}_j$ setting $CF_j = 0$ if plant j has retired or stopped producing. Markers weight observations (i.e. plants) using the product of $k = 3$ nearest-neighbor weights and total 2007–15 coal deliveries.

First, suppose that coal mines and rail carriers coordinate, behaving as vertically integrated monopolists. This may occur without formal profit-sharing, as geographically isolated mines depend heavily on rail carriers to transact their coal, while rail carriers with sunk track investments stand to gain considerable profits by cooperating with mines. In this case, I could rewrite the rail carrier's profit function replacing coal quantity (q_{ioj}) with the mine's production function, and replacing constant marginal cost (C_j) with total mining input costs. This would cause coal mines and rail carriers to jointly share the burden of a carbon tax.

Second, suppose that coal mines can exert market power at the mine-mouth when selling to rail carriers. This may occur if a few large firms dominate mining operations (as in the Powder River Basin; Atkinson and Kerkvliet (1986)), if multiple rail carriers compete to purchase coal from a single mining firm, or if mines can sell to non-rail intermediaries (e.g. coal-by-barge firms). In this case, double marginalization would shift delivered coal prices even further from the competitive benchmark (ignoring externalities).[38] This would create a second opportunity for incomplete pass-through of a cost shock (or a carbon tax), but adjustments in markups would not be coordinated along the coal supply chain. If mines responded by reducing mine-mouth markups, this would reduce the tax burden borne by rail carriers. Mines earning market power rents would also stand to lose under a carbon tax.

Third, suppose that rail carriers can exert monopsony power at the mine-mouth. This may occur if coal mines are captive to a single rail carrier with strong bargaining power: whereas these mines depend on revenue a single product, diversified rail carriers may divert resources (e.g. locomotives, labor) to other profitable shipping opportunities. In this case, rail carriers could adjust prices both at the mine-mouth and at the power plant; they would likely bear an even greater tax burden for having extracted rents on both sides of the market. If a downstream carbon tax caused rail carriers to raise mine-mouth prices (i.e. incomplete pass-through at the mine-mouth), the effects on coal mining profits would be theoretically ambiguous—depending on whether mine-mouth prices increased by enough to offset the reduction in coal quantity.

38. Alexandrov, Pittman, and Ukhaneva (2018) find no empirical evidence of double marginalization between railroads, in cases where multiple rail carriers own segments along the same shipping route.

D CO_2 emissions counterfactuals

Here, I describe how I estimate CO_2 emissions counterfactuals in Section 7.3. I start by estimating the relationship between fuel costs and generation for each coal unit. Then, I use the fitted models to predict generation under two counterfactual scenarios: (i) if the fracking boom never happened; and (ii) if the fracking boom happened, but coal markups remained fixed (i.e., $\frac{d\mu_j}{dZ} = 0$ and $\rho_j = 1$). Comparing across these two scenarios, I can calculate CO_2 abatement due to short-run coal-to-gas substitution, both with and without changes to coal markups.

D.1 Counterfactuals algorithm

Step 1: I construct counterfactual Henry Hub gas prices in the absence of the fracking boom, using historic NYMEX monthly futures prices as of December 2008. This follows Holladay and LaRiviere (2017), who estimate a structural break in Henry Hub spot prices on December 5, 2008. The left panel of Figure D1 compares actual prices to these futures prices, which are close to pre-Recession levels (and would be roughly constant absent seasonal variation).[39] To construct a counterfactual price Z_{gd}^{CF} for each gas plant g on day d, I adjust for the daily wedge between the price at each plant's nearest trading hub and the (factual) Henry Hub price.[40]

