

Fracking, Coal, and Air Quality

Reid Johnsen, Jacob LaRiviere, Hendrik Wolff

Abstract: This paper estimates indirect benefits of improved air quality induced by hydraulic fracturing, or “fracking” in the continental United States. The recent increase in natural gas supply led to displacement of coal-fired electricity by cleaner natural gas-fired generation. Using detailed spatial panel data comprising the near universe of air quality monitors merged with US power plant locations, we find that coal generation decreased by 28% attributable to lower natural gas prices. Using an IV identification strategy to isolate fracking’s impact on natural gas prices, we identify a 4% decrease in average PM_{2.5} levels due to decreased coal generation. These benefits vary geographically; air pollution levels decreased most in parts of Alabama by 35%. Back of the envelope calculations imply accumulated health benefits of roughly \$17 billion annually.

JEL Codes: I18, Q41, Q53

Keywords: fracking, coal-fired power plants, air pollution, health, electricity

HOW DO NEW TECHNOLOGIES IMPACT the environment? Many innovations are surrounded by extensive discussions about their direct impacts for the natural environment, potentially leading to the ban of the new technology. In many settings of the economy, new technologies replace existing dirty technologies, leading to important indirect environmental benefits. In the wake of new innovations, these indirect effects are sometimes overlooked, although they are required in full cost-benefit analysis (Prest and Turvey 1968). Furthermore, to estimate these effects can be challenging when new technologies are introduced at times of major macroeconomic shifts.

Reid Johnsen is in the Department of Agricultural and Resource Economics, University of California, Berkeley (rlkj@berkeley.edu). Jacob LaRiviere is at Microsoft (jlariv@microsoft.com) and the Department of Economics at University of Tennessee. Hendrik Wolff is in the Department of Economics, Simon Fraser University (Hendrik_wolff@sfu.ca). We thank Werner Antweiler, Catherine Hausman, Joshua Linn, Erin Mansur, Joseph Shapiro, and Nathan Tefft for particularly helpful criticism and suggestions. We are also grateful for comments from seminar participants at the University of British Columbia, the University of Hawaii, and the AERE meetings. Hendrik Wolff acknowledges financial support by Social Sciences and Humanities Research Council grant 20586. All errors are ours.

Received January 3, 2017; Accepted October 30, 2018; Published online July 30, 2019.

JAERE, volume 6, number 5. © 2019 by The Association of Environmental and Resource Economists. All rights reserved. 2333-5955/2019/0605-0006\$10.00 <http://dx.doi.org/10.1086/704888>

This paper estimates the short-run causal effect of hydraulic fracturing (“fracking”) on ambient air quality attributable to natural gas’s displacement of coal in the electricity sector. In particular, we quantify indirect air quality and health benefits due to displacement of coal-fired electricity generation by cheaper natural gas-fired generation. The 2009 advance in horizontal drilling, and fracking technology is arguably the most important change to US energy markets since the OPEC crisis and has had vast implications for the American economy (Hausman and Kellogg 2015), substantially increasing the US natural gas supply.

The conventional wisdom is that fracking induces negative localized environmental externalities, many of which have been recently discussed in the growing economics literature on fracking.¹ Some regulators have limited or prohibited fracking. New York State, for example, has banned fracking entirely, and many US municipalities and European Union countries have limited its scope.

But because fracking has dramatically decreased US natural gas prices, it has potentially decreased coal consumption and concomitantly improved US air quality. According to the EIA (1999), coal-fired power plants, as compared to natural gas-fired power plants, emit 392 times as many units of particulate matter (PM)² per megawatt-hour (MWh) of electrical generation.³ Between 1950 and 2008 about half of US electricity generation came from coal. Between 2009 and 2016 this proportion dropped to 35%–40%. Indeed, one contribution of our paper is to estimate the relative impact of coal versus natural gas on air quality econometrically using the (near) universe of all natural gas and coal-fired power plants. As a result of its effect on natural gas prices, fracking may have introduced indirect nonmarket benefits through air quality improvements. Air quality benefits are especially important because a large share of the US air quality monitors are in noncompliance with recently tightened EPA air quality limits, which were imposed in lieu of Pigouvian taxes to mitigate PM’s negative health externalities. Figure 1 illustrates the distribution of air quality monitors that are out of compliance during our sample period.

Estimating these indirect effects, however, poses major econometric challenges because the fracking boom occurred simultaneously with fundamental shifts in the macroeconomy. During our period of study, 2007 through 2012, changing electricity market structures, growing Chinese demand for coal, and most of all the Great Recession and

1. Damages discussed in the economics literature include toxic leaks into groundwater supplies (Muehlenbachs et al. 2013, 2015), chemicals exposing surface water (Olmstead et al. 2013), traffic accidents (Muehlenbachs and Krupnick 2014; Graham et al. 2015), earthquakes (Koster and van Ommeren 2015), and price shocks to local nontradable goods adversely affecting individuals living near fracked wells (Allcott and Keniston 2015). Natural gas leaks are discussed in Brandt et al. (2014), and Jackson et al. (2014) provide an overview of the costs and benefits of fracking.

2. PM is linked to a number of serious health and other externalities, well documented in the economics literature; see Graff Zivin and Neidell (2013) for a recent review of the literature.

3. The cited ratio of 392 is based on engineering studies (EIA 1999).

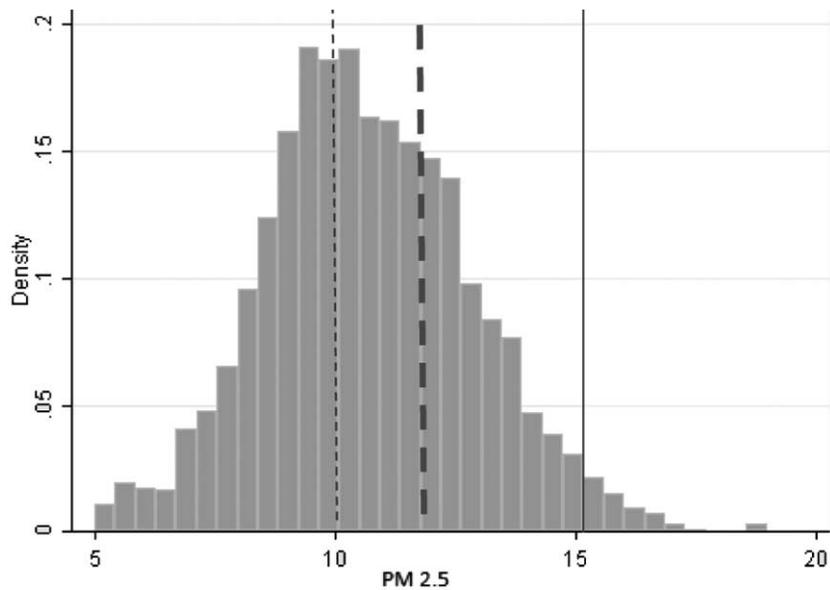


Figure 1. Histogram of $PM_{2.5}$ average daily sample values. The mean daily concentration of $PM_{2.5}$ at the monitors in our sample is $10.69 \mu\text{g}/\text{m}^3$, with a standard deviation of $2.17 \mu\text{g}/\text{m}^3$. The light dotted line displays the WHO $PM_{2.5}$ standard of $10 \mu\text{g}/\text{m}^3$. Since December 14, 2012, the EPA has set the standard to $12 \mu\text{g}/\text{m}^3$ (bold striped line). From 1997 to 2012, the federal standard for compliance with the Clean Air Act was $15 \mu\text{g}/\text{m}^3$. Color version available as an online enhancement.

subsequent recovery had the potential to affect patterns of electricity generation.⁴ For example, if the recession led to lower electricity demand and a corresponding decrease of fossil fuel consumption for electrical generation, air quality might have improved anyway.

To estimate the causal impacts of fracking on electricity sector pollution we therefore must construct the appropriate counterfactual levels of coal and natural gas consumption for electricity generation, in the absence of fracking's impact on natural gas prices, *ceteris paribus*. In this way we estimate what would have happened to electricity sector pollution if fracking had been banned in the United States. Aside from adequately modeling fuel substitution, constructing the counterfactual presents three principal endogeneity challenges, each of which our research design must address. The first is accounting for the recession's impact on electricity demand. Our research design must condition on observed electricity generation when determining what generation would have been

4. We begin our data set in 2007, 2 years before the fracking revolution that started around 2009. We end our study in 2012 to mitigate concerns about increases in natural gas capacity attributable to the natural gas price decrease.

in 2012 but for the decrease in natural gas price attributable to fracking. Second, some of the decrease in natural gas prices that occurred over our sample period was due to decreased demand attributable to the recession. We must isolate fracking's impact on natural gas prices within our research design. Third, over our sample period there were various EPA regulatory changes. Our research design must also account for EPA regulatory changes that occur in our sample period. Note that by ending our sample period before 2013 we avoid the more complicated question of how cheaper natural gas impacted investment decisions for new natural gas-fired capacity. As a result, our study focuses on short to medium run impacts of fracking on US air quality.

To overcome these identification problems, we propose a three-step instrumental variable (IV) empirical strategy. First, we isolate the effects of relative fuel prices from other industry forces by constructing regional least cost electricity dispatch models. Using hourly data comprising the electrical generation, fuel input prices, and boiler-specific efficiency levels of the near universe of US power plants,⁵ we construct an IV from the simulated power plant dispatch order. The simulated dispatch orders indicate how much power each US power plant produces during each hour as a function of the relative input prices of coal, oil, and natural gas as well as regional power plant capacities. Macroeconomic impacts during this period are therefore held constant in our counterfactual scenario. While dispatch models are common in the industrial organization literature (Wolfram [1999] and Borenstein et al. [2002] are two examples),⁶ to our knowledge we are the first to use a dispatch model to econometrically estimate the causal impacts of changing market conditions on environmental outcomes.

Second, we evaluate how changes in electricity dispatch affect local air pollution. This analysis combines precise information from the universe of EPA-operated air quality monitors with the exact location of all power plants. Rather than conduct an ex ante simulation of air quality changes using atmospheric chemistry models, we look at how observed changes in electricity generation at power plants affect the observed patterns of air quality. We show that changes in relative prices substantially alter generation patterns across power plants and significantly change local air quality. Not all price changes, however, can be attributed to fracking (e.g., coal price increases because of surging coal demand in China are unrelated to fracking).

Our third and final step isolates the portion of the natural gas price change that is related to fracking. Due to limitations of the global natural gas transportation

5. The heat rate of a boiler is measured by the quantity of fossil fuel burned per unit of electricity generated. The higher the heat rate the more inefficient the production of electricity. While typically marginal costs cannot be observed in most industries, the public information of heat rates together with the public information of input fuel costs allows us to estimate the marginal costs for each power plant.

6. As in those papers, we maintain the assumption that load is inelastic with respect to wholesale electricity prices since the marginal user rarely pays wholesale prices and instead pays retail prices.

network, increases in the US natural gas supply during our sample period were largely consumed domestically, leading to a decrease in US natural gas prices relative to international prices. Hausman and Kellogg (2015) attribute a US natural gas price decrease of \$3.41/mmBtu to fracking,⁷ which represents a decline of roughly 50% from the 2007 price. We also perform various robustness checks around this number, including varying the number of air quality monitors in the analysis, adding operations and maintenance costs, and allowing for ramping constraints. Using our IV estimate of the impact of the relative price change on air quality, we simulate air quality in a counterfactual scenario in which no fracking had occurred, *ceteris paribus*, by adding the Hausman and Kellogg (2015) estimate of \$3.41 to the 2012 price of natural gas. Because the dispatch model is constructed using relative prices, instrumented changes in generation are orthogonal to changes in EPA regulation. Hence, in our 2012 counterfactual we let the economy evolve from 2007 to 2012 as it actually developed, except that we set the price of natural gas to the estimated price from a scenario in which fracking had been banned since 2007.

