

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

Louis Preonas*

November 21, 2017

JOB MARKET PAPER

[MOST RECENT VERSION](#)

Abstract

Economists have widely endorsed pricing CO₂ emissions to internalize climate change-related externalities. Doing so would significantly affect coal, which is the most carbon-intensive major energy source. However, U.S. coal markets exhibit an additional distortion, as the railroads that transport coal to power plants can exert market power. This upstream distortion can mute the price signal of a corrective tax, due to changes in markups or incomplete tax pass-through. In this paper, I provide the first empirical estimates of how coal-by-rail markups respond to changes in coal demand. I find that rail carriers reduce coal markups when downstream power plant demand changes, due to a decrease in the price of a natural gas (a competing fuel). I estimate markup changes that vary substantially across coal plants, resulting from a combination of heterogeneous transportation market structure and plant-specific demand shocks. Since low natural gas prices and a CO₂ emissions tax similarly disadvantage coal, observed decreases in coal markups imply that pass-through of a federal carbon tax to coal power plants may be heterogeneous and incomplete. This could substantially erode the environmental benefits of a price-based climate policy. My results suggest that decreases in coal markups have increased recent climate damages by \$2.4 billion, compared to a counterfactual where markups do not change.

*University of California, Berkeley, Department of Agricultural and Resource Economics, and Energy Institute at Haas. Email: lpreonas@berkeley.edu. I thank Meredith Fowlie, Severin Borenstein, Maximilian Auffhammer, Lint Barrage, Cyndi Berck, Peter Berck, Susanna Berkouwer, Joshua Blonz, Fiona Burlig, Lucas Davis, Karl Dunkle Werner, Don Fullerton, James Gillan, Stephen Jarvis, Akshaya Jha, Jeremy Magruder, Dave McLaughlin, Elisabeth Sadoulet, James Sallee, Leo Simon, James Stock, Sofia Villas-Boas, Reed Walker, Matt Woerman, Catherine Wolfram, and Derek Wolfson. I also thank seminar participants at the Energy Institute at Haas, UC Berkeley Dept. of Agricultural and Resource Economics, and the Heartland ERE Workshop at University of Illinois Urbana-Champaign. All remaining errors are mine. The online appendix for this paper is available at http://www.louispreonas.com/s/preonas_jmp_appendix.pdf.

1 Introduction

Economists have widely advocated policies that price carbon dioxide emissions to reflect their marginal external cost (Nordhaus (1993)). While such policies are efficient under perfect competition (Pigou (1932)), additional distortions such as market power reduce the efficiency of a Pigouvian tax (Buchanan (1969); Barnett (1980)). Economists have long understood that firms with market power may adjust prices in response to taxation (Cournot (1838)). However, there exists surprisingly little empirical research on how market power impacts the pass-through of an environmental tax, or the transmission of the desired price signal to market participants.

This paper investigates market power in the transportation of coal, and analyzes its potential impacts on the efficacy and efficiency of U.S. climate policy. Coal is likely the most environmentally damaging and carbon-intensive industry in the U.S. economy (Muller, Mendelsohn, and Nordhaus (2011)). Many geographically concentrated mines supply coal to many geographically dispersed power plants, and the railroads that transport coal from mines to plants can exercise market power (Busse and Keohane (2007)). If a carbon tax causes these oligopolist railroads to reduce coal markups, this could mute the carbon price signal received by power plants and erode the environmental benefits of the tax. While previous research has studied environmental and economic outcomes under a carbon tax, I provide the first estimates of how upstream market power in coal supply might impact climate policy outcomes.

I begin by estimating the size of the market power distortion in coal transportation, or the average markup levels faced by coal power plants. Then, I estimate how markups change due to changes in the demand for coal. Theory suggests that a shift in coal demand should cause a profit-maximizing railroad to reoptimize coal markups. Recent decreases in the price of natural gas, coal’s primary competitor in electricity markets, represent a negative shock to power plants’ coal demand. Since a carbon tax would induce a similar shift in coal demand, observed changes in coal markups due to decreases in the gas price may predict how railroads would reoptimize markups under a carbon tax.¹

To identify markup levels, I exploit predetermined cross-sectional heterogeneity in market power. Some coal plants are “captive” to a single rail carrier and face an effective transportation monopoly; other plants may purchase coal from multiple railroads, and these

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

“non-captive” plants face more competitive coal shipping. I implement a nearest-neighbor matching strategy to compare the price of coal delivered to captive vs. non-captive plants, which takes advantage of plants’ inability to arbitrage spatial price differences. By flexibly controlling for coal commodity value and railroad freight costs, I recover the average differential markup faced by captive plants.

I use a difference-in-differences design to identify changes in coal markups caused by changes in the price of natural gas. This leverages two sources of cross-sectional heterogeneity: (i) geographic variation in transport market power, and (ii) variation in coal plants’ sensitivity to gas price changes, which I predict from microdata on U.S. electricity generation. Using a simple oligopoly model as a guide, I combine these two sources of variation into a single cross-sectional predictor of markup changes, and interact this variable with the time series of gas prices. Regressing the delivered price of coal on this interaction in a panel fixed effects framework, I estimate the extent to which gas price changes cause differential changes in coal markups. Given that natural gas is the primary substitute for coal in electricity supply, negative shocks to the gas price disadvantage coal generation in a manner similar to a tax on CO₂ emissions (Cullen and Mansur (2017)). Hence, observed gas price shocks mimic the variation of a carbon tax, and my estimates of markup changes can help predict the pass-through of such a tax.

I find that coal plants facing the most market power in transportation pay \$2–5 per ton higher average markups, compared to plants facing the least market power. This translates to an average markup of 4–15 percent of delivered coal prices, explaining 13–42 percent of the average spatial gap between mines’ sales prices and plants’ delivered prices. I also find robust and statistically significant *changes* in markups for approximately 42 percent of plants—the subset of plants that face the most market power *and* are sensitive to competition from gas-fired generation. For these “markup-sensitive” plants, a \$1/MMBTU drop in gas price causes coal markups to fall by \$1 per ton. I find no evidence that markups change for the remaining 58 percent of coal plants, which are less sensitive to gas-fired competition *or* face relatively little market power.

I demonstrate that rail carriers reoptimize markups to effectively insulate some coal plants against shocks to their competitiveness. As decreasing gas prices reduce the marginal cost of gas-fired generation, “markup-sensitive” plants see their coal prices decrease, thereby reducing these plants’ *own* marginal costs. Such offsetting changes in markups help this subset of coal plants to remain competitive with their gas-fired rivals. By contrast, over half

of coal plants experience the full gas price shock, as their markups do not adjust. These heterogeneous impacts across plants are qualitatively consistent with the predictions of the static oligopoly model that I develop, implying that rail carriers indeed reoptimize markups heterogeneously to maximize profits in coal shipping.

Falling gas prices have given gas-fired plants a competitive advantage over coal-fired plants. This is similar to what might occur under a carbon tax, which would penalize coal as the more carbon-intensive fuel. Therefore, I can convert my estimated markup changes into the pass-through rates of an implicit carbon tax, or the rates at which rail carriers would have passed a mine-mouth carbon tax on to delivered coal prices.² For the subset of plants whose markups do not change, this translates to full pass-through, or implied pass-through rates statistically indistinguishable from 1. By contrast, this translates to incomplete pass-through for “markup-sensitive” plants, with plant-specific pass-through rates ranging from 0.98 to 0.41. This suggests that market imperfections in coal shipping are likely to distort the price signal of a federal carbon tax, such that certain coal plants may experience as little as 41 percent of the desired cost increase.

This paper contributes to four different literatures. First, my results contribute to the literature on market power in intermediaries. Atkin and Donaldson (2015) develop techniques to identify markups separately from transportation costs, and several recent papers estimate how oligopolistic intermediaries influence both upstream and downstream outcomes (e.g., Startz (2016); Ganapati (2017)). While these studies typically focus on differentiated product markets, coal is a globally traded commodity that is relatively homogeneous. I leverage a unique feature of coal markets—limited spatial arbitrage between power plants—to credibly identify transport markups while invoking relatively few structural assumptions on coal demand. My results have important implications for many commodities with high geographic specificity and high transportation costs, including crude oil, cement, and metals.

My analysis also contributes to the literature on coal intermediaries, which has largely focused on the railroads’ interactions with upstream mines (e.g., Kolstad and Wolak (1983); Wolak and Kolstad (1988)), rather than downstream power plants. A notable exception is Busse and Keohane (2007), who provide the first evidence of price discrimination due to geographic variation in coal shipping during the 1990s. My results demonstrate that

2. The physical location of the tax along the coal supply chain should not change the economic interpretation of pass-through, in the absence of additional market distortions beyond rail market power (Weyl and Fabinger (2013)). “Forward” pass-through of a mine-mouth tax (i.e. a cost shock to rail shipping) follows the standard formulation of a cost shock passed through to final goods prices. However, in practice, carbon taxes are typically levied on electricity sales.

heterogeneous coal markups have persisted through recent years, and I am the first to show that markup changes have led to economically meaningful impacts on CO₂ emissions.

Second, my results contribute to a growing empirical literature on environmental policy in the presence of market power. Given widespread evidence of market power in major polluting industries (e.g., Bushnell, Mansur, and Saravia (2008) on electricity markets; Hastings (2004) on gasoline markets), surprisingly few studies have empirically estimated the theoretically ambiguous interactions of these two market failures. Mansur (2007) finds that market power in electricity markets can increase pollution abatement under environmental regulation. On the other hand, Ryan (2012) and Fowlie, Reguant, and Ryan (2016) find that emissions regulation exacerbates market power distortions in the cement industry. I find that changes in coal markups may significantly erode the environmental benefits of a carbon tax, as incomplete pass-through would mute the price signal felt by a subset of coal plants. My results suggest that incomplete pass-through increased CO₂ emissions damages during my sample period by roughly \$2.4 billion, compared to a full pass-through counterfactual. Hence, the magnitude of this effect would likely be economically meaningful, despite the fact that incomplete pass-through would only impact a fraction of coal plants.

Incomplete pass-through of a carbon tax could increase or decrease welfare, depending on the size of the tax relative to marginal external costs. If the tax were equal to the social cost of carbon, then the presence of markups would restrict coal consumption past the socially optimal quantity. In this case, incomplete pass-through would reduce the markup distortion and likely increase welfare.³ However, real-world carbon prices are typically much smaller than the estimated social cost of carbon (Carl and Fedor (2016); Revesz et al. (2017)). Under a suboptimally low carbon tax, coal markups would increase welfare by raising coal prices closer to their marginal social cost. In this case, incomplete pass-through would lower coal markups and reduce welfare.

Third, my analysis contributes to the literature on estimating tax pass-through in energy markets. Previous studies have often found heterogeneous pass-through of energy taxes, due to variation in market structure both across and within industries (e.g., Ganapati, Shapiro, and Walker (2016) in manufacturing; Pouliot, Smith, and Stock (2017) in transport fuels). Muehlegger and Sweeney (2017) find that pass-through in petroleum refining also varies by whether cost shocks are firm-specific or common across all firms. I find heterogeneous pass-

3. Coal also emits harmful local air pollutants such as SO₂, NO_x, and particulate matter. A tax greater than the social cost of carbon may partially internalize damages from these other pollutants. Hence, incomplete pass-through could reduce welfare even under a tax equal to marginal CO₂ damages.

through due to a combination of these two effects: spatial variation in the competitiveness of coal shipping, and variation in coal plants' sensitivity to relative fuel price shocks. To my knowledge, this is the first evidence that pass-through of a carbon tax in the U.S. electricity sector may be heterogeneous and incomplete.

Heterogeneous pass-through implies that the economic incidence of a carbon tax would likely vary across coal plants. I apply the theoretical tools of Weyl and Fabinger (2013) to translate pass-through to incidence, which reveals substantial heterogeneity in the share of the implied tax burden borne by plants. While most plants bear the full decline in profits from a gas price drop, a subset of plants bear less than half, with the remainder coming out of railroad oligopoly rents. This finding contributes to the literature on environmental tax incidence, which has shown that imperfect competition and heterogeneous pass-through can shift the tax burden towards producers and make climate policy less regressive (Ganapati, Shapiro, and Walker (2016); Stolper (2016b)). In my setting, shifting a share of the tax burden upstream from coal power plants may also benefit electricity consumers, potentially reducing the regressivity of a carbon tax.

Finally, my results contribute to the literature on fuel-switching between coal and natural gas. Recent decreases in the gas price have crowded out coal-fired generation, thereby reducing CO₂ emissions from the U.S. electricity sector. While several previous studies have estimated the magnitude of these environmental benefits (e.g., Holladay and LaRiviere (2017); Fell and Kaffine (forthcoming)), I show that decreasing coal markups may have attenuated this shift away from coal. A simple counterfactual exercise suggests that short-run fuel substitution could have yielded 8 percent greater CO₂ abatement, if coal markups had not changed. This suggests that previous retrospective analyses may have understated the potential environmental benefits of a carbon tax, if the tax is large enough to drive coal markups close to zero and eliminate the countervailing effect of incomplete pass-through.

This paper proceeds as follows. Section 2 describes the institutions of U.S. coal markets and the recent boom in natural gas production. Section 3 develops a static oligopoly model to predict how railroads reoptimize markups as gas prices change. Section 4 outlines the data I use to implement my empirical strategy, detailed in Section 5. Section 6 reports results from estimating levels and changes of coal markups. Section 7 analyzes the implications of these results for climate policy. Section 8 concludes.

2 U.S. Coal Markets

U.S. coal markets feature three primary types of agents: mining firms, power plants, and transport intermediaries.⁴ Mines are spatially concentrated in regions with productive coal deposits, most notably the Powder River Basin in northeastern Wyoming, and the Appalachian Basin in West Virginia and Kentucky. By contrast, coal power plants are spatially dispersed across the country, due to the regionally fragmented nature of electricity markets. Coal is heavy relative to its commodity value, and plants located far from mines incur substantial coal shipping costs (Joskow (1985)). Railroads are the dominant transportation mode, and a few large rail carriers deliver over 70 percent of coal shipments. Figure 1 maps the geographic configuration of coal producing regions, coal power plants, and major rail lines. Plants located on navigable waterways may also receive coal shipments via barge, a more competitive outside option with low barriers to entry. Barges contribute roughly 17 percent of coal deliveries.⁵

Four firms control most of the coal shipping industry, with two large rail carriers dominating both the western and eastern U.S. (see Figure 1). Ever since the Staggers Act of 1980 substantially weakened rail price regulations, railroads have been able to set freight shipping rates with limited government oversight (MacDonald (1989)).⁶ In cases where a single rail carrier exhibits “market dominance” along a given route, regulators may intervene to prevent rail revenues from exceeding 180 percent of total variable costs.⁷ This means that rail oligopolists have significant leeway to exercise market power and negotiate complex long-term contracts with power plants (Joskow (1988)). By allowing carriers to extract oligopoly rents and exploit economies of scale, the Staggers Act also spurred a series of railroad mergers; the 33 “Class I” railroads of 1980 have consolidated into the 7 Class I railroads of today (Schmidt (2001); Prater, Sparger, and O’Neil (2014)).⁸

4. Over 90 percent of U.S. coal consumption occurs in the electric power sector. My analysis does not include other industrial consumers of coal, such as steel, cement, and paper manufacturers. I also ignore coal imports (less than 2 percent of U.S. consumption) and exports (roughly 3 percent of U.S. production).

