Fracking, Coal and Air Quality

Reid Johnsen, Jacob LaRiviere and Hendrik Wolff

January 2017

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 air quality monitors merged with US power plants 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.

Key words: Fracking, Coal-Fired Power Plants, Air Pollution, Health, Electricity

JEL Codes:	Q41	Demand and Supply • Prices
	Q53	Air Pollution
	I18	Government Policy • Regulation • Public Health

Johnsen, Reid: Department of Agricultural and Resource Economics, University of California at Berkeley, (rlkj@berkeley.edu); Jacob LaRiviere: Microsoft (jlariv@microsoft.com) and Department of Economics at University of Tennessee and Hendrik Wolff, Department of Economics, Simon Fraser University (Hendrik_wolff@sfu.ca). Acknowledgments: We thank Werner Antweiler, Catherine Hausman, Joshua Linn, Erin Mansur, Joseph Shapiro and Nathan Tefft for particularly helpful criticisms 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. All errors are ours.

Introduction

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 macro-economic shifts.

This paper estimates the short run causal effect of 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 hydraulic fracturing ("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 U.S. municipalities and EU countries have strictly limited its scope.

But because fracking has dramatically decreased U.S. natural gas prices, it has potentially decreased coal consumption and concomitantly improved U.S. air quality. According the EIA (1998), 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 percent. Indeed, one contribution of our paper is to estimate the relative impact of coal vs. natural gas on air quality econometrically using the (near) universe of all natural gas and coal-fired power plants. As a

[1] Damages discussed in the economics literature include: toxic leaks into groundwater supplies (Muehlenbachs, Spiller and Timmins, 2013, 2015); chemicals exposing surface water (Olmstead et al., 2013); traffic accidents (Muehlenbachs and Krupnick, 2014, Graham et al., 2015), earthquakes (Koster, 2015), 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 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.

result of its effect on natural gas prices, fracking may have introduced indirect non-market benefits through air quality improvements. Air quality benefits are especially important, as a large share of the U.S. air quality monitors are in noncompliance with the current EPA air quality limits, which were imposed in lieu of Pigouvian taxes to mitigate PM's negative health externalities. Figure 1 shows this explicitly by showing the distribution of air quality monitors that are out of compliance during our sample period.

To estimate these indirect effects, however, poses major econometric challenges because the fracking boom occurred simultaneously with fundamental shifts in the macro-economy. During our period of study, 2007 through 2012, changing structures of electricity markets, growing Chinese demand for coal, and most of all the Great Recession and subsequent recovery all had the potential to affect patterns of electricity generation.[4] For example if the recession led to lower electricity demand and fossil fuel 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 of what electricity generation would have been in the absence of fracking's impact on natural gas prices, ceteris paribus. In this way, we estimate what would have occurred to electricity sector pollution if fracking had been banned in the U.S. Aside from adequately modeling fuel substitution, there are three principal endogeneity challenges to do so. The first is the accounting for the recession's impact on electricity demand. As a result, we must have a research design which conditions on observed electricity generation. Second, some of the decrease in natural gas prices which occurred over our sample was due to decreased demand attributable to the recession. As a result, we must isolate fracking's impact on natural gas prices specifically within our research design. Third, over our time period there were various EPA regulatory changes. As a result, our research design must also account for EPA regulatory changes which occur in our sample period. Note that we do not attempt to answer the more complicated question of how cheaper natural gas impacted investment decisions for new natural gas fired capacity, thus ending our sample period before 2013 when natural gas capacity began to increase. As a result, our study focuses on short to medium run impacts of fracking on U.S. air quality.

To overcome this identification problem, we propose a three step IV empirical strategy to quantify how fracking indirectly affected air quality via impacts on the electricity sector. 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

[4] We begin our dataset 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.

efficiency levels[5] of the near universe of all US power plants, we construct an instrumental variable (IV) from the dispatch order. The simulated dispatch orders indicate how many MWhs each US power plant produces each hour as a function of the relative input prices of coal, oil, and natural gas as well as the regional composition of 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 (1998) and Borenstein, Bushnell and Wolak (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 (for example, 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 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[7] to fracking, which represents a decline of roughly 50% from the 2007 price. We also perform robustness checks around this number. 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 we focus on relative fuel prices with a dispatch model, instrumented changes in generation are orthogonal to changes in EPA regulation. Hence, for our 2012 counterfactual, we let the economy

[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 pay retail prices.

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

evolve from 2007 to 2012 as it actually developed, except that we set the price of natural gas to the level as if 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 non-market benefit account 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* non-market environmental benefit attributable to fracking using standard VSL estimates. However, this estimate ignores other non-market 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 pollution. We find that shutting down all U.S. 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 US, 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, 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 NOx 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 2011) or observed policies with sharp variation that is conducive to difference-in-differences designs (Chay and Greenstone

[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), that estimate that fracking produces an upper bound damage on groundwater of $250 million per year.

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 Deschenes, Greenstone, and Shapiro (2012) 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 use a dispatch model to econometrically identify impacts on ambient pollution levels. In addition, recent papers by Linn et al. (2014), Cullen and Mansur (2015), Knittel et al. (2015), and Holladay and LaRiviere (2016) analyze the mechanisms of the fuel switching behavior and discuss implication 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 2.1 describes our data sources and 2.2 provides a first intuition of the correlation between electricity production by coal, plant emissions, and ambient air quality. Section 3 describes the construction of our dispatch model that we use to simulate our instrumental variable. Section 4 outlines our econometric framework and Section 5 presents regression results. Conclusions and further thoughts on research are offered in Section 6. The Appendix lists details of our data and several alternative specifications.

2. Data

2.1 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 which 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. As a result, the study period takes as the counterfactual what would have occurred had there been no endogenous investment decisions by electricity generators. Studying the effects of fracking while allowing endogenous investment would require a dynamic model and is outside the scope of our study. Third, in late December 2012 the EPA announced more stringent air quality standards for PM 2.5, which is a by-product of coal-fired electricity generation. Stopping the sample in 2012 avoids the confounding effects of this regulatory change.

We collect ambient air pollution data from the Environmental Protection Agency's (EPA) 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 5000 monitors across the United States. Following the example of Deschenes, Greenstone, and Shapiro (2012), we restrict our observations to those monitors that report a minimum of one time during each of at least 47 weeks out of every year from 2007 to 2012. This restriction is imposed to eliminate biased data from monitors that were either decommissioned or taken out of service during the period of interest, or were operated on a seasonal basis. Additionally, we further restrict the sample of monitors to those designated as having "population exposure" to be consistent with EPA guidelines. 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. In total, 537 $PM_{2.5}$ monitors in 363 counties and 193 SO_2 monitors in 154 counties remain in the final dataset. AQS provides hourly pollution measurements for some monitors. For these we compute daily averages. Those daily averages are subsequently averaged again to aggregate up to the monthly, quarterly, and annual levels. Figure 2 shows the locations of the $PM_{2.5}$ and SO_2 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. Triangles indicate $PM_{2.5}$ monitors, circles indicate SO_2 monitors, and crosses indicate coal-fired power plants. Figure 2 shows a significant number of air quality monitors proximate to coal-fired power plants, especially in the eastern US. We also investigate natural gas power plants proximate to air quality monitors below. Given that the eastern US has more coal generation than the western US we are able to include the vast majority of coal-fired power plants in our study.

We obtain generation data from the EPA Air Markets Program Data (AMPD) via the Continuous Emissions Monitoring System (CEMS), which contains data on the near universe of all US power plants equipped with generators with rated capacity of 10 MW or greater.[10] In addition to hourly generation and hourly SO_2 and NO_x emissions data, AMPD also contains data on primary and secondary fuel type, and exact power plant location. 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 is a reason to suspect data entry error.[11] We take fuel input price data from

[10] Holladay and LaRiviere (2016) 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.

[11] For example, if a plant reported that a large amount of electricity was generated on a given day but no SO_2 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 SO_2 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.

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. [12]

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 US. Of these, we consider only those stations that reported data for at least 25 days out of every month during the period of 2007-2012, eliminating those stations that did not consistently report. The resulting dataset 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.

2.2 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 – 2012, the compliance standard was set to an average daily concentration (over the course of one year) of 15 $\mu g/m^3$ 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 14th of 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 μg/m3. 553 monitors in 328 counties reported average annual PM2.5 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 – 2012. The counties that experienced the largest decrease, shown in black and dark grey, 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 grey.

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 US, while those counties in the Midwest tend to show a smaller decrease or slight increase. The visual correlation between reduced coal-fired electrical

[12] Chu, Holladay and LaRiviere (2016) 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 purchases prices of power plants.

generation and reduced levels of $PM_{2.5}$ is stark. In this paper we are interested in quantifying the relationship between coal generation and air quality.

Figure 5 illustrates one mechanism that may have contributed to the regional dichotomy. 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 "Small trend" represents the average of the lowest third of states, ordered by change in coal share from 2007 to 2012. This suggests that some regions of the US 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 which is itself a function of macroeconomic and regulatory activity. The price of natural gas is similarly a function of national demand, which is itself a function of macroeconomic activity. Since a sharp increase 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.