We find large and precisely estimated effects of fracking: coal generation declined by 28% on average and ambient air quality increased an average of 4% due to the displacement of coal by 2012. Our spatially differentiated impact analysis, however, shows substantial heterogeneity in the geographic distribution of air quality benefits. Local air quality increased by 35% in the area of greatest coal displacement. Back of the envelope calculations imply that fracking produced health benefits of roughly \$17 billion annually. To put this benefit into context, note that Hausman and Kellogg (2015) estimate direct benefits of \$25 billion annually for the electricity market due to fracking. Hence our indirect estimated nonmarket benefit accounts for an additional 68% of this annual market-based surplus.⁸ Note that this indirect environmental benefit from fracking is much larger compared to the environmental costs heretofore estimated.⁹ As a result, we find evidence of a significant indirect nonmarket environmental benefit attributable to fracking using standard value of statistical life (VSL) estimates. However, this estimate ignores other nonmarket costs like methane leakage, damage to local roads, earthquakes, and any other externalities.

In addition, this paper contributes to the knowledge of atmospheric pollution conditions. The reduction in coal-fired generation provides us with a unique opportunity to econometrically estimate the contribution of coal-fired power plants to air

7. All prices in this paper are in 2012 US dollars, using CPI according to the Historical Chained Consumer Price Index for All Urban Consumers, US city average, all items (C-CPI-U); mmBtu stands for 1 million British thermal units and is the standard measure for one unit of natural gas.

8. This paper does not estimate the direct negative effects of fracking. These would need to be added for a full cost-benefit analysis, which is beyond the scope of this paper.

9. The largest upper bound monetary estimate we could find is from Ames et al. (2012), who estimate that fracking produces an upper bound damage on groundwater of \$250 million per year.

pollution. We find that shutting down all US coal-fired power plants would on average decrease local air pollution by 16% (confidence interval from 9% to 23%). There is substantial heterogeneity, however—in the most coal-intensive area of the United States, a complete shutdown of coal-fired generation would decrease local $PM_{2.5}$ levels by 89%. While the atmospheric chemistry literature continues to debate the source apportionment and spatial modeling of $PM_{2.5}$ (i.e., Yu et al. 2013; Hodan and Barnard 2004; Crawford et al. 2015; Pirovano et al. 2015), our study—to our knowledge—is the first to empirically estimate the apportionment for coal on a nationwide level. We also find estimates for NO_x and SO_2 , but these estimates are less precisely estimated.

This paper builds on a growing literature that uses quasi-experimental research designs to quantify how environmental regulation, manufacturing production, transportation, and other forms of economic activity affect air quality and human health. This body of work uses either local variation in air quality conditional on detailed fixed effects (Currie and Neidell 2005; Schlenker and Walker 2016) or observed policies with sharp variation that is conducive to difference-in-differences designs (Chay and Greenstone 2005; Isen et al. 2014). Fabra and Reguant (2014) use emission prices to instrument for emission costs and identify pass-through in the electricity market, and Deschênes et al. (2017) use a triple difference design to study the effects of a NO_x regulation on defensive health expenditures. A conceptually closer research design to ours is Mansur (2007), which develops an electricity dispatch model to simulate emission rates of firms with different levels of market power but does not econometrically identify impacts on ambient pollution levels. In addition, recent papers by Linn et al. (2014), Cullen and Mansur (2017), Knittel et al. (2015), and Holladay and LaRiviere (2017) analyze the mechanisms of fuel-switching behavior and discuss implications on carbon emissions profiles. We expand on this work by closely linking an economic model of electricity markets together with detailed data on air quality to quantify the indirect environmental consequences of a general equilibrium economic shock—a change to natural gas extraction technology.

The rest of the paper is organized as follows. Section 1.1 describes our data sources, and section 1.2 provides a first intuition for the correlation between electricity production by coal, plant emissions, and ambient air quality. Section 2 describes the construction of our dispatch model that we use to construct our instrumental variable. Section 3 outlines our econometric framework, and section 4 presents regression results. Conclusions and further thoughts on research are offered in section 5. The appendix, available online, lists details of our data and several alternative specifications.

1. DATA

1.1. Description of Data

We study the period between 2007 and 2012 for several reasons. First, the period of study captures the nationwide decrease in natural gas prices, which begin in late 2008 and early 2009. Second, by ending the sample in 2012, we avoid any changes

to the stock of natural gas-fired generation capacity that resulted from decreased natural gas prices. In order to extend the analysis to 2013 and beyond we would need a dynamic model of natural gas investment for the appropriate no-fracking counterfactual, which is outside the scope of our study.

We collect ambient air pollution data from the EPA's Air Quality System (AQS), which compiles ambient levels of $PM_{2.5}$, SO_2 , and a variety of other air pollutants as measured by a network of approximately 5,000 monitors across the United States. Following the example of Deschênes et al. (2017), we restrict our observations to those monitors that report a minimum of one time during each of at least 47 weeks during every year from 2007 to 2012. This restriction eliminates biased data from monitors that were either taken out of service or were operated on a seasonal basis during the period of study. Additionally, we further restrict the sample of monitors to those designated as having "population exposure." This restriction is intended to reduce noise from monitors that are located in industrial areas, although in practice very few monitors were dropped due to this criterion. Finally, we restrict the sample to monitors that are within 70 miles of a coal-fired plant in our main specification, also performing robustness checks by using a 40-mile radius and a 100-mile radius.¹⁰ In total, 537 $PM_{2.5}$ monitors in 363 counties remain in the final data set. Observed $PM_{2.5}$ levels are averaged and aggregated up to the monthly, quarterly, and annual levels. Figure 2 shows the locations of the $PM_{2.5}$ monitors that we include in our analysis, as well as the locations of all power plants that have at least one boiler for which the primary fuel is coal. A significant number of air quality monitors are proximate to coal-fired power plants, especially in the eastern United States, and we are able to include the vast majority of US coal-fired power plants in our study. We do an identical analysis for SO_2 and NO_x monitors to identify effects of SO_2 and NO_x but focus on $PM_{2.5}$ in this paper.

We obtain electrical generation data from the EPA Air Markets Program Data (AMPD) via the Continuous Emissions Monitoring System (CEMS), which contains data on all US power plants equipped with generators rated at 10 megawatts (MW) or greater.¹¹ In addition to hourly generation at the generator level and hourly SO_2 and NO_x emissions data, AMPD also contains data on primary and secondary fuel type and exact power plant location. The CEMS data call generation "gross load," and we thus use load and generation interchangeably in the paper. It is not uncommon for some plants to cease electricity generation either seasonally or during periods of low demand, and so we interpret missing values for generation or stack emissions as zeros unless there

10. Having a small radius decreases possible measurement error but also reduces sample size. Because measurement error has been shown to be important in causal studies of air pollution on economic outcomes (Sullivan 2017), we prefer a smaller vs. larger radius.

11. Holladay and LaRiviere (2017) describe in more detail the subset of AMPD data that is collected by Continuous Emissions Monitoring Systems (CEMS), which are required to be installed on any power plant with a capacity of 25 MW or greater.



Figure 2. Locations of coal-fired power plants and air pollution monitors. The map displays all power plants that have at least one boiler for which coal is the primary fuel and all air pollution monitors that detect either PM_{2.5} or SO₂ and meet the reporting criteria as explained in the text. Color version available as an online enhancement.

is a reason to suspect data entry error.¹² We take fuel input price data from the State Energy Systems (SEDS) database housed by EIA. The database reports average fuel input price used by electricity generators using a given fuel in a given year by state.¹³

We obtain weather data from the National Oceanic and Atmospheric Administration's (NOAA) Quality Controlled Local Climatological Data (QCLCD), which catalogs daily weather information as recorded at approximately 1,600 weather stations across the United States. Of these, we consider only those stations that reported data for at least 25 days out of every month during the period of 2007–12, eliminating those stations that

12. For example, if a plant reported that a large amount of electricity was generated on a given day but no SO₂ was emitted, this suggests that there was a data entry error. The converse is not necessarily true—many power plants have the secondary function of generating steam for municipal heating. Since steam generation produces air pollution but does not produce electricity, it is possible for daily SO₂ emissions to be positive while electricity generation is zero. Any observation that indicates generation with missing data for pollution is dropped, but those observations that are missing generation data but have positive pollution data are assumed to have zero generation.

13. Chu et al. (2017) show that using fuel input spot prices can cause misleading inference when constructing dispatch models since coal spot prices are a poor proxy for actual coal purchase prices of power plants.

did not consistently report. The resulting data set contains monthly, quarterly, and annual averages for 886 weather stations in the lower 48 states. Each power plant is assumed to experience the weather conditions that are reported by the station nearest to the plant.

Finally, in order to control for changes in economic activity that affect ambient $PM_{2.5}$ and are correlated with electricity generation, we draw from the Bureau of Labor Statistics' Local Area Unemployment (LAU) data set. LAU contains monthly observations of the number of employed and unemployed workers for each US county.¹⁴ By accounting for employment level in our empirical analysis we mitigate the confounding effects of changes in activity, including changes related to the Great Recession, in industrial sectors that emit $PM_{2.5}$, such as construction or trucking.

1.2. First Empirical Results

Figure 1 shows the distribution of average daily $PM_{2.5}$ measurements for the monitors in our sample. The federal government sets the standard for allowable concentrations of $PM_{2.5}$ through regular updates and amendments to the Clean Air Act of 1976. During the period of 2007–12, the compliance standard was set to an average daily concentration (over the course of one year) of 15 micrograms per cubic meter ($\mu g/m^3$) and a peak concentration of 35 $\mu g/m^3$ per day. The average daily concentration limit was revised to 12 $\mu g/m^3$ on December 14, 2012, putting 24% of our sample monitors out of compliance. The mean $PM_{2.5}$ concentration in our sample is 10.69 $\mu g/m^3$, with a standard deviation of 2.17 $\mu g/m^3$; 553 monitors in 328 counties reported average annual $PM_{2.5}$ levels that exceeded the new standard in at least one year in our study period.

Preliminary analysis of the AQS $PM_{2.5}$ data suggests that the change in ambient $PM_{2.5}$ levels over time is not uniform across all regions of the United States. Figure 3 shows the change in county average daily ambient $PM_{2.5}$ levels from 2007 to 2012. The counties that experienced the largest decrease, shown in darker red, tend to be clustered around Appalachia and the Great Lakes region. The Midwest and New England appear to have experienced a smaller decrease or an increase, shown in shades of light blue.

Figure 4 shows a similar trend in the change in county average coal-fired electrical generation over the same time period. Those counties that show the greatest decrease in coal-fired generation most often appear in the Appalachian or Great Lakes regions of the United States, while those counties in the Midwest tend to show a smaller decrease or slight increase. The visual correlation between reduced coal-fired electrical generation and reduced levels of $PM_{2.5}$ is stark.

14. We make use of the LAU data because they provide a high level of temporal and spatial resolution. One alternative data set is the US Census County Business Patterns (CBP) survey, which provides county-level annual measures of economic activity by sector. Robustness checks at the annual level using CBP data on specific industries that have large impacts on air pollution, such as construction, transportation, and manufacturing, are quantitatively similar to LAU results and are available from the authors by request.

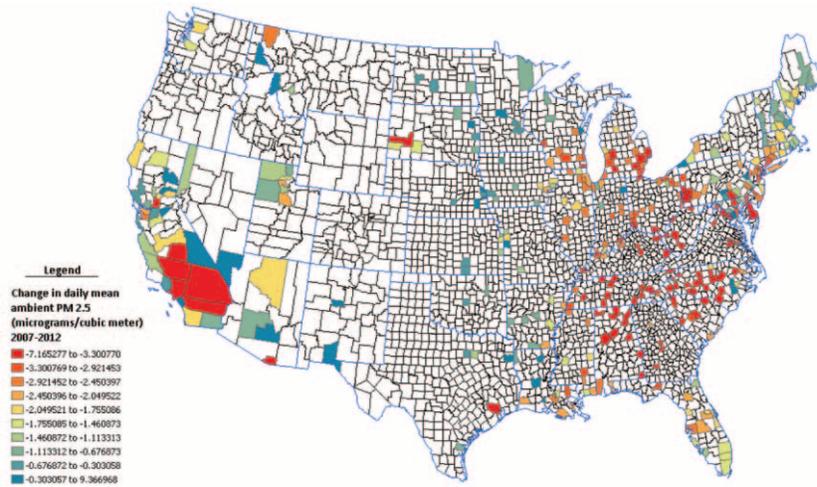


Figure 3. Change in ambient PM_{2.5} levels 2007–12. There is a regional dichotomy in the change in ambient PM_{2.5} levels. Most counties in the Midwest and Mid-Atlantic regions show a decrease in average daily pollution level, but the decrease is more pronounced in Appalachia and near the Great Lakes than in the Midwest or New England.