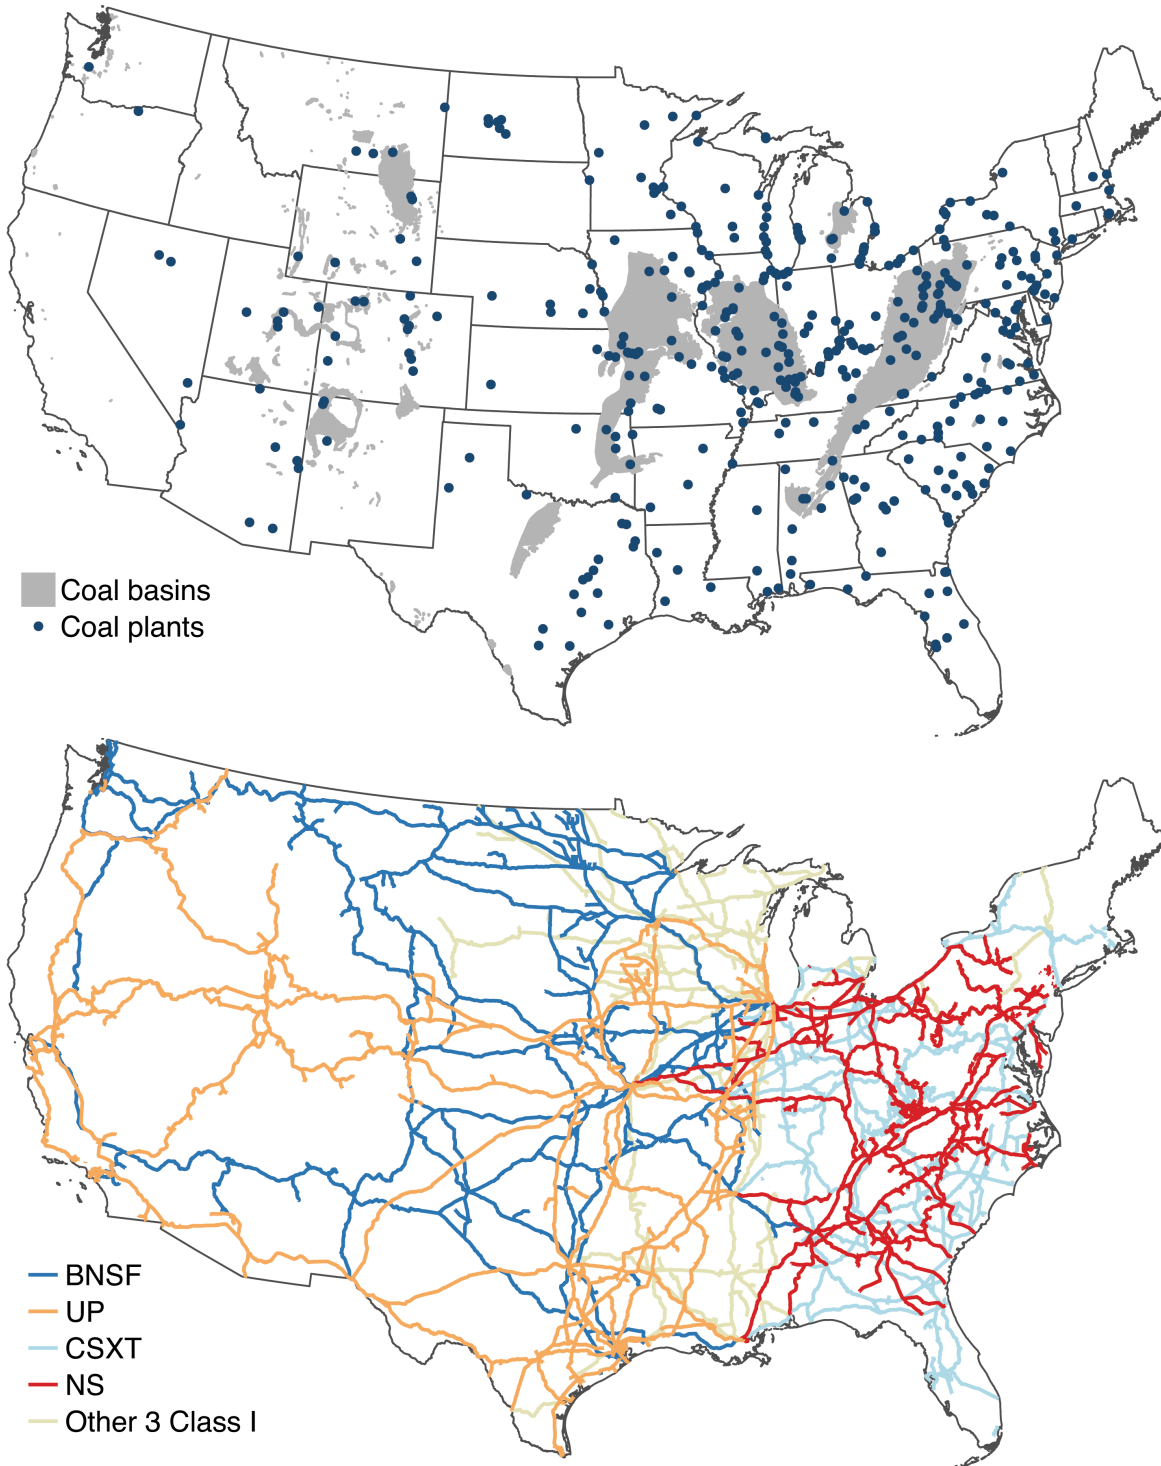
5. Trucks also transport a small share of coal deliveries. However, trucking is relatively costly and likely cannot compete directly with rail and water (Busse and Keohane (2007)).

6. MacDonald (2013) offers a comprehensive history of U.S. railroad regulation.

7. In practice, regulators loosely interpret this threshold such that railroads may earn an adequate return on investment (Wilson (1996)). While the Surface Transportation Board reviews only 1–7 rate challenges each year, rate cases for coal shipping occur more frequently than for all other commodities combined (https://www.stb.gov/stb/industry/Rate_Cases.htm).

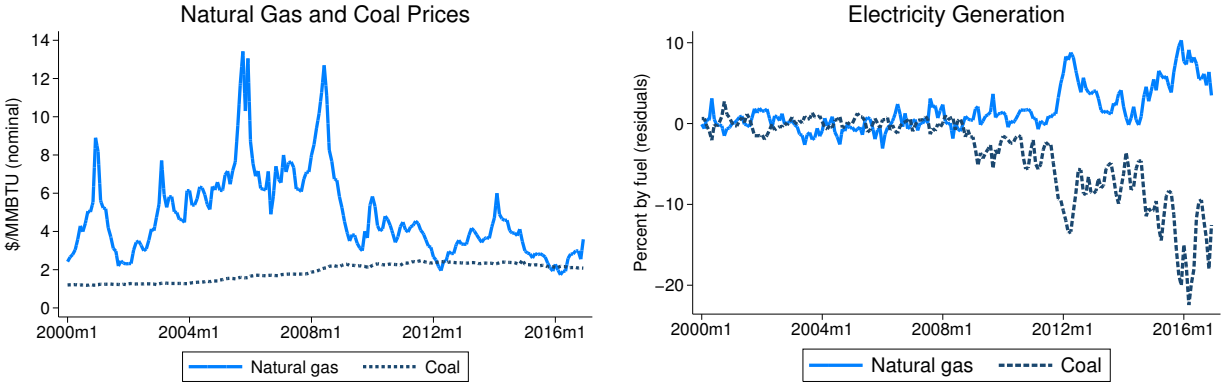
8. The Class I designation includes carriers with annual operating revenues exceeding \$453 million. These seven firms account for approximately 69 percent of rail mileage and 94 percent of rail freight revenues.

Figure 1: U.S. Coal Geography



Notes: This top panel maps all productive deposits of power plant grade coal (i.e. bituminous and sub-bituminous) in the contiguous U.S. The vast majority of coal production occurs in three regions: the Powder River Basin in northeastern Wyoming; the Appalachian Basin in West Virginia and eastern Kentucky; and the Illinois Basin in southern Illinois and western Kentucky. Dots denote all 430 large coal-fired electric power plants that operated between 2002–2015. The bottom panel maps major rail lines owned and operated by the seven Class I rail carriers. Two rail carriers each dominate the West (BNSF, Union Pacific) and East (CSX Transportation, Norfolk Southern). I combine the remaining three carriers (Canadian National, Canadian Pacific, Kansas City Southern) into a single color, as these rail carriers cover smaller territories far from most major coal deposits.

Figure 2: U.S. Fuel Prices and Electricity Generation



Notes: The left panel reports the monthly average Henry Hub spot prices for natural gas, and the monthly average cost of coal delivered to electric generating plants. This average cost of coal does not account for heterogeneous coal attributes, including BTU content by weight, sulfur content, or ash content. The right panel plots U.S. monthly electricity generation by fuel as a percent of total monthly generation, controlling for month fixed effects and a 2000–2008 time trend for each fuel.

Three factors have led to substantial spatial dispersion in coal-by-rail markups. First, unlike most commodities, coal consumption must occur in precise geographic locations with potentially limited access to transportation networks. While some power plants have the option to purchase coal from multiple rail carriers or via barge (by virtue of their locations), other plants must rely on a single rail carrier for all coal deliveries. Second, many plants are constrained to buy a particular type of coal, produced in only one mining region (Joskow (1987)). This further restricts plants’ shipping options, as mines may also have limited access to rail and water networks.⁹ Third, the resale of coal is cost-prohibitive, because infrastructure is built to carry coal *to* (not *away from*) plants (Busse and Keohane (2007); Jha (2015)). Hence, plants are unable to arbitrage away spatial price differences, allowing railroads to charge higher markups to plants with fewer shipping options.¹⁰

U.S. coal consumption has declined over the past decade, largely due to decreases in the price of natural gas. Technological advances in hydraulic fracturing (“fracking”) have led to

9. Coal’s physical characteristics vary across coal regions, and even across mines within a region. Plants typically value coal with high energy content (or BTUs per ton), and with low sulfur and ash content (which create local air pollution). Plants self-calibrate to a pre-specified mix of coal attributes, and deviations can reduce the efficiency of boilers and pollution-control devices (Kerkvliet and Shogren (1992)). Also, many plants comply with SO₂ regulations by burning low-sulfur coal from Wyoming’s Powder River Basin (Schmalensee and Stavins (2013)). If such a plant has access to two rail carriers and the Ohio River, but only one rail carrier connects to the Powder River Basin, then it has effectively one shipping option.

10. Power plants may purchase coal directly from rail carriers; alternatively, plants may purchase freight services from railroads and separately purchase coal from upstream mines. This distinction does not affect the economic interpretation of delivered coal markups. My analysis treats rail intermediaries as both the owners of the commodity and the providers of freight services.

a boom in natural gas extraction, causing a historic drop in U.S. gas prices.¹¹ Because coal plants compete directly with natural gas generation in electricity markets, low gas prices have crowded out coal-fired electricity generation. The left panel of Figure 2 shows how the fracking boom has led to a decline in U.S. gas prices since 2008, and the right panel illustrates how the electricity sector has shifted towards gas and away from coal. The corresponding decrease in coal demand has likely caused rail oligopolists to reoptimize coal markups. These observed changes in markups can predict what might occur under a carbon tax, which would similarly disadvantage coal relative to low-carbon natural gas (Cullen and Mansur (2017)). If coal markups decrease (increase), this would dampen (magnify) the carbon tax price signal as it passes along the coal supply chain. This effect would likely be heterogeneous across plants, due to variation in pre-existing markups *and* variation in plants’ exposure to gas-fired competition.¹²

3 Theoretical Framework

In this section, I develop a simple Cournot oligopoly model of railroad intermediaries who sell coal to power plants. This allows me to predict how markups should respond to gas price changes, which depends on three key dimensions of plant heterogeneity: (1) number of potential rail carriers; (2) availability of water transport as a more competitive outside option; and (3) price elasticity of demand for coal as an input to electricity production.

3.1 Symmetric Rail Oligopoly

Consider power plant j that is a price-taker in the market for coal. This plant consumes a specific type of coal from origin o , which is produced at constant marginal cost by a perfectly competitive mining sector.¹³ Plant j is fully captive to N_{oj} identical rail carriers for its coal deliveries from origin o , and each rail carrier i chooses the best-response quantity of coal

11. Two separate technological innovations have facilitated the “fracking boom”: horizontal drilling and hydraulic fracturing. Fitzgerald (2013) provides a comprehensive overview of these technological advances and their effect on the costs of gas extraction. The physical properties of natural gas make it expensive to export, which is why a domestic supply glut has depressed U.S. gas prices.

12. Low gas prices have not impacted all coal plants equally. For a coal plant located in an electricity market with many gas-fired competitors, a negative gas price shock will likely cause its coal demand to decrease. For a coal plant in a market without any gas-fired competitors, the same gas price shock may have no effect on its coal demand. Coal plants also vary in their productive efficiency, and low gas prices should disproportionately hurt relatively inefficient plants.

13. This assumption greatly simplifies my theoretical framework. In reality, coal supply may be upward-sloping, and mining need not be perfectly competitive, especially in Wyoming’s Powder River Basin where

q_{ioj} that maximizes its profits on route oj . In equilibrium, plant j consumes $N_{oj}q_{ioj} = Q_{oj}$ units of coal at price P_{oj} . Plant j cannot resell its purchased coal, meaning that P_{oj} is not restricted by a binding arbitrage condition and rail carriers may effectively treat each plant as its own isolated coal market.

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

$$\pi_{ioj}(q_{ioj}) = q_{ioj} \left[P_{oj}(Q_{oj}; \mathbf{Z}_{oj}) - C_o - S(\mathbf{T}_{oj}) \right] - F_{oj} \quad (1)$$

where P_{oj} is the plant j 's inverse demand for coal shipped from origin o , as a function of Q_{oj} and a vector of parameters \mathbf{Z}_{oj} . C_o is the exogenous mine-mouth coal price in origin o . The function $S(\mathbf{T}_{oj})$ denotes the average per-unit cost of shipping coal on route oj , where \mathbf{T}_{oj} is a vector of transportation cost parameters, including rail mileage, diesel costs, and the opportunity cost of a rail car. Finally, the oligopolist incurs F_{oj} , a fixed cost of entry on shipping route oj .¹⁴ In reality, carrier i maybe also be subject to regulatory oversight, but I abstract from rail regulation in this simple model.

Taking rail carrier i 's first-order condition, and rearranging in terms of a price-cost markup μ_{oj} :

$$\mu_{oj} \equiv P_{oj} - C_o - S(\mathbf{T}_{oj}) = - \left(\frac{\theta_{oj}}{N_{oj}} \right) \frac{\partial P_{oj}}{\partial Q_{oj}} Q_{oj} \quad (2)$$

where the ‘‘conduct parameter’’ θ_{oj} equals 1 under a pure Cournot oligopoly and 0 under perfect competition.¹⁵ Plant j 's markup depends on both its coal transportation options and its demand for type- o coal. If plant j is captive to a single rail carrier (i.e., $N_{oj} = 1$), it should face higher markups than if multiple carriers were competing on route oj . At the same time, if plant j is located on a navigable waterway and can receive type- o coal via barge, this should limit railroads' ability to set high markups. Since waterways have less restricted usage rights and lower barriers to entry, I treat barge shipments as a competitive outside

a few large firms dominate mining operations (Atkinson and Kerkvliet (1986)). Appendix A.3 discusses the welfare implications of alternative market structures.

14. I use a symmetric oligopoly model for tractability. In reality, each firm's shipping routes are constrained by track ownership and trackage rights, implying non-identical costs $S(\mathbf{T}_{oj})$ and F_{oj} . For simplicity, I assume that quantity q_{ioj} does not enter into $S(\mathbf{T}_{oj})$, which ignores rail capacity constraints or increasing returns to scale in shipping. My empirical analysis relaxes this assumption.

15. $\theta_{oj} = \frac{\partial Q_{oj}}{\partial q_{ioj}}$ is identical for all i , by symmetry. I use this ‘‘conduct parameter’’ formulation for notational convenience (following Atkin and Donaldson (2015); Bergquist (2017)), and I treat θ_{oj} only as a continuous heuristic for distance from perfect competition. Calibrating θ_{oj} as a structural parameter can be problematic, as it only takes on a well-defined interpretation at a few values (Corts (1999)).

option (i.e., $\theta_{oj} = 0$). Finally, if plant j 's inverse demand for coal is relatively inelastic, it should face relatively higher markups, all else equal.

3.2 Comparative Statics for Coal Markups

Coal demand depends on the price of natural gas, because the two fuels compete in electricity dispatch. If the gas price decreases (increases), a coal plant may supply less (more) electricity at a given coal price. The gas price also influences the elasticity of coal demand, by determining the range of coal prices over which a coal plant is marginal in electricity supply. A marginal plant has (locally) elastic coal demand, because its coal consumption responds to small changes in coal price. At lower coal prices, the coal plant will be inframarginal and its strict capacity constraint will bind. This translates to (locally) inelastic coal demand, as small changes in coal price will not change its coal consumption.

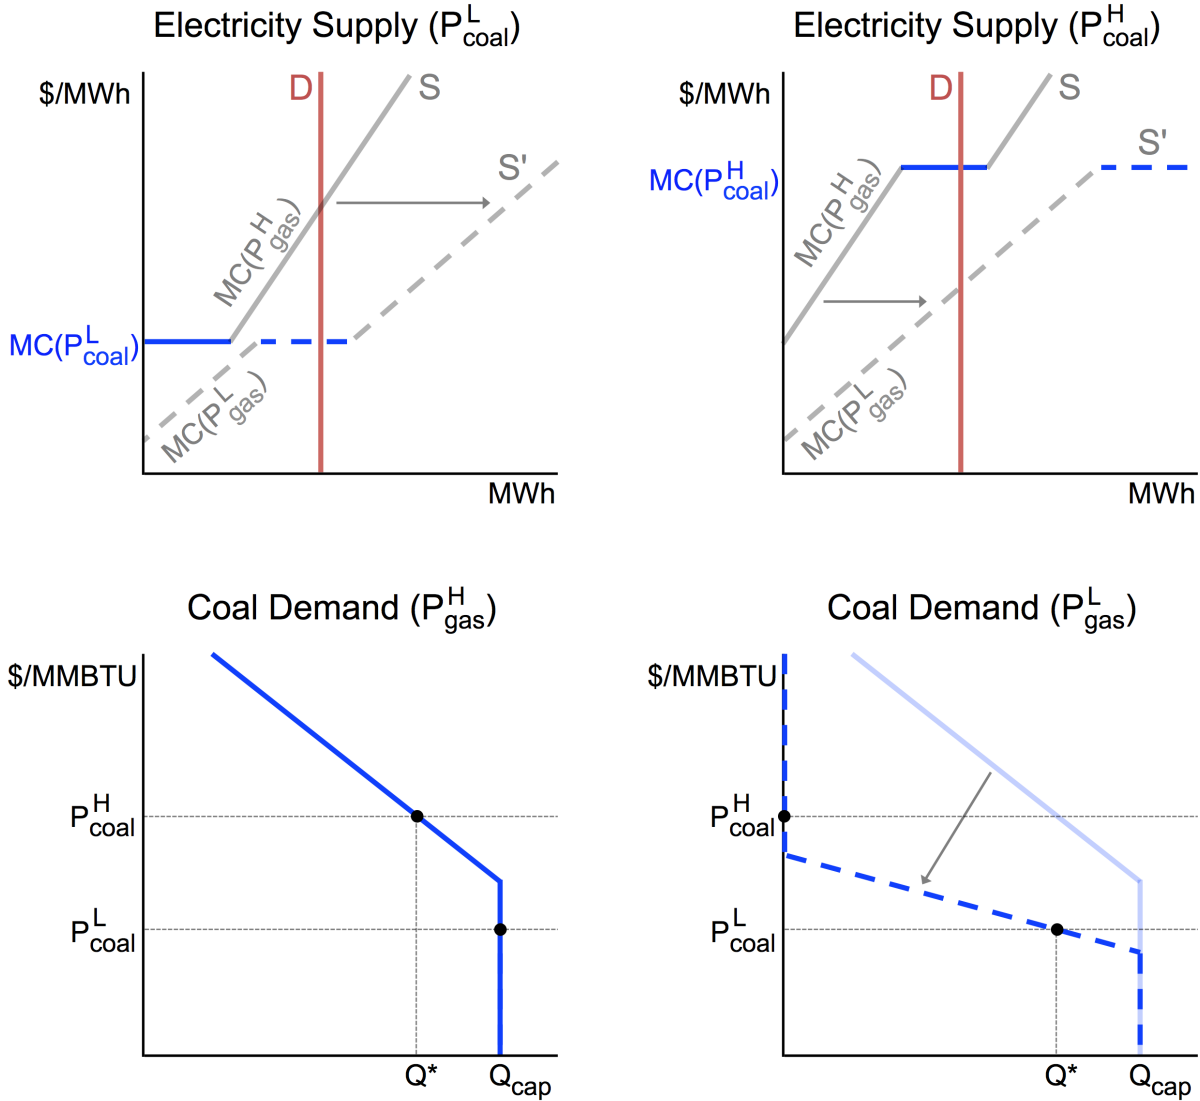
Figure 3 presents a stylized electricity market to illustrate how a negative gas price shock impacts both the level and slope of coal demand. There is a single coal plant with constant marginal cost, and an upward-sloping supply of gas-fired generation. Each technology's marginal costs scale with its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