3. 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 three). 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, Bushnell and Wolak (2002). The purpose of the dispatch model is to predict the total amount of megawatt hours (MWhs) 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.

3.1 The Wholesale Electricity Market

The U.S. electricity market accounts for roughly 2.2% of GDP.[13] 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 frying of transmission lines. As a result, each National Energy Regulatory Commission (NERC) region is governed by one or more an Independent System Operators (ISO) which balance electricity supply and electricity demand at every point in time in a given region.

Second, the vast majority of cases the marginal retail end consumer of electricity does not pay the wholesale price of electricity but rather a constant average price of electricity. As a result, electricity demand is highly inelastic at every point in time (Borenstein 2002).[14] Determinants of electricity demand include weather and time of day. For example, during hot summer days, peak demand often occurs when temperatures are highest as 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.

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.[15] The ISO then uses a complicated linear programming mechanism to select the cost minimizing dispatch order given a particular day's forecasted composition of electricity demand.[16] In regulated markets, the ISO instead dictates to generators who will produce based upon inferred costs of electricity generators. The ISO mandates significant data sharing processes to help ensure a cost minimizing composition of generation.

---

[13] 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).

[14] There are two exceptions. First, large industrial users sometimes have bilateral bargains with electricity producers that include them 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).

[15] There is a large literature which shows how market power in the bidding process can lead to departures from least costs generation in determining wholesale electricity costs (Wolfram 1998 and Borenstein, Bushnell and Wolak 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.

[16] 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.

3.2 The Dispatch Model

Our dispatch model attempts to exploit the cost minimization feature of the wholesale electricity market in order to construct an instrument for total generation by each electricity generating unit directly related to input prices changes from 2007-2012. In the dispatch model we calculate the costs 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.

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 and Borenstein, Bushnell and Wolak 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 MW. 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.[17] This allows us to construct the average observed heat rate over the sample period for each generator for each hour.[18] We similarly use observed maximums of generation levels by boilers to identify capacity. Specifically, we construct heat rate for a boiler *i* over time period in which they are operating for *T* hours:

$$heat\_rate_i \ = \sum\nolimits_{t=1}^{T} \frac{mmBtu_{it}}{MW_{it}}.$$

Input prices are provided at the state- fuel type- year level in dollars per mmBtu: $\frac{\$}{mmBtu_{sfy}}$. We match this SEDS data to boilers in the CEMS database. We choose the SEDS database for two reasons. First, other EIA databases which 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 which report no input prices for an entire year. Because SEDS is aggregated to state averages, there is reporting for states in both regulated and unregulated markets. Second, we are most concerned with long run prices 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 of fossil fuel inputs and we both want to control for that variation.

[17] 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.

[18] Using observed heat rate rather than reported from EIA forms as others have done effectively weights by time spent generating for the envelope of heat rates. Davis and Hausman (2016) use this technique. We similarly use observed maximum capacity from the CEMS data rather than reported maximum capacity from EIA forms. In both cases we restrict to taking observed averages and maximums to days in which a boiler produced for more than 1.5 hours to eliminate rounding errors.

We construct the cost of generation for a generating unit by taking the product of the heat rate and the input price. We order generating units by costs within a NERC region and create a running sum of plant capacities creating a NERC level marginal cost curve for electricity generation.[19] Figure 6 displays the marginal cost curve for the southeast US (SERC region) for 2007 and 2012. We show fuel type by plotting different colors. Figure 6 shows that while in 2007 the low MC electricity was mainly supplied by coal generators, by 2012 this dispatch got replaced by natural gas generators. This change is purely attributable to the change in the relative prices of the cheaper natural gas compared to the price of coal.

Using observed NERC hourly generation allows us to construct instrumented hours of generation for each boiler in a NERC region. Figure 7 displays the observed and predicted data of our first stage IV approach below for three different NERC regions (SERC, MRO and RFC) for the first and last year of the study period. We restrict the figure to include only coal-fired generation. 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, we would observe a 45 degree line. We take this as evidence that the dispatch model is doing an adequate job of predicting generation for many generators. There are, though, a significant number of individual plants across both regions in 2012 for which the dispatch model predicts zero generation in contrast to positive observed generation. What are the causes for the lower 2012 predicted generation from the dispatch model? First, coal-fired power plants may have long term contracts with coal mines which may require them to purchase coal at pre-specified prices which could impact input prices we use versus those faced by firms. Second, geographically isolated power plants providing electricity 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 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 in constructing the instrument. 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, like previous usage of the dispatch model in the literature (Borenstein et. al. 2002 and Fabrizio et. al. 2008).

[19] We have also constructed state level MC 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.

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.

4. Econometric Framework

4.1 OLS Model

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

(1) $$Ambient_{it} = \beta_0 + \beta_1 Generation_{it} + \beta_2 Weather_{it} + \beta_3 FixedEffects_{it} + \varepsilon_{it}$$

where $Ambient_{it}$ is a measure of ambient $PM_{2.5}$ or SO_2, $Generation_{it}$ 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, $Weather_{it}$ is a vector of weather variables including average wind speed, temperature, humidity, barometric pressure, and precipitation and $FixedEffects_{it}$ is a vector of various spatial and time fixed effects as outlined below. The variables are indexed by the spatial unit *i* and are indexed over time by *t*. We explore three separate specifications with *t* either aggregated to the annual, quarterly, or monthly level.

We estimate two specifications over the spatial scale (indexed by *i*). In the first case, $Ambient_{it}$ and $Generation_{it}$ are aggregated to the county level. This simple specification has the advantage of partitioning the United States into discrete, non-overlapping areas in which every unit of analysis contains at least one coal-fired power plant and at least one ambient air pollution monitor. Figure 8 displays the county observations (in red) that contain both a $PM_{2.5}$ monitor and a coal-fired power plant. The disadvantage of the county aggregated model is that it implicitly assumes that air pollution from the generation of electricity does not cross county borders. While this assumption is clearly unrealistic, this first specification serves us as our simple baseline model to which we compare subsequent models.

Our second 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 more realistically modelling 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 single power plant is located within 70 miles of six pollution monitors, the influence of that plant will be reported in six separate observations. We perform robustness checks creating a 40 mile and 100 mile buffer as well and find that the 70 mile specification generates the most robust results. In all regressions we cluster standard errors by state, accounting for both temporal autocorrelation within each spatial unit of observation and for spatial correlation of ambient air pollution over across monitors within each state.

The final number of observations used in each regression is restricted by the combined constraints of the various datasets. For those regressions for which the unit of observation is a county, the total number of observations is restricted to those counties that contain both a pollution monitor and a power plant: 166 in the case of $PM_{2.5}$ monitors and 66 in the case of SO_2 monitors[20]. Similarly, for those regressions that rely on a 70 mile circle as the unit of observation, the total number of observations is limited to the number of pollution monitors that have at least one coal-fired power plant within a 70 mile radius. Here the total number of observations is 387 in the case of $PM_{2.5}$ monitors and 73 in case of the SO_2 monitors. Figure 9 shows the monitors (in red) that are within 70 miles of at least one coal-fired power plant (indicated by black crosses). The blue 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.

4.2 Instrumental Variables Approach

The main contribution of this section is identifying the causal effect of the relative input price change on ambient air quality. We use observed input prices for fossil fuel-fired powers plants to construct 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. Thus, predicted levels of generation are taken as the instrument for changes in coal-fired generation.[21] Specifically, for each region we estimate:

---

[20] We find that using ambient $PM_{2.5}$ as a dependent variable tends to produce more stable results than using ambient SO_2. This may be a result of a smaller sample size—more than twice as many $PM_{2.5}$ monitors meet our selection criteria as do SO_2 monitors. Details of the regressions on ambient SO_2 can be found in appendix 1.

[21] Note that our instrument becomes necessary because fracking is not the only macroeconomic shift in the US energy landscape in this time. 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 footnote 2 as well as the particular market conditions described in footnote 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 regional US electricity market to construct an instrumental variable which identifies the causal impact of the changes in natural gas prices.

$$Generation_{it} = \alpha + Z_{it}\theta + Inst.Generation_{it}\delta + \varepsilon_{it}$$

(2)

$$Ambient_{jt} = \alpha + X_{it}\varphi + \sum_{j} \widehat{Generation}_{jit}\,\beta + \varepsilon_{it}$$

The top equation represents the first stage regression. It estimates how observed generation relates to predicted generation due to relative marginal costs of fossil fuel generators by the dispatch model. The bottom equation takes generator predicted values from the first equation 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.[22]

Our instrument is valid only if fracking did not affect $PM_{2.5}$ through any channels other than electricity generation. In the US 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 $PM_{2.5}$ levels.

Figure 7 shows our first stage relationship between instrumented coal and observed coal. In theory, without any ramping, maintenance, etc., the relationship should be a 45 degree line. Figure 7 shows, however, that the generation level is over-predicted at some of the power plants and under-predicted 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 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.