Figure D1: Counterfactual fuel price scenarios

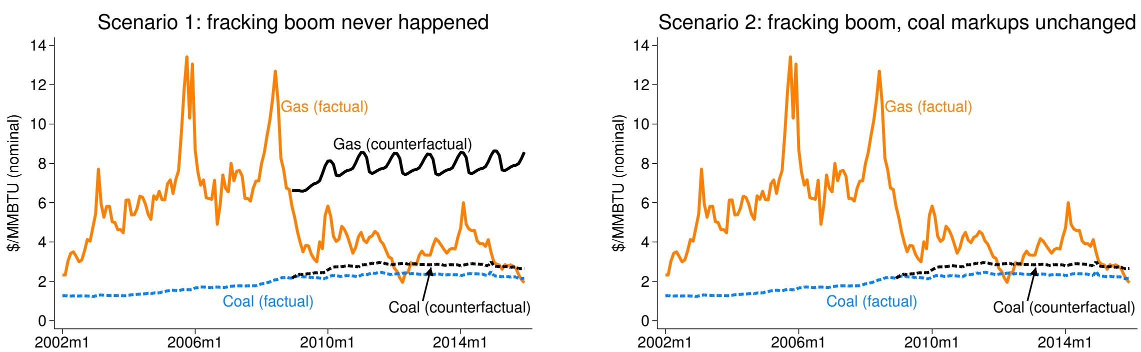

Notes: The left panel plots my first scenario: gas prices absent the fracking boom, and coal markups that did not adjust to the (now absent) coal demand shock. The right panel plots my second scenario: observed gas prices, but coal markups that did not adjust to the coal demand shock. I plot Henry Hub futures prices as of December 2008; in this figure, higher coal prices are illustrative.

Step 2: Figure D1 illustrates average coal prices absent the fracking boom—had rail carriers not decreased coal markups. I predict counterfactual coal prices P_{ojms}^{CF} using the fitted model estimating Equation (5) with $TREAT_j = \hat{M}_j$, and replacing factual Henry Hub prices (Z_m) with counterfactual Henry Hub prices (i.e. the solid black line in the left panel of Figure D1).

My DD regressions include only the subset of coal plants with nearest-neighbor matches and unmasked coal prices. Here, I predict CO_2 abatement across *all* coal plants. I regress P_{ojms}^{CF} on the interaction of factual coal prices (P_{ojms}) and predicted markup changes ($\hat{M}_j$), estimating a separate coefficient for each sample month.[41] Taking predicted values from this regression, I populate P_{jm}^{CF} for the 41% of utility-owned coal-by-rail plants that aren't in Column (2) of Table 3. Taking predicted values from the same regression removing P_{ojms} from the right-hand

39. Importantly, my counterfactual analysis begins *after* the Recession-related price spike in gas prices.

40. This preserves network congestion and pipeline costs in Z_{gd}^{CF}. I also add the average gap between hub prices and delivered prices, to account for last-mile pipeline charges (see discussion in Appendix G.5).

41. I include fixed effects for plant regions, as well as $\mathbf{C}_{ojms}$, $S(\mathbf{T}_{ojms})$, $\mathbf{X}_{jm}$, and η_o from Equation (4).

side, I populate P^{CF}_{jm} for the 23% of non-utility-owned coal-by-rail plants with masked prices. I populate $P^{CF}_{jm} = P_{ojms}$ for non-rail coal shipments, assuming no markup changes.[42]

Step 3: I construct factual and counterfactual coal-to-gas cost ratios following Step 1 in Appendix A.1. However, I now average *both* coal and gas marginal costs across all generating units of each fuel type within each PCA, replacing Equation (A1) with:

$$MC^{coal}_{um} \equiv \sum_{j \in \text{PCA}_u} \left(\frac{Q^{elec}_{jm} \cdot HR_{jm} \cdot (P_{jm} + MC^{env}_{jm})}{\sum_{j \in \text{PCA}_u} Q^{elec}_{jm}} \right) \tag{D1}$$

This facilitates counterfactuals where *many* plants' coal prices change. I construct counterfactual cost ratios replacing factual with counterfactual prices for each fuel:

$$MC^{coal,CF}_{um} \equiv \sum_{j \in \text{PCA}_u} \left(\frac{Q^{elec}_{jm} \cdot HR_{jm} \cdot (P^{CF}_{jm} + MC^{env}_{jm})}{\sum_{j \in \text{PCA}_u} Q^{elec}_{jm}} \right) \tag{D2}$$

$$MC^{gas,CF}_{ud} \equiv \sum_{g \in \text{PCA}_u} \left(\frac{Q^{elec}_{gm} \cdot HR_{gm} \cdot (Z^{CF}_{gd} + MC^{env}_{gm})}{\sum_{g \in \text{PCA}_u} Q^{elec}_{gm}} \right) \tag{D3}$$