At a given gas price, the plant's coal demand is the 1-to-1 mapping between coal price and coal consumption. Under a high coal price and high gas price (i.e., the solid supply curve in the top-right panel), the coal plant is marginal in the electricity market and generates at 70 percent capacity. Hence, it demands 70 percent of its throughput capacity for coal, or Q^* in the bottom-left panel. Comparing the bottom two panels, the gas price governs the range of coal prices for which the plant is marginal, and coal demand is not vertical. A negative gas price shock causes inverse coal demand to shift down *and* become less steep.¹⁶

Using my symmetric oligopoly model, I can derive how rail carriers should reoptimize coal markups in response to gas price changes. Let Z_j denote the gas price of coal plant j 's

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

Figure 3: Coal Demand and Natural Gas Prices



Notes: This figure presents a stylized electricity market to illustrate the relationship between gas prices and coal demand. There is one coal generator with a fixed capacity, and constant marginal cost at a given coal price ($MC(P_{\text{coal}})$, in blue). There is also an upward-sloping supply of many small natural gas generators, with marginal costs that scale multiplicatively with the gas price ($MC(P_{\text{gas}})$, in grey). Electricity demand (D) is perfectly inelastic, and deterministic (for simplicity). The top panels show four electricity supply curves, each for a given combination of coal price (low in the left panel, high in the right panel) and gas price (high for solid lines, low for dashed lines). The bottom panels translate the coal plant's electricity production into its corresponding demand for coal (MWh out is a function of MMBTU of coal in, given the plant's fixed production technology). Under a high gas price (P_{gas}^H), the coal plant consumes at full capacity (Q_{cap}) given a low coal price (P_{coal}^L) and at Q^* given a high coal price (P_{coal}^H). If the gas price decreases to P_{gas}^L , the coal plant is now marginal given P_{coal}^L (where it had been inframarginal) and above the margin given P_{coal}^H (where it had been marginal). The decrease in gas price has caused inverse coal demand at Q^* to shift down and become less steep.

competitors, which enters plant j 's inverse demand function as an element of the parameter vector \mathbf{Z}_{oj} . The change in markup μ_{oj} that results from a small change in gas price Z_j is:¹⁷

$$\frac{d\mu_{oj}}{dZ_j} = \frac{\frac{\partial P_{oj}}{\partial Z_j}(2 + E_{D_{oj}} - N_{oj}) - \frac{\partial^2 P_{oj}}{\partial Q_{oj} \partial Z_j} Q_{oj} \theta_{oj}}{2 + E_{D_{oj}}} \quad (3)$$

where $E_{D_{oj}}$ is the elasticity of the slope of inverse demand scaled by the degree of competitiveness $\frac{\theta_{oj}}{N_{oj}}$:

$$E_{D_{oj}} \equiv \left(\frac{\partial^2 P_{oj}}{\partial Q_{oj}^2} \right) \left(\frac{\partial P_{oj}}{\partial Q_{oj}} \right)^{-1} Q_{oj} \left(\frac{\theta_{oj}}{N_{oj}} \right) \quad (4)$$

Equation (3) depends on the level, slope, and curvature of plant j 's inverse demand. $\frac{\partial P_{oj}}{\partial Z_j}$ captures how gas price affects the *level* of inverse coal demand. If a negative gas price shock (i.e. $dZ_j < 0$) causes plant j 's inverse coal demand to shift down as in Figure 3, then $\frac{\partial P_{oj}}{\partial Z_j} > 0$. The cross-partial $\frac{\partial^2 P_{oj}}{\partial Q_{oj} \partial Z_j}$ captures how gas price affects the *slope* of inverse coal demand. If lower gas prices make inverse coal demand less steep (i.e. if $dZ_j < 0$ causes $\frac{\partial P_{oj}}{\partial Q_{oj}}$ to become less negative), then $\frac{\partial^2 P_{oj}}{\partial Q_{oj} \partial Z_j} < 0$. Finally, the change in markup depends on the degree to which inverse demand is concave ($E_{D_{oj}} > 0$) or convex ($E_{D_{oj}} < 0$). More concave demand will tend to increase $\frac{d\mu_{oj}}{dZ_j}$, while more convex demand will tend to decrease $\frac{d\mu_{oj}}{dZ_j}$.¹⁸

These three features of coal demand interact with route oj 's rail market size (N_{oj}) and structure (θ_{oj}) to jointly determine how railroads should reoptimize markups when the gas price changes. The sign of $\frac{d\mu_{oj}}{dZ_j}$ is theoretically ambiguous and depends on the relative sizes of $\frac{\partial P_{oj}}{\partial Z_j}$, $\frac{\partial^2 P_{oj}}{\partial Q_{oj} \partial Z_j}$, and $E_{D_{oj}}$, which may vary considerably across heterogeneous coal plants. Rail carrier behavior may also depart from the predictions of this simple model, especially if regulatory constraints bind or if markups are not truly independent across plants.¹⁹ Below, I econometrically estimate the degree to which observed gas price changes have caused changes in coal markups. I directly estimate plant-specific demand parameters $\frac{\partial P_{oj}}{\partial Z_j}$, $\frac{\partial^2 P_{oj}}{\partial Q_{oj} \partial Z_j}$, and $E_{D_{oj}}$, which I use to construct an empirical approximation of $\frac{d\mu_{oj}}{dZ_j}$. This allows me to test the theoretical predictions and empirical magnitudes of Equation (3).

17. Appendix A.1 provides a full derivation of this comparative static.

18. This is a standard finding from the pass-through literature on imperfectly competitive product markets, where the pass-through rate is closely related to the curvature of demand (Weyl and Fabinger (2013)).

19. In reality, rate regulation prevents carriers from extracting full (unconstrained) oligopoly rents. This simple model does not account for multiple-market negotiations between rail carriers or dynamic interactions between carriers and plants.

4 Data

My analysis combines several publicly available datasets, published by U.S. government agencies. This section highlights the core datasets for my empirical analysis, including data on coal shipments, power plants, and the U.S. rail network.²⁰ It also describes how I use GIS data to construct a measure of plants’ rail captiveness.

4.1 Data Sources

The Energy Information Administration’s (EIA) Form 923 collects detailed data on coal deliveries to power plants, for all large U.S. coal plants. These data are at the month-shipment level, where “shipments” aggregate deliveries received on a single purchase order or contract, in a given month, with the same supplier, county of origin, and coal rank (i.e. bituminous vs. sub-bituminous). For each observation, EIA reports the total tons of coal delivered; their average BTU, sulfur, and ash content by weight; and the primary modes of transportation (e.g., rail, barge, truck). EIA also classifies each shipment as either a long-term contract or a spot market transaction.

Coal plants must report the average prices for each observation, inclusive of commodity costs, transportation costs, and markups. EIA redacts price data for independent power producers, and I only observe prices for utility-owned plants. My empirical analysis focuses on this subset of plants, which represent 77 percent of coal deliveries since 2002.²¹ Because coal is a heterogeneous commodity without a uniform price index, I control for average mine-mouth prices at the county-year level, published in EIA’s Annual Coal Report.

I merge coal shipment data with several EIA datasets on power plant characteristics (Form 860), operations (Forms 906 and 923), and pollution abatement (Forms 767, 860, and 923). The EPA eGRID database reports each plant’s power control area (PCA), or its region on the electricity transmission grid. To estimate plant-specific coal demand parameters, I leverage data from the EPA’s Continuous Emissions Monitoring System (CEMS), which

20. Appendix B describes each of these datasets in further detail, while also describing how I merge datasets and construct key variables.

21. Beginning in the late 1990s, electricity market restructuring forced many vertically integrated utilities to sell their coal plants. Most of these divestments were in just four states (Pennsylvania, Illinois, Ohio, and New York), and the vast majority occurred before the start of the fracking boom. Previous research has focused directly on the effects of coal plant divestment (Cicala (2015); Chan et al. (2017)), and these studies have obtained non-disclosure agreements with EIA to unmask prices for non-utility plants.

report hourly generation and emissions for all large fossil fuel generating units.²² This allows me to estimate each coal unit’s probability of generating in a given hour, conditional on the relative prices of coal and natural gas.

I use detailed GIS data on the U.S. rail network published by the Bureau of Transportation Statistics (BTS). I apply a graph algorithm to find the shortest path along the rail network connecting each coal-producing county to each power plant destination.²³ Then, I calculate the proportion of each shortest route owned or operated by each of the 7 Class I rail carriers, and assign a “dominant” (modal) carrier to each route. BTS also reports the average traffic density of rail lines, which I integrate over the full length of each route as a proxy for rail network congestion. To control for time-series variation in shipping costs, I use the Association of American Railroads’ (AAR) monthly fuel price index, which compiles survey data on actual diesel prices paid by railroad operators.²⁴ I also calculate each plant’s proximity to a navigable river, Great Lake, or coastline. This allows me to identify the subset of plants with the option to receive waterborne coal deliveries.

4.2 Defining Rail Captiveness

I treat each power plant’s location on the rail network as pre-determined. Plant geographic locations are obviously fixed, and I exclude the few plants constructed during my 2002–2015 sample period. More importantly, the rail network was largely static throughout this period, with very few changes in the ownership or trackage rights of individual rail lines.²⁵ This means that each coal plant has faced the same set of potential rail carriers.

I partition plants into two time-invariant groups, “captive” and “non-captive”, based on their locations on the rail network and the counties from which they purchase coal. I define the “captive” group as plants that either (i) become unconnected from the rail network after removing any single Class I carrier, or (ii) become unconnected from *all* observed trading partners after removing the dominant carrier along each origin-destination route. For example, consider a plant that only purchases coal from two counties in Wyoming. I

22. Most coal plants comprise multiple generating units (or boilers), each with different operating constraints and variable costs.

23. Hughes (2011) applies a similar algorithm to calculate the shortest rail distance, and finds that GIS-derived shortest distances closely approximate (yet slightly understate) actual rail shipping distances. Appendix C describes this shortest-distance algorithm in detail.

24. Diesel purchases represent roughly half of railroads’ total variable transportation costs.

25. The last merger between Class I carriers occurred in 1999. 99.3 percent of Class I track mileage maintained constant ownership since 2006, the earliest year of available BTS data.

classify this plant as captive if a single Class I carrier controls all terminal nodes within a 7-mile radius. This plant is also captive if, after removing the dominant carrier on its shortest route to *each* Wyoming county, the new shortest routes *both* increase by over 300 miles.²⁶

5 Empirical Strategy

The goal of my empirical analysis is to estimate how coal-by-rail markups vary across heterogeneous power plants. First, I estimate how markup *levels* vary across captive vs. non-captive plants, using a nearest-neighbor matching strategy. Next, I estimate how gas price shocks cause markups to *change* differentially across two sources of heterogeneity: (i) variation in market power, and (ii) variation in plants' sensitivity to competition from natural gas generation. These difference-in-differences estimates combine data on plants' transportation options with plant-specific coal demand estimates, using the comparative static $\frac{d\mu}{dZ}$ from my oligopoly model as a guide.

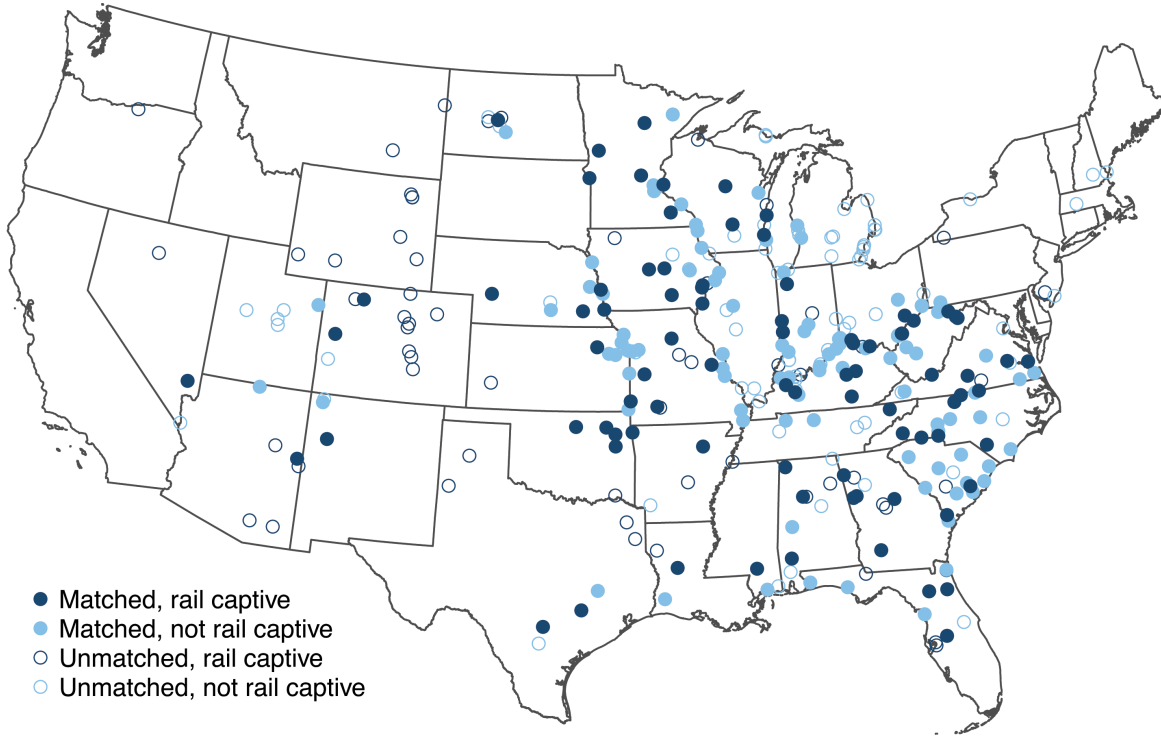
5.1 Matching Captive vs. Non-Captive Plants

Rail captiveness is not randomly assigned, and we might expect captive and non-captive plants to differ systematically. Because captiveness depends on geography, plants of each group might be spatially concentrated and burn similar types of coal, have similar operating characteristics, or face similar market conditions. Any observed or unobserved differences that are correlated with rail markups would lead to biased estimates of the markup differential between captive and non-captive groups.

To address this identification challenge, I apply a nearest-neighbor matching strategy in the tradition of Heckman, Ichimura, and Todd (1997). I match each plant in the “captive” group to k plants in the “non-captive” group with the closest geographic proximity, with a maximum distance of 200 miles between matched plants. I also force exact matches on plants' preferred coal type from the pre-fracking period (2002–2006). I omit plants that rely exclusively on non-rail shipping modes (i.e. barges and trucks), and plants that are not utility-owned (for which I do not observe coal price data). I assign nearest-neighbor weights

26. Each of these thresholds is quite conservative. 7 miles is the 95th percentile of plants' distance to the closest rail node. A 300-mile increase in distance implies a 20 percent increase over the median delivered coal price. Appendix C discusses my choice of thresholds, while Appendix E reports sensitivity analysis.

Figure 4: Nearest Neighbor Matching



Notes: This map plots the location of all 324 utility-owned coal power plants in the full 2002–2015 sample. Rail captive plants are in navy and non-captive plants are in light blue; plants in the matched sample are filled, while unmatched plants are hollow. Match criteria: up to k nearest neighbors ($k = 3$), with a maximum distance of 200 miles; exact matches on coal rank; and removing non-utility and non-rail plants.

as the inverse of the number of matches. For example, if a non-captive plant is one of 3 matches for captive plant A and one of 2 matches for captive plant B, it receives a weight of $\frac{1}{3} + \frac{1}{2} = \frac{5}{6}$. Matched captive plants receive a weight of 1, while all unmatched plants receive a weight of 0. This ensures that weights sum to twice the number of matched captive plants.²⁷

Figure 4 maps the full sample of captive and non-captive plants. This reveals broad geographic overlap, except for certain regions where plants tend to be either only captive (i.e. the western Great Plains) or only non-captive (i.e. Michigan). This map also shows the outcome of my matching algorithm with $k = 3$ nearest neighbors. Matched plants tend to be located near multiple plants of the opposite group. While I allow matches up to 200 miles apart, most matched plants have a nearest neighbor within 100 miles.²⁸

Table 1 presents summary statistics for both groups of plants from 2002–2006, including plant and coal characteristics. In the full sample, captive plants systematically purchase more

27. My approach closely resembles that of Cicala (2015), who also imposes a maximum match distance of 200 miles. He and Lee (2016) similarly match coal plants that sell gypsum byproduct vs. non-gypsum plants.