4.3 Stage Three

[22] As discussed above, despite the observed positive relationship between instrumented coal-fired generation in the dispatch model and observed generation, there are reasons why the instrument could cause concern. First, we implicitly assume that firms do not exercise market power over time. Market power would not be a large problem for the instrument per se so long as the composition of market power stayed constant over time. Second, we ignore operation and maintenance costs and pollution costs incurred by firms in constructing the instrument in this iteration of the paper. It is plausible that operation and maintenance costs are constant over our study window and it thus controlled for by fixed effects. Third, pollution costs for coal-fired power plants decrease in regulated areas over our study period as total emissions are falling due to inexpensive natural gas. This biases our instrument toward predicting too little coal. Finally, coal power plants sometimes stay on even when they are not the lowest-cost provider since ramping costs are large. We ignore ramping with our dispatch model, like previous usage of the dispatch model in the literature (Borenstein et. al. 2002 and Fabrizio et. al. 2008). This neglect of ramping has implications for our instrument.

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 10 displays the Henry Hub natural gas spot and futures prices for 2005 – 2014. If a futures contract displays the same price as a spot contract then the market predicts no price change. Figure 10 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 the unanticipated price effects due to fracking. The 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 ΔP^{Medium} = \$3.41/mmBtu, relative to a counterfactual scenario of no fracking, with upper and lower bounds of ΔP^{Upper} =4.11 and ΔP^{Lower} =2.16 in 2012 USD. We calculate the counterfactual natural gas price by adding the Hausman and Kellogg (2015) estimates to 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 (column 5) links changes in generation under the Hausman and Kellogg counterfactuals to predicted changes in air quality.[23] 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 $G_{ih}(P_{2012})$ in each hour h and boiler i.
b. Run dispatch model at observed 2012 prices plus decrease in NG prices attributable to fracking and record generation as $G_{ih}(P_{2012} + \Delta P^k)$.
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 :

[23] Given a valid estimate of the effect of coal generation on $PM_{2.5}$ 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.

(3) $$\Delta PM^{j} = \widehat{\boldsymbol{\beta}_{IV}} * \Sigma_{h,i\epsilon j}\, G_{ih}(P_{2012} + \Delta P^{k}) - \widehat{\boldsymbol{\beta}_{IV}} * \Sigma_{h,i\epsilon j}\, G_{ih}(P_{2012})$$

with k ϵ {Medium, Upper, Lower} denoting a Hausman and Kellogg scenario. Accounting for differential impacts across space is important because generation is determined by the composition of installed power plant efficiency by fuel type, which varies spatially (Holladay and LaRiviere 2016).[24]

5. 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 5[25].

5.1 OLS Estimates

Table 1 reports the results of 30 different linear regressions, each of which seeks to quantify the relationship between power plant activity and ambient $PM_{2.5}$. Results are presented in two panels: Panel A reports results from the county-level unit of observation, and Panel B contains those results obtained by defining the unit of observation as the 70-mile radius circle around a pollution monitor. Each regression in Panel A is based on a sample of 166 counties, each of which contain both a $PM_{2.5}$ monitor and a coal-fired power plant. Reported electrical load and emitted pollution are the sums of the respective quantities from all coal-fired power plants within a given county, and measured ambient $PM_{2.5}$ is the simple mean of all monitors with that county. Each regression in Panel B is based on 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 power plant. Within each panel, 5 models are specified for each of three measures of plant activity: gross electrical load, emitted SO_2, and emitted NOx. In columns (1) – (2) we employ the annual time unit of observation, progressively adding weather controls and county fixed effects and year fixed effects. Columns (3) reports results of our quarterly aggregation, and columns (4) and (5) report results of monthly aggregation, while (5) also includes county as well as US Census Region by year by month fixed effects.

[24] 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 FRCC, MRO, SPP, TRE or WECC. In RFC, 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 NG generation is identical, the percentage change in coal generation in smaller than the percentage change in NG generation in these regions because the installation bases are so large.

[25] We also find that a log-linear specification produces similar results. See appendix 2 for details.

Throughout this and the following tables, our preferred model is specification 5 of Panel B. 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 off of variation in generation across monitors within region within a 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 specification (column 5 of Panel B of Table 1), suggests that an increase of daily generation by one terawatt-hour of coal-fired electrical generation corresponds to a 17.39 $\mu g/m^3$ increase in ambient $PM_{2.5}$ (which is equivalent of one gigawatt hour increasing $PM_{2.5}$ by 0.0174 $\mu g/m^3$). This result is robust to a variety of time unit specifications--estimates using annual and quarterly timeframes return values ranging from 16.11 $\mu g/m^3$ to 18.10 $\mu g/m^3$ per terawatt-hour.

Our county-level analyses of Panel A produces similar but less precise results, returning values ranging from 12.08 $\mu g/m^3$ to 36.61 $\mu g/m^3$. The larger range for county level estimates is attributable to imprecision in the aggregation of the left hand side variable: county level ambient $PM_{2.5}$ levels. While there is no theory to guide the appropriate aggregation technique, we focus on the 70 mile radius around air quality monitors as our primary specification.

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 $\mu g/m^3$. In our sample, the 2012 mean daily $PM_{2.5}$ is 9.47 $\mu g/m^3$. 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 11 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 $\mu g/m^3$, 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$\mu g/m^3$. This implies that in such a region, the shut down of all plants would lower $PM_{2.5}$ levels by 74%. Note that not all of this change would be attributable to fracking.

Rows A-2 and B-2 of Table 1 summarize the results, for the county level and monitor level respectively, that are obtained from using emitted SO_2 as the independent variable of interest. At the county level (row A-2) and the monitor level (row B-2), we find small but statistically significant coefficients in all specifications, indicating a strong relationship between emitted SO_2 and measured $PM_{2.5}$. We note that the coefficient of interest is smaller at the monitor level in all cases, which suggests that SO_2 may dissipate at a distance of less than 70 miles from the point at which it is emitted. Although significant, the absolute effect of SO_2 on $PM_{2.5}$ is very small. Shutting down all SO_2 emitted from power plants would reduce PM2.5 by only up to 0.4%.[26] This is because only some of the SO_2 transforms to sulfate, and it is sulfate which is ultimately captured by the PM2.5 monitors (Hodan and Barnard (2004).

Rows A-3 and B-3 indicate inconclusive results when emitted NOx is substituted as the regressor of interest. At the monitor level the coefficients for NOx are significant in some specifications, but are insignificant at our preferred specification (5). These mixed results may be attributable to the action of the NOx Budget Program (NBP), a seasonal cap-and-trade system which regulates NOx emissions in many of the counties in our sample. Because the NBP is active only during summer months, its effects may introduce additional noise to the model, especially when aggregated to the annual level. When we instrument NOx by the relative prices of the dispatch model, orthogonal to NBP, then we would expect our results to have the correct sign. We will explore this in the next section.

5.2 IV Estimates

Table 2 presents the $PM_{2.5}$ results using the instrument. We focus the discussion on the monitor level 70-mile buffer estimates rather than the county level estimates because they are measured with less noise.[27] We first note that the IV estimates tend to have the same sign as the OLS estimates and in some cases are larger by an order of magnitude.[28]

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] $\mu g/m^{3,}$ with results in brackets accounting for the 95% confidence interval. On average, then, our estimates suggest that coal-

[26] This calculation is based upon that the 70 mile buffer with the highest SO_2 emissions in our dataset in the U.S. emits 2513 tons of SO_2 per day.

[27] As in the OLS specifications, the county level IV estimates tell a similar story but have more noise associated with them.

[28] 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 the at the yearly level the estimates are almost identical in magnitude as the OLS estimates. As in the OLS specifications, county level linear time trends and state year time trends reduce the level and significance of the IV estimates. This is expected: much of the variation in average yearly $PM_{2.5}$ levels are absorbed by these fixed effects. For example, the state year fixed effect specification implies that we identify the effect of coal generation on different $PM_{2.5}$ readings within a state within a year attributable to different coal-fired being in the 70 mile radius of those monitors. Given that we observe input prices at the state level as well, this absorbs a significant amount of variation in the data.

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$\mu g/m^3$. This implies that in such a region, the shutdown of all plants[29] would lower $PM_{2.5}$ levels by 89% +/- [39%].

In Panel 2, the independent variable of interest is instrumented SO_2 emission at coal-fired power plants. The IV estimates are about 3 to 4 times larger compared to the corresponding OLS estimates. Still, even in the most aggressive specification (4), a complete ban of SO_2 at powerplants would only reduce ambient $PM_{2.5}$ by only up to 1.9% in the dirtiest area and by a tiny 0.13% in the average 70 mile buffer zone.

Panel 3 of Table 2 displays the results by regressing ambient $PM_{2.5}$ on the sum of the instrumented NOx emissions from coal-fired power plants within a 70 mile buffer. While the corresponding OLS results (Table 1) produced very unstable results, our IV results are significand and have the correct sign. This is important in terms of verifying our identification strategy. While the OLS results were inconclusive and likely biased because of the regional and temporal NOx cap and trade programs, in our the IV specifications NOx emission are instrumented by the predicted coal quantities from our dispatch model. Hence, in our IV specification, the relative price change between fossil fuels only (and not the NOx Budget Program) explain changes in PM. When we single out these price changes as the causal channel our regressions produce the correct sign throughout all of our IV specifications of Panel 3 in Table 2. However, the OLS specification presented in Panel 3 of Table 1 yielded inconclusive results; one possible explanation is the confounding effect of the NOx Budget Program.