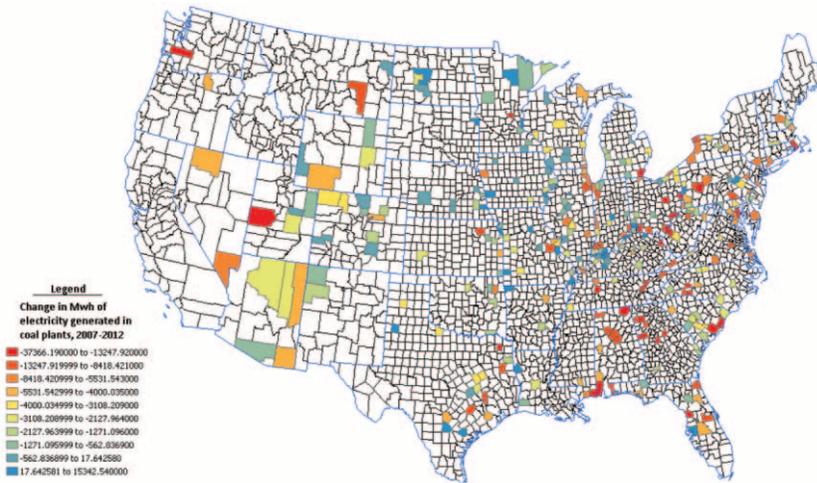


Figure 4. Change in county average daily electrical generation (MWh) by coal-fired power plants, from 2007 to 2012. Counties in the Appalachian and in the Great Lakes regions tended to decrease their coal-fired electrical generation by the largest amounts. Counties in the Midwest generally show smaller decreases or slight increases.

Figure 5 illustrates one mechanism that may have contributed to the regional dichotomy in coal-fired generation. The plot shows the price of natural gas and the percentage of electricity generated by coal. Figure 5 suggests that prior to 2009 there was not a strong correlation between the price of natural gas and the percentage of coal-fired electricity generation. However, when the price of gas dropped below \$6/mmBtu, the correlation strengthened significantly. The big trend represents the average of the top third of states, ordered by change in coal share from 2007 to 2012 and the small trend represents the average of the lowest third of states, ordered by the change in coal share from 2007 to 2012. This suggests that some regions of the United States are more able to take advantage of low natural gas prices than others and are therefore more able to substitute away from coal in electricity generation.

There are several issues with figure 5, though, which lead us to use a dispatch model of the electricity wholesale market. Demand for electricity from coal is primarily an outcome of (1) relative coal and natural gas prices and (2) electricity demand, each of which are functions of macroeconomic and regulatory activity. Since a sharp increase

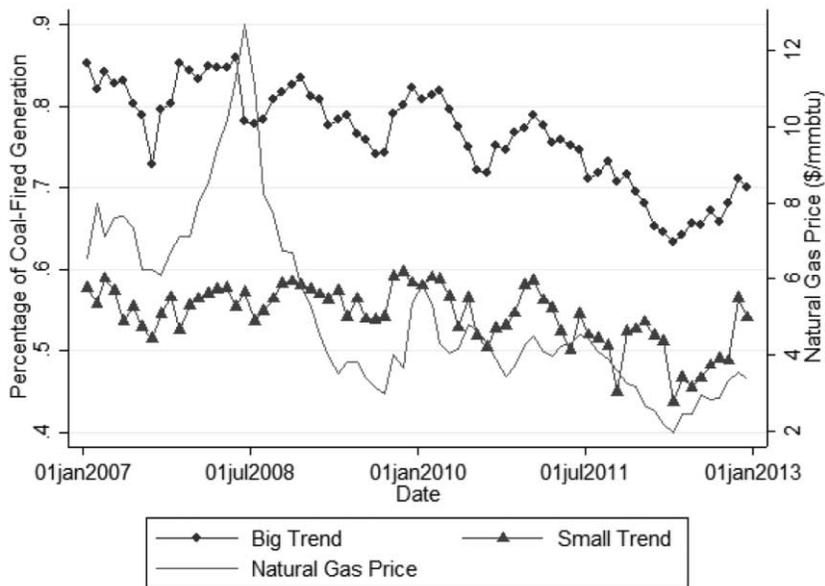


Figure 5. Percentage of electricity that is generated by coal and natural gas price. During the period 2007–9, there does not appear to be a strong correlation between the real price of natural gas and the percentage of electricity that is generated by coal-fired power plants. In 2009 the price of natural gas drops below \$6/mmBtu, and the correlation appears to strengthen substantially. Big trend represents the average of the top third of states, ordered by change in coal share from 2007 to 2012. Small trend represents the average of the lowest third of states, ordered by change in coal share from 2007 to 2012. Color version available as an online enhancement.

in the supply of natural gas, due to fracking, occurred during our study period, we employ a dispatch model to attribute changes in coal generation to relative changes in the price of coal and natural gas.

2. CONSTRUCTING THE INSTRUMENT AND THE DISPATCH MODEL

In order to estimate the indirect benefits of fracking on air quality through the reduction of coal-fired generation, our identification strategy proceeds in three broad stages. This section describes the first stage (which is similarly repeated in stage 3). In our first stage, we develop regional specific dispatch models in order to isolate the effect of changes in relative prices on changes in observed electricity generation. The dispatch model is motivated by Borenstein et al. (2002). The purpose of the dispatch model is to predict the total amount of electricity generated by fossil fuels at each hour and boiler. Since our primary goal in the model is to isolate the change in generation exclusively due to changes in relative prices, we make several assumptions detailed in this section.

2.1. The Wholesale Electricity Market

The US electricity market accounts for roughly 2.2% of GDP.¹⁵ There are two important characteristics of the wholesale electricity market in the context of our paper. First, the total generation of electricity must equal the total demand of electricity at every point in time. If there is an imbalance, it leads to blackouts or damage to transmission lines. As a result, each National Energy Regulatory Commission (NERC) region is governed by one or more Independent System Operators (ISO), which balance electricity supply and electricity demand at every point in time in a given region.

Second, in the vast majority of cases the retail end consumer of electricity does not pay the wholesale price of electricity but rather a constant average price. As a result, electricity demand is highly inelastic at every point in time (Borenstein 2002).¹⁶ Determinants of electricity demand include weather and time of day. For example, during hot summer days, peak demand often occurs when temperatures are highest and households run air conditioners intensely. Because demand is exogenous and supply must equal demand at every point in time, the order in which electricity-generating units are dispatched is the main determinant of electricity costs.

15. In 2014, total electricity sales to end customers was \$393 billion, equivalent to 2.2% of the \$17.9 trillion GDP (see http://www.eia.gov/electricity/annual/html/epa_01_01.html).

16. There are two exceptions. First, large industrial users sometimes have bilateral bargains with electricity producers that include their paying wholesale electricity prices. Second, some residential consumers are beginning to pay real time wholesale electricity prices. These fractions of demand are, however, still very small, making electricity demand effectively exogenous at every point in time. In addition, our results rest on the assumption that the demand function for natural gas in 2012 is not impacted by the price of electricity. In the short run over our time span from 2007 to 2012 this assumption likely holds (Quistorff 2015).

There are two types of wholesale electricity markets in the United States: deregulated and regulated. In both deregulated and regulated markets, the ISO attempts to minimize the cost of needed generation. In deregulated electricity markets private electricity producers bid for the right to produce electricity at every point in time during the day.¹⁷ The ISO then uses a complicated linear programming mechanism to select the cost-minimizing dispatch order given a particular day's forecasted electricity demand.¹⁸ In one important regulated market for our study, the Southeastern Electric Reliability Council (SERC), the ISO handles most dispatch decisions based upon inferred costs of electricity generators. The ISO mandates data-sharing processes to help ensure cost-minimizing generation.

2.2. The Dispatch Model

Our dispatch model exploits the cost-minimization feature of the wholesale electricity market in order to construct an instrument for generation that is directly related to input price changes. The model takes as inputs the calculated cost of electricity generation for every boiler with a capacity of over 10 MW in the United States. We then simulate the total generation of each boiler had each generating unit been dispatched, in order of our constructed cost measure, until the electricity demand is met.

As in prior literature on the wholesale electricity market, we construct marginal cost curves for a NERC region by calculating the average cost of generation by boiler (Wolfram 1999; Borenstein et al. 2002). To do so, we construct two important plant-specific characteristics: the heat rate and input prices. The heat rate is the average of mmBtus used per MWh. The heat rate is a measure of a boiler's efficiency. To construct the heat rate, we use the CEMS database containing hourly generation and fossil fuel input (measured in mmBtu) for all generators.¹⁹ This allows us to construct the average observed heat rate over the sample period for each generator.²⁰ We similarly use observed

17. There is a large literature that shows how market power in the bidding process can lead to departures from least costs generation in determining wholesale electricity costs (Wolfram 1999; Borenstein et al. 2002). The main identifying assumption of our model is simply that the order of dispatch is determined by cost rather than the level of wholesale costs.

18. Most of this bidding process takes place on a "day ahead" market. There is also a "real time" market in which the ISO can purchase additional electricity as needed. There are additional types of electricity producer contracts such as reserve contracts, in which producers are paid to withhold generation capacity should it be needed in real time.

19. Although there are multiple possible fuel types (e.g., various types of coal, natural gas, and oil), the CEMS data converts each fuel input into mmBtus so that they are directly comparable by heat content.

20. Using observed heat rate rather than that reported from EIA forms follows Davis and Hausman (2016). Similarly, we use observed maximum capacity from the CEMS data rather than reported maximum capacity from EIA forms. In both cases we restrict our sample to days in which a boiler produced for more than 1.5 hours to eliminate rounding errors.

maximums of generation levels by boilers to identify capacity. Specifically, we construct heat rate for a boiler i over a period in which it is in operation for T hours by

$$\text{heat_rate}_i = \sum_{t=1}^T \frac{\text{mmBtu}_{it}}{\text{MW}_{it}}.$$

Input prices are provided by SEDS at the state/fuel type/year level in dollars per mmBtu: $\$/\text{mmBtu}_{sfy}$, which we match to boilers in the CEMS database. We choose the SEDS database for three reasons. First, other EIA databases that have input price data, like the EIA 923 form, have asymmetric reporting requirements for firms in regulated versus unregulated markets. As a result, there are many generators that report no input prices for an entire year. Because SEDS is aggregated to state averages, prices are reported for states in both regulated and unregulated markets. Second, we are most concerned with long-run price changes due to increases in the supply of natural gas, making the year level an appropriate level of analysis. Third, different states face different transportation costs for fossil fuel inputs, and we want to control for that variation.

We construct the cost of generation for each generating unit by taking the product of the unit's heat rate and input price. NERC-level marginal cost curves for electricity generation are created by ordering generating units by cost and creating a running sum of plant capacities within each NERC region.²¹ Figure 6 displays the marginal cost curve for the southeast United States (SERC region) for 2007 and 2012. Figure 6 shows that while in 2007 the low marginal cost of electricity was mainly supplied by coal-fired generators, by 2012 many of the high-cost coal-fired generators were displaced by natural gas-fired ones. This change is purely attributable to the decrease in the price of natural gas relative to that of coal.

Using observed NERC hourly generation allows us to construct instrumented hours of generation for each boiler in a given NERC region. Figure 7 displays the observed and predicted coal-fired generation from our first-stage IV approach for three different NERC regions (SERC, Midwest Reliability Organization [MRO], and Reliability First Corporation [RFC]) in 2007 and 2012. Figure 7 shows a strong positive correlation between observed MWs and instrumented MWs from the least cost dispatch model. If the dispatch model perfectly predicted market events, all points in the figure would land on a 45 degree line. We take the positive correlation shown in the figure as evidence that the dispatch model is doing an adequate job of predicting generation for many generators. There are, though, a significant number of plants across all three regions for which the 2012 dispatch model predicts zero generation but observed generation is positive.