28. Appendix E.1 reports sensitivity analysis on the number of matches (k) and maximum match distance.

Table 1: Summary Statistics – Captive vs. Non-Captive Coal Plants (2002–2006)

A. Plant characteristics	Full sample			Matched sample ($k = 3$)		
	Rail Captive	Not Rail Captive	Difference in means	Rail Captive	Not Rail Captive	Difference in means
Total plant capacity (MW)	908.19 (780.72)	865.73 (742.77)	42.46 [0.57]	908.59 (69.45)	912.04 (87.77)	-3.45 [0.98]
Coal-fired capacity (MW)	806.13 (738.72)	760.84 (703.91)	45.29 [0.52]	812.92 (67.15)	797.67 (85.43)	15.26 [0.89]
Number of coal units	2.36 (1.32)	2.62 (1.64)	-0.26 [0.08]*	2.56 (0.15)	2.57 (0.14)	-0.01 [0.97]
Coal unit vintage (year)	1968.85 (13.90)	1962.88 (13.34)	5.97 [0.00]***	1966.60 (1.40)	1962.35 (1.43)	4.25 [0.03]**
Annual capacity factor	0.63 (0.17)	0.60 (0.17)	0.03 [0.04]**	0.63 (0.01)	0.63 (0.01)	0.00 [0.92]
Heat rate (MMBTU/MWh)	11.09 (1.40)	11.06 (1.52)	0.03 [0.86]	10.96 (0.14)	10.74 (0.13)	0.22 [0.25]
Scrubber installed	0.36 (0.48)	0.29 (0.45)	0.07 [0.12]	0.27 (0.05)	0.28 (0.06)	-0.01 [0.91]
Market participant	0.49 (0.50)	0.71 (0.46)	-0.22 [0.00]***	0.47 (0.05)	0.51 (0.06)	-0.04 [0.62]

B. Coal deliveries	Full sample			Matched sample ($k = 3$)		
	Rail Captive	Not Rail Captive	Difference in means	Rail Captive	Not Rail Captive	Difference in means
Annual deliveries (million MMBTU)	48.82 (47.90)	44.00 (43.70)	4.82 [0.29]	47.28 (4.13)	43.68 (4.93)	3.60 [0.58]
Sulfur content (lbs/MMBTU)	0.87 (0.61)	1.02 (0.79)	-0.15 [0.03]**	0.80 (0.06)	0.84 (0.07)	-0.05 [0.58]
Ash content (lbs/MMBTU)	8.46 (4.21)	8.96 (8.24)	-0.50 [0.46]	7.99 (0.36)	7.75 (0.31)	0.24 [0.61]
Share spot market	0.19 (0.29)	0.19 (0.25)	-0.00 [0.87]	0.19 (0.03)	0.17 (0.03)	0.02 [0.56]
Share of contracts expiring ≤ 2 years	0.22 (0.25)	0.24 (0.26)	-0.01 [0.61]	0.20 (0.02)	0.19 (0.02)	0.01 [0.80]
Share sub-bituminous	0.41 (0.47)	0.31 (0.42)	0.10 [0.03]**	0.42 (0.05)	0.40 (0.06)	0.02 [0.80]
Average rail distance (miles)	554.91 (385.90)	620.34 (417.90)	-65.43 [0.12]	569.77 (40.51)	585.04 (41.22)	-15.27 [0.79]

C. Number of plants	Full sample			Matched sample ($k = 3$)		
	Rail Captive	Not Rail Captive	Total	Rail Captive	Not Rail Captive	Total
Preferred coal rank: bituminous	94	149	243	49	60	109
Preferred coal rank: sub-bituminous	77	76	153	35	35	70
Non-rail plants	17	14	31	0	0	0
Utility plants	148	176	324	86	96	182
Total plants	190	240	430	86	96	182

Notes: This table compares coal plants captive to a single rail carrier to non-captive plants. The left three columns include all CEMS electric power plants with at least 1 coal generating unit in 2002–2015, and with at least 1 reported coal delivery in both 2002–2006 and 2007–2015. The right three columns are weighted by nearest-neighbor matches, with unmatched plants receiving weight 0, matched captive plants receiving weight 1, and matched non-captive plants weighted by the inverse of the number of matches. Matching criteria: up to k nearest neighbors ($k = 3$), with a maximum distance of 200 miles; exact matches on preferred coal rank; and removing non-utility and non-rail plants. Standard deviations are in parentheses, and p -values [in brackets] are clustered at the plant level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

low sulfur, sub-bituminous coal, and are more likely to participate in wholesale electricity markets. However, nearest-neighbor weights adjust the distributions of both captive and non-captive plants such that they are no longer statistically different. While geographic proximity alone does not ensure that matched non-captive plants can serve as plausible controls, distance-based matching yields covariate balance across a wide range of observables.

5.2 Markup Levels

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

$$P_{ojms} = \tau D_j + \beta_C \mathbf{C}_{ojms} + S(\mathbf{T}_{ojms}; \beta_T) + \beta_X \mathbf{X}_{jm} + \eta_o + \delta_m + \varepsilon_{ojms} \quad (5)$$

P_{ojms} is the average delivered price of coal, for shipment s from county o to plant j in month m . D_j is an indicator for rail captiveness, and the coefficient τ estimates the average differential markup faced by captive plants, relative to non-captive plants. Since I do not directly observe coal markups, I use price as an outcome variable and control for shipment-level variation in both commodity value and shipping costs (i.e. $C_o + S(\mathbf{T}_{oj})$ in Equation (2)). I also include nearest-neighbor weights and plant-specific controls (\mathbf{X}_{jm}), in the style of a doubly robust estimator (Wooldridge (2007)). After controlling for both county fixed effects (η_o) and month-of-sample fixed effects (δ_m), the remaining variation in P_{ojms} is close to the variation I would use in the ideal experiment: comparing the price of two identical coal shipments to two otherwise identical coal plants, where only one plant is rail captive.

The matrix \mathbf{C}_{ojms} controls for determinants of commodity value, including average heat, sulfur, and ash content; coal rank; and the average annual mine-mouth price for coal produced in county o . \mathbf{C}_{ojms} also includes dummies for spot market transactions and contracts expiring within 2 years, since plants pay higher prices for (longer) contracts that minimize the risk of supply disruptions.²⁹ The matrix \mathbf{T}_{ojms} controls for the two primary factors affecting the cost of transporting coal: the shortest rail distance between coal county o and plant j , and the average diesel price paid by rail carriers in month m . \mathbf{T}_{ojms} also includes the

29. Wolak (1996) finds that coal plants simultaneously purchase on long-term contracts and the spot market, even as contract purchases tend to have higher delivered prices. Jha (2017) estimates that the average regulated coal plant is willing to trade a \$1.66 increase in expected delivered coal price for a \$1.00 decrease in the standard deviation of delivered coal price.

log of shipment size, as marginal freight costs are likely decreasing in tons shipped. Finally, \mathbf{T}_{ojms} includes the proportion of the shortest oj route on rail lines with high traffic density, to allow for higher costs along more congested routes. The function $S(\cdot)$ flexibly models shipping costs as the four-way interaction of the components of \mathbf{T}_{ojms} .³⁰ The matrix \mathbf{X}_{jm} controls for both predetermined and time-varying plant characteristics, including each of the variables in panel A of Table 1. I cluster standard errors at the plant level, which allows for arbitrary within-plant serial correlation. In order to adjust the distribution of captive vs. non-captive plants, while also inflating each observation by the quantity of coal transacted, I weight by the product of the nearest-neighbor weights and shipment size.³¹

To interpret $\hat{\tau}$ as causal, D_j must be uncorrelated with plant unobservables, after nearest-neighbor matching and conditioning on observable characteristics in \mathbf{X}_{jm} . Captiveness is geographically predetermined, and 85 percent of matched plants predate the 1980 Staggers Act, which effectively legalized rail price discrimination. These plants likely could not have strategically influenced their degree of captiveness. It is also unlikely that coal plant unobservables impacted railroads' decisions to consolidate and increase market power, given that individual coal plants are small relative to the rail network. Spurious correlations could violate this identifying assumption if, for example, captive plants tended to have less sophisticated managers. As such a violation is unlikely, I interpret $\hat{\tau}$ as the causal effect of rail captiveness on markups. My subsequent results do not hinge on this interpretation.

5.3 Coal Demand Estimation

My theoretical framework illustrates how changes in markups likely depend on coal plants' sensitivity to the gas price. In order to account for this additional source of variation, I estimate plant-specific coal demand curves. In most settings with detailed data on production functions, this demand would follow from applying the Envelope Theorem to the firm's profit function at different factor prices, and inverting its production function to solve for the corresponding input quantity. However, four features of electricity markets make this approach infeasible for deriving coal demand.

30. Appendix E demonstrates that Equation (5) is not sensitive to the components of \mathbf{C}_{ojms} and \mathbf{T}_{ojms} , which supports my interpretation of $\hat{\tau}$ as an average differential *markup*. Misspecification of these cost controls would mean that $\hat{\tau}$ might confound markups *and* cost differences.

31. Observations in EIA's coal delivery data vary substantially by size, and this enables me to estimate the differential markup for the average ton of coal shipped, giving each ton of coal equal weight. Equation (5) also includes the log of shipment size as a linear control, as part of \mathbf{T}_{ojms} .

First, regulated coal plants do not behave as short-run profit-maximizers; their profits are calibrated to a fixed rate-of-return, and they are not the residual claimants on marginal costs of coal purchases or marginal revenues from electricity production. Second, for plants that do not participate in wholesale electricity markets, production decisions depend on complex engineering algorithms designed to balance the electricity grid; such algorithms may not define a marginal electricity price. Third, it would be extremely difficult to characterize the spatial and temporal constraints of the grid, where supply must respond to instantaneous changes in electricity demand, subject to available transmission capacity. Fourth, coal plants face their own dynamic operating constraints; delayed startup/shutdown decisions, ramping constraints, and maintenance outages would imply a unique state-dependent objective function for each plant.³²

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

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

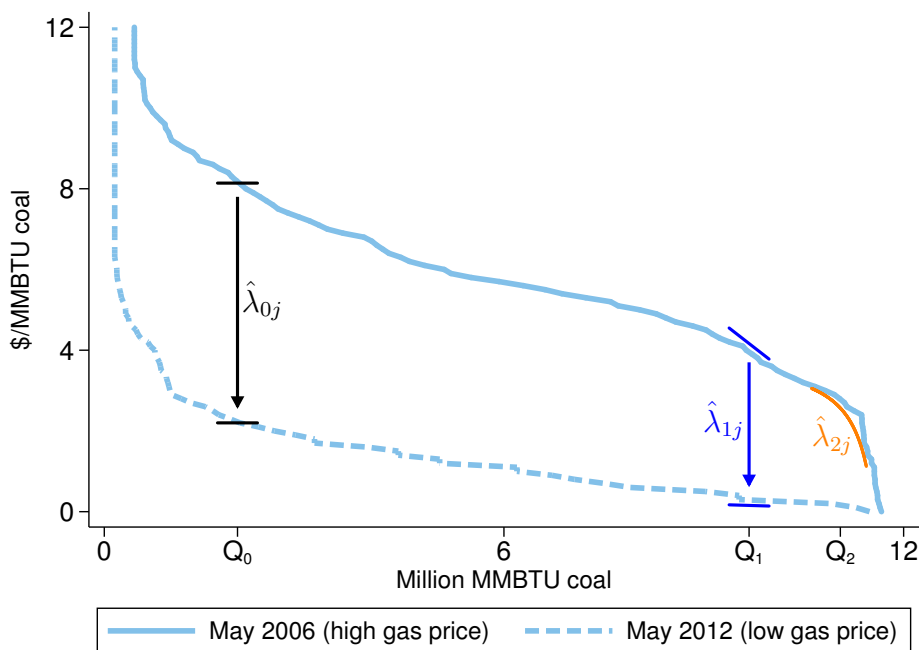
Each unit-specific regression predicts unit u 's generation conditional on *aggregate* fossil electricity generation in u 's market region (G_{uh} , in discrete bins b), the daily ratio of u 's marginal costs relative to the marginal costs of gas generators (CR_{ud}), and a matrix of controls (\mathbf{G}_{uh}).³⁴

32. A large body of research addresses each of these issues. Fowlie (2010) finds that rate-of-return regulation distorts coal plants' incentives to adopt least-cost pollution abatement strategies, while Cicala (2015) finds that regulated plants do not minimize coal purchase costs. Cicala (2017) demonstrates that the lack of a marginal electricity price signal leads to large allocative inefficiencies under non-market dispatch. Borenstein, Bushnell, and Stoft (2000); and Davis and Hausman (2016) demonstrate how transmission constraints directly impact electricity market outcomes. Mansur (2008) and Reguant (2014) focus on power plants' dynamic production constraints and non-convex operating costs.

33. I use the term "policy function" because I treat coal demand estimation as a prediction problem, rather than estimating an optimized function of price or quantity.

34. $CF_{uh} \in [0, 1]$ by construction, and $CF_{uh} = 0$ if unit u does not operate in hour h . G_{uh} sums hourly CEMS generation across all units in u 's market region. The cost ratio CR_{ud} divides unit u 's marginal cost (including coal price and marginal environmental costs) by the average marginal cost of gas units in the same PCA. Generation bins b allow for a flexible relationship between CF_{uh} , G_{uh} , and CR_{ud} . \mathbf{G}_{uh} includes the daily sum, maximum, minimum, and standard deviation of G_{uh} ; the daily maximum temperature; hour-of-day, quarter-of-year, and year-of-sample fixed effects; and year dummies interacted with the daily sum of G_{uh} . Appendix D describes my demand estimation procedure in further detail, and conducts sensitivity analysis on Equation (6).

Figure 5: Coal Demand Estimation Example



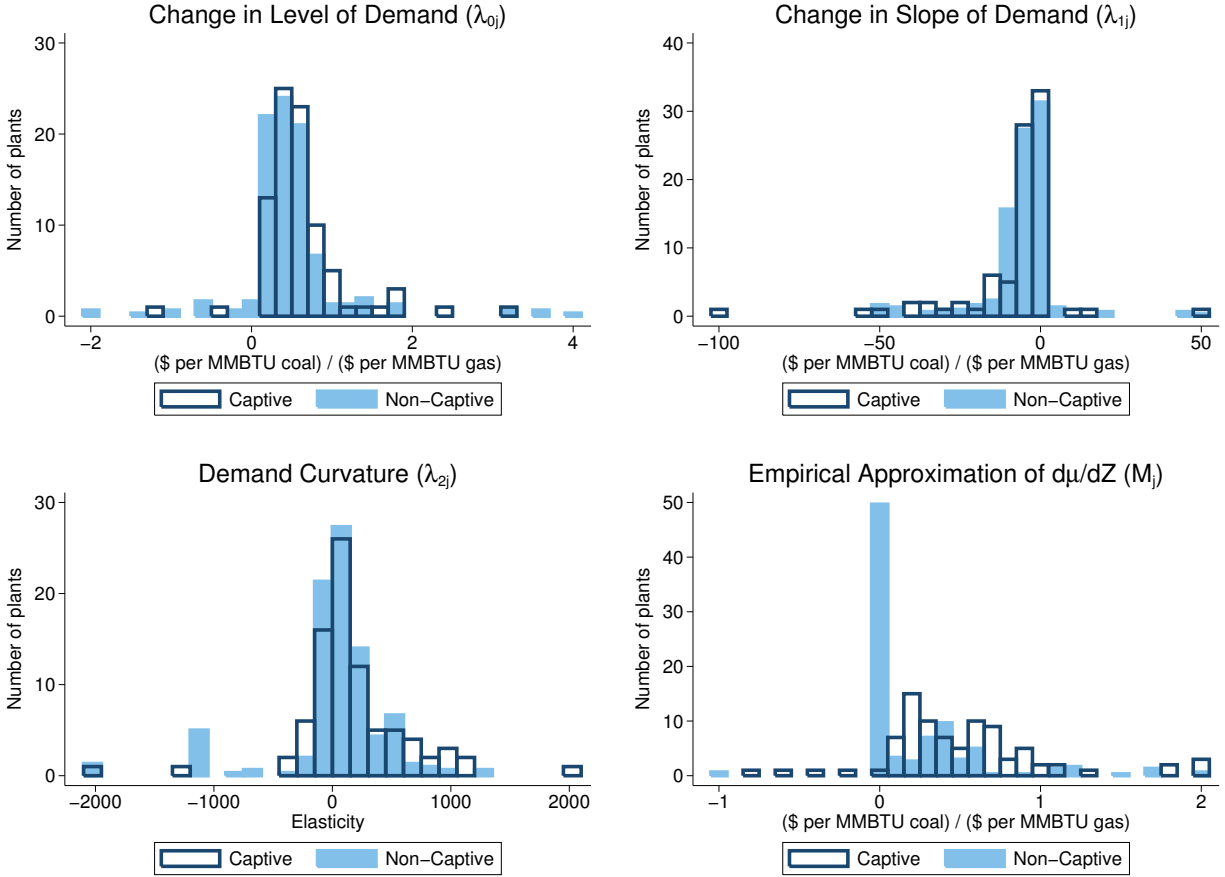
Notes: This figure plots two estimated coal demand curves for a representative plant j . The solid curve shows plant j 's monthly coal demand for May 2006, when the gas price was high; the dashed curve shows its demand for May 2012, when the gas price was low. I estimate how a change in gas price affects the level ($\hat{\lambda}_{0j}$) and slope ($\hat{\lambda}_{1j}$) of plant j 's coal demand. In this simplified example, a decrease in gas price decreases the level of inverse demand (i.e., $\hat{\lambda}_{0j} > 0$ at Q_0) and makes its slope less negative (i.e., $\hat{\lambda}_{1j} < 0$ at Q_1). The third parameter $\hat{\lambda}_{2j}$ estimates the average (local) curvature of coal demand, which is concave (i.e., $\hat{\lambda}_{2j} > 0$) at Q_2 . My algorithm estimates all three parameters at plant j 's observed coal quantities, using estimated demand curves and gas prices for all months. Appendix D.1 describes this algorithm in further detail.