Throughout this paper we assume that coal-fired generation produces more air pollution than natural gas-fired generation. To test to this assumption[30], the above set of regressions are repeated by including the production of electricity from natural gas-fired power plants in addition to the coal-fired generation as two separate regressors.[31] Table 4 displays the results for several of the above OLS and IV specifications. 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 and 2.

[29] 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.

[30] According 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.

[31] 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).

Finally, note again, that not all of the air quality changes in the above IV regressions are attributable to fracking, but are attributable to the changes in the relative prices. To isolate the effect of fracking, we proceed in stage three.

5.3. Third Stage

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 U.S. 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 of the United States. Figure 12 shows the results of the counterfactual analysis for generation of both coal (panel a) and natural gas (panel b) power plants from the median case of Hausman and Kellogg (2015) with price differential ΔP^{Medium}. The figure shows the changes in generation which result from changes in the NERC level dispatch order as a result of NG prices which were uniformly \$3.41/mmBtu higher across the U.S. The implication is that these changes would not have occurred if fracking had been banned. Each black singular dot refers to the location of a power plant in our sample. Figure 12 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 timeframe of the analysis. While region level patterns across NG generation increases and coal generation decreases are commensurate, there is clear within region variation. This is clear in the northeast where coal decreased across the entire region but increases in NG 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 12 displays information regarding generation level impacts, we can perform the same task at the air quality monitors. In doing so, we can match levels to percentage changes relative to observed 2012 ambient air pollution levels. In Panel (a), the shaded areas around each power plant in Figure 12 hence also display the corresponding additional $PM_{2.5}$ emitted in the counterfactual 2012 scenario. The shaded areas correspond to a percentage change in air pollution ranging from a 0% (white) to 35% (black).

The darkest spot is an Alabama which today would observe an additional 35% in PM2.5 if fracking had been banned.[32]

Figure 13 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 NG. This is an outlier, though, because it is nearby not only several coal-fired power plants, but inefficient coal-fired power plants which according to the dispatch model would be generating much more intensely with high NG prices.

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 no qualitative differences in generation. In Appendix Table A3 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 to model endogenous capacity decisions which would be complicated by increase renewable generation in addition to lower natural gas prices from 2012 going forward. We conclude that due to the fixed stock of generating capacity the change in dispatch does not scale linearly with changes in NG prices. This is important for extrapolating these findings to larger wedges in coal and NG prices which could be caused, for example, by carbon taxes in which a full investment model like Gowrisankaran et. al. (2016) or Cullen and Reynolds (2016) is needed. Our findings highlight the need for those approaches.

5.4 Monetarization of Indirect Benefits

To put the change in air pollution levels into context, using the monetary externality cost measurers derived by Muller, Mendelsohn and Nordhaus (2011), we find that fracking provided health benefits by $17 billion per year, with a lower bound externality benefit of $5.4 billion and an upper bound of $43 billion per year, due to the displacement of coal. The “Medium”, “Lower” and “Upper” benefits are calculated under the assumptions as described in the note of Table 5. 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

[32] The percentage is calculated by $\Delta PM^{j} / [\widehat{\boldsymbol{\beta}_{IV}} * \Sigma_{h,i\epsilon j} G_{ih}(P_{2012} + \Delta P^{k}) + PM_{2.5}$ *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, $\widehat{\boldsymbol{\beta}_{IV}} * \Sigma_{h,i\epsilon j} G_{ih}(P_{2012})$.

natural gas price supply function. In comparison, the magnitude of the changes across the rows is considerably, due to the different assumptions in calculating the health costs, as outlined in the note below Table 5 and detailed in Muller, Mendelsohn, and Nordhaus (2011).

As an alternative, singling out only $PM_{2.5}$ and analyzing lung cancer cases only, the medical literature reports that a decrease in $PM_{2.5}$ of .25-.64 μg/m^3 would correspond to a decrease in lung cancer rates of 0.9 to 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 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 129-354 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 U.S. dollars.[34]

6. Conclusion

Our main contribution is to offer a new method 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. In this way, the longer run impacts are likely to be larger than what we find here making this a lower bound.

Assuming this coal displacement was uniformly distributed, our IV estimates imply a 95% confidence interval for decreases in $PM_{2.5}$ due to fracking of 2-6%. We show, however, that these benefits vary spatially, with the largest gains in Eastern United States and most pronounced 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]

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

[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 powerplants 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).

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 form the US. As a result, US natural gas prices capitalize US ‘fracking’ more directly. Secondly, 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, Metaxoglou and Trindade (2015) as well as Linn, Muehlenbachs, and Wang (2014) are two recent promising working papers that provide methods to analyze CO_2 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 which 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 disjointed energy policy which has characterized the US in recent decades. Furthermore, states like New York, which have banned fracking, are able to enjoy lower natural gas prices without suffering negative non-market impacts or positive market impacts through leasing revenue. We are not aware of studies which answer regulatory federalism questions about the efficiency of this type of policy in nationwide input markets.

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.

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.

Borenstein, S. (2002). The Trouble with Electricity Markets Understanding California's Restructuring Disaster. *Journal of Economic Perspectives,* 88(4): 669-685.

Borenstein, S., Bushnell, J. and Wolak, F. (2002). Measuring Market Inefficiencies in Calfornia's Wholesale Electricity Market. *American Economic Review*, 92(5):1376-1405.

Brandt, A. R., G. A. Heath, E. A. Kort, F. O'Sullivan, and others. 2014. "Methane Leaks from North American Natural Gas Systems." Science 343, no. 6172: 733–35.

Chay, K. Y., and Greenstone, M. (2005). Does Air Quality Matter? Evidence from the Housing Market. *Journal of Political Economy*, *113*(2).

Crawford, Jagoda, et al. "Using Radon-222 as an Indicator of Atmospheric Mixing Depth in ME-2 for PM2. 5 Source Apportionment. *Aerosol and Air Quality Research* 15 (2015): 611-624.

Cullen, J and Reynolds, S (2016). The Long Run Impact of Environmental Policies on Wholesale Electricity Markets: A Dynamic Competitive Analysis. *University of Arizona Working Paper.*

Currie, J., and Neidell, M. (2005). Air pollution and infant health: what can we learn from California's recent experience? *The Quarterly Journal of Economics*,120(3), 1003-1030.

Davis, L. and Hausman, C. (2016). Market Impacts of a Nuclear Power Plant Closure. *American Economics Journal: Applied Economics, 8(2): 92-122.*

Deschenes, O., Greenstone, M., and Shapiro, J. (2013) Defensive Investments and the Demand for Air Quality: Evidence from the NOx Budget Program and Ozone Reductions. *Yale University working paper.*

EIA (1998): Natural Gas 1998: Issues and Trends, Office of Oil and Gas of the Energy Information Administration (EIA).

EIA (2016): Today In Energy. Monthly Reports, available online at https://www.eia.gov/todayinenergy/ (Last accessed April 24, 2016).https://www.eia.gov/todayinenergy/ (Last 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, K., Rose, N., Wolfram, C. (2008) Do Markets Reduce Costs? Assessing the Impact of Regulatory Restructuring on U.S. Electric Generation Efficiency. *American Economic Review*, 97(4): 1250-1277.

Graff Zivin, Joshua and Matthew Neidell. "Environment, health, and human capital." *Journal of Economic Literature* 51.3 (2013): 689-730.

Gowrissankaran, G., Reynolds, S., and Samano, M (2016). Intermittency and the Value of Renewable Energy. *Journal of Political Economy.* 124(4): 1187-1234.

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: 203–09.

Hausman, C., & Kellogg, R. (2015). Welfare and Distributional Implications of Shale Gas. *Brookings Papers on Economic Activity*.

Hodan, W. M., and Barnard, W. R. (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 LaRiviere, J. (2016). The Impact of Cheap Natural Gas on Marginal Emissions and Implications for Energy Policy. *University of Tennessee working paper.*

Isen, A., Rossin-Slater, M., and Walker, W. R. (2014). *Every Breath You Take–Every Dollar You'll Make: The Long-Term Consequences of the Clean Air Act of 1970* (No. w19858). National Bureau of Economic Research

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. MIT Working Paper, October 1, 2015.

Koster, Hans R.A. and Jos van Ommeren (2015): A shaky business: Natural gas extraction, earthquakes and house prices. *European Economic Review*. Vol 80(November), pp. 120-139.

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, Washington: Resources for the Future.

Mansur, E. T. (2007). Do oligopolists pollute less? Evidence from a restructured electricity market. *The Journal of Industrial Economics*, *55*(4), 661-689.

Muehlenbachs, L., Spiller, E., and Timmins, C. (2013). Shale Gas Development and the Costs of Groundwater Contamination Risk. *Resources for the Future Discussion Paper*, (12-40).

Muehlenbachs, L., E. Spiller, and C. Timmins, 2015, "The Housing Market Impacts of Shale Gas Development," *American Economic Review*, 105(12): 3633-59

Muehlenbachs, Lucija Anna and Alan J. Krupnick (2014): Infographic: Shale Gas Development Linked to Traffic Accidents in Pennsylvania, Resources for the Future, (185).