21. We have also constructed state-level marginal cost curves. Results from using state-level curves in constructing instrumented hours are available from the authors upon request. These instruments tend to be noisier in the East where states are physically smaller and work better in states that are physically larger. We take observed maximum capacity from the CEMS data for each boiler.

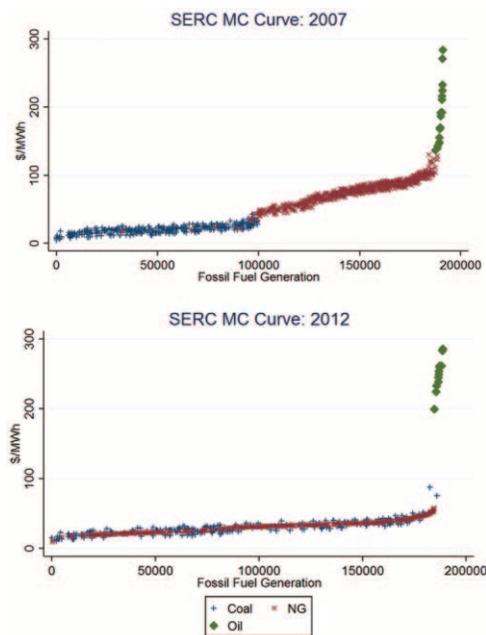


Figure 6. Marginal cost curves of electricity generation: 2007 and 2012 constructed SERC supply curves. Data are taken from CEMS and SEDS data sets. Capacity measure constructed from max average hourly generation observed within a day conditional on boiler generating. Red dots represent natural gas-fired capacity; the 2012 curve displays significantly more mixing at lower levels of generation.

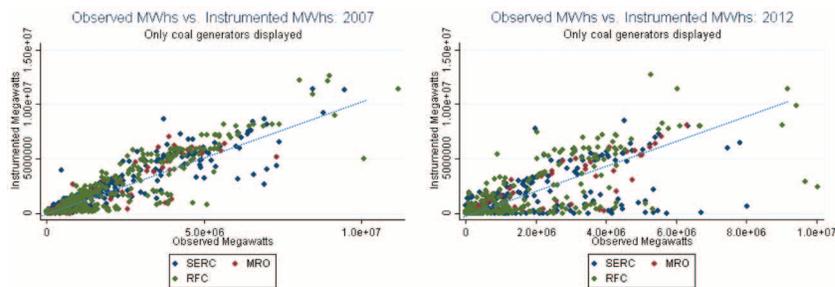


Figure 7. Observed and instrumented megawatt-hours 2007 and 2012: 2007 and 2012 instrumented hours from dispatch model (y -axis) against observed generation in CEMS data (x -axis). Capacity measure constructed from max average hourly generation observed within a day conditional on boiler generating. Color indicates NERC region. Only coal generation displayed. 45 degree line indicated by dotted line.

There are several possible explanations for this underprediction. First, coal-fired power plants may have long-term contracts with coal mines that require them to purchase coal at pre-specified prices instead of at the contemporaneous market price. Second, geographically isolated power plants without nearby competition are likely to stay on to serve local populations. Our dispatch model ignores spatial heterogeneity in demand within the NERC region. More generally, NERC level dispatch models impose a no-trading assumption across NERC regions. We have performed the same analysis with state-level dispatch models, but larger dispatch models rather than smaller ones better serve the medium-run analysis that we perform here. Third, we implicitly assume that firms do not exercise market power asymmetrically over load levels and time. Fourth, we ignore operation and maintenance costs and pollution costs incurred by firms. If operation and maintenance costs are constant over our study window, they would be controlled for by the fixed effects in our empirical specifications below. Finally, coal power plants sometimes stay on even when they earn a loss since ramping costs are large. We ignore ramping with our dispatch model, following previous usage of the dispatch model in the literature (Borenstein et al. 2002; Fabrizio et al. 2008).

Finally, we again note that the overall price decrease of natural gas (from a maximum \$8 to a minimum of \$2 from the mid-2000s to 2012) is not attributable entirely to fracking. According to Hausman and Kellogg (2015), the price reduction due to fracking is roughly \$3.41. Hence, in our third stage we calculate fracking's contribution to reduction in $PM_{2.5}$ using various price scenarios from Hausman and Kellogg.

3. ECONOMETRIC FRAMEWORK

3.1. OLS Model

The purpose of this section is to describe our econometric model relating ambient air quality to electricity generated by coal-fired power plants. The model takes the following form:

$$\text{Ambient}_{jt} = \beta_0 + \beta_1 \text{Generation}_{jt} + \beta_2 \text{Weather}_{jt} + \beta_3 \text{Fixed Effects}_{jt} + \epsilon_{jt}, \quad (1)$$

where Ambient_{jt} is a measure of ambient $PM_{2.5}$ or SO_2 , Generation_{jt} is a measure of power plant output defined as either (i) megawatt-hours of electricity generated, (ii) tons of SO_2 emitted by the plant, or (iii) tons of NO_x emitted by the plant, Weather_{jt} is a vector of weather variables including average wind speed, temperature, humidity, barometric pressure, and precipitation, and $\text{Fixed Effects}_{jt}$ is a vector of spatial and temporal fixed effects as outlined below. The variables are indexed over space by j and over time by t . We explore three separate specifications, with t aggregated either to the annual, quarterly, or monthly level.

Our spatial specification sums the total generation for all coal-fired power plants within a 70-mile radius circle centered at each pollution monitor. This specification has the advantage of modeling the area of influence of a given smokestack but has the potential disadvantage of weighting some power plants more heavily than others. For instance, if a

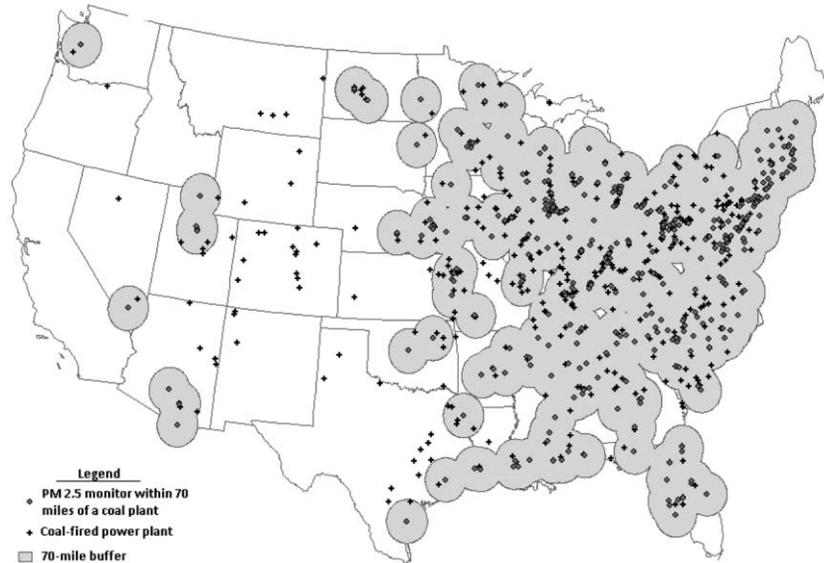


Figure 8. 70-mile radius circles centered on $PM_{2.5}$ monitors. Color version available as an online enhancement.

single power plant is located within 70 miles of six pollution monitors, the influence of that plant will be reported in six separate observations.²² We perform robustness checks by limiting (extending) the radius to 40 (100) miles in the empirical section below.

The total number of observations is limited to the number of pollution monitors that are located within 70 miles of at least one coal-fired power plant: 387 in the case of $PM_{2.5}$ monitors and 73 in case of the SO_2 monitors. Figure 8 maps the monitors that are within 70 miles of a coal-fired power plant (indicated by black crosses). The shaded areas denote the buffers around the $PM_{2.5}$ monitors, within which the power plants are assumed to affect ambient air pollution levels. Electrical generation and plant emissions are summed within each circle, and those sums are treated as discrete observations.

3.2. Instrumental Variables Approach

This section identifies the causal effect of the relative input price change on ambient air quality. We use observed input prices for fossil fuel-fired power plants to construct

22. In earlier versions of this paper we tried to incorporate wind speeds into the analysis as well to account for dispersion of $PM_{2.5}$ for coal stacks. However, there were at least two problems: the monitors take readings only every 5 days, meaning that we would need a model of how to aggregate pollution at air quality monitors. Further, getting the appropriate wind speed and direction for the appropriate altitude at the precise needed location (e.g., the top of the smoke stack) was impossible. As a result, we ended up choosing the circles around air quality monitors rather than a more complicated structure.

the marginal cost of generation for each power plant. We then use the dispatch model to construct a predicted level of generation at each coal-fired power plant. Predicted levels of generation are subsequently taken as an instrument for changes in coal-fired generation.²³ Specifically, for each region we estimate

$$\text{Generation}_{it} = \alpha + Z_{it}\theta + \text{Inst. Generation}_{it}\delta + \varepsilon_{it}, \quad (2a)$$

$$\text{Ambient}_{jt} = \alpha + X_{it}\varphi + \sum_j \widehat{\text{Generation}}_{jit}\beta + \varepsilon_{it}. \quad (2b)$$

Equation (2a) estimates how observed generation relates to predicted generation due to relative marginal costs of fossil fuel generators by the dispatch model. Equation (2b) takes predicted generation values from equation (2a) for each generator i at time t and associates it with a level of spatial aggregation j as discussed in the previous section. The estimated coefficient β describes the change in air quality attributable to relative input price changes over our sample period.

Our instrument is valid only if fracking did not affect $\text{PM}_{2.5}$ through any channels other than electricity generation, and it is not correlated with other things impacting coal-fired power plants over this time period. In the United States, natural gas is used primarily for electricity generation, heating, and industrial production. Since heating via natural gas did not vary greatly over our study window (EIA 2016), electricity generation is likely the only channel through which fracking affected ambient $\text{PM}_{2.5}$ levels. There might be concern that pollution stringency increased over our study period. However, previous work shows that pollution permits that disproportionately impact coal generators decreased in prices significantly over our sample due to increase natural gas electricity production (Holladay and LaRiviere 2017).

Figure 7 shows our first-stage relationship between instrumented and observed coal-fired generation. In theory, without any ramping, maintenance, and so forth, the relationship should be a 45 degree line. Figure 7 shows, however, that the generation level is overpredicted at some of the power plants and underpredicted at others. In particular, note that the attenuation of the instrument increases from 2007 to 2012, which is likely due to increased ramping of coal-fired plants and today natural gas replaces previous

23. Note that our instrument is necessary because fracking is not the only macroeconomic shift in the US energy landscape during our period of study. Changing structures of electricity markets and policies (renewable portfolio standards), growing Chinese demand for raw materials, increasing Chinese supply of energy-intensive manufactured goods, and most of all the Great Recession all had potential to affect patterns of electricity generation between 2007 and 2012. The type of potential data error in our data discussed in n. 2 as well as the particular market conditions described in n. 16 are additional reasons why OLS could lead to attenuated results. Our IV strategy makes our estimates robust against these issues. To isolate the effects of natural gas price shocks from other industry forces, we use a dispatch model of a regional US electricity market to construct an instrumental variable that identifies the causal impact of the changes in natural gas prices.

coal-fired baseload generators. Over the entire nation, the percentage difference between the simulated dispatch and observed coal generation is -14.4% in 2007 and $+3.7\%$ in 2012. We conclude that while the instrument is noisy it is an adequate representation of observed generation behavior. In the empirical results section we also perform robustness checks for the dispatch model, including adding in operations and maintenance (O&M) costs and allowing for coal and natural gas power plants to have ramping constraints.

3.3. Stage 3

Our third and final step isolates the effect of the fracking-related natural gas price change from any other confounder. In order to attribute the changes in observed generation to fracking, we must be certain that fracking caused the long-term decrease in natural gas prices. Figure 9 displays the Henry Hub natural gas spot and futures prices for 2005–14. If a futures contract displays the same price as a spot contract then the market predicts no price change. Figure 9 shows that commodities markets did not anticipate the natural gas price decrease in late 2008 due to the recession. More importantly for our purposes, in early 2009 the market expected natural gas prices to increase rather than decrease. We attribute this to fracking's unexpected impact on the price of



Figure 9. Natural gas spot and futures prices: Henry Hub natural gas prices over time. Figure 9 displays prices for contracts on the date they were written. Therefore, if a futures contract is the same price as a spot contract then the market predicts no price change. In early 2009 the market expected natural gas prices to increase rather than decrease.

natural gas. The high prices observed before the great recession of 2008 have not been observed since then, despite increases in total natural gas consumption since that time. While we attribute a portion of the price change to fracking, we point out that other factors also influenced natural gas prices during our study period. In this third and final stage, we simulate the counterfactual scenario using the price change that Hausman and Kellogg (2015) attribute to fracking.