After estimating Equation (6) for each unit, I solve each fitted model in terms of the coal price (a component of CR_{ud}), and derive the counterfactual prices for which $\widehat{CF}_{uh} = 0.5$. These are the predicted coal prices at which unit u would have been exactly marginal in electricity supply, in each hour h . I integrate the distribution of counterfactual coal prices across all hours in each month and across each of plant j 's units, transforming unit-specific capacity factors into their corresponding coal quantities. This yields price-quantity mappings, or coal demand curves for each plant-month. Finally, I use monthly gas price variation to estimate three parameters that map directly to my comparative static in Equation (3):

$$\hat{\lambda}_{0j} \sim \frac{\partial P_{oj}}{\partial Z_j} \quad \hat{\lambda}_{1j} \sim \frac{\partial^2 P_{oj}}{\partial Q_{oj} \partial Z_j} Q_j \quad \hat{\lambda}_{2j} \frac{\theta_{oj}}{N_{oj}} \sim E_{D_{oj}} \quad (7)$$

$\hat{\lambda}_{0j}$ and $\hat{\lambda}_{1j}$ estimate how gas price affects the *level* and *slope* of plant j 's inverse demand, respectively. $\hat{\lambda}_{2j}$ estimates the average *curvature* of plant j 's inverse demand, where $E_{D_{oj}}$ is

Figure 6: Coal Demand Estimation Results



Notes: These histograms report the distributions of estimated demand parameters ($\hat{\lambda}_{0j}$, $\hat{\lambda}_{1j}$, $\hat{\lambda}_{2j}$ from Equation (7)), and the empirical approximation of the comparative static $\frac{d\mu}{dZ}$ from Equation (8). Each histogram includes one observation per plant and applies nearest neighbor weights ($k = 3$). Matching criteria: up to k nearest neighbors, with a maximum distance of 200 miles; exact matches on coal rank; and removing non-utility and non-rail plants. The outermost bins are bottom-coded and top-coded, for ease of presentation. These outliers likely reflect idiosyncratic factors (other than the gas price) that have affected coal demand. Each histogram includes 86 captive plants and 96 matched non-captive plants.

the elasticity of the slope of inverse demand (as defined in Equation (4)). Figure 5 describes these parameters graphically, using two estimated demand curves from one plant.

Figure 6 presents the estimated distributions of these parameters, separately for captive and non-captive plants. The distribution of $\hat{\lambda}_{0j}$ has a median of 0.46, which implies that for a \$1 per MMBTU decrease in gas prices, coal prices would need to fall by \$0.46 per MMBTU to maintain the median plant’s baseline coal consumption.³⁵ These estimates reveal considerable heterogeneity across coal plants, with $\hat{\lambda}_{0j} \in [0, 1]$ for 86 percent of plants,

35. Unlike my regressions on coal shipments where the rail carrier’s relevant price is in dollars per ton, these demand parameters use dollars per MMBTU, in order to denominate coal in terms of its energy content (i.e. its value to power plants as a fuel input). BTU content varies substantially across coal shipments, with a mean (standard deviation) of 19.7 (3.4) MMBTU/ton in my estimation sample.

implying reasonable elasticities of substitution between coal and gas. $\hat{\lambda}_{1j} < 0$ for 88 percent of plants, implying that low gas prices have caused most plants' inverse coal demand to become less steep. As gas prices have fallen, most coal plants have become more marginal in electricity markets, which has made them more coal-price-elastic. Finally, $\hat{\lambda}_{2j} > 0$ for 67 percent of plants, suggesting that coal demand tends to be (locally) concave.

Importantly, $\hat{\lambda}_{0j}$, $\hat{\lambda}_{1j}$, and $\hat{\lambda}_{2j}$ are the outcome of a linear prediction algorithm that imposes no assumptions on plant j 's objective function. I do assume that plants with multiple generating units operate these units independently, that Equation (6) is not misspecified, and that counterfactual coal consumption in each hour is either zero or at maximum capacity. Counterfactual coal prices also hold the rest of the electricity market constant, including the coal prices faced by other plants.³⁶ This means that my demand estimates could not predict the effects of a common shock to coal commodity prices. However, they *can* predict variation in plant j 's idiosyncratic opportunity cost of coal — the very type of price changes that occur when a rail carrier reoptimizes plant-specific markups.

5.4 Markup Changes

Changes in plant j 's markups will depend on both its coal demand and the extent to which it faces market power in coal shipping. Recall the comparative static from Equation (3):

$$\frac{d\mu_{oj}}{dZ_j} = \frac{\frac{\partial P_{oj}}{\partial Z_j}(2 + E_{D_{oj}} - N_{oj}) - \frac{\partial^2 P_{oj}}{\partial Q_{oj} \partial Z_j} Q_{oj} \theta_{oj}}{2 + E_{D_{oj}}}$$

Using this theory as a guide, I can construct an empirical approximation for the change in plant j 's average markup, μ_j , that should result from a change the gas price, Z_j :

$$\Rightarrow M_j \equiv \frac{\hat{\lambda}_{0j} \left[D_j + \hat{\lambda}_{2j}(1 - W_j)(2 - D_j)^{-1} \right] - \hat{\lambda}_{1j}(1 - W_j)}{2 + \hat{\lambda}_{2j}(1 - W_j)(2 - D_j)^{-1}} \quad (8)$$

This variable M_j combines my estimated demand parameters ($\hat{\lambda}_{0j}$, $\hat{\lambda}_{1j}$, and $\hat{\lambda}_{2j}$) with data on plant j 's transport market structure. I can translate the captiveness indicator D_j into a binary version of N_{oj} , where $\hat{N}_j = 2 - D_j$: for plants captive to a single carrier, $\hat{N}_j = 1$;

36. In reality, rail carriers may jointly reoptimize markups across multiple coal plants selling into the same electricity market. If plants j markups move in the same direction as other plants' markups, then my estimates for plant j 's coal demand may be too large (small) at low (high) coal prices.

for non-captive plants, $\hat{N}_j = 2$.³⁷ Likewise, the option to receive waterborne deliveries can serve as a (crude) empirical proxy for the conduct parameter $\hat{\theta}_j = 1 - W_j$, where W_j is an indicator of water access. For plants without a water option, $\hat{\theta}_j = 1$, consistent with Cournot oligopoly; for plants with the ability to receive barge deliveries that bypass rail carriers, $\hat{\theta}_j = 0$, consistent with a competitive fringe.

The bottom-right panel in Figure 6 reports the distributions of M_j , separately for captive and non-captive plants. This underscores two potentially important shortcomings of a difference-in-differences strategy that would split plants based on captiveness alone. First, M_j varies considerably across captive plants, with a median of 0.43 and an interquartile range of [0.23, 0.70]. This suggests that a captive plant at the 75th percentile of this distribution should have experienced three times larger changes in markups, compared to a captive plant at the 25th percentile. The binary indicator $D_j = 1$ obscures this key heterogeneity.³⁸ Second, while the distribution of M_j for non-captive plants has a large mass at 0, M_j is positive for 43 percent of non-captive plants. In fact, for many non-captive plants, M_j is larger than for their captive counterparts, implying that W_j , $\hat{\lambda}_{0j}$, $\hat{\lambda}_{1j}$, and $\hat{\lambda}_{2j}$ combine to outweigh $D_j = 0$. This suggests that non-captive plants likely also experienced decreases in markups as gas prices fell, even though these plants face less rail market power.

Using M_j as cross-sectional variation, I estimate markup changes using a lagged difference-in-differences model:

$$P_{ojms} = \tau M_j \cdot Z_{m-L}^{HH} + \sum_{\ell=0}^{L-1} \tau_{\ell} M_j \cdot \Delta Z_{m-\ell}^{HH} \dots + \beta_C \mathbf{C}_{ojms} + S(\mathbf{T}_{ojms}; \beta_T) + \beta_X \mathbf{X}_{jm} + \eta_{oj} + \delta_m + \varepsilon_{ojms} \quad (9)$$

Z_m^{HH} is the average Henry Hub spot price in month m , and $\Delta Z_m^{HH} = Z_m^{HH} - Z_{m-1}^{HH}$. The coefficient of interest τ captures the cumulative effect of a \$1/MMBTU change in gas price, over $L = 36$ months. Each lagged coefficient τ_{ℓ} captures the cumulative effect after ℓ months, for a plant with $M_j = 1$ relative to a plant with $M_j = 0$. I accommodate delayed effects of gas prices on markups because most coal deliveries occur on long-term contracts, which may

37. Given that the rail network is close to a duopoly in both the western and eastern U.S., I assign $\hat{N}_j = 2$ for non-captive plants. Using this formulation, $E_{D_j} \sim \hat{\lambda}_{2j}(1 - W_j)(2 - D_j)^{-1}$.

38. In the absence of a theory model, one might interact D_j with the gas price to estimate a reduced-form difference-in-differences version of Equation (5). Appendix E.3 reports results from this model, which yields imprecise point estimates close to zero. This is unsurprising, given that these regressions rely on captiveness alone, while ignoring heterogeneity in coal demand.

be slow to adjust to changing market conditions.³⁹ Equation (9) includes origin-destination fixed effects η_{oj} , which control for the average markup of all shipments from county o to plant j . As in Equation (5), \mathbf{C}_{ojms} and $S(\mathbf{T}_{ojms})$ control for the costs of each coal shipment, while \mathbf{X}_{jm} controls for time-varying plant characteristics. I cluster standard errors by plant, and weight observations by the product of nearest-neighbor weights and shipment size.

I interpret $\hat{\tau}$ as the causal effect of gas price changes on coal-by-rail markups. The key identifying assumption is that gas price changes are uncorrelated with other factors affecting the differential trajectory of coal markups, after controlling for time-varying plant characteristics in \mathbf{X}_{jm} . Technological advances of the fracking boom were unrelated to coal mining costs; the Henry Hub spot price is also uncorrelated with U.S. diesel prices, which drive coal shipping costs.⁴⁰ A violation of parallel trends would occur if a coal plant unobservable that is correlated with coal prices (e.g., how electricity regulators monitor plants' coal purchase costs) changed differentially for plants with high vs. low M_j .⁴¹

6 Results

6.1 Markup Levels

Table 2 reports results from estimating Equation (5), which demonstrate that captive plants indeed face higher markups than their matched non-captive counterparts. Point estimates of \$2–3 translate to average differential markups of 4–7 percent, on an average delivered price of \$37–41 per ton for non-captive plants. This implies that markups for captive plants contribute 13–23 percent of the spatial gap between mine-mouth prices (averaging \$22–24

39. It is common to allow for delayed pass-through, in settings where price changes may not be instantaneous. Pouliot, Smith, and Stock (2017) use the same differenced lag structure to estimate delayed pass-through in the market for renewable fuel credits. This is algebraically equivalent to the standard (non-differenced) distributed lag model, $\sum_{\ell=0}^L \beta_{\ell} D_j \cdot Z_{m-\ell}^{HH}$, where $\sum_{\ell=0}^L \beta_{\ell} = \tau$. Coal prices for long-term contracts should be quite sticky, even though many coal contracts include flexible price-adjustment provisions that enable rail carriers to partially re-optimize markups before these contracts expire (Joskow (1988); Kosnik and Lange (2011)). I estimate Equation (9) separately for contract and spot market shipments.

40. The fracking boom may have impacted local labor markets in certain coal mining regions, and oil-by-rail shipping increased congestion in certain portions of the rail network. However, my results are largely consistent across geographic regions, making these violations of parallel trends unlikely. During 2002–2015, the correlation between Henry Hub and U.S. average monthly diesel prices was -0.01 . If these two price series were correlated, I would worry about multicollinearity between Z_m^{HH} and diesel prices in \mathbf{T}_{ojms} .

41. Christian and Barrett (2017) show that even spurious time trends can induce bias for a difference-in-differences treatment variable that interacts a cross-sectional characteristic with an exogenous time series. In Appendix E.4.2, I show that my data exhibit parallel pre-trends in M_j . I also use interacted time fixed effects to help rule out potential time-varying confounders.

Table 2: Markup Levels – Captive vs. Non-Captive Coal Plants

	(1)	(2)	(3)	(4)	(5)	(6)
$1[\text{Captive}]_j$	2.174*** (0.743)	1.797*** (0.609)	1.656*** (0.580)	2.819*** (0.799)	2.270*** (0.650)	2.006*** (0.607)
k nearest neighbors	1	3	5	1	3	5
Balanced panel				Yes	Yes	Yes
Coal county FEs	Yes	Yes	Yes	Yes	Yes	Yes
Month-of-sample FEs	Yes	Yes	Yes	Yes	Yes	Yes
Mean of dep var	40.74	37.97	39.42	40.42	37.04	38.43
Plants	141	182	193	96	126	133
Captive plants	86	86	86	60	60	60
Plant-county-months	32,995	44,888	47,985	27,338	36,984	39,548
Observations	66,473	88,511	93,950	57,914	76,831	81,670

Notes: Each regression estimates Equation (5) at the coal shipment level, with delivered coal price (\$ per short ton) as the dependent variable. I control for shipping costs using the 4-way interaction of rail distance, diesel price, tons shipped, and rail traffic density. Plant- and delivery-specific controls are listed in panels A and B of Table 1, respectively. I also control for the average annual coal price from the originating county, each plant’s distance to its closest rail terminal, and baseload natural gas capacity in each plant’s PCA. Matching criteria: up to k nearest neighbors within a 200-mile radius; exact matches on coal rank; and removing non-utility and non-rail plants. Regressions apply nearest-neighbor weights, and also weight each observation by the quantity of coal transacted. Balanced panels include plants receiving at least 1 shipment in each sample year (2002–2015). I report means of the dependent variable for non-captive plants only. Standard errors are clustered by plant. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

per ton) and delivered prices. Given that the indicator D_j applies an arbitrary threshold to define captiveness, and non-captive plants likely face nonzero markups, these estimated *differentials* likely understate the average markup *level* faced by captive plants.

These results are robust to the number of nearest-neighbor matches, which I vary from $k = 1$ to $k = 5$. I construct my estimation sample to exclude the (very few) coal plants constructed after 2001, but many plants retired during my 2002–2015 sample period. If these plants’ exit decisions were correlated with their delivered coal prices, which affected their ability to compete with low-cost natural gas, endogenous exit could bias my estimates in an unbalanced panel. Columns (4)–(6) in Table 2 restrict the sample to plants receiving at least 1 delivery in each year.⁴² While this removes 31 percent of plants, point estimates remain statistically significant and increase slightly in magnitude.

42. This is not a fully balanced panel. Coal shipments are lumpy, and many active plants do not report deliveries in each month. I “balance” the panel to mitigate any confounding effects from plant exit, not to take a stand on the timing of coal deliveries. Olley and Pakes (1996) demonstrate that bias due to endogenous exit may remain even after balancing the panel, if exit is correlated with unobserved firm productivity. This is not a concern in my setting, as I directly control for each plant’s productivity (i.e. its heat rate).