Muller, N. Z., Mendelsohn, R., and Nordhaus, W. (2011). Environmental accounting for pollution in the United States economy. The American Economic Review, 101(5), 1649-1675.

Olmstead, Sheila M., Lucija A. Muehlenbachs, Jhih-Shyang Shih, Ziyan Chu, and Alan J. Krupnick (2013): Shale gas development impacts on surface water quality in Pennsylvania. *Proceedings of the National Academy of Sciences*, March 2013, Vol. 110, No. 13, pp. 4962-4967

Quistorff, Brian (2015): The Effects of Energy Prices on Manufacturing Employment, University of Maryland Job Market Paper.

Pirovano, G., et al. "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, The Economic Journal, Vol. 75(300), pp. 683-735

Raaschou-Nielson, Ole, 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). *The Lancet Oncology*, 14(9): 813-822.

Schlenker, W., and Walker, W. R. (2016). Airports, air pollution, and contemporaneous health Review of Economic Studies, 83(2): 768-809 .

Wolfram, C. (1999). Measuring Duopoly Power in the British Electricity Spot Market. *American Economic Review*, 89(September): 805-826.

Yu, Lingda, et al. (2013): Characterization and source apportionment of PM2. 5 in an urban environment in Beijing. *Aerosol and air quality research* 13.2 574-583.

Figure 1: Histogram of $PM_{2.5}$ Average Daily Sample Values

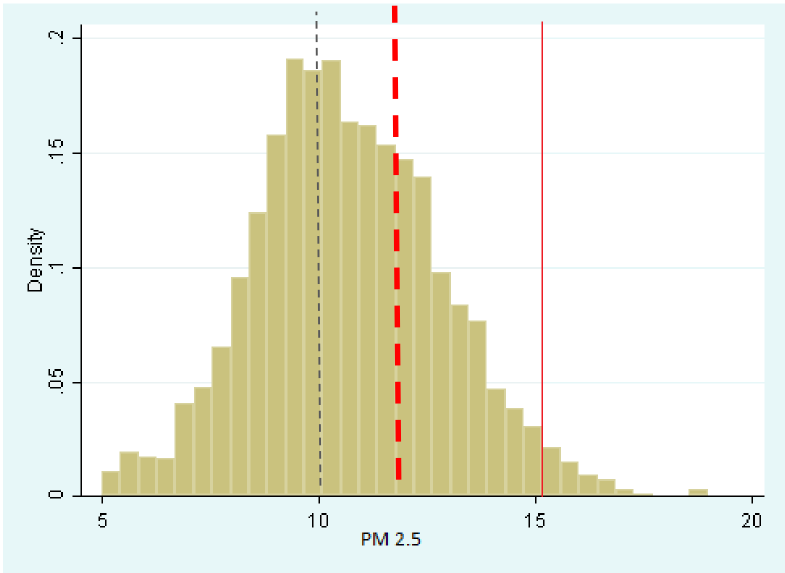

Note: The mean daily concentration of $PM_{2.5}$ at the monitors in our sample is 10.69 $\mu g/m^3$, with a standard deviation of 2.17 $\mu g/m^3$. The light dotted line displays the WHO $PM_{2.5}$ standard of 10 $\mu g/m^3$. Since December 14th of 2012, the EPA has set the standard to 12 $\mu g/m^3$ (bold striped line). From 1997-2012, the federal standard for compliance with the Clean Air Act was 15 $\mu g/m^3$.

Figure 2: Locations of Coal-Fired Power Plants and Air Pollution Monitors

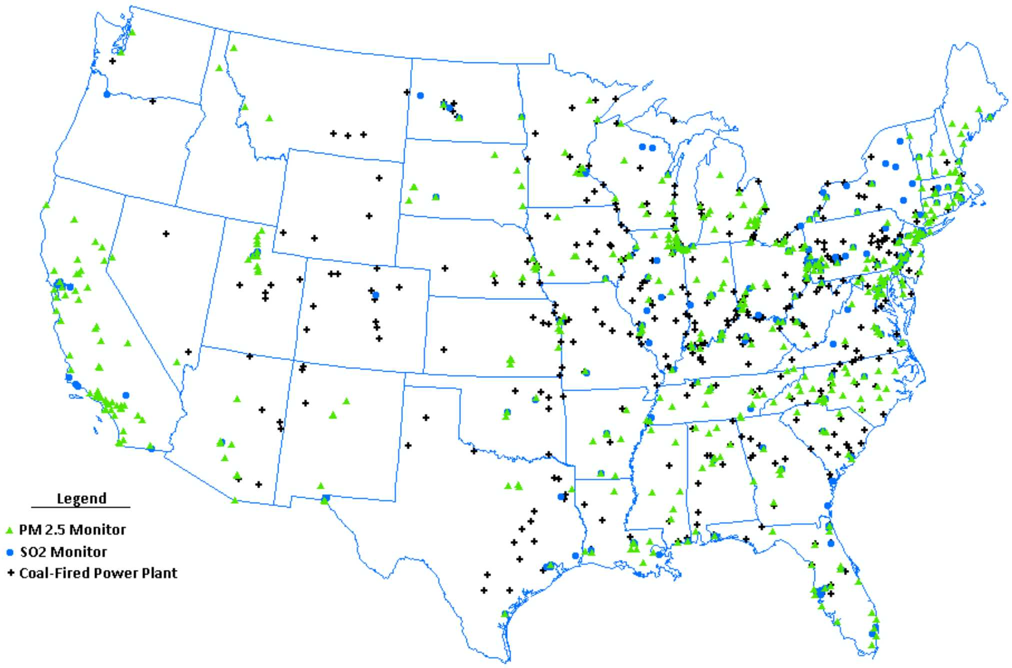

Note: The map displays all power plants that have at least one boiler for which coal is the primary fuel and all air pollution monitors that detect either $PM_{2.5}$ or SO_2 and meet the reporting criteria as explained in the text.

Figure 3: Change in Ambient $PM_{2.5}$ Levels 2007-2012

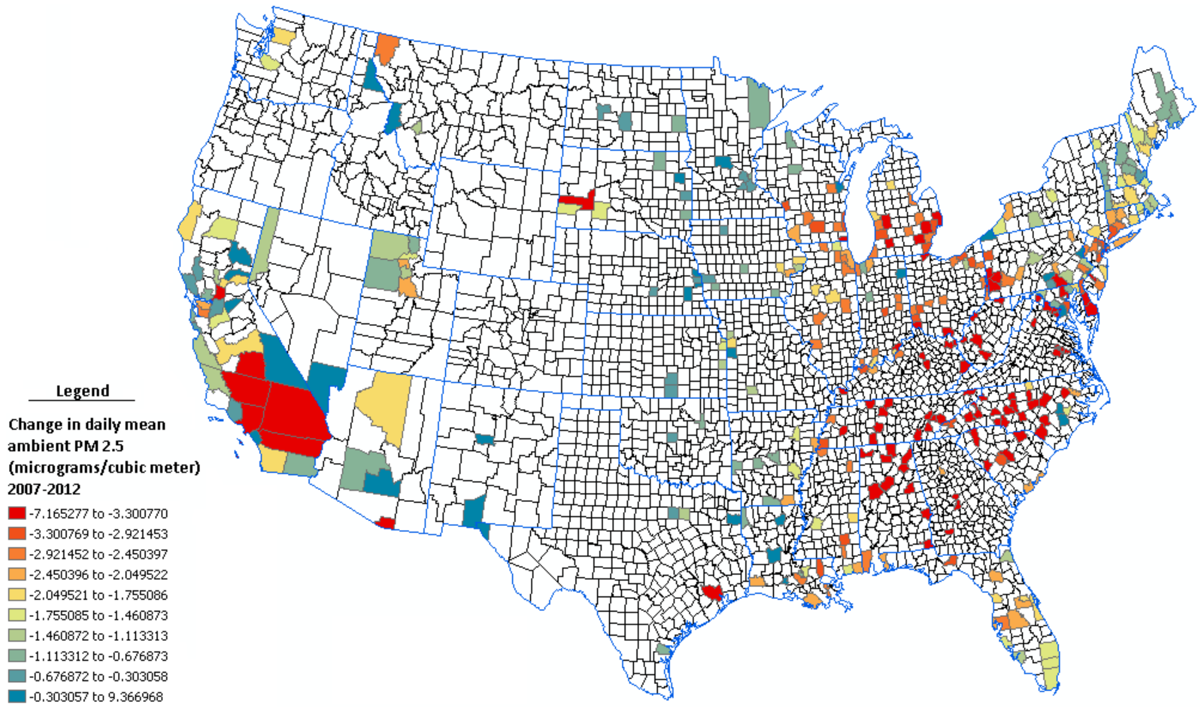

Note: There is a regional dichotomy in the change in ambient $PM_{2.5}$ levels. Most counties in the Midwest and Mid-Atlantic regions show a decrease in average daily pollution level, but the decrease is more pronounced in Appalachia and the near the Great Lakes than in the Midwest or New England.