Hausman and Kellogg (2015) estimate the expected price decrease of natural gas attributable to fracking in their “medium” case to be $\Delta P^{\text{Medium}} = \$3.41/\text{mmBtu}$, relative to a counterfactual scenario of no fracking, with upper and lower bounds of $\Delta P^{\text{Upper}} = 4.11$ and $\Delta P^{\text{Lower}} = 2.16$ in 2012 USD. We calculate the counterfactual natural gas price by adding the Hausman and Kellogg (2015) estimates to the observed 2012 price of natural gas. For any given natural gas price, our dispatch model yields a dispatch order for all power plants in each NERC region. Using counterfactual natural gas prices, we simulate a counterfactual dispatch order. In section 4.2 we presented an IV approach for estimating the impact of a natural gas price change on air quality. Using the estimate from our preferred IV specification (table 1, col. 5) links changes in generation under the Hausman and Kellogg counterfactuals to predicted changes in air quality.²⁴ In summary, we estimate the change in air quality attributable to fracking using the following technique:

- a. Run dispatch model at observed 2012 input prices and record generation as $G_{ib}(P_{2012})$ in each hour b and boiler i .
- b. Run dispatch model at observed 2012 prices plus decrease in natural gas prices attributable to fracking and record generation as $G_{ib}(P_{2012} + \Delta P^k)$.
- c. Predict change in 2012 air quality at location j using IV coefficient estimates for the sum of the i boilers in the j th buffer:

$$\Delta PM^j = \widehat{\beta}_{IV} \cdot \sum_{b,i \in j} G_{ib}(P_{2012} + \Delta P^k) - \widehat{\beta}_{IV} \cdot \sum_{b,i \in j} G_{ib}(P_{2012}), \quad (3)$$

with $k \in \{\text{Medium, Upper, Lower}\}$ denoting a Hausman and Kellogg scenario. Accounting for differential impacts across space is important because generation is

24. Given a valid estimate of the effect of coal generation on $\text{PM}_{2.5}$ concentrations we can identify the effect of fracking on air quality. To do so we take the instrumented change in coal generation and aggregate it to the monitor level. We then multiply the estimated coefficient by the predicted generation. This creates a map with the spatial distribution of air quality post-fracking. We then perform the same exercise with a counterfactual input price schedule which assumes that natural gas prices return to their pre-recession level (adding the Hausman-Kellogg price of \$3.41). This creates a map with the spatial distribution of air quality assuming that fracking had not caused a decrease in natural gas prices. The difference between these two maps produces the spatial distribution of air quality changes attributable to fracking.

determined by the composition of installed power plant efficiency by fuel type, which varies spatially (Holladay and LaRiviere 2017).²⁵

4. RESULTS

This section describes the reduced form relationship between electricity generation at coal-fired power plants and the level of ambient air pollution using our data from 2007 to 2012 and the econometric framework described in section 4.²⁶

4.1. OLS Estimates

Table 1 reports the results of 15 different linear regressions, each of which seeks to quantify the relationship between power plant activity and ambient $PM_{2.5}$. The electrical load and emitted pollution are summed over all coal-fired power plants that fall within the circle surrounding a given power plant. Within each panel, five models are specified for each of three measures of plant activity: gross electrical load, emitted SO_2 , and emitted NO_x . In columns 1–2 we employ the annual time unit of observation, progressively adding weather controls and county fixed effects and year fixed effects. Column 3 reports results of our quarterly aggregation, and columns 4 and 5 report results of monthly aggregation, while column 5 also includes county as well as US Census Region by year by month fixed effects.

Our preferred model is specification 5. Column 5 takes the unit of observation as the 70-mile circle around a pollution monitor and uses monthly average daily generation levels as a right-hand-side variable while including region-by-month-by-year fixed effects. In this specification we identify the coefficient through variation in generation across monitors within a region-month-year. Put another way, observed differences in $PM_{2.5}$ monitors in those regions are identified by coal-fired power plants in one area of a region that are less efficient than in another and therefore decreasing production due to inexpensive natural gas.

Our OLS estimates show that there is a statistically significant positive relationship between electricity generated at coal-fired power plants and the level of ambient $PM_{2.5}$ as measured at pollution monitors in nearby populated areas. Our preferred model

25. This shows up in our counterfactuals asymmetrically across regions: adding the lower bound for fracking's impact on natural gas prices (\$2.16/mmBtu) relative to upper bound for fracking's impact (\$4.11/mmBtu) leads to no large differences in coal generation in the Florida Reliability Coordinating Council (FRCC), MRO, Southwest Power Pool (SPP), Texas Reliability Entity (TRE), or Western Electricity Coordinating Council (WECC). In RFC, Northeast Power Coordinating Council (NPCC), and SERC, however, natural gas generation increases by between roughly 40%, 8%, and 130%, respectively. Note that while the level change in coal and natural gas generation is identical, the percentage change in coal generation is smaller than the percentage change in natural gas generation in these regions because the installation bases are so large.

26. We also find that a log-linear specification produces similar results. See app. 2 for details.

Table 1. OLS at Monitor Level (70-Mile Radius)

| | (1) | (2) | (3) | (4) | (5) |
|-----------------------------------|---------------------|---------------------|---------------------|---------------------|---------------------|
| Measure of plant activity: | | | | | |
| 1. Average daily gross load (TWh) | 17.80*** (3.572) | 16.29*** (3.362) | 16.12*** (3.084) | 18.11*** (2.902) | 17.40*** (2.495) |
| R-squared | .853 | .862 | .693 | .591 | .696 |
| Adjusted R-squared | .832 | .842 | .682 | .586 | .689 |
| Time Unit of Observation | | | | | |
| | Annual | Annual | Quarterly | Monthly | Monthly |
| Employment controls | yes | yes | yes | yes | yes |
| Weather controls | | yes | yes | yes | yes |
| Fixed effects: | | | | | |
| County | yes | yes | yes | yes | yes |
| Year | yes | yes | | | |
| Year × quarter | | | yes | | |
| Year × month | | | | yes | |
| Region × year × month | | | | | yes |
| Number of monitors | 387 | 387 | 387 | 387 | 387 |
| Observations | 2,316 | 2,316 | 9,255 | 27,763 | 27,763 |

Note. Table 1 reports the results of 15 different linear regressions, each of which seeks to quantify the relationship between power plant activity and ambient $PM_{2.5}$. Five models are specified. The spatial unit of observation is a 70-mile radius circle centered at each of 387 $PM_{2.5}$ monitors. The electrical load and emitted pollution are summed over all coal-fired power plants that fall within the circle surrounding a given air quality monitor. In cols. 1–2 we employ the annual time unit of observation, progressively adding weather controls and county fixed effects and year fixed effects. Column 3 reports results of our quarterly aggregation, and cols. 4 and 5 report results of monthly aggregation, while col. 5 also includes county as well as US Census Region by year by month fixed effects. Standard errors clustered by state in parentheses.

* $p < .1$.

** $p < .05$.

*** $p < .01$.

specification (col. 5 of table 1) suggests that an increase of daily generation by one terawatt-hour of coal-fired electrical generation corresponds to a $17.39 \mu\text{g}/\text{m}^3$ increase in ambient $PM_{2.5}$ (which is the equivalent of one gigawatt-hour increasing $PM_{2.5}$ by $0.0174 \mu\text{g}/\text{m}^3$). This result is robust to a variety of time unit specifications—estimates using annual and quarterly time frames return values ranging from $16.11 \mu\text{g}/\text{m}^3$ to $18.10 \mu\text{g}/\text{m}^3$ per terawatt-hour.

To put these numbers into context: the largest coal-fired power plant in our sample (W. A. Parish Power Station near Houston, TX) is capable of producing 95 GWh per day. The effect of such a power plant on $PM_{2.5}$ would be an increase of $1.65 \mu\text{g}/\text{m}^3$. In

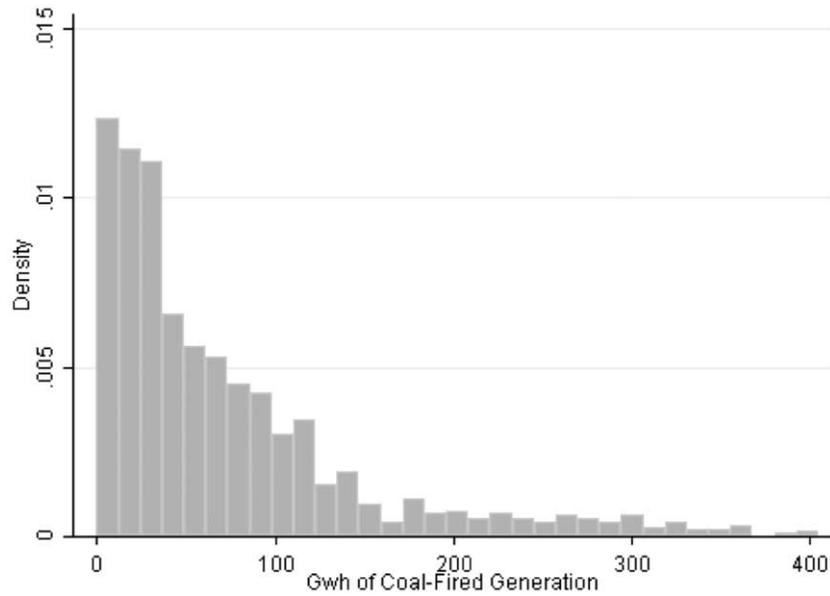


Figure 10. Histogram of total generation (GWh) in 70-mile radius observations. Under the 70-mile radius unit of observation, we sum the electrical generation of all plants within that radius of each pollution monitor. The mean total generation is 73.7 GWh; the maximum is 405. Color version available as an online enhancement.

our sample, the 2012 mean daily $PM_{2.5}$ is $9.47 \mu\text{g}/\text{m}^3$. Hence, if W. A. Parish Power Station were to go from being shut down to operating at full capacity in a typical county, we would expect it to increase the ambient $PM_{2.5}$ level by approximately 17.4% based on these OLS regressions. This effect is amplified if we assume that all power plants were to shut down around a given pollution monitor within 70 miles. Figure 10 shows the histogram of total generation under this assumption. The mean total generation of 73.7 GWh corresponds to $PM_{2.5}$ level changes of $1.3 \mu\text{g}/\text{m}^3$, and in the most coal-intensive area of the United States—with a generation of 405 GWh in one 70 buffer—we expect $PM_{2.5}$ level changes of $7.04 \mu\text{g}/\text{m}^3$. This implies that in such a region, the shutdown of all plants would lower $PM_{2.5}$ levels by 74%. Note that not all of this change would be attributable to fracking.

Finally, we used an OLS regression to test the primary assumption in the paper that coal-fired generation contributes to higher $PM_{2.5}$ than natural gas-fired generation. To test this assumption,²⁷ the above set of regressions is repeated by including the production of electricity from natural gas-fired power plants in addition to the coal-fired generation

27. According to the EIA (1998), coal-fired power plants emit 392 times as many units of particulate matter per unit of electricity compared to natural gas-fired power plants.

as two separate regressors.²⁸ Similar to table 1, we find that in our preferred specification (monthly data with full controls) the point estimate on coal is 20.54 (standard error 5.4) but that the coefficient on natural gas is imprecisely estimated: -101 (standard error 110). Thus, natural gas-fired generation does not have a significant impact on local air pollution in any specification, while the effect of coal-fired generation is qualitatively very similar to our previous regressions of table 1. A full output table is available upon request.