Table 3: Markup Levels – Captiveness and Water Access

	(1)	(2)	(3)	(4)	(5)	(6)
1[Captive, No Water] _j	3.694*** [1.5, 5.9]	2.368*** [0.7, 4.0]	2.338*** [0.8, 3.9]	4.846*** [2.7, 7.0]	3.089*** [1.3, 4.9]	2.910*** [1.3, 4.5]
1[Captive, Water] _j	2.110** [0.0, 4.2]	0.844 [-0.7, 2.4]	0.819 [-0.6, 2.3]	3.038*** [1.0, 5.1]	1.480* [-0.2, 3.2]	1.326* [-0.2, 2.9]
1[Non-Captive, No Water] _j	2.135 [-0.5, 4.8]	0.293 [-1.6, 2.2]	0.512 [-1.3, 2.4]	2.878** [0.0, 5.7]	0.700 [-1.4, 2.8]	0.860 [-1.2, 2.9]
<i>k</i> nearest neighbors	1	3	5	1	3	5
Balanced panel				Yes	Yes	Yes
Coal county FEs	Yes	Yes	Yes	Yes	Yes	Yes
Month-of-sample FEs	Yes	Yes	Yes	Yes	Yes	Yes
Mean of dep var	33.91	34.31	34.81	33.31	33.91	34.53
Plants	141	182	193	96	126	133
Captive plants	86	86	86	60	60	60
Water-access plants	56	84	88	39	62	64
Plant-county-months	32,995	44,888	47,985	27,338	36,984	39,548
Observations	66,473	88,511	93,950	57,914	76,831	81,670

Notes: Each regression estimates Equation (5) at the coal shipment level, with delivered coal price (\$ per short ton) as the dependent variable, and interacting rail captiveness with an indicator for water access (i.e. plants having the option to purchase coal via barge). I control for shipping costs using the 4-way interaction of rail distance, diesel price, tons shipped, and rail traffic density. Plant- and delivery-specific controls are listed in panels A and B of Table 1, respectively. I also control for the average annual coal price from the originating county, each plant’s distance to its closest rail terminal, and baseload natural gas capacity in each plant’s PCA. Matching criteria: up to *k* nearest neighbors within a 200-mile radius; exact matches on coal rank; and removing non-utility and non-rail plants. Regressions apply nearest-neighbor weights, and also weight each observation by the quantity of coal transacted. Balanced panels include plants receiving at least 1 shipment in each sample year (2002–2015). I report means of the dependent variable for the omitted group (non-captive plants with a water option). Standard errors are clustered by plant, and 95 percent confidence intervals are in brackets. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 3 interacts rail captiveness with another pre-determined plant characteristic likely to affect markups: an indicator for access to waterborne shipments. These results reveal differential markups of \$2–5 per ton for captive plants with no coal-by-barge option, relative to plants with the most competitive shipping regimes (i.e. the omitted group with multiple rail carriers and barges). While these point estimates are more sensitive to the number of nearest neighbors and to balancing the panel, they show that the markup distortion may be as large as \$4–5 per ton, or 11–15 percent of the average delivered coal price. This implies markups of up to 30–42 above rail carriers’ marginal shipping costs. My point estimates are consistent with the magnitudes of Busse and Keohane (2007), who estimate differential markups of \$4 per ton for coal shipped from Wyoming in the late 1990s. My analysis

demonstrates that coal-by-rail price discrimination, due to geographic variation in market power, has persisted through recent years and across multiple coal-producing regions.⁴³

6.2 Markup Changes

Having established that markup *levels* vary across coal plants, I now estimate heterogeneous *changes* in markups due to changes in the gas price. I identify markup changes using cross-sectional variation in M_j , a predictor of how railroads reoptimize markups due to variation in both transport market power and plant-specific coal demand. Table 4 reports cumulative effects across 36 months, from estimating Equation (9). I find positive, statistically significant point estimates, which are qualitatively consistent with the predictions of my oligopoly model. As gas prices fell during the fracking boom, rail carriers reoptimized markups *heterogeneously* across plants, causing markups to decrease *more* for plants with greater M_j .

I estimate separate regressions for long-term contracts vs. spot-market shipments, as the timing of markup changes will likely differ by transaction type. Rail carriers should be able to reoptimize spot markups more quickly than contract markups; however, markup changes for relatively less flexible contracts should be more persistent, yielding larger cumulative effects. Figure 7 plots lagged coefficients $\hat{\tau}_\ell$ for each regression in Table 4, where each coefficient represents the cumulative effect through ℓ months. This reveals that contract markups begin to adjust 6 months after a gas price shock, with delayed effects that accumulate until month 18 and persist through month 36. By contrast, the effects for spot market shipments attenuate and lose significance after 36 months. This likely reflects a difference in transaction costs, as bilateral contract negotiations facilitate greater opportunity for price discrimination than posted spot shipping rates.⁴⁴ Table 5 reveals that differential changes in markups were more pronounced for bituminous coal and for plants in the South and East, even though my results remain statistically significant across coal ranks and plant regions.

Table 4 implies that for a \$1/MMBTU decrease in gas price, markups decreased by \$1.21–1.57/ton for a plant with $M_j = 1$, compared to plants with $M_j = 0$. However, a literal interpretation of $M_j = 1$ would suggest an effect size of \$1/MMBTU of coal, equivalent to

43. These results are consistent if I restrict the sample to by origin or destination region. Appendix E.2 reports these results, along with additional sensitivity analyses.

44. Railroads likely invest more resources in reoptimizing less flexible, longer-lived contracts. Whereas pooled and contract results are robust to alternative nearest-neighbor weights, fixed effects, and controls, spot market results are not. Appendix E.4 reports these sensitivities, and additional robustness checks.

Table 4: Markup Changes – Demand Parameters Interacted with Gas Price

	Both Types		Contracts		Spot Market	
	(1)	(2)	(3)	(4)	(5)	(6)
$(\Delta \widehat{\text{Markup}})_j \times (\text{Gas Price})_m$	1.411*** (0.403)	1.205*** (0.407)	1.574*** (0.298)	1.543*** (0.288)	1.620 (1.116)	0.491 (0.920)
k nearest neighbors	1	3	1	3	1	3
Balanced panel	Yes	Yes	Yes	Yes	Yes	Yes
Plant \times county FEs	Yes	Yes	Yes	Yes	Yes	Yes
Month-of-sample FEs	Yes	Yes	Yes	Yes	Yes	Yes
Mean of dep var	33.27	33.86	32.07	32.67	41.76	41.19
Plants	93	123	91	121	84	114
Plant-county-months	25,876	35,467	21,848	29,653	6,753	9,586
Observations	55,941	74,805	44,431	58,954	11,430	15,738

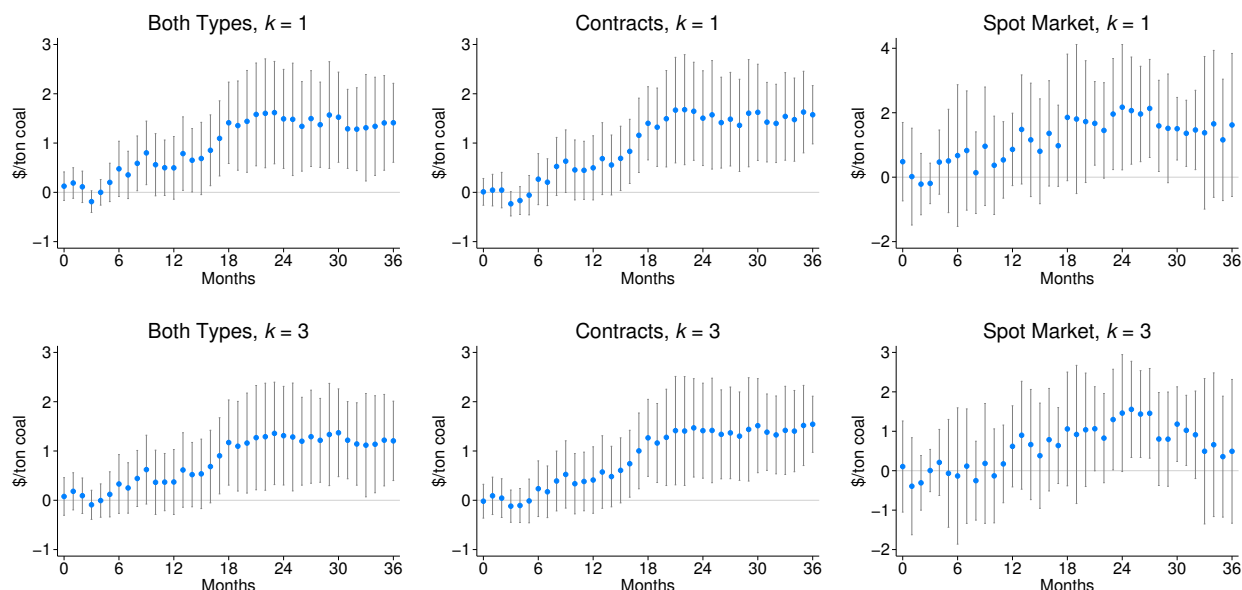
Notes: Each regression estimates Equation (9) at the coal shipment level, with delivered coal price (\$ per ton) as the dependent variable. The first 2 columns pool long-term contracts and spot market shipments, while the middle and right columns split the sample by transaction type. The DD treatment variable interacts the empirical approximation M_j of the comparative static $\frac{d\mu}{dZ}$ (from Equation (8)) with the Henry Hub average monthly spot price for natural gas, using a lagged-differenced model with $L = 36$ lags. This table reports estimates for $\hat{\tau}$, or the cumulative effects over $L = 36$ months. Figure 7 plots each lagged coefficient $\hat{\tau}_\ell$, which reports the cumulative effect through ℓ months. M_j is in units of \$ per MMBTU of coal, and BTU content ranges from 14–30 MMBTU/ton. I control for shipping costs using the 4-way interaction of rail distance, diesel price, tons shipped, and rail traffic density. Plant- and delivery-specific controls are listed in panels A and B of Table 1, respectively. I also control for the average annual coal price from the originating county, and baseload natural gas capacity in each plant’s PCA. Matching criteria: up to k nearest neighbors within a 200-mile radius; exact matches on coal rank; and removing non-utility and non-rail plants. Regressions apply nearest-neighbor weights, and also weight each observation by the quantity of coal transacted. Balanced panels include plants receiving at least 1 shipment in each sample year (2002–2015). I report means of the dependent variable for plants with $M_j = 0$. Standard errors are clustered by plant. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

roughly \$20/ton of coal.⁴⁵ This mismatch in magnitudes underscores the shortcomings of my oligopoly model, which does not account for railroad regulation or the threat of arbitrage. Suppose that binding regulation limited the maximum markup to \$5/ton, but unconstrained markups would have been \$14/ton before the fracking boom and \$4/ton after the fracking boom. In this case, I would only observe a \$1/ton decrease in markups, rather than the \$10/ton decrease predicted by an unconstrained model. This explains why M_j does not generate accurate *quantitative* predictions of how markups change.⁴⁶ However, my results in Table 4 demonstrate that M_j can generate accurate *qualitative* predictions, by capturing cross-sectional heterogeneity in both transport market structure and coal demand.

45. Recall that M_j is in units of \$/MMBTU of coal, but the dependent variable in Tables 4–6 is the coal price in \$/ton of coal. BTU content varies across coal shipments, with an average of 19.7 MMBTU/ton.

46. My model also assumes that markups are independent across plants. If rail carriers jointly optimize markups across multiple markets, this would attenuate my estimates. M_j linearly extrapolates to a finite change in gas price, which may overstate markup changes if markups approach their lower bound of zero. Measurement error in M_j may also attenuate my estimates of $\hat{\tau}$, and I address this issue in Appendix E.4.3.

Figure 7: Markup Changes – Cumulative Effects



Notes: This figure plots 36 lag-differenced DD coefficient estimates ($\hat{\tau}_0, \dots, \hat{\tau}_{35}$) and $\hat{\tau}$, from 6 separate regressions of Equation (9) with $L = 36$ lags. Each panel corresponds to a column in Table 4, which reports $\hat{\tau}$ only (i.e. the rightmost point in each graph). Each coefficient estimates the interaction of M_j with the ℓ -month lagged difference in natural gas prices ($\Delta Z_{m-\ell}^{HH}$), such that each dot represents the cumulative effect through ℓ months. Whiskers denote 95 percent confidence intervals for each point estimate, with standard errors clustered by plant. See the notes below Table 4 for further details on the estimation.

To better capture how magnitudes vary across the full range of coal plants, I discretize M_j into five indicator variables corresponding to the quintiles of its positive support.⁴⁷ Table 6 reports results for quintiles 2–5, revealing magnitudes that increase monotonically in M_j . The omitted category is the 42 percent of plants with $M_j \leq 0.22$, most of which are non-captive and located on navigable waterways. If I assume no markup changes for these omitted plants, the point estimates in Table 6 represent the average change in markups for plants in each quintile.⁴⁸ I find statistically significant decreases in contract markups, for the 42 percent of plants in quintiles 3–5 (i.e., $M_j > 0.35$). For the 14 percent of plants in quintile 5 (i.e., $M_j > 0.70$), a \$1/MMBTU gas price decrease caused average markups to fall by \$1.07–1.34/ton, and caused contract markups to fall by \$1.35–1.49/ton. Given that gas prices fell by \$4/MMBTU during the fracking boom, and that average markup *levels*

47. Each “quintile” includes 14–16 percent of plants, because $M_j \leq 0$ for 28 percent of plants. Tables 4–6 omit 3 plants with extremely low/high M_j (i.e. $|M_j| > 2$), which almost certainly reflect errors in estimating these plants’ demand parameters. Appendix E.4.3 reports results including these outliers, and my point estimates attenuate slightly but largely retain statistical significance.

48. Non-captive plants with a water delivery option benefit from the most competitive shipping regimes. These plants likely faced markups close to zero, prior to the fracking boom. Hence, if gas price changes caused any markup decreases for omitted plants, such changes should have been relatively small.

Table 5: Markup Changes – Contract Shipments, Split Samples

	Split by Coal Grade		Removing Plants in Region		
	Bituminous (1)	Sub-bituminous (2)	West (3)	Midwest (4)	South/East (5)
$(\widehat{\Delta \text{Markup}})_j \times (\text{Gas Price})_m$	2.137*** (0.303)	0.899** (0.417)	1.472*** (0.271)	1.808*** (0.235)	1.045*** (0.343)
k nearest neighbors	3	3	3	3	3
Balanced panel	Yes	Yes	Yes	Yes	Yes
Plant \times county FEs	Yes	Yes	Yes	Yes	Yes
Month-of-sample FEs	Yes	Yes	Yes	Yes	Yes
Mean of dep var	48.27	22.02	40.01	29.54	28.98
Plants	76	68	85	64	93
Plant-county-months	18,692	10,758	23,181	17,487	18,720
Observations	31,591	26,995	44,497	35,105	38,394

Notes: Each regression estimates Equation (9) at the coal shipment level for contract shipments only, with delivered coal price (\$ per ton) as the dependent variable. Columns (1)–(2) include only shipments of a given coal grade. I partition all plants into three regions, and Columns (3)–(5) each exclude plants from a given region. The DD treatment variable interacts the empirical approximation M_j of the comparative static $\frac{d\mu}{dZ}$ (from Equation (8)) with the Henry Hub average monthly spot price for natural gas, using a lagged-differenced model with $L = 36$ lags. This table reports estimates for $\hat{\tau}$, or the cumulative effects over $L = 36$ months. M_j is in units of \$ per MMBTU of coal, and BTU content ranges from 14–30 MMBTU/ton. I control for shipping costs using the 4-way interaction of rail distance, diesel price, tons shipped, and rail traffic density. Delivery-specific controls are listed in panel B of Table 1. I omit plant-specific controls (except for ISO and scrubber indicators), in order to reduce the number of covariates to less than the number of plant clusters (pooled results in Table 4 are not sensitive to inclusion of these controls). I also control for the average annual coal price from the originating county, and baseload natural gas capacity in each plant’s PCA. Matching criteria: up to k nearest neighbors within a 200-mile radius; exact matches on coal rank; and removing non-utility and non-rail plants. Regressions apply nearest-neighbor weights, and also weight each observation by the quantity of coal transacted. Balanced panels include plants receiving at least 1 shipment in each sample year (2002–2015). I report means of the dependent variable for plants with $M_j = 0$. Standard errors are clustered by plant. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

were \$2–5/ton, these magnitudes imply that rail carriers have heterogeneously reoptimized markups, to eliminate most of the market power distortion for a subset of plants.

My results demonstrate that market power exists in coal transportation, and that rail carriers strategically reoptimize coal markups in a manner consistent with profit maximization. Rail market power arises primarily from coal’s geographic specificity, as the production and consumption of coal are both highly locationally constrained. This affords rail intermediaries substantial bargaining power, and coal’s low value-to-weight ratio increases the premium on transportation access. To identify market power, I exploit price dispersion due to the lack of spatial arbitrage in coal deliveries. While this feature is likely unique to coal markets, the features that generate market power in coal shipping—geographic specificity and high freight costs—exist in many other commodity markets (e.g., Covert and Kellogg (2017) on crude oil; Hortaçsu and Syverson (2007) on cement).