Figure 4: Change in the county average daily electrical generation (MWh) by coal-fired power plants, from 2007-2012

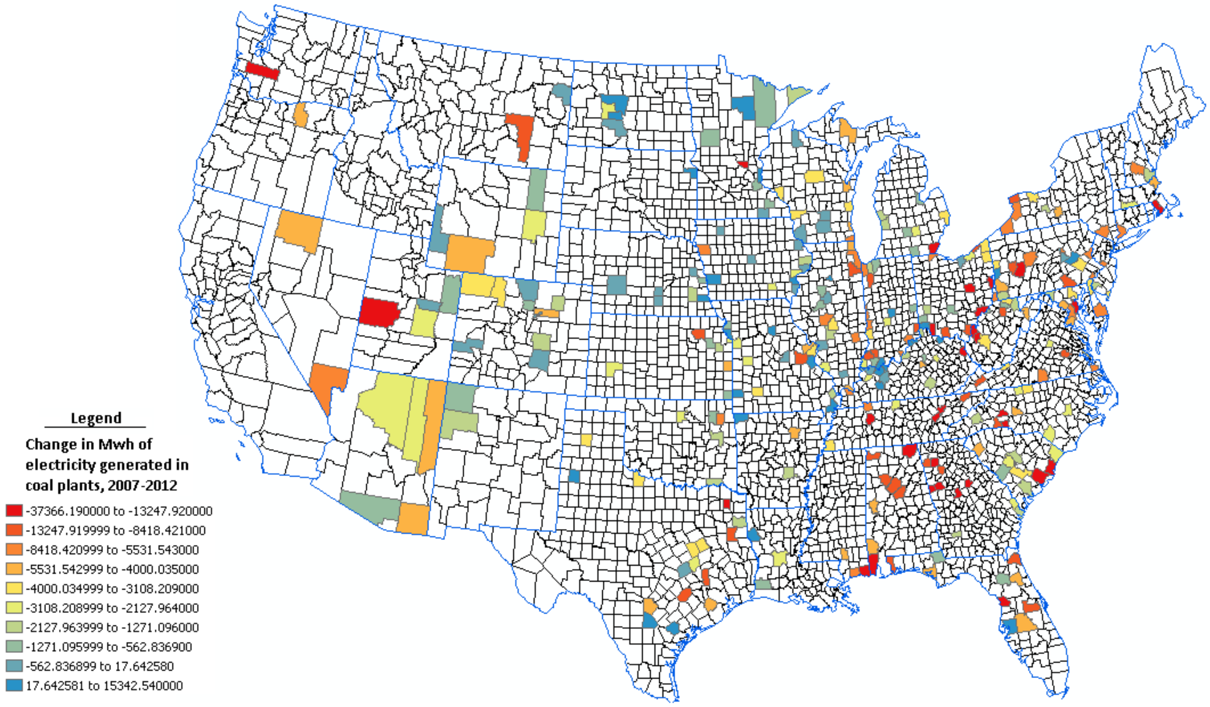

Note: Counties in the Appalachian and in the Great Lakes regions tended to decrease their coal-fired electrical generation by the largest amounts. Counties in the Midwest generally show smaller decreases or slight increases.

Figure 5: Percentage of Electricity that is Generated by Coal and Natural Gas Price

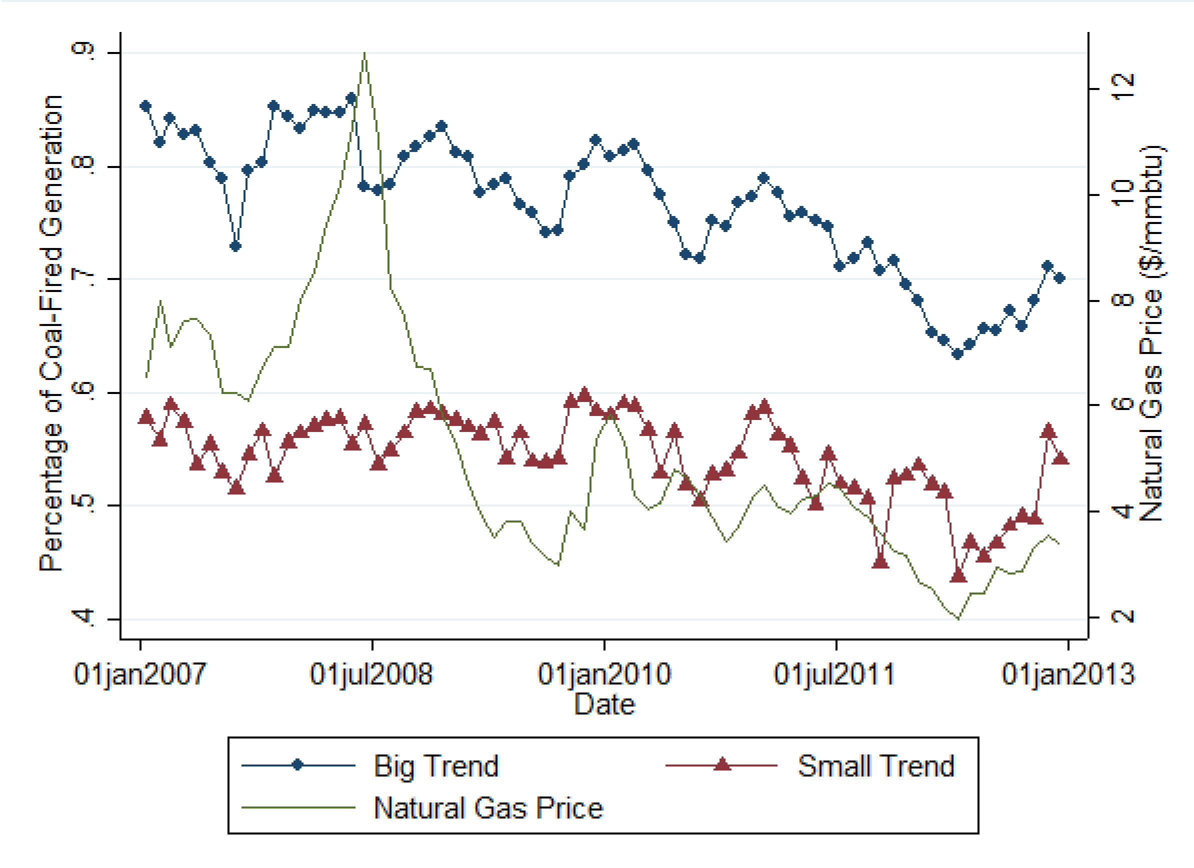

Note: During the period of 2007 – 2009, there does not appear to be a strong correlation between the real price of natural gas and the percentage of electricity that is generated by coal-fired power plants. In 2009 the price of natural gas drops below $6/mmBtu, and the correlation appears to strengthen 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.

Figure 6: Marginal Cost Curves of Electricity Generation

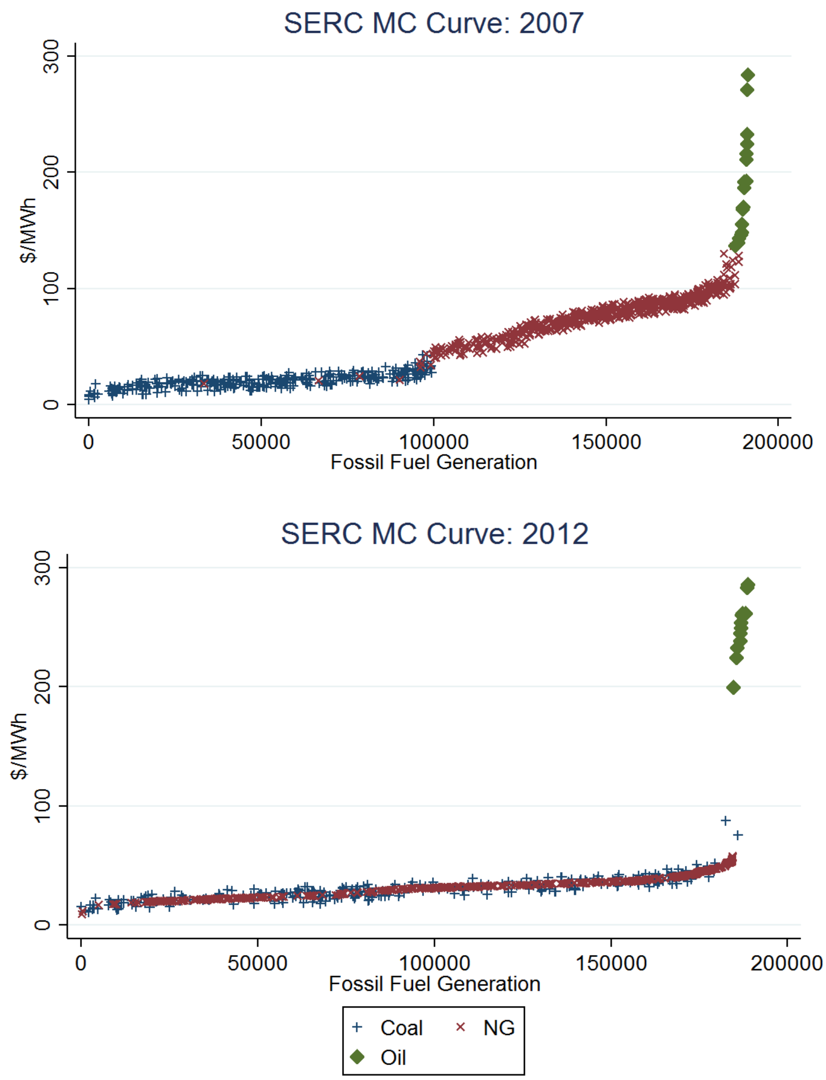

Note: 2007 and 2012 constructed SERC supply curves. Data are taken from CEMS and SEDS datasets. Capacity measure constructed from max average hourly generation observed within a day conditional on boiler generating. Red dots represent natural gas-fired capacity; the 2012 curve displays significantly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