4.2. IV Estimates

Table 2 presents the PM_{2.5} results using the instrument. We first note that the IV estimates tend to have the same sign as the OLS estimates and in some cases are larger.²⁹ The coefficient estimate on instrumented coal generation seems plausible: the mean amount of daily coal-fired generation at monitors is 73.7 GWhs per day. According to our IV preferred estimate in specification 5, that corresponds to ambient PM_{2.5} concentration of 1.53 +/- [.67] $\mu\text{g}/\text{m}^3$, with results in brackets accounting for the 95% confidence interval. On average, then, our estimates suggest that coal-fired generation's contribution to ambient PM_{2.5} levels is 16% +/- [7%] in areas where coal-fired generation is present. In the most coal-intensive area of the United States—with a generation level of 405 GWh in a single 70 buffer—our IV method predicts PM_{2.5} changes of 8.41 $\mu\text{g}/\text{m}^3$. This implies that in such a region, the shutdown of all plants would lower PM_{2.5} levels by 89% +/- [39%].³⁰

4.3. IV Robustness Checks

We performed four robustness checks around our main IV results. The first was to evaluate and justify the 70-mile radius used in Deschênes et al. (2017) in our context. Table 3 shows the results using 40- and 100-miles radiuses around power plant locations to select air quality monitors. The main challenge with using only a 40-mile radius is that the sample size is greatly reduced since 40% of air quality monitors do not have any coal-fired power plants within 40 miles. The main reason to not use a larger radius is

28. In the IV specification 6, natural gas is instrumented by the dispatch models boiler-specific hourly natural gas predictions (in the same fashion as we instrumented for coal-fired generation).

29. Focusing on panel 1 (with the right-hand-side variable of coal in TWhs), the IV estimate is 19% larger in our preferred specification 5, 63% larger in specification 4, and 72% larger in specification 3 compared to the OLS specifications. Only at the yearly level are the estimates almost identical in magnitude as the OLS estimates.

30. To put these numbers into context: the largest coal-fired power plant in our sample (W. A. Parish Power Station near Houston, TX) is capable of producing 95 GWh per day. If W. A. Parish Power Station were to go from being shut down to operating at full capacity in a typical county, we would expect it to increase the ambient PM_{2.5} level by 21% +/- [9%] based on these IV regressions.

Table 2. IV Regression of Average PM_{2.5} on Instrumented Average Daily Coal Generation

| Variables | 70 Mile | | | | |
|------------------------------------|--------------------------|---------------------|---------------------|---------------------|---------------------|
| | (1) | (2) | (3) | (4) | (5) |
| 1. Average daily gross load (TWhs) | 17.61*** (6.140) | 15.06*** (5.001) | 27.64*** (7.112) | 29.43*** (6.499) | 20.76*** (4.614) |
| R-squared | .853 | .862 | .694 | .591 | .696 |
| Adjusted R-squared | .832 | .842 | .683 | .586 | .689 |
| | Time Unit of Observation | | | | |
| | Annual | Annual | Quarterly | Monthly | Monthly |
| Employment controls | yes | yes | yes | yes | yes |
| Weather controls | | yes | yes | yes | yes |
| Fixed effects: | | | | | |
| County | yes | yes | yes | yes | yes |
| Year | yes | yes | | | |
| Year × quarter | | | yes | | |
| Year × month | | | | yes | |
| Region × year × month | | | | | yes |
| First-stage <i>F</i> -statistic | 33.5244 | 33.6892 | 45.7696 | 53.4829 | 53.4625 |
| Prob > <i>F</i> | .0000 | .0000 | .0000 | .0000 | .0000 |
| Observations | 2,316 | 2,316 | 9,255 | 27,763 | 27,763 |

Note. This table reports the results of 15 different IV regressions, each of which seeks to quantify the relationship between power plant activity and ambient PM_{2.5}. The unit of observation is an air pollution monitor with a 70-mile radius circle. The electrical load and emitted pollution are summed over all coal-fired power plants that fall within the 70-mile circle surrounding a given monitor. In cols. 1–2 we employ the annual time unit of observation, progressively adding weather controls and county fixed effects and year fixed effects. Column 3 reports results of our quarterly aggregation, and cols. 4 and 5 report results of monthly aggregation, while col. 5 also includes county as well as US Census Region by year by month fixed effects. Standard errors clustered by state in parentheses.

* $p < .1$.

** $p < .05$.

*** $p < .01$.

that additional coal power plants are included for each air quality monitor, which may or may not be ideal.

Table 3 shows the result of the IV regression for the 40- and 100-mile circle specifications. The point estimates are somewhat smaller in magnitude for the 40-mile specification (roughly 60% of the 70-mile point estimates) but roughly the same for the 100-mile specification. The 40-mile specification point estimates are sometimes not statistically significant as well due to the reduced number of observations. As a result, we view that following the earlier empirical work and using a 70-mile radius is reasonable to use in our case.

Table 3. IV Regression of Average PM_{2.5} on Instrumented Average Daily Coal Generation

| | (1) | (2) | (3) | (4) | (5) |
|---------------------------|---------------------|-------------------|---------------------|---------------------|---------------------|
| A. 40 Mile | | | | | |
| Average daily coal (TWhs) | 8.821 (11.32) | 6.935 (9.938) | 14.82 (11.20) | 18.26* (10.40) | 12.20 (9.253) |
| Constant | 12.51*** (.270) | -35.72 (38.14) | -10.08 (40.74) | 96.80*** (34.98) | 79.16** (30.89) |
| R-squared | .842 | .852 | .697 | .602 | .694 |
| Adjusted R-squared | .820 | .830 | .687 | .596 | .687 |
| First-stage F-statistic | 20.2252 | 19.5422 | 24.0902 | 28.0278 | 25.8021 |
| Prob > F | .0001 | .0001 | .0000 | .0000 | .0000 |
| Observations | 1,980 | 1,980 | 7,899 | 26,942 | 26,942 |
| B. 100 Mile | | | | | |
| Average daily coal (TWhs) | 18.50 (11.72) | 16.12 (10.92) | 26.63*** (9.830) | 26.64*** (7.272) | 19.70*** (6.397) |
| Constant | 10.22*** (1.530) | 39.63 (48.47) | 61.97 (46.10) | 139.4*** (32.94) | 121.7*** (27.60) |
| R-squared | .848 | .860 | .683 | .602 | .694 |
| Adjusted R-squared | .826 | .840 | .672 | .596 | .687 |
| First-stage F-statistic | 27.4496 | 25.9687 | 49.0441 | 74.7078 | 84.3424 |
| Prob > F | .0000 | .0000 | .0000 | .0000 | .0000 |
| Observations | 2,472 | 2,472 | 9,888 | 26,942 | 26,942 |
| Time unit of observation | Annual | Annual | Quarterly | Monthly | Monthly |
| Weather controls | | yes | yes | yes | yes |
| Fixed effects: | | | | | |
| County | yes | yes | yes | yes | yes |
| Year | yes | yes | | | |
| Year × quarter | | | yes | | |
| Year × month | | | | yes | |
| Region × year × month | | | | | yes |

Note. This table provides robustness checks for the IV 70-mile specification of table 2. Panel A shows the results for a 40-mile specification. The coefficient of interest is not statistically significant in most specifications. Panel B gives the results for a 100-mile specification. The coefficient of interest is statistically significant in three of the five specifications and similar in magnitude to the 70-mile specifications shown in table 2. A larger mileage range implies more viable coal-fired power plants be included in the analysis, explaining differences in number of observations. Standard errors clustered by state in parentheses.

* $p < .1$.

** $p < .05$.

*** $p < .01$.

The second robustness check was to include nonfuel variable operations and maintenance (O&M) costs in the dispatch model. Nonfuel variable O&M costs were taken from the FERC form 1 data and include all nonfuel O&M costs, including pollution allowances. The FERC form 1 data are often not reported at the unit ID level in the CEMS data we use for the dispatch model but instead at the more aggregate Office of the Regulatory Information System Plant (ORISPL) code level. As a result, some type of allocation of O&M costs to a unit ID in our data must take place. Where there is not a direct match we use the allocations used by S&P Global Market Intelligence (GMI, formerly SNL). The appendix details the specifics of this allocation process.

The main insight of the O&M robustness exercise is that coal has higher O&M costs than natural gas units. For example, combined cycle coal units have an O&M cost of \$8.05/MWh whereas combined cycle natural gas units are \$3.45/MWh. Since these costs do not change with fuel prices, however, it is unclear that including them will impact the IV results. Table 4 shows that the results are almost identical to the results from not including O&M costs and that r^2 changes by only a small amount ($\sim .005$). As a result, we conclude that using O&M costs does not impact our qualitative findings. We exclude them from our main specification.

The third robustness check uses an algorithm to allow for possible ramping constraints of coal and natural gas units (Reguant 2014). We include a detailed description of the algorithm in the appendix, which we adopt from Linn and McCormack (2018). The main assumption is that any coal (natural gas) power plant that produces during the peak hour within a day must also generate 30% of its capacity for the 6 (3) hours on either side of the peak hour. The 30% and 6 (3) hours are certainly ad hoc but informed by reduced form evidence in Linn and McCormack (2018).

Table 5 shows the results from the model with ramping constraints. It shows that there are significant differences in the point estimates from adding ramping constraints in the way we have done it; point estimates basically double as do standard errors, although r^2 is effectively unchanged. We interpret this in two ways. First, the estimates that do not include ramping constraints could be underestimates: the IV approach might be dispatching too much coal and therefore attributing not enough $PM_{2.5}$ to coal generation. This would imply that our point estimates without ramping constraints are too low and conservative estimates. More importantly, without a more principled way of choosing parameters for the ramping algorithm, it is unclear that it can be relied upon. As a result we note that our estimates are likely to be conservative and encourage additional research to use other approaches, including machine learning approaches, to tackle the challenging ramping issue. The appendix also includes detailed discussion of first-stage results of this approach.

Finally, we performed a fourth robustness check in which we allow the weather variables to enter the second-stage regression relating instrumented generation to ambient $PM_{2.5}$ levels nonlinearly. We do not observe any changes to the qualitative results when doing so. The one exception is in the model with ramping constraints when the point estimates decrease in the monthly and quarterly specifications by roughly 15% and r^2

Table 4. IV Regression of Average PM_{2.5} on Instrumented Average Daily Coal Generation with O&M

| Variables | 70 Mile | | | | |
|------------------------------------|--------------------------|---------------------|---------------------|---------------------|---------------------|
| | (1) | (2) | (3) | (4) | (5) |
| 1. Average daily gross load (TWhs) | 18.71*** (5.804) | 16.48*** (4.639) | 27.65*** (6.646) | 29.67*** (6.366) | 21.05*** (5.364) |
| R-squared | .854 | .863 | .691 | .589 | .696 |
| Adjusted R-squared | .832 | .843 | .680 | .584 | .689 |
| | Time Unit of Observation | | | | |
| | Annual | Annual | Quarterly | Monthly | Monthly |
| Employment controls | yes | yes | yes | yes | yes |
| Weather controls | | yes | yes | yes | yes |
| Fixed effects: | | | | | |
| County | yes | yes | yes | yes | yes |
| Year | yes | yes | | | |
| Year × quarter | | | yes | | |
| Year × month | | | | yes | |
| Region × year × month | | | | | yes |
| First-stage <i>F</i> -statistic | 33.6692 | 33.6243 | 44.6538 | 52.1944 | 52.1995 |
| Prob > <i>F</i> | .0000 | .0000 | .0000 | .0000 | .0000 |
| Observations | 2,316 | 2,316 | 9,255 | 27,763 | 27,763 |

Note. This table reports the results of 15 different IV regressions including O&M costs, each of which seeks to quantify the relationship between power plant activity and ambient PM_{2.5}. The unit of observation is an air pollution monitor with a 70-mile radius circle. The electrical load and emitted pollution are summed over all coal-fired power plants that fall within the 70-mile circle surrounding a given monitor. In cols. 1–2 we employ the annual time unit of observation, progressively adding weather controls and county fixed effects and year fixed effects. Column 3 reports results of our quarterly aggregation, and cols. 4 and 5 report results of monthly aggregation, while col. 5 also includes county as well as US Census Region by year by month fixed effects. Standard errors clustered by state in parentheses.

* $p < .1$.

** $p < .05$.

*** $p < .01$.

increases by between .02 and .06, making it more in line with the other specifications. We do not present the nonlinear weather results here, but they are available upon request.