Then, I construct three cost ratios:

$$CR_{ud} = \frac{MC^{coal}_{um}}{MC^{gas}_{ud}} \quad , \qquad CR^{NO\Delta Z}_{ud} = \frac{MC^{coal,CF}_{um}}{MC^{gas,CF}_{ud}} \quad , \qquad CR^{NO\Delta\mu}_{ud} = \frac{MC^{coal,CF}_{um}}{MC^{gas}_{ud}} \tag{D4}$$

CR_{ud} uses factual gas prices and factual coal prices. $CR^{NO\Delta Z}_{ud}$ aligns with my first scenario, using both counterfactual gas and coal prices (i.e., if the fracking boom never happened, or "$NO\Delta Z$"). $CR^{NO\Delta\mu}_{ud}$ aligns with my second scenario, using factual gas prices and counterfactual coal prices (i.e., if the fracking boom happened but markups never changed, or "$NO\Delta\mu$").

Step 4: For each coal unit u, I estimate the following time-series regression, for each day d, from 2002 to 2015:

$$MWH_{ud} = \sum_b \alpha_{ub} \mathbf{1}[G_{ud} \in b] + \sum_b \gamma_{ub} \mathbf{1}[G_{ud} \in b] \cdot CR_{ud} + \mathcal{SP}(CR_{ud}; \zeta_u) + \xi_u \mathbf{G}_{ud} + \omega_{ud} \tag{D5}$$

This specification is similar to Equation (6), but differs in several key ways:

- It uses daily (not hourly) observations, since I no longer seek to integrate over hours.[43]
- It uses daily net generation (MWH_{ud}) as the dependent variable, instead of capacity factor. This is largely a normalization, and facilitates translating into predicted emissions.
- Its cost ratio (CR_{ud}) averages across all coal units in unit u's PCA, as I now want to accommodate price changes across *many* coal plants (rather than idiosyncratic price changes).

42. These predicted counterfactual prices are likely too mismeasured to estimate markup changes, but are better suited for estimating the conditional probability of a coal unit operating (using Equation (D5) below).

43. For estimating coal demand, hourly observations allow me to more accurately discretize each unit's capacity factor. For counterfactuals, estimating Equation (D5) at the daily level reduces computation time without meaningfully changing the relationship between fuel prices and unit u's predicted generation.

- Following Cullen and Mansur (2017), it includes a cubic spline of the cost ratio, $\mathcal{SP}(CR_{ud}; \zeta_u)$. This more flexibly models the effect of relative fuel price changes on unit u's generation.[44]

- It includes only month-of-year fixed effects. This avoids removing year-on-year variation in fuel prices, which is important for these two counterfactual scenarios.

Step 5: I store predicted values ($\widehat{MWH}_{ud}$) from Equation (D5), estimated using CR_{ud}. Then, I predicted counterfactual generation ($\widehat{MWH}_{ud}^{NO\Delta Z}$, $\widehat{MWH}_{ud}^{NO\Delta\mu}$), by plugging two counterfactual cost ratios ($CR_{ud}^{NO\Delta Z}$, $CR_{ud}^{NO\Delta\mu}$) into this fitted model.

Step 6: I convert $\widehat{MWH}_{ud}$, $\widehat{MWH}_{ud}^{NO\Delta Z}$, and $\widehat{MWH}_{ud}^{NO\Delta\mu}$ into $\widehat{CO2}_{ud}$, $\widehat{CO2}_{ud}^{NO\Delta Z}$, and $\widehat{CO2}_{ud}^{NO\Delta\mu}$, multiplying by unit u's monthly CO_2 emissions rate.