Figure 7: Observed and Instrumented Megawatt hours 2007 and 2012

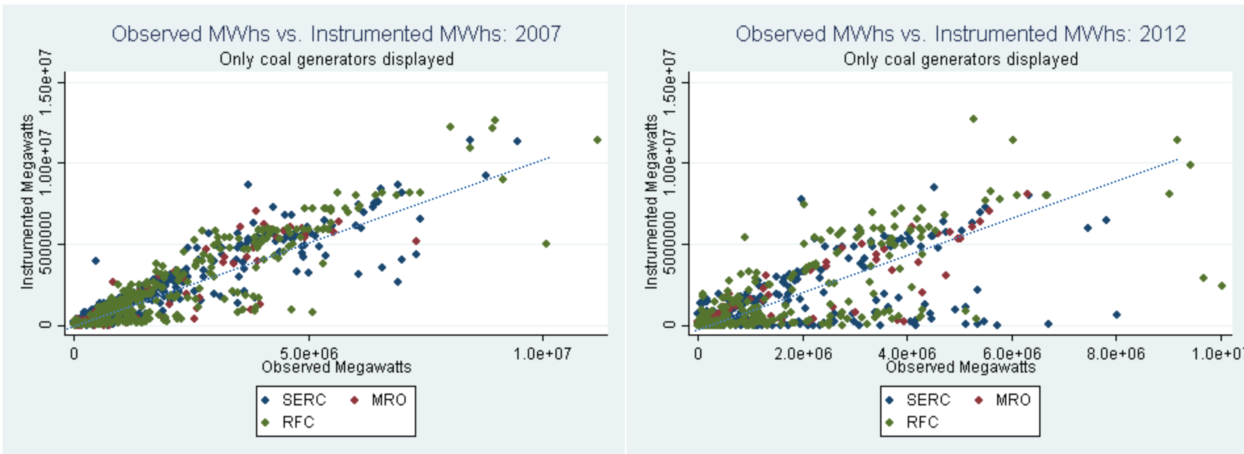

Note: 2007 and 2012 instrumented hours from dispatch model (y-axis) against observed generation in CEMS data (x-axis). Capacity measure constructed from max average hourly generation observed within a day conditional on boiler generating. Color indicates NERC region. Only coal generation displayed. 45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

Figure 8: Counties that Contain Both a $PM_{2.5}$ Monitor and a Coal-Fired Power Plant

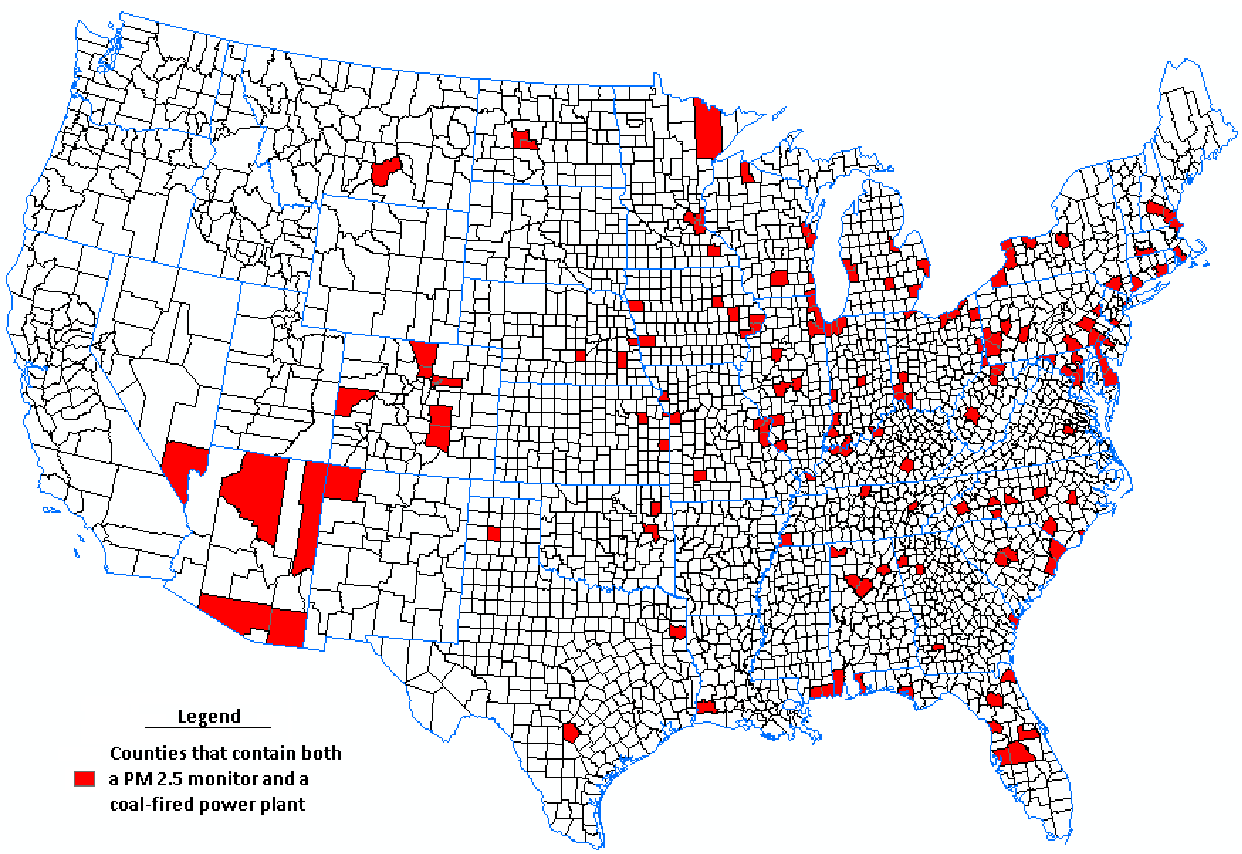

Note: For regressions using county as the unit of observation, only those counties that contain both a $PM_{2.5}$ monitor and a coal-fired power plant are included.

Figure 9: 70 Mile Radius Circles Centered on $PM_{2.5}$ Monitors

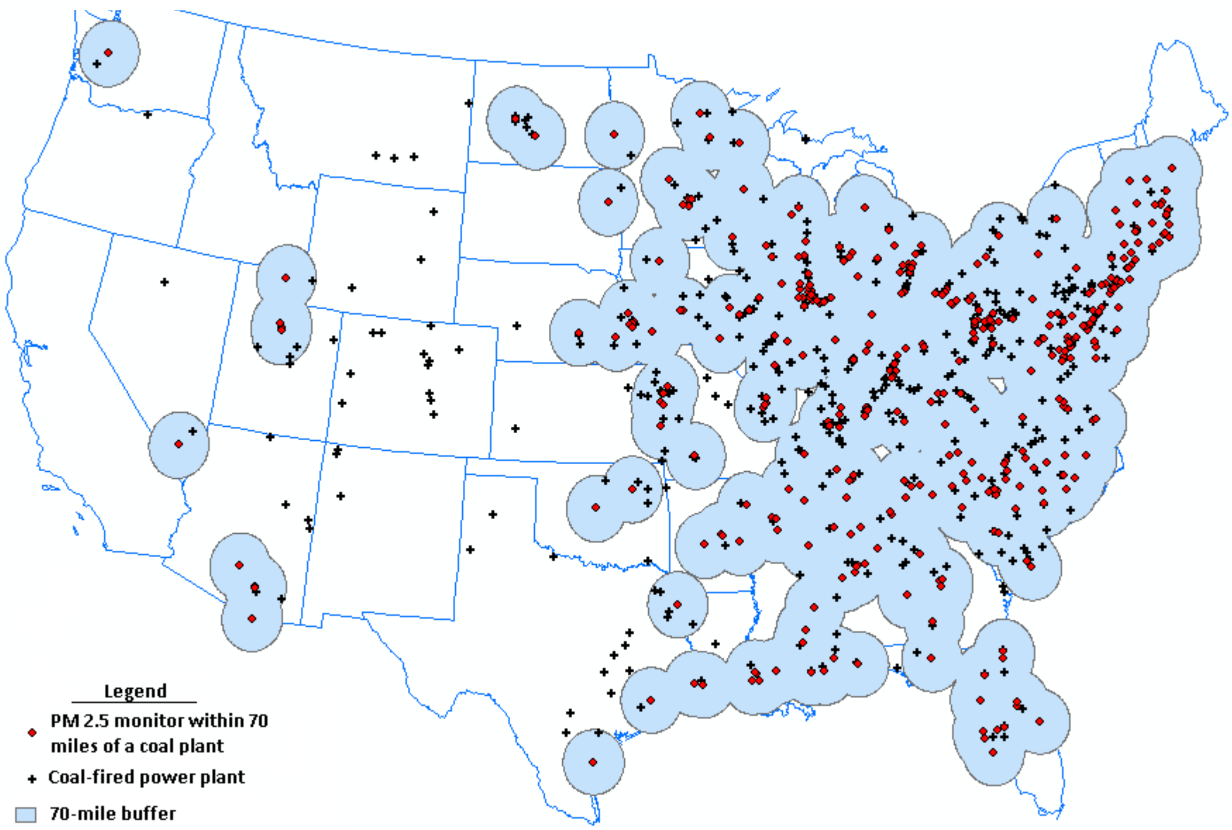

Note: In regressions using the 70 mile radius circle as the unit of observation, a circle is drawn around each $PM_{2.5}$ Monitor. Within each circle, the total amount of electricity generated and pollution emitted is summed. Each circle is then treated as a single observation.