In sum, our robustness checks show the method to be robust to a variety of model permutations. The one exception is the model with ramping constraints which implies that our estimates are possibly conservative. We thus err on the side of conservatism. One possible extension beyond the scope of this paper would be to evaluate the impact of different types of coal (high vs. low sulfur coal) on ambient air pollution, although

Table 5. IV Regression of Average PM_{2.5} on Instrumented Average Daily Coal Generation with O&M and Ramping Constraints

| Variables | 70 Mile | | | | |
|------------------------------------|--------------------------|--------------------|---------------------|---------------------|---------------------|
| | (1) | (2) | (3) | (4) | (5) |
| 1. Average daily gross load (TWhs) | 27.58** (11.68) | 28.56** (11.67) | 63.90*** (13.04) | 69.43*** (12.87) | 52.22*** (11.89) |
| R-squared | .850 | .860 | .685 | .584 | .693 |
| Adjusted R-squared | .828 | .839 | .674 | .579 | .686 |
| | Time Unit of Observation | | | | |
| | Annual | Annual | Quarterly | Monthly | Monthly |
| Employment controls | yes | yes | yes | yes | yes |
| Weather controls | | yes | yes | yes | yes |
| Fixed effects: | | | | | |
| County | yes | yes | yes | yes | yes |
| Year | yes | yes | | | |
| Year × quarter | | | yes | | |
| Year × month | | | | yes | |
| Region × year × month | | | | | yes |
| First-stage <i>F</i> -statistic | 27.8119 | 30.1356 | 47.6743 | 72.5875 | 86.5714 |
| Prob > <i>F</i> | .0000 | .0000 | .0000 | .0000 | .0000 |
| Observations | 2,316 | 2,316 | 9,255 | 27,763 | 27,763 |

Note. This table reports the results of 15 different IV regressions including O&M costs, each of which seeks to quantify the relationship between power plant activity and ambient PM_{2.5}. The unit of observation is an air pollution monitor with a 70-mile radius circle. The electrical load and emitted pollution are summed over all coal-fired power plants that fall within the 70-mile circle surrounding a given monitor. In cols. 1–2 we employ the annual time unit of observation, progressively adding weather controls and county fixed effects and year fixed effects. Column 3 reports results of our quarterly aggregation, and cols. 4 and 5 report results of monthly aggregation, while col. 5 also includes county as well as US Census Region by year by month fixed effects. Standard errors clustered by state in parentheses.

* $p < .1$.

** $p < .05$.

*** $p < .01$.

the correlation of coal-burning technologies and proximity to different types of coal would require a different research design.

4.4. Third Stage

Not all of the air quality changes in the above IV regressions are attributable to fracking; some are attributable to the changes in the relative prices. To isolate the effect of fracking, we proceed in stage 3 where we leverage the dispatch model and the natural gas price change attributable to fracking. Our third and final step isolates the effect

of the fracking-related change in the price of natural gas from any other confounder. Using formula (3) to identify the predicted $PM_{2.5}$ reduction using the price difference by Hausman and Kellogg (2015), we find a US national average decrease in air pollution of 4% within 70 miles of a coal plant.

One advantage of our approach is that we can study the spatial incidence of fracking at each air quality monitor in the United States. Figure 11 shows the results of the counterfactual analysis for generation of both coal (panel *a*) and natural gas (panel *b*) power plants based on the medium case of Hausman and Kellogg (2015) with price differential ΔP^{Medium} . The figure shows the changes in generation that result from changes in the NERC level dispatch order as a result of natural gas prices that were uniformly \$3.41/mmBtu higher across the United States. The implication is that these changes would not have occurred if fracking had been banned. Figure 11 shows the spatial substitution patterns of coal and natural gas generation. These patterns are dictated by the NERC level supply model and the embedded generation capacity within each region during the time frame of the analysis. While regional patterns across natural gas-fired generation increases and coal-fired generation decreases are commensurate, there is clear within-region variation. In the Northeast, coal-fired generation decreased across the entire region, but increases in natural gas-fired generation were concentrated close to the New York City metro area. A similar pattern holds for the Southeast. It is also clear that the WECC region was not dramatically impacted by a change in the dispatch order.

While figure 11 maps changes in generation, we can apply our section 4.2 estimate of the relationship between coal-fired generation and $PM_{2.5}$ to interpret the figure in terms

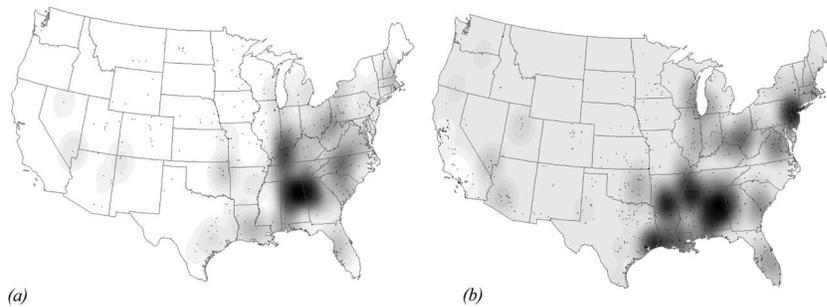


Figure 11. Regionally differentiated incidence of fracking on ambient air quality. *a*, Counterfactual increase in coal generation. *b*, Increase in natural gas generation. Panel *a* displays the simulated counterfactual situation in 2012 if fracking had not occurred in the United States from 2007 to 2012. In panel *a*, darkness indicates more dramatic increase in levels of coal generation if fracking had not occurred. The darkest spot in Alabama corresponds to a causal decrease in $PM_{2.5}$ levels of 35% due to fracking. Panel *b* displays the causal spatial increase in electricity generation by natural gas due to fracking in 2012 relative if fracking had not occurred. In panel *a* each black dot represents a coal-fired power plant. In panel *b* each dot represents a natural gas-fired power plant. Color version available as an online enhancement.

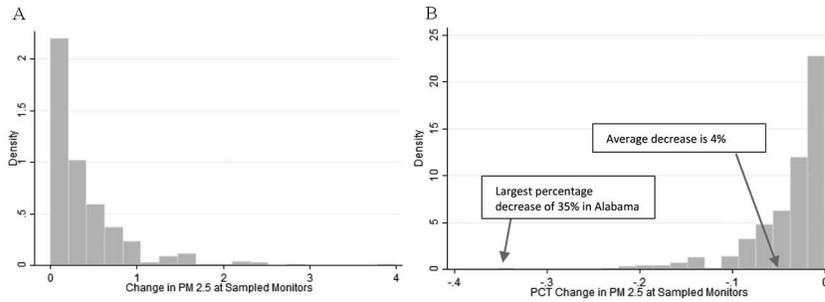


Figure 12. Histogram of level and percentage change of $PM_{2.5}$ at the monitor level. *A*, Counterfactual increase in $PM_{2.5}$ without fracking. *B*, Percentage change due to fracking. Panel *A* displays 2012 monitor level changes in ambient air quality relative to “no fracking” counterfactual. Panel *B* displays these changes in terms of percentage points. The largest percentage decrease in $PM_{2.5}$ is 35% in Alabama, but that was an outlier; the next closest was a decrease of 23%. The national average decrease in air pollution is 4% around coal-fired power plants caused by fracking. Color version available as an online enhancement.

of the difference in ambient air pollution between the 2012 observed levels and the counterfactual scenario. The shaded areas of figure 11*a* correspond to a change in air pollution levels ranging from 0% (white) to 35% (black). Figure 12 shows two histograms of monitor level outcomes from the counterfactual dispatch model, one in levels and one in percentages. The largest percentage change in air quality was in Alabama which would have observed 35% higher levels of $PM_{2.5}$ had fracking not decreased the price of natural gas.³¹

We have performed the same analysis with the Hausman and Kellogg (2015) lower and upper scenarios. While there are small level changes in the results if fracking impacted the price drop at the upper bound (\$4.11) or lower bound (\$2.16) rather than the expected impact \$3.41, we found little difference in dispatch order. In table A3 (tables A1–A4 are available online) we argue that the lack of difference between the counterfactual scenarios is the result of the limited ability of combined cycle boilers to exploit lower natural gas prices. This is consistent with our research design; we chose to stop our sample in 2012 to avoid modeling endogenous capacity decisions. We conclude that due to the fixed stock of generating capacity the change in dispatch composition does not scale linearly with changes in natural gas prices. This finding cautions against extrapolation to larger wedges between coal and natural gas prices, for example, as a result of carbon taxes. Our findings highlight the need for full investment

31. The percentage is calculated by $\Delta PM_{2.5}^j [\widehat{\beta}_{IV} \cdot \sum_{b,i \in j} G_{ib}(P_{2012} + \Delta P^k) + PM_{2.5}^j]$ from other sources at location j , whereby “ $PM_{2.5}$ from other sources” is the difference between the observed $PM_{2.5}$ in 2012 at location j and the $PM_{2.5}$ contribution from the surrounding power plants, $\widehat{\beta}_{IV} \cdot \sum_{b,i \in j} G_{ib}(P_{2012})$.

Table 6. Gross Externality Benefits Measured in USD from Improved Ambient Air Quality due to the Displacement of Coal from the Fracking Revolution

| Health Cost Scenarios | Scenario of Natural Gas Price Reduction due to Fracking | | |
|-----------------------|---|----------------|----------------|
| | Lower | Medium | Upper |
| Lower bound | 5,403,056,856 | 5,918,083,939 | 5,947,291,829 |
| Medium | 15,239,391,133 | 16,692,031,624 | 16,774,412,851 |
| Upper bound | 38,998,987,309 | 42,716,426,382 | 42,927,247,432 |

Note. The columns reflect the three calculated scenarios of the reduction in the price of natural gas due to fracking. The lower, medium, and upper scenarios correspond to a decrease in the price of natural gas of \$2.16, \$3.41, and \$4.11 per mmBtu, respectively. The rows represent the gross externality measures that relate coal generation to health costs: the lower, medium, and upper assume an age-adjusted value of statistical life of 2.8 million, age-adjusted VSL of 7.7 million and age-unadjusted (uniform) VSL of 7.7 million, respectively, as described in Muller et al. (2011). All prices are in 2012 USD.

models like Gowrisankaran et al. (2016) or Cullen and Reynolds (2016) when analyzing counterfactual scenarios with large price differentials.

4.5. Monetization of Indirect Benefits

To put the change in air pollution levels into context we use the monetary externality cost measures of the local air pollutants derived by Muller et al. (2011).³² We find that displacement of coal due to fracking provided health benefits of \$17 billion per year by reducing PM_{2.5}, SO₂, and NO_x emissions, with a lower bound externality benefit of \$5.4 billion and an upper bound of \$43 billion per year. The “medium,” “lower,” and “upper” benefits are calculated under the assumptions as described in the note to table 6. Note again that values between the columns do not change substantially, as the medium and upper bound scenarios of the natural gas price difference have little impact on coal-fired generation relative to the lower bound scenario, due to the nonlinear natural gas price supply function. In comparison, the magnitude of the changes across the rows is considerable, due to the different assumptions in calculating the health costs, as outlined in the note to table 6 and detailed in Muller et al. (2011).

As an alternative measure of health externalities, the medical literature reports that a decrease in PM_{2.5} of .25–.64 μg/m³ would correspond to a decrease in lung cancer rates of 0.9%–2.3% (Raaschou-Nielson et al. 2013).³³ The American Lung Association reports that there are 160,000 deaths related to lung cancer each year and that roughly

32. This monetary externality cost accounts for local air pollutants only. It excludes the social cost of carbon.

33. These calculations assume linear dose-response functions. Raaschou-Nielson et al. (2013) experimented with nonlinear models as well but conclude that the results do not deviate from their linear dose response model.

85% of these are due to smoking. Assuming that 10% are due to external air quality conditions, our estimates suggest that fracking then is responsible for saving between 129 and 354 lung cancer lives each year, the equivalent of \$1.2 to \$3.3 billion annually using current VSL measures. Overall, these back of the envelope estimates suggest that the increased air quality attributable to fracking, via the natural gas price decrease, is on the order of the lower billions of annual US dollars.³⁴

5. CONCLUSION

Our main contribution is to leverage a structural dispatch model with exogenous natural gas price changes to identify the causal effect of natural gas price decreases attributable to fracking on air quality. We find that fracking displaced 28% of coal-fired generation in the short term during our sample period. A full investigation of long-run impacts would require a dynamic model of natural gas capacity investments, but as natural gas-fired generation capacity increases coal-fired generation should be further displaced. We can therefore interpret our medium-run findings as a lower bound on the long-run impacts of fracking on air pollution.