Step 7: I sum factual and counterfactual coal generation and coal emissions across all units in each month, for all months between December 2008 and December 2015:

$$\widehat{\mathbf{MWH}}_m^{coal} \equiv \sum_{d\in m}\sum_u \widehat{MWH}_{ud} \;, \qquad \widehat{\mathbf{CO2}}_m^{coal} \equiv \sum_{d\in m}\sum_u \widehat{CO2}_{ud} \tag{D6}$$

$$\widehat{\mathbf{MWH}}_m^{coal,NO\Delta Z} \equiv \sum_{d\in m}\sum_u \widehat{MWH}_{ud}^{NO\Delta Z} \;, \qquad \widehat{\mathbf{CO2}}_m^{coal,NO\Delta Z} \equiv \sum_{d\in m}\sum_u \widehat{CO2}_{ud}^{NO\Delta Z} \tag{D7}$$

$$\widehat{\mathbf{MWH}}_m^{coal,NO\Delta\mu} \equiv \sum_{d\in m}\sum_u \widehat{MWH}_{ud}^{NO\Delta\mu} \;, \qquad \widehat{\mathbf{CO2}}_m^{coal,NO\Delta\mu} \equiv \sum_{d\in m}\sum_u \widehat{CO2}_{ud}^{NO\Delta\mu} \tag{D8}$$

Step 8: I predict counterfactual natural gas emissions in each month by replacing changes in coal generation with gas generation on a 1-for-1 basis, and multiplying by the average CO_2 emissions rate for combined-cycle gas plants from December 2008 to December 2015 ($\mathbf{E}_m^{gas}$):

$$\widehat{\mathbf{CO2}}_m^{gas,NO\Delta Z} \equiv \mathbf{E}_m^{gas} \times \left[\mathbf{MWH}_m^{gas} - \left(\widehat{\mathbf{MWH}}_m^{coal} - \widehat{\mathbf{MWH}}_m^{coal,NO\Delta Z}\right)\right] \tag{D9}$$

$$\widehat{\mathbf{CO2}}_m^{gas,NO\Delta\mu} \equiv \mathbf{E}_m^{gas} \times \left[\mathbf{MWH}_m^{gas} - \left(\widehat{\mathbf{MWH}}_m^{coal} - \widehat{\mathbf{MWH}}_m^{coal,NO\Delta\mu}\right)\right] \tag{D10}$$

Step 9: I sum total CO_2 emissions for both fuels, from December 2008 to December 2015:

$$\widehat{\mathbf{CO2}} \equiv \sum_m \left[\widehat{\mathbf{CO2}}_m^{coal} + \mathbf{CO2}_m^{gas}\right] \tag{D11}$$

$$\widehat{\mathbf{CO2}}^{NO\Delta Z} \equiv \sum_m \left[\widehat{\mathbf{CO2}}_m^{coal,NO\Delta Z} + \widehat{\mathbf{CO2}}_m^{gas,NO\Delta Z}\right] \tag{D12}$$

$$\widehat{\mathbf{CO2}}^{NO\Delta\mu} \equiv \sum_m \left[\widehat{\mathbf{CO2}}_m^{coal,NO\Delta\mu} + \widehat{\mathbf{CO2}}_m^{gas,NO\Delta\mu}\right] \tag{D13}$$

44. I use cubic splines with 6 knots, however the number of knots does not affect the estimation results.

Step 10: I calculate realized abatement under the fracking boom as the percent reduction in realized CO_2 emissions, compared to the no-fracking counterfactual:

$$\mathbf{ABATE}^{REALIZED} \;=\; \frac{\widehat{\mathbf{CO2}}^{NO\Delta Z} - \widehat{\mathbf{CO2}}}{\widehat{\mathbf{CO2}}^{NO\Delta Z}} \;=\; 0.0470 \tag{D14}$$

I calculate potential abatement under the fracking boom as the percent reduction counterfactual CO_2 emissions for $\Delta\mu = 0$, compared to the no-fracking counterfactual:

$$\mathbf{ABATE}^{POTENTIAL} \;=\; \frac{\widehat{\mathbf{CO2}}^{NO\Delta Z} - \widehat{\mathbf{CO2}}^{NO\Delta\mu}}{\widehat{\mathbf{CO2}}^{NO\Delta Z}} \;=\; 0.0521 \tag{D15}$$

Based on these predictions and calculations, decreasing coal markups eroded roughly 10% of the potential CO_2 abatement of the fracking boom:

$$1 - \frac{\mathbf{ABATE}^{REALIZED}}{\mathbf{ABATE}^{POTENTIAL}} \;\approx\; 1 - \frac{0.0470}{0.0521} \;=\; 0.0965 \tag{D16}$$