Table 6: Markup Changes – Quantiles of $\widehat{\Delta \text{Markup}}$

	Both Types		Contracts		Spot Market	
	(1)	(2)	(3)	(4)	(5)	(6)
$\mathbf{1}[M_j \in (0.22, 0.35]] \times (\text{Gas Price})_m$	0.044 (0.206)	-0.017 (0.178)	0.245 (0.208)	0.280 (0.172)	-0.757 (0.487)	-1.259** (0.621)
$\mathbf{1}[M_j \in (0.35, 0.52]] \times (\text{Gas Price})_m$	0.264 (0.230)	0.144 (0.201)	0.498** (0.237)	0.475** (0.209)	-0.243 (0.454)	-0.954* (0.572)
$\mathbf{1}[M_j \in (0.52, 0.70]] \times (\text{Gas Price})_m$	0.804*** (0.276)	0.641** (0.247)	0.800*** (0.255)	0.733*** (0.218)	1.085 (0.852)	0.345 (0.953)
$\mathbf{1}[M_j \in (0.70, 2.00]] \times (\text{Gas Price})_m$	1.336*** (0.493)	1.065*** (0.376)	1.494*** (0.470)	1.347*** (0.370)	1.367 (1.126)	0.092 (0.979)
<i>k</i> nearest neighbors	1	3	1	3	1	3
Balanced panel	Yes	Yes	Yes	Yes	Yes	Yes
Plant \times county FEs	Yes	Yes	Yes	Yes	Yes	Yes
Month-of-sample FEs	Yes	Yes	Yes	Yes	Yes	Yes
Mean of dep var	31.46	34.85	30.83	34.67	33.55	35.45
Plants	93	123	91	121	84	114
Plant-county-months	25,876	35,467	21,848	29,653	6,753	9,586
Observations	55,941	74,805	44,431	58,954	11,430	15,738

Notes: Each regression estimates a modified Equation (9) at the coal shipment level, with delivered coal price (\$ per ton) as the dependent variable, and $L = 36$ lags. Instead of interacting the L -month lagged gas price with a continuous M_j to estimate the coefficient of interest $\hat{\tau}$, I estimate four $\hat{\tau}$'s using indicator variables for quintiles of M_j 's positive support. The omitted group is the first quintile, plus all plants with $M_j \leq 0$. This table reports the average cumulative change in markups caused by a \$1/MMBTU change in gas price, for plants in a given quintile relative to the omitted group. Lag-differenced coefficients $\hat{\tau}_\ell$ still use a continuous M_j interaction. The first 2 columns include long-term contracts and spot market shipments, while the middle and right columns split the sample by transaction type. I control for shipping costs using the 4-way interaction of rail distance, diesel price, tons shipped, and rail traffic density. Plant- and delivery-specific controls are listed in panels A and B of Table 1, respectively. I also control for the average annual coal price from the originating county, and baseload natural gas capacity in each plant's PCA. Matching criteria: up to k nearest neighbors within a 200-mile radius; exact matches on coal rank; and removing non-utility and non-rail plants. Regressions apply nearest-neighbor weights, and also weight each observation by the quantity of coal transacted. Balanced panels include plants receiving at least 1 shipment in each sample year (2002–2015). I report means of the dependent variable for the omitted group of plants, with $M_j \leq 0.22$. Standard errors are clustered by plant. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

7 Implications for Climate Policy

7.1 Markup Size vs. External Costs of Coal

Given that coal intermediaries exercise market power, a carbon tax has the potential to restrict aggregate coal consumption below the social optimum. Buchanan (1969) demonstrates that if a market power distortion is sufficiently large relative to external costs, a Pigouvian tax could actually reduce welfare. Based on my estimates in Table 3, I can reject differential coal markups greater than \$7 per ton, relative to the “most-competitive” omitted category (i.e. plants with multiple rail carriers and a coal-by-barge option). \$7 per short ton of coal is roughly equivalent to \$2–5 per metric ton of CO₂, which is far below recent social cost of carbon estimates of \$50 per metric ton (Interagency Working Group (2016); Revesz et al. (2017)). Hence, coal markups are an order of magnitude smaller than the carbon externality.

This means that the welfare gains from Pigouvian taxation would likely dwarf any welfare loss from exacerbating the market power distortion (echoing Oates and Strassmann (1984)).⁴⁹ However, real-world carbon prices typically range from \$3–30 per metric ton of CO₂, which is far below estimated climate damages of \$50 per metric ton (Carl and Fedor (2016)). Under such a suboptimally low carbon tax, the presence of coal markups should increase welfare by internalizing an additional fraction of marginal damages. Even under a tax equal to marginal climate damages, markups could potentially *still* increase welfare by internalizing local air pollution damages from coal combustion (Levy, Baxter, and Schwartz (2009); Muller, Mendelsohn, and Nordhaus (2011)).

7.2 Pass-Through of Implicit Carbon Tax

A negative gas price shock makes coal plants less competitive in electricity supply, and a tax on CO₂ emissions similarly disadvantages coal, the more carbon-intensive fuel. Cullen and Mansur (2017) argue that under reasonable assumptions, the coal-to-gas price ratio is a sufficient statistic for CO₂ emissions from the electricity sector under a counterfactual carbon tax. If electricity demand is perfectly inelastic, and only coal or gas generators can be marginal in electricity supply, then a short-run change in relative fuel prices should yield the same emissions outcomes as the equivalent carbon tax.⁵⁰

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

50. Appendix A.2 provides further detail on the assumptions underlying this section, along with derivations of implied pass-through rates.

Using Cullen and Mansur’s framework, a gas price change ΔZ yields the same fuel cost ratio (CR) as the carbon tax t (suppressing plant j subscripts):

$$CR = \frac{MC_{coal}}{MC_{gas}} = \frac{P}{Z + \Delta Z} = \frac{P + t E_{coal}}{Z + t E_{gas}} \quad (10)$$

where P is the coal price paid by power plants, MC_{fuel} are marginal costs per MMBTU, and E_{fuel} are fuel-specific CO₂ emissions factors (in metric tons CO₂/MMBTU). My empirical results demonstrate that P is not fixed, and I can rewrite this expression to allow coal markups (μ) to endogenously respond to ΔZ :

$$CR = \frac{MC_{coal}}{MC_{gas}} = \frac{P + \Delta\mu}{Z + \Delta Z} = \frac{P + \rho t E_{coal}}{Z + t E_{gas}} \quad (11)$$

ρ is the pass-through rate of the implicit tax t . If markups do not change ($\Delta\mu = 0$), the cost ratio reflects full pass-through of the carbon tax, or $\rho = 1$ as in Equation (10). If markups decrease in response to a negative gas price shock ($\Delta Z < 0$ causing $\Delta\mu < 0$, consistent with Table 6), then pass-through of t is incomplete, and $\rho < 1$. By reoptimizing markups during the fracking boom, rail carriers effectively lowered the coal-to-gas cost ratio, which led to incomplete pass-through of the negative shock to coal demand.

I can rearrange Equation (11) to translate my point estimates from Table 6 into implied tax pass-through rates, setting $\Delta Z = 1$ and using average fuel prices from the start of the fracking boom. Table 7 reports pass-through rates for the five quintiles of M_j , assuming full pass-through ($\rho_j = 1$) for the omitted group in Table 6. Panel A reveals substantial heterogeneity both across and within plant groups. While most plants experience full pass-through, the 14 percent of plants with the highest M_j have pass-through ranging from $\rho_j = 0.41$ to $\rho_j = 0.90$, with an average rate of $\rho_j = 0.80$. Isolating long-term contracts implies even lower pass-through rates, due to larger changes in markups for contract shipments. Panel B calculates pass-through rates for a cost ratio inclusive of marginal environmental costs, to account for marginal costs of pollution abatement already incurred by coal and gas plants.⁵¹ While pre-existing environmental policies reduce the size of the implicit carbon tax, implied pass-through rates increase only slightly. To my knowledge, this is the first

51. For plants covered by SO₂, NO_x, or CO₂ allowance trading regimes, I monetize each generating unit’s emissions rates using prevailing allowance prices. I also include variable costs of operating pollution control devices, such as scrubbers.

Table 7: Heterogeneous Pass-Through of Implied Carbon Tax

	Plant Group (Quintile of M_j)				
	(1)	(2)	(3)	(4)	(5)
$\Delta \widehat{\text{Markup}} (M_j)$	(-2.00, 0.22]	(0.22, 0.35]	(0.35, 0.52]	(0.52, 0.70]	(0.70, 2.00]
Share of plants	0.42	0.16	0.14	0.14	0.14
A. Fuel prices only					
ρ_j , all shipments	1.00	1.00 [1.00, 1.01] <i>(0.90, 1.11)</i>	0.97 [0.91, 0.98] <i>(0.88, 1.06)</i>	0.89 [0.79, 0.93] <i>(0.80, 0.98)</i>	0.80 [0.41, 0.90] <i>(0.67, 0.93)</i>
ρ_j , contract shipments	1.00	0.93 [0.83, 0.96] <i>(0.82, 1.03)</i>	0.89 [0.71, 0.94] <i>(0.80, 0.99)</i>	0.87 [0.76, 0.92] <i>(0.79, 0.95)</i>	0.75 [0.27, 0.87] <i>(0.62, 0.88)</i>
B. Fuel + environmental costs					
ρ_j , all shipments	1.00	1.00 [1.00, 1.01] <i>(0.91, 1.10)</i>	0.97 [0.93, 0.99] <i>(0.89, 1.05)</i>	0.90 [0.82, 0.95] <i>(0.83, 0.98)</i>	0.83 [0.51, 0.91] <i>(0.72, 0.94)</i>
ρ_j , contract shipments	1.00	0.93 [0.85, 0.96] <i>(0.84, 1.02)</i>	0.91 [0.75, 0.95] <i>(0.82, 0.99)</i>	0.89 [0.78, 0.95] <i>(0.82, 0.95)</i>	0.78 [0.39, 0.88] <i>(0.67, 0.90)</i>

Notes: This table converts point estimates from Table 6 into pass-through rates of an implied carbon tax, for $k = 3$ nearest neighbors. I rearrange Equation (11) to solve for ρ as a function $\Delta\mu$; assign P and Z their average prices from 2007–08, the beginning of the fracking boom; and set $E_{coal} = 0.095$ and $E_{gas} = 0.053$, their average emissions rates in metric tons CO₂ per MMBTU. For a \$1/MMBTU change in gas price (i.e. $\Delta Z = 1$) and assuming full pass-through ($\rho_j = 1$) in the omitted group, Table 6 point estimates the average change in markups for each group (i.e. $\Delta\mu$, after converting from \$/ton to \$/MMBTU of coal). Average pass-through rates are in bold, and square brackets report the minimum and maximum pass-through rates for plants in each group. I report the 95 percent confidence intervals for the *average* (bolded) pass-through rates in parentheses and italics (calculated from the confidence interval of each $\hat{\tau}$ estimate). Whereas Panel A follows Equation (11) by only including fuel prices, Panel B includes environmental costs following Equation (A21) in Appendix A.2.

empirical evidence that predicts heterogeneous and incomplete pass-through of a carbon tax in either U.S. coal markets or the U.S. electricity sector.⁵²

My results contribute to a growing body of research finding heterogeneous pass-through of price-based climate policies. Previous work has shown that variation in market structure either across industries (Ganapati, Shapiro, and Walker (2016)), or across space within an industry (Pouliot, Smith, and Stock (2017)), can generate substantial heterogeneity in pass-

52. Chu, Holladay, and LaRiviere (2017) estimate incomplete pass-through from coal spot prices to delivered coal prices; the authors caution that their analysis is not predictive of long-term price changes that would occur under a carbon tax. Kim, Chattopadhyay, and Park (2010) conceptually illustrate how variation in power plants' costs may lead to incomplete carbon tax pass-through.

through rates.⁵³ Similarly, I find that heterogeneous pass-through of a carbon tax in U.S. coal markets would arise largely from spatial variation in the competitiveness of coal shipping. However, coal markups also adjust heterogeneously to plant-specific demand shocks; I am able to detect incomplete pass-through only after accounting for this second dimension of heterogeneity.

Muehlegger and Sweeney (2017) estimate incomplete pass-through of firm-specific cost shocks in petroleum refining, but full pass-through of cost shocks that are common across firms. Given that CO₂ emissions rates vary substantially across refineries, this implies that a carbon tax would likely lead to heterogeneous pass-through by inducing variation in refinery-specific costs. Coal power plants exhibit similar variation in CO₂ emissions rates, and I likewise find incomplete pass-through resulting from plant-specific shocks. My analysis is the first to show that pass-through of a carbon tax in the U.S. electricity sector may be heterogeneous and incomplete, in part due to variation in plants' sensitivity to relative cost shocks. By contrast, Fabra and Reguant (2014) estimate full pass-through of carbon prices in the Spanish wholesale electricity market, which they attribute to highly correlated emissions cost shocks among marginal plants. My results demonstrate that variation in upstream market power may weaken the correlation in cost shocks across plants, potentially leading to heterogeneous pass-through despite an *average* pass-through rate close to 1.⁵⁴

7.3 Heterogeneous Tax Incidence

Weyl and Fabinger (2013) show how pass-through under imperfect competition is closely linked to economic incidence. In fact, the pass-through rate (ρ), conduct parameter (θ), and number of symmetric firms (N) are sufficient to characterize the incidence (I) of a tax (t):

$$I = \frac{dCS/dt}{dPS/dt} = \frac{\rho}{1 - (1 - \theta/N)\rho} \quad (12)$$

53. Ganapati, Shapiro, and Walker (2016) find heterogeneous energy cost pass-through for manufacturing industries under imperfect competition. Pouliot, Smith, and Stock (2017) find lower pass-through rates of renewable fuel credits in less integrated market regions (see also Knittel, Meiselman, and Stock (2017); Li and Stock (2017)). Spatial and temporal variation in production capacity can also lead to heterogeneous pass-through in petroleum refining (Marion and Muehlegger (2011)); however incomplete pass-through caused by capacity constraints does not necessarily reflect market power (Borenstein and Kellogg (2014)).

54. Fabra and Reguant (2014) also attribute their finding of full pass-through to inelastic aggregate electricity demand and high-frequency uniform-price auctions. Nazifi (2016) similarly predicts full pass-through of a carbon tax in the Australian electricity market. Stolper (2016a) shows how the jurisdictional borders of energy taxes can create also induce variation in firm-specific costs, resulting in heterogeneous pass-through.

where CS and PS are consumer and producer surplus. Lower pass-through rates imply that consumers (i.e. coal plants) bear relatively less of the tax burden than producers (i.e. rail carriers). For a given pass-through rate ρ , a less competitive market structure (i.e. greater θ/N) implies that rail oligopolists bear a relatively greater tax burden.

Given the range of pass-through estimates in Table 7, the incidence of a carbon tax would likely vary substantially across coal plants. During the fracking boom, plants in the least competitive shipping regimes that were most sensitive to gas prices paid only 44 percent of the burden of low gas prices (i.e., $\rho_j = 0.80$, $\theta_j = 1$, $N_j = 1$). By contrast, plants with full pass-through and a water delivery option paid 100 percent of the lost surplus in coal shipping (i.e., $\rho_j = 1$, $\theta_j = 0$).⁵⁵ The complete incidence calculation would also include lost profits in electricity markets, which would depend in part on plants' ability to pass marginal emissions costs through to wholesale electricity prices (Fabra and Reguant (2014)).

My results add to a nascent body of evidence that the assumption of homogeneous incidence can obscure the true distributional impacts of energy taxes. Stolper (2016b) uncovers heterogeneous tax incidence for Spanish transportation fuels, which renders a seemingly regressive tax unambiguously progressive. Similarly, Ganapati, Shapiro, and Walker (2016) show that a carbon tax appears less regressive after accounting for variation in the competitiveness of intermediate product markets. In my setting, heterogeneous incidence suggests that under a carbon tax, certain coal plants would stand to lose relatively less than others.⁵⁶ By shifting a share of the tax burden further upstream from electricity consumers, market imperfections in coal shipping may also reduce the regressivity of a carbon tax.

7.4 Counterfactuals

Figure 2 illustrates how U.S. electricity generation has shifted away from coal as gas prices have fallen, and several previous studies have sought to quantify the environmental benefits of coal-to-gas switching induced by the fracking boom. For example, Knittel, Metaxoglou, and Trindade (2015) estimate that the 70-percent drop in gas prices between 2008 and

55. The share of the burden borne by plant j is $\frac{I_j}{1+I_j} = \frac{\rho_j}{1+(\theta_j/N_j)\rho_j}$. Appendix A.3 contains a more detailed discussion of implied carbon tax incidence as it pertains to my theoretical framework.

56. All coal plants would likely see profits decrease under a carbon tax, yet some plants would likely bear *relatively* less tax burden in the short run. Muehlegger and Sweeney (2017) find that a carbon tax on petroleum refiners would imply heterogeneous firm-specific cost shocks, also creating relative winners and losers.

2012 caused CO₂ emissions from electricity generation to fall by up to 19–33 percent.⁵⁷ My analysis is the first to show that coal markups have adjusted to partially offset this change in relative fuel prices. This suggests that if coal markups had not changed, the fracking boom could have yielded even greater reductions in CO₂ emissions from electricity generation.

To quantify how decreasing coal markups may have limited coal-to-gas switching in the short run, I first estimate a time series regression similar to Equation (6) for each coal generating unit. This calibrates a semi-parametric relationship between each unit’s electricity generation and the coal-to-gas cost ratio. Next, I use this fitted model to infer predicted generation under two counterfactual cost ratios: (1) if the fracking boom never happened and gas prices had remained high; and (2) if the fracking boom did happen but coal markups had remained fixed. Converting predicted generation into predicted CO₂ emissions and summing across all coal units, I can calculate short-run CO₂ abatement from the fracking boom both with and without changes to coal markups.⁵⁸

This exercise suggests that decreases in coal markups eroded roughly 8 percent of the fracking boom’s short-run abatement potential. Based on these calculations, CO₂ emissions fell by 4.5 percent during the fracking boom, as a result of short-run coal-to-gas substitution alone. However, if coal markups had not changed, this would have been a 4.9-percent emissions reduction. These numbers capture only short-run changes on the intensive margin of fossil generation, and several other margins have contributed to the 20–25 percent decrease in CO₂ emissions from electricity.⁵⁹ Even so, they suggest that falling coal markups meaningfully reduced the environmental benefits of low natural gas prices, with unrealized CO₂ abatement equal to \$2.4 billion in climate damages.⁶⁰

57. Holladay and LaRiviere (2017) estimate short-run changes in the marginal CO₂ emissions rates that vary substantially across electricity market regions. Fell and Kaffine (forthcoming) attribute the decline in coal generation to a combination of low natural gas prices and increased wind generating capacity. Wolak (2016) applies a general equilibrium framework to estimate the fracking boom’s impact on global coal markets, and finds that increased U.S. coal exports have not led to increases in global CO₂ emissions.

58. This short-run exercise abstracts from changes to electric generating capacity. I assume that electricity demand is perfectly inelastic, with gas generation crowding out coal generation 1-for-1. Following Cullen and Mansur (2017), I include a cubic spline in the average cost ratio across all fossil generators in unit u ’s PCA. Unlike in Equation (6), I use the *average* cost ratio because I now want to allow unit u ’s generation to respond to changes in *other* units’ coal prices. Appendix F discusses these counterfactuals in further detail.