Assuming that coal displacement was uniformly distributed, our IV estimates imply a 95% confidence interval of 2%–6% for decreases in $PM_{2.5}$ due to fracking. We show, however, that coal displacement varies spatially, with the most pronounced air quality improvements in Alabama, where air pollution decreased by 35% due to fracking. As a result, we find evidence of a significant environmental benefit attributable to fracking, with an average estimate of \$17 billion in annual health benefits.³⁵ Identifying this lower bound from short-run impacts is a key contribution of this paper, and longer-run impacts are almost surely larger as natural gas capacity changes.

In this paper we limit our study to natural gas fracking (although oil-related fracking also fundamentally changed international energy markets in significant ways). Focusing on natural gas provides methodological advantages over oil because very little natural gas is exported from the United States. As a result, US natural gas prices capitalize US fracking more directly. Second, we look at local air pollutants only, and not at CO_2 , which is a global air pollutant. Even though CO_2 emissions decreased substantially in the US electricity-generation sector during our study period, US coal continues to be mined for export. As a result, fracking's impact on global CO_2 emissions is ambiguous. Locally measured $PM_{2.5}$ does not suffer from this leakage problem. Knittel et al. (2015)

34. For any of the above damage calculations, note that these are likely lower bounds. If the damage function from air pollution on health is convex, the monetary benefits would be larger because most coal-fired power plants with the largest decreases in production due to fracking are located in the most $PM_{2.5}$ polluted areas of the eastern United States.

35. We interpret, our analysis to be short run to medium run. For a discussion on the interpretation of short-, medium-, vs. long-run elasticities, see the comment and discussion section in Hausman and Kellogg (2015).

as well as Linn et al. (2014) are two recent promising working papers that provide methods to analyze CO₂ in this context.

Finally, we once again point out that this study by no means presents a full cost-benefit analysis of fracking. Rather, it contributes to cost-benefit analyses by developing a novel three-stage methodology to estimate indirect partial air pollution benefits. Our results highlight the importance of incidence for developing policies that maximize national welfare. Politicians motivate bans on fracking by pointing out negative externalities, localized at fracking sites. However, our results highlight large positive indirect impacts in areas in which coal-fired electricity production has been replaced by cleaner natural gas. Similar political issues occur in free trade debates, where local job loss receives enormous political attention while marginal decreases in consumer prices at the national level have disparate benefits. Because fracking policy is created at the state level (rather than trade policy being created at the national level), this discord highlights the costs of the disjointed energy policy that has characterized the United States in recent decades. Furthermore, states like New York, which have banned fracking, are able to enjoy lower natural gas prices without suffering negative nonmarket impacts or positive market impacts through leasing revenue. Holland et al. (2016) addresses a topic with similar regulatory federalism environmental externalities issues regarding state versus federal laws for electric cars about the efficiency of this type of policy in nationwide input markets. More work along these lines might provide guidance for policy makers.

REFERENCES

- Allcott, Hunt, and Daniel Keniston. 2015. Dutch disease or agglomeration? The local economic effects of natural resource booms in modern America. NBER Working paper 20508, National Bureau of Economic Research, Cambridge, MA.
- Ames, Robert, Anthony Corridore, Joel N. Ephross, Edward A. Hirs III, Paul W. MacAvoy, and Richard Tavelli. 2012. The arithmetic of shale gas. Yale Graduates in Energy Study Group, Yale University.
- Borenstein, Severin. 2002. The trouble with electricity markets: Understanding California's restructuring disaster. *Journal of Economic Perspectives* 88 (4): 669–85.
- Borenstein, Severin, James B. Bushnell, and Frank A. Wolak. 2002. Measuring market inefficiencies in California's wholesale electricity market. *American Economic Review* 92 (5): 1376–1405.
- Brandt, Adam R., Garvin Heath, E. A. Kort, Francis O'Sullivan, Gabrielle Pétron, Sarah M. Jordaan, P. Tans, et al. 2014. Methane leaks from North American natural gas systems. *Science* 343 (6172): 733–35.
- Chay, Kenneth Y., and Michael Greenstone. 2005. Does air quality matter? Evidence from the housing market. *Journal of Political Economy* 113 (2): 376–424.
- Chu, Y., J. S. Holladay, and J. LaRiviere. 2017. Opportunity cost pass-through from fossil fuel market prices to procurement costs of the US power producers. *Journal of Industrial Economics* 65 (4): 842–71.
- Crawford, Jagoda, Scott D. Chambers, David D. Cohen, Alan Griffiths, Alastair G. Williams, and E. Stelcer. 2015. Using Radon-222 as an indicator of atmospheric mixing depth in ME-2 for PM_{2.5} source apportionment. *Aerosol and Air Quality Research* 15:611–24.

- 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): 106–33.
- 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, University of Arizona.
- Currie, Janet, and Matthew Neidell. 2005. Air pollution and infant health: What can we learn from California's recent experience? *Quarterly Journal of Economics* 120 (3): 1003–30.
- Davis, Lucas, and Catherine Hausman. 2016. Market impacts of a nuclear power plant closure. *American Economics Journal: Applied Economics* 8 (2): 92–122.
- Deschênes, Olivier, Michael Greenstone, and Joseph S. Shapiro. 2017. Defensive investments and the demand for air quality: Evidence from the NOx budget program. *American Economic Review* 107 (10): 2958–89.
- EIA (Energy Information Administration). 1999. Natural gas 1998: Issues and trends. Energy Information Administration, Office of Oil and Gas, Washington, DC. http://inis.iaea.org/search/search.aspx?orig_q=RN:30046498.
- . 2016. Today in energy: Monthly reports, Energy Information Administration, Office of Oil and Gas, Washington, DC. <https://www.eia.gov/todayinenergy> (accessed April 24, 2016).
- Fabra, Natalia, and Mar Reguant. 2014. Pass-through of emissions costs in electricity markets. *American Economic Review* 104 (9): 2872–99.
- Fabrizio, Kira R., Nancy L. Rose, and Catherine D. Wolfram. 2008. Do markets reduce costs? Assessing the impact of regulatory restructuring on U.S. electric generation efficiency. *American Economic Review* 97 (4): 1250–77.
- Gowrissankaran, Gautam, Stanley S. Reynolds, and Mario Samano. 2016. Intermittency and the value of renewable energy. *Journal of Political Economy* 124 (4): 1187–1234.
- Graff Zivin, Joshua, and Matthew Neidell. 2013. Environment, health, and human capital. *Journal of Economic Literature* 51 (3): 689–730.
- Graham, Jove, Jennifer Irving, Xiaoqin Tang, Stephen Sellers, Joshua Crisp, Daniel Horwitz, Lucija Muehlenbachs, Alan Krupnick, and David Carey. 2015. Increased traffic accident rates associated with shale gas drilling in Pennsylvania. *Accident Analysis and Prevention* 74:203–9.
- Hausman, Catherine, and Ryan Kellogg. 2015. Welfare and distributional implications of shale gas. *Brookings Papers on Economic Activity*, Washington, DC.
- Hodan, William M., and William R. Barnard. 2004. Evaluating the contribution of PM2.5 precursor gases and re-entrained road emissions to mobile source PM2.5 particulate matter emissions. MACTEC Federal Programs, Research Triangle Park, NC.
- Holladay, J. S., and J. 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* 85:205–27.
- Holland, Stephen P., Erin T. Manser, Nicholas Z. Muller, and Andrew J. Yates. 2016. Are there environmental benefits from driving electric vehicles? The importance of local factors. *American Economic Review* 106 (12): 3700–3729.
- Isen, Adam, Maya Rossin-Slater, and W. Reed Walker. 2014. Every breath you take—every dollar you'll make: The long-term consequences of the Clean Air Act of 1970. NBER Working paper 19858, National Bureau of Economic Research, Cambridge, MA.

- Jackson, Robert B., Avner Vengosh, J. William Carey, Richard J. Davies, Thomas H. Darrah, Francis O'Sullivan, and Gabrielle Petron. 2014. The environmental costs and benefits of fracking. *Annual Review of Environment and Resources* 39:327–62.
- Knittel, Christopher R., Konstantinos Metaxoglou, and Andre Trindade. 2015. Natural gas prices and coal displacement: Evidence from electricity markets. NBER Working paper 21627, National Bureau of Economic Research, Cambridge, MA.
- Koster, Hans R. A., and Jos van Ommereen. 2015. A shaky business: Natural gas extraction, earthquakes and house prices. *European Economic Review* 80 (C): 120–39.
- Linn, Joshua, and Kristen McCormack. 2018. The roles of energy markets and environmental regulation in reducing coal-fired plant profits and electricity sector emissions. Resources for the Future, Washington, DC.
- Linn, Joshua, Lucija Muehlenbachs, and Yushuang Wang. 2014. How do natural gas prices affect electricity consumers and the environment? RFF Discussion paper no. 14-19, Resources for the Future, Washington, DC.
- Mansur, Erin T. 2007. Do oligopolists pollute less? Evidence from a restructured electricity market. *Journal of Industrial Economics* 55 (4): 661–89.
- Muehlenbachs, Lucija, and Alan J. Krupnick. 2014. Infographic: Shale gas development linked to traffic accidents in Pennsylvania. RFF Resources no. 185, Resources for the Future, Washington, DC.
- Muehlenbachs, Lucija, Elisheba Spiller, and Christopher Timmins. 2013. Shale gas development and the costs of groundwater contamination risk. RFF Discussion paper no. 12-40, Resources for the Future, Washington, DC.
- . 2015. The housing market impacts of shale gas development. *American Economic Review* 105 (12): 3633–59.
- Muller, Nicholas Z., Robert Mendelsohn, and William Nordhaus. 2011. Environmental accounting for pollution in the United States economy. *American Economic Review* 101 (5): 1649–75.
- Olmstead, Sheila M., Lucija A. Muehlenbachs, Jhib-Shyang Shih, Ziyang Chu, and Alan J. Krupnick. 2013. Shale gas development impacts on surface water quality in Pennsylvania. *Proceedings of the National Academy of Sciences* 110 (13): 4962–67.
- Quistorff, Brian. 2015. The effects of energy prices on manufacturing employment. University of Maryland Job Market paper.
- Pirovano, Guido, Cristina Colombi, Alessandra Balzarini, Gieseppe Riva, V. Gianelle, and G. Lonati. 2015. PM2.5 source apportionment in Lombardy (Italy): Comparison of receptor and chemistry-transport modelling results. *Atmospheric Environment* 106 (2015): 56–70.
- Prest, A. R., and R. Turvey. 1968. Cost-benefit analysis: A survey. *Economic Journal* 75 (300): 683–735.
- Raaschou-Nielson, Ole, Zorana J. Andersen, Rob Beelen, Evangelia Samoli, Massimo Stafoggia, Gudrun Weinmayr, Barbara Hoffmann, et al. 2013. Air pollution and lung cancer incidence in 17 European cohorts: Prospective analyses from the European Study of Cohorts for Air Pollution Effects (ESCAPE). *Lancet Oncology* 14 (9): 813–22.
- Reguant, M. 2014. Complementary bidding mechanisms and startup costs in electricity markets. *Review of Economic Studies* 81 (4): 1708–42.
- Schlenker, Wolfram, and Reed Walker. 2016. Airports, air pollution, and contemporaneous health. *Review of Economic Studies* 83 (2): 768–809.
- Sullivan, Daniel. 2017. The true cost of air pollution: Evidence from the housing market. RFF Working paper, Resources for the Future, Washington, DC.

- Wolfram, Catherine. 1999. Measuring duopoly power in the British electricity spot market. *American Economic Review* 89 (4): 805–26.
- Yu, Lingda, Guangfu Wang, Renjian Zhang, Leiming Zhang, Yu Song, Bingbing Wu, Xufang Li, Kun An, and Junhan Chu. 2013. Characterization and source apportionment of PM_{2.5} in an urban environment in Beijing. *Aerosol and Air Quality Research* 13:574–83.