I monetize this eroded abatement (i.e., $\left[\widehat{\mathbf{CO2}}^{NO\Delta Z} - \widehat{\mathbf{CO2}}^{NO\Delta\mu}\right] - \left[\widehat{\mathbf{CO2}}^{NO\Delta Z} - \widehat{\mathbf{CO2}}\right]$) using the most up-to-date estimates of the social cost of carbon. Rennert et al. (2022) report a central estimate of \$185 per metric ton of CO_2, which the EPA has recently incorporated into its framework for regulatory impact analysis of greenhouse gas regulations (U.S. EPA (2022)). Previously, the EPA had assumed a value closer to \$50 per metric ton of CO_2, which aligned with pre-2016 integrated assessment modeling (Interagency Working Group on Social Cost of Greenhouse Gases (2016)).

D.2 Sensitivities and interpretation

I estimate several alternate version of Equation (D5), and I report these counterfactual sensitivities in Table D1. First, for more flexibility, I add cubic splines in $G_d \times CR_{ud}$, G_d, and daily maximum temperature; this has little effect on my counterfactual predictions.[45] Next, I replace month-of-year fixed effects with quarter-of-year fixed effects, to match Equation (6); this yields similar counterfactual predictions. I also add year fixed effects to control for medium-to-long run changes in plant operations; this absorbs most of the identifying gas price variation, yielding much smaller estimates of CO_2 abatement (2.6% vs. 4.7%) but implying a greater share of abatement eroded (11.3% vs. 9.7%). Finally, I estimate an hourly version of Equation (D5) to align with my demand estimation; this yields quite similar results.

How do these magnitudes compare to the existing literature on the fracking boom? Using similar time-series design at the interconnection level, Cullen and Mansur (2017) estimate that a tax of \$22/metric ton CO_2 would have yielded 4.9% reductions in daily CO_2 emissions.[46] Taking my derived expression for the implicit tax (Equation (C4)) and plugging in $\Delta Z = -4$

45. This matches Cullen and Mansur (2017)'s main specification, which includes cubic splines in the coal-to-gas cost ratio, total system load, and temperature.

46. Cullen and Mansur (2017) use short tons, while I use metric tons: \$20/ton CO_2 = \$22/metric ton CO_2.

Table D1: CO_2 counterfactuals under alternate versions of Equation (D5)

	Realized abatement		Potential abatement		Percent
	M tonnes	Percent	M tonnes	Percent	eroded
Preferred Equation (D5)	661.5	4.7%	732.1	5.2%	9.7%
Spline of $G_d \times CR_{ud}$	681.0	4.8%	750.9	5.3%	9.3%
Splines of $G_d \times CR_{ud}$, G_d, temp.	678.0	4.8%	747.3	5.3%	9.3%
Quarter-of-year fixed effects	654.3	4.7%	724.4	5.2%	9.7%
Adding year fixed effects	350.2	2.6%	395.0	2.9%	11.3%
Hourly (not daily) observations	651.3	4.6%	722.0	5.1%	9.8%

Notes: The top row reports counterfactual results using my preferred version of Equation (D5). The second row replaces the interacted sum (i.e. the second term in Equation (D5)) with a cubic spline in $G_d \times CR_{ud}$. The third row includes this spline and two additional cubic splines in G_d and daily maximum temperature. The fourth row replaces month-of-year fixed effects with quarter-of-year fixed effects. The fifth row uses both month-of-year and year fixed effects. The last row estimates Equation (D5) at the hourly level, adding hour-of-day fixed effects. "Realized abatement" reports the share of counterfactual no-fracking CO_2 emissions avoided. "Potential abatement" reports the share that would have been avoided if coal markups did not change. I also report abatement in million metric tons (a.k.a. tonnes), which corresponds to the numerators of Equation (D14)–(D15). The right-most column reports Equation (D16).

(i.e. the observed drop in Henry Hub prices), the avearge coal plant in my sample faced an implicit tax of \$24/metric ton CO_2. My prediction of 4.7% realized CO_2 abatement is quite close to Cullen and Mansur (2017)'s results from interconnection-wide reduced-form time series regressions, despite coming from plant-specific time-series regressions for coal units only. As with Cullen and Mansur (2017), my use of time fixed effects (for econometric identification) absorbs some of the time series variation in gas prices (and implicit carbon prices). While this could render a \$20–32 carbon tax effectively out-of-sample (giving the remaining variation), Equation (D5) uses only month-of-year fixed effects—removing seasonal gas price variation but preserving year-on-year trends.