59. Many coal plants have invested in medium-run efficiency improvements (Linn, Mastrangelo, and Burtraw (2014)). On the capacity margin, the fracking boom has spurred investment in new gas plants (Brehm (2017)), while accelerating coal plant retirements (Meng (2016)). Increases in non-fossil generation have also crowded out conventional fossil generation. This counterfactual exercise also ignores the long-run dynamic effects of relative fuel prices changes (Cullen and Reynolds (2016)).

60. I monetize the difference between 4.5 and 4.9 percent abatement at \$50 per metric ton CO₂. Accounting for medium- and long-run margins would increase the value of unrealized abatement. Importantly, these calculations only consider the electricity sector. Low gas prices have increased CO₂ emissions from other end

Extrapolating to future climate policy, decreases in coal markups may similarly erode the environmental benefits of a carbon tax. However, this countervailing effect would likely disappear if the tax were sufficiently large, as markups should not decrease below zero. This suggests that existing retrospective analyses may underestimate CO₂ abatement under future climate policy. By not accounting for incomplete pass-through in coal markets, these studies likely understate the amount of coal displacement that would occur if a sufficiently stringent climate policy pushed delivered coal prices down to marginal cost.

8 Conclusion

This paper demonstrates that decreases in natural gas prices have caused decreases in coal markups. These effects vary substantially across coal-fired power plants, due to the interaction of heterogeneous transport market structures and plant-specific shocks to coal demand. While previous studies have documented market power in coal shipping, my analysis is the first to show that oligopolist rail carriers reoptimize markups due to heterogeneous changes in plants' coal demand. I also show that pass-through of a carbon tax in the electricity sector may be heterogeneous and incomplete, as railroads would likely reduce markups to effectively buffer a subset of coal plants against the tax. This has the potential to significantly erode the environmental benefits of climate policy, and incomplete pass-through would likely reduce welfare under a carbon tax smaller than marginal climate damages.

My analysis highlights the need for further research estimating pass-through of environmental taxes under imperfect competition. Markets for energy or energy-intensive products tend to be highly concentrated, and the assumption of perfect competition can generate both misguided welfare estimates and biased policy counterfactuals. In order to more fully characterize welfare under climate policy, future research should incorporate market imperfections in both upstream fuel markets (Gillingham et al. (2016)) and downstream electricity markets (Bushnell, Mansur, and Saravia (2008)). My analysis also underscores how heterogeneous market imperfections can generate heterogeneous pass-through of environmental taxes. If pass-through varies across polluting firms, then a uniform carbon price may not incentivize an efficient allocation of CO₂ abatement (Montgomery (1972)), and the optimal second-best climate policy may feature a non-uniform carbon tax.

uses (e.g. residential space heating) and methane leaks from gas drilling. These factors may have combined to outweigh fracking-induced CO₂ abatement from electricity generation (Hausman and Kellogg (2015)).

Future research should also investigate how coal-by-rail market power impacts climate policy outcomes in the medium-to-long run. For example, a carbon tax may incentivize investment in new coal transportation infrastructure, which would mitigate market power and reduce dispersion in delivered coal prices. My analysis largely ignores the coal mining sector, and it is important to consider how a carbon tax might disproportionately hurt labor markets in coal communities (Lobao et al. (2016)). Finally, similar market imperfections in coal transportation likely exist outside the U.S., due to coal’s geographic specificity and high transportation costs. Hence, market power in coal shipping may impact climate policy outcomes in the developing world, where coal consumption continues to rise (Wolfram, Shelef, and Gertler (2012)).

References

- 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.
- Barnett, A. H. 1980. “The Pigouvian Tax Rule under Monopoly.” *American Economic Review* 70 (5): 1037–1041.
- Bergquist, Lauren Falcao. 2017. “Pass-Through, Competition, and Entry in Agricultural Markets: Experimental Evidence from Kenya.” Working Paper.
- Borenstein, Severin, James Bushnell, and Steven Stoft. 2000. “The Competitiveness Effects of Transmission Capacity in a Deregulated Electricity Industry.” *RAND Journal of Economics* 31 (2): 294–325.
- Borenstein, Severin, and Ryan Kellogg. 2014. “The Incidence of an Oil Glut: Who Benefits from Cheap Crude Oil in the Midwest?” *Energy Journal* 35 (1): 15–33.
- Brehm, Paul. 2017. “Natural Gas Prices, Electric Generation Investment, and Greenhouse Gas Emissions.” Working Paper.
- Buchanan, James M. 1969. “External Diseconomies, Corrective Taxes, and Market Structure.” *American Economic Review* 59 (1): 174–177.
- Bushnell, James B., Erin T. Mansur, and Celeste Saravia. 2008. “Vertical Arrangements, Market Structure, and Competition: An Analysis of Restructured U.S. Electricity Markets.” *American Economic Review* 98 (1): 237–266.
- 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.

- Carl, Jeremy, and David Fedor. 2016. "Tracking Global Carbon Revenues: A Survey of Carbon Taxes Versus Cap-and-Trade in the Real World." *Energy Policy* 96: 50–77.
- Chan, H. Ron, Harrison Fell, Ian Lange, and Shanjun Li. 2017. "Efficiency and Environmental Impacts of Electricity Restructuring on Coal-fired Power Plants." *Journal of Environmental Economics and Management* 81: 1–18.
- Christian, Paul, and Christopher B. Barrett. 2017. "Revisiting the Effect of Food Aid on Conflict: A Methodological Caution." Working Paper.
- Chu, Yin, Scott Holladay, and Jacob LaRiviere. 2017. "Opportunity Cost Pass-Through from Fossil Fuel Market Prices to Procurement Costs of the U.S. Power Producers." Working Paper 2017-02. University of Tennessee, Department of Economics.
- Cicala, Steve. 2015. "When Does Regulation Distort Costs? Lessons from Fuel Procurement in U.S. Electricity Generation." *American Economic Review* 105 (1): 411–444.
- . 2017. "Imperfect Markets Versus Imperfect Regulation in U.S. Electricity Generation." NBER Working Paper 23053.
- Corts, Kenneth S. 1999. "Conduct Parameters and the Measurement of Market Power." *Journal of Econometrics* 88 (2): 227–250.
- Cournot, Augustin. 1838. "Of the Influence of Taxation on Commodities Produced under a Monopoly." In *Researches into the Mathematical Principles of the Theory of Wealth*, translated by Nathaniel T. Bacon, 67–78. London: Macmillan, 1897.
- Covert, Thomas R., and Ryan Kellogg. 2017. "Crude by Rail, Option Value, and Pipeline Investment." NBER Working Paper 23855.
- 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. 2016. "The Long Run Impact of Environmental Policies on Wholesale Electricity Markets: A Dynamic Competitive Analysis." Working Paper.
- Davis, Lucas W., and Erich Muehlegger. 2010. "Do Americans Consume Too Little Natural Gas? An Empirical Test of Marginal Cost Pricing." *RAND Journal of Economics* 41 (4): 791–810.
- 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.
- Fell, Harrison, and Daniel T. Kaffine. Forthcoming. "The Fall of Coal: Joint Effects of Fuel Prices and Renewables on Generation and Emissions." *American Economic Journal: Economic Policy*.
- Fitzgerald, Timothy. 2013. "Frackonomics: Some Economics of Hydraulic Fracturing." *Case Western Reserve Law Review* 63 (4): 1337–1362.

- Fowlie, Meredith. 2010. "Emissions Trading, Electricity Restructuring, and Investment in Pollution Abatement." *American Economic Review* 100 (3): 837–869.
- Fowlie, Meredith, Mar Reguant, and Stephen P. Ryan. 2016. "Market-Based Emissions Regulation and Industry Dynamics." *Journal of Political Economy* 124 (1): 249–302.
- Ganapati, Sharat. 2017. "The Modern Wholesaler: Global Sourcing, Domestic Distribution, and Scale Economies." Working Paper.
- Ganapati, Sharat, Joseph S. Shapiro, and Reed Walker. 2016. "The Incidence of Carbon Taxes in U.S. Manufacturing: Lessons from Energy Cost Pass-Through." NBER Working Paper 22281.
- Gillingham, Kenneth, James Bushnell, Meredith Fowlie, Michael Greenstone, Charles Kolstad, Alan Krupnick, Adele Morris, Richard Schmalensee, and James Stock. 2016. "Reforming the U.S. Coal Leasing Program." *Science* 354 (6316): 1096–1098.
- Hastings, Justine S. 2004. "Vertical Relationships and Competition in Retail Gasoline Markets: Empirical Evidence from Contract Changes in Southern California." *American Economic Review* 94 (1): 317–328.
- Hausman, Catherine, and Ryan Kellogg. 2015. "Welfare and Distributional Implications of Shale Gas." *Brookings Papers on Economic Activity* Spring: 71–125.
- He, Qingxin, and Jonathan M. Lee. 2016. "The Effect of Coal Combustion Byproducts on Price Discrimination by Upstream Industries." *International Journal of Industrial Organization* 44: 11–24.
- Heckman, James J., Hidehiko Ichimura, and Petra E. Todd. 1997. "Matching as an Econometric Evaluation Estimator: Evidence from Evaluating a Job Training Programme." *Review of Economic Studies* 64 (4): 605–654.
- 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.
- Hortaçsu, Ali, and Chad Syverson. 2007. "Cementing Relationships: Vertical Integration, Foreclosure, Productivity, and Prices." *Journal of Political Economy* 115 (2): 250–301.
- Hughes, Jonathan E. 2011. "The Higher Price of Cleaner Fuels: Market Power in the Rail Transport of Fuel Ethanol." *Journal of Environmental Economics and Management* 62 (2): 123–139.
- Interagency Working Group on Social Cost of Greenhouse Gases. 2016. "Technical Update of the Social Cost of Carbon for Regulatory Impact Analysis." United States Government, Technical Support Document.
- Jha, Akshaya. 2015. "Dynamic Regulatory Distortions: Coal Procurement at U.S. Power Plants." Working Paper.
- . 2017. "Regulatory Induced Risk Aversion: Coal Procurement at U.S. Power Plants." Working Paper.

- Joskow, Paul L. 1985. "Vertical Integration and Long-term Contracts: The Case of Coal-burning Electric Generating Plants." *Journal of Law, Economics, & Organization* 1 (1): 33–80.
- . 1987. "Contract Duration and Relationship-Specific Investments: Empirical Evidence from Coal Markets." *American Economic Review* 77 (1): 168–185.
- . 1988. "Price Adjustment in Long-term Contracts: The Case of Coal." *The Journal of Law and Economics* 31 (1): 47–83.
- Kerkvliet, Joe, and Jason F. Shogren. 1992. "The Impacts of Environmental Regulation on Coal Procurement Strategies: Design Coal and Multi-attributed Quality." *Journal of Environmental Management* 35 (2): 83–91.
- Kim, Wook, Deb Chattopadhyay, and Jong-bae Park. 2010. "Impact of Carbon Cost on Wholesale Electricity Price: A Note on Price Pass-through Issues." *Energy* 35 (8): 3441–3448.
- Knittel, Christopher R., Ben S. Meiselman, and James H. Stock. 2017. "The Pass-Through of RIN Prices to Wholesale and Retail Fuels under the Renewable Fuel Standard." *Journal of the Association of Environmental and Resource Economists* 4 (4): 1081–1119.
- Knittel, Christopher R., Konstantinos Metaxoglou, and Andre Trindade. 2015. "Natural Gas Prices and Coal Displacement: Evidence from Electricity Markets." NBER Working Paper 21627.
- Kolstad, Charles D., and Frank A. Wolak. 1983. "Competition in Interregional Taxation: The Case of Western Coal." *Journal of Political Economy* 91 (3): 443–460.
- Kosnik, Lea, and Ian Lange. 2011. "Contract Renegotiation and Rent Re-distribution: Who Gets Raked over the Coals?" *Journal of Environmental Economics and Management* 62 (2): 155–165.
- Levy, Jonathan I., Lisa K. Baxter, and Joel Schwartz. 2009. "Uncertainty and Variability in Health-Related Damages from Coal-Fired Power Plants in the United States." *Risk Analysis* 29 (7): 1000–1014.
- Li, Jing, and James H. Stock. 2017. "Cost Pass-Through to Higher Ethanol Blends at the Pump: Evidence from Minnesota Gas Station Data." Working Paper.
- 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.
- Lobao, Linda, Minyu Zhou, Mark Partridge, and Michael Betz. 2016. "Poverty, Place, and Coal Employment across Appalachia and the United States in a New Economic Era: Poverty, Place, and Coal Employment." *Rural Sociology* 81 (3): 343–386.
- MacDonald, James M. 1989. "Railroad Deregulation, Innovation, and Competition: Effects of the Staggers Act on Grain Transportation." *Journal of Law and Economics* 32 (1): 63–95.

- MacDonald, James M. 2013. "Railroads and Price Discrimination: The Roles of Competition, Information, and Regulation." *Review of Industrial Organization* 43 (1–2): 85–101.
- Mansur, Erin T. 2007. "Do Oligopolists Pollute Less? Evidence from a Restructured Electricity Market." *Journal of Industrial Economics* 55 (4): 661–689.
- . 2008. "Measuring Welfare in Restructured Electricity Markets." *Review of Economics and Statistics* 90 (2): 369–386.
- Marion, Justin, and Erich Muehlegger. 2011. "Fuel Tax Incidence and Supply Conditions." *Journal of Public Economics* 95 (9–10): 1202–1212.
- Meng, Kyle C. 2016. "Estimating Path Dependence in Energy Transitions." NBER Working Paper 22536.
- Montgomery, W. David. 1972. "Markets in Licenses and Efficient Pollution Control Programs." *Journal of Economic Theory* 5 (3): 395–418.
- Muehlegger, Erich, and Richard L. Sweeney. 2017. "Competition and Pass-Through of Input Cost Shocks: Evidence from the U.S. Fracking Boom." Working Paper.
- Muller, Nicholas Z., Robert Mendelsohn, and William Nordhaus. 2011. "Environmental Accounting for Pollution in the United States Economy." *American Economic Review* 101 (5): 1649–1675.
- Nazifi, Fatemeh. 2016. "The Pass-Through Rates of Carbon Costs on to Electricity Prices within the Australian National Electricity Market." *Environmental Economics and Policy Studies* 18 (1): 41–62.
- Nordhaus, William D. 1993. "Reflections on the Economics of Climate Change." *Journal of Economic Perspectives* 7 (4): 11–25.
- Oates, Wallace E., and Diana L. Strassmann. 1984. "Effluent Fees and Market Structure." *Journal of Public Economics* 24 (1): 29–46.
- Olley, G. Steven, and Ariel Pakes. 1996. "The Dynamics of Productivity in the Telecommunications Equipment Industry." *Econometrica* 64 (6): 1263–1297.
- Pigou, Arthur C. 1932. *The Economics of Welfare, 1920*. London: Macmillan.
- Pouliot, Sebastien, Aaron Smith, and James H. Stock. 2017. "RIN Pass-Through at Gasoline Terminals." Working Paper.
- Prater, Marvin, Adam Sparger, and Daniel O’Neil Jr. 2014. *Railroad Concentration, Market Shares, and Rates*. Technical report. U.S. Department of Agriculture, Agricultural Marketing Service.
- Reguant, Mar. 2014. "Complementary Bidding Mechanisms and Startup Costs in Electricity Markets." *Review of Economic Studies* 81 (4): 1708–1742.
- Revesz, R., M. Greenstone, M. Hanemann, M. Livermore, T. Sterner, D. Grab, P. Howard, and J. Schwartz. 2017. "Best Cost Estimate of Greenhouse Gases." *Science* 357 (6352): 655.

- Ryan, Stephen P. 2012. "The Costs of Environmental Regulation in a Concentrated Industry." *Econometrica* 80 (3): 1019–1061.
- Schmalensee, Richard, and Robert N Stavins. 2013. "The SO₂ Allowance Trading System: The Ironic History of a Grand Policy Experiment." *Journal of Economic Perspectives* 27 (1): 103–122.
- Schmidt, Stephen. 2001. "Market Structure and Market Outcomes in Deregulated Rail Freight Markets." *International Journal of Industrial Organization* 19 (1): 99–131.
- Startz, Meredith. 2016. "The Value of Face-to-face: Search and Contracting Problems in Nigerian Trade." Working Paper.
- Stolper, Samuel. 2016a. "Competition and Incidence: Automotive Fuel Tax Pass-Through at State Borders." Working Paper.
- . 2016b. "Who Bears the Burden of Energy Taxes? The Role of Local Pass-Through." Working Paper.
- 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. 1996. "Why Do Firms Simultaneously Purchase in Spot and Contract Markets? Evidence from the United States Steam Coal Market." In *Agricultural Markets: Mechanisms, Failures, and Regulations*, edited by David Martimort, 109–168. Elsevier Science.
- . 2016. "Assessing the Impact of the Diffusion of Shale Oil and Gas Technology on the Global Coal Market." Working Paper.
- Wolak, Frank A., and Charles D. Kolstad. 1988. "Measuring Relative Market Power in the Western U.S. Coal Market Using Shapley Values." *Resources and Energy* 10 (4): 293–314.
- Wolfram, Catherine, Ori Shelef, and Paul Gertler. 2012. "How Will Energy Demand Develop in the Developing World?" *Journal of Economic Perspectives* 26 (1): 119–138.
- Wooldridge, Jeffrey M. 2007. "Inverse Probability Weighted Estimation for General Missing Data Problems." *Journal of Econometrics* 141 (2): 1281–1301.