My estimates of short-run abatement from fuel switching do not capture the full extent of fracking-induced decreases in CO_2 emissions from U.S. electricity generation. Linn, Mastrangelo, and Burtraw (2014) show that coal plants increase their thermal efficiency (i.e. lower their heat rates) in response to competitive pressure; the fracking boom has likely contributed meaningful medium-run CO_2 abatement through this channel. Low gas prices have also led to medium-to-long-run abatement on the capacity margin, by incentivizing investments in new combined-cycle gas plants and accelerating coal plant retirements (Brehm (2019)). Finally, in the long-run, even a small carbon tax could have a large effect on coal capacity (Cullen and Reynolds (2017)). Equation (D5) ignores each of these sources of fracking-induced CO_2 abatement, which is likely why my counterfactual exercise finds only 4.7% abatement.

A simple event-study analysis suggests that CO_2 emissions from the U.S. electricity sector have fallen by 20–25% during the fracking boom. While this does not establish the *causal* effect of low gas prices, it does suggest that total abatement was potentially much greater than 4.7%. It also suggests that the unrealized environmental benefits of fracking may have been much larger, since decreasing coal markups likely also impacted each of the above medium- and long-run abatement channels *in addition to* their impact on the short-run coal-to-gas switching margin. Importantly, CO_2 is a global pollutant, and my analysis focuses on U.S. emissions only. The fracking boom also impacted global energy markets, with theoretically ambiguous implications for global CO_2 emissions (Knittel, Metaxoglou, and Trindade (2016); Wolak (2016)).

Appendix references

Alexandrov, Alexei, Russell W. Pittman, and Olga Ukhaneva. 2018. "Pricing of Complements in the U.S. Freight Railroads: Cournot versus Coase." Working Paper.

Atkin, David, and Dave Donaldson. 2015. "Who's Getting Globalized? The Size and Implications of Intra-national Trade Costs." NBER Working Paper 21439.

Atkinson, Scott E., and Joe Kerkvliet. 1986. "Measuring the Multilateral Allocation of Rents: Wyoming Low-Sulfur Coal." *RAND Journal of Economics* 17 (3): 416–430.

Brehm, Paul. 2019. "Natural Gas Prices, Electric Generation Investment, and Greenhouse Gas Emissions." *Resource and Energy Economics* 58:101106.

Busse, Meghan R., and Nathaniel O. Keohane. 2007. "Market Effects of Environmental Regulation: Coal, Railroads, and the 1990 Clean Air Act." *RAND Journal of Economics* 38 (4): 1159–1179.

Chu, Yin, J. Scott Holladay, and Jacob LaRiviere. 2017. "Opportunity Cost Pass-Through from Fossil Fuel Market Prices to Procurement Costs of the U.S. Power Producers." *Jounal of Industrial Economics* 65 (4): 842–871.

Cicala, Steve. 2022. "Imperfect Markets versus Imperfect Regulation in U.S. Electricity Generation." *American Economic Review* 112 (2): 409–41.

Corts, Kenneth S. 1999. "Conduct Parameters and the Measurement of Market Power." *Journal of Econometrics* 88 (2): 227–250.

Cullen, Joseph A. 2015. "Dynamic Response to Environmental Regulation in the Electricity Industry." Working Paper.

Cullen, Joseph A., and Erin T. Mansur. 2017. "Inferring Carbon Abatement Costs in Electricity Markets: A Revealed Preference Approach Using the Shale Revolution." *American Economic Journal: Economic Policy* 9 (3): 106–133.

Cullen, Joseph A., and Stanley S. Reynolds. 2017. "Market Dynamics and Investment in the Electricity Sector." Working Paper.

Davis, Lucas, and Catherine Hausman. 2016. "Market Impacts of a Nuclear Power Plant Closure." *American Economic Journal: Applied Economics* 8 (2): 92–122.

Fabra, Natalia, and Mar Reguant. 2014. "Pass-Through of Emissions Costs in Electricity Markets." *American Economic Review* 104 (9): 2872–2899.

Fabrizio, Kira R., Nancy L. Rose, and Catherine D. Wolfram. 2007. "Do Markets Reduce Costs? Assessing the Impact of Regulatory Restructuring on U.S. Electric Generation Efficiency." *American Economic Review* 97 (4): 1250–1277.

Holladay, J. Scott, and Jacob LaRiviere. 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.

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.

InterVISTAS. 2016. *An Examination of the STB's Approach to Freight Rail Rate Regulation and Options for Simplification.* Report to Surface Transportation Board. Project FY14-STB-157.

Jha, Akshaya. 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 (1): 51–78.

Katz, Michael L. 1987. “The Welfare Effects of Third-Degree Price Discrimination in Intermediate Good Markets.” *American Economic Review* 77 (1): 154–167.

Knittel, Christopher, Konstantinos Metaxoglou, and Andre Trindade. 2016. “Are We Fracked? The Impact of Falling Gas Prices and the Implications for Coal-to-Gas Switching and Carbon Emissions.” *Oxford Review of Economic Policy* 32 (2): 241–259.

Linn, Joshua, Erin Mastrangelo, and Dallas Burtraw. 2014. “Regulating Greenhouse Gases from Coal Power Plants under the Clean Air Act.” *Journal of the Association of Environmental and Resource Economists* 1 (1/2): 97–134.

Linn, Joshua, and Lucija Muehlenbachs. 2018. “The Heterogeneous Impacts of Low Natural Gas Prices on Consumers and the Environment.” *Journal of Environmental Economics and Management* 89:1–28.

MacDonald, James M. 1987. “Competition and Rail Rates for the Shipment of Corn, Soybeans, and Wheat.” *RAND Journal of Economics* 18 (1): 151–163.

———. 2013. “Railroads and Price Discrimination: The Roles of Competition, Information, and Regulation.” *Review of Industrial Organization* 43 (1–2): 85–101.

Mansur, Erin T. 2013. “Prices versus Quantities: Environmental Regulation and Imperfect Competition.” *Journal of Regulatory Economics* 44 (1): 80–102.

Mayo, John W., and David E. M. Sappington. 2016. “Regulation in a ‘Deregulated’ Industry: Railroads in the Post-Staggers Era.” *Review of Industrial Organization* 49 (2): 203–227.

Rennert, Kevin, Frank Errickson, Brian C. Prest, Lisa Rennels, Richard G. Newell, William Pizer, Cora Kingdon, Jordan Wingenroth, Roger Cooke, Bryan Parthum, David Smith, Kevin Cromar, Delavane Diaz, Frances C. Moore, Ulrich K. Müller, Richard J. Plevin, Adrian E. Raftery, Hana Ševčíková, Hannah Sheets, James H. Stock, Tammy Tan, Mark Watson, Tony E. Wong, and David Anthoff. 2022. “Comprehensive Evidence Implies a Higher Social Cost of CO_2.” *Nature* 610:687–692.

Schmalensee, Richard. 1981. “Output and Welfare Implications of Monopolistic Third-Degree Price Discrimination.” *American Economic Review* 71 (1): 242–247.

U.S. Environmental Protection Agency. 2008. *Optimal Emissions from Commuting, Business Travel and Product Transport.* Technical report EPA430-R-08-006.

———. 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.

Varian, Hal R. 1985. “Price Discrimination and Social Welfare.” *American Economic Review* 75 (4): 870–875.

Wetzstein, Brian, Raymond Florax, Kenneth Foster, and James Binkley. 2021. “Transportation Costs: Mississippi River Barge Rates.” *Journal of Commodity Markets* 21:100123.

Weyl, E. Glen, and Michal Fabinger. 2013. “Pass-Through as an Economic Tool: Principles of Incidence under Imperfect Competition.” *Journal of Political Economy* 121 (3): 528–583.

Wilson, Wesley W. 1996. “Legislated Market Dominance in Railroad Markets.” *Research in Transportation Economics* 4:49–67.

Wolak, Frank A. 2016. “Assessing the Impact of the Diffusion of Shale Oil and Gas Technology on the Global Coal Market.” Working Paper.