Figure 10: Natural gas spot and futures prices

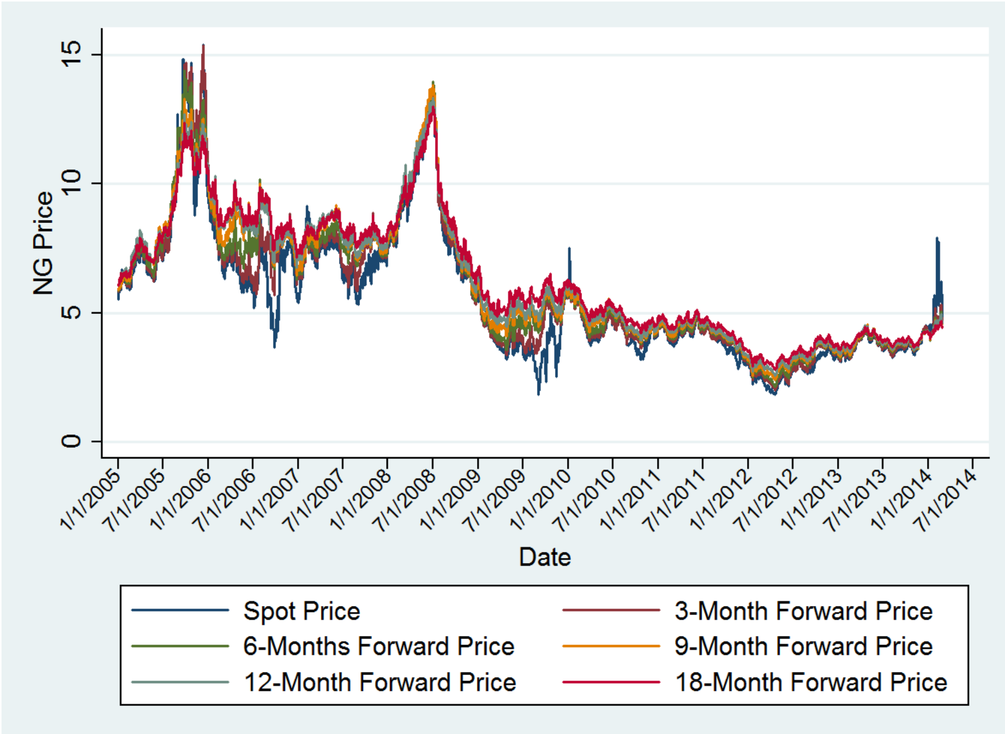

Note: Hunry Hub natural gas prices over time. Figure 11 displays prices for contracts on the date they were written. Therefore, if a futures contract is the same prices as a spot contract then the market predicts no price change. In early 2009 the market expected natural gas prices to increase rather than decrease.

Figure 11: Histogram of Total Generation (GWh) in 70-mile Radius Observations

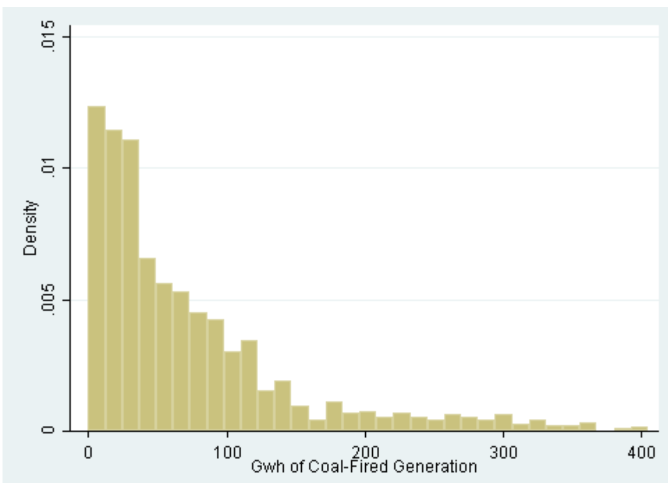

Note: Under the 70-mile radius unit of observation, we sum the electrical generation of all plants within that radius of each pollution monitor. The mean total generation is 73.7 GWh, the maximum is 405.

Figure 12: Regionally differentiated incidence of fracking on ambient air quality

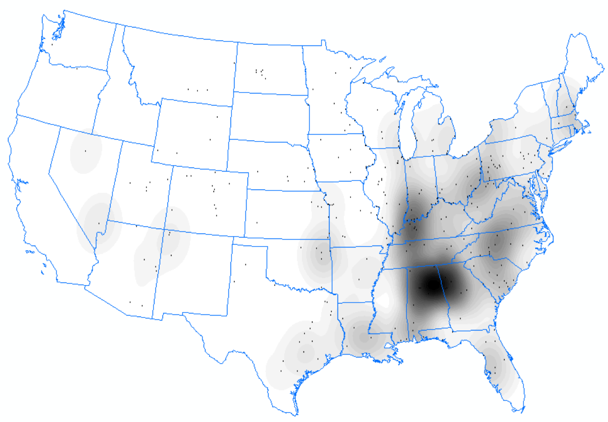

Panel (a): Counterfactual increase in coal generation *Panel (b): Increase in NG Generation*

Note: 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.

Figure 13: Histogram of level and percentage change of PM 2.5 at the monitor level.

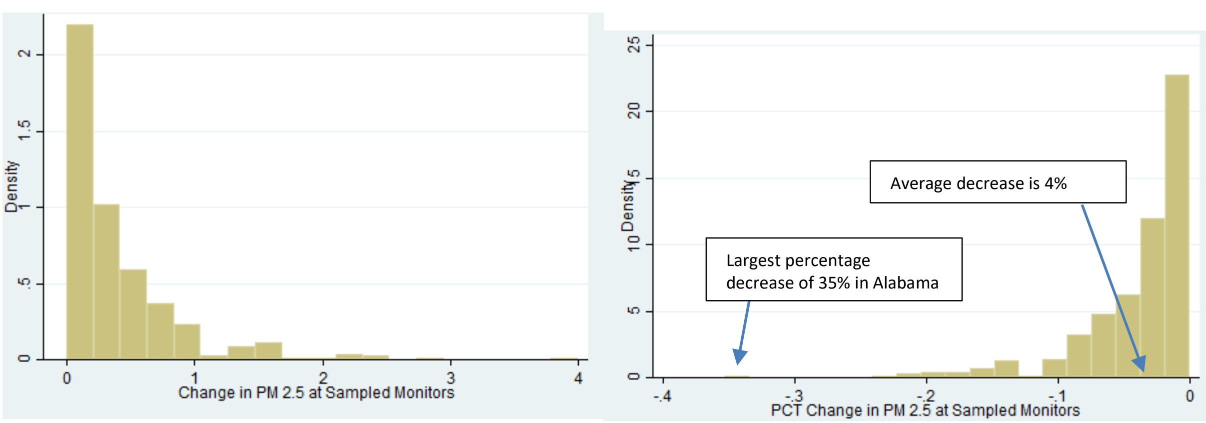

Panel (a): Change in Levels *Panel (b): Percentage change*

Note: 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.

Table 1: Regression results for $PM_{2.5}$ on various specifications of power plant activity.

Panel A: County Level Regressions

	(1)	(2)	(3)	(4)	(5)
Measures of Plant Activity					
A-1: Avg Daily Gross Load (TWh)	36.61**	25.98	23.06**	23.97***	12.08*
	(16.84)	(16.93)	(10.29)	(7.634)	(6.940)
R-squared	0.917	0.923	0.759	0.637	0.719
Adjusted R-squared	0.897	0.904	0.744	0.629	0.705
A-2: Avg Daily SO_2 (Million Tons)	1.795***	1.599***	1.789***	1.884***	1.222***
	(0.497)	(0.356)	(0.417)	(0.383)	(0.324)
R-squared	0.913	0.920	0.752	0.634	0.719
Adjusted R-squared	0.894	0.901	0.738	0.626	0.707
A-3: Avg Daily NOx (Million Tons)	0.754	0.208	-6.084***	-5.130***	-5.448***
	(3.328)	(3.186)	(1.612)	(1.419)	(1.462)
R-squared	0.915	0.921	0.752	0.634	0.719
Adjusted R-squared	0.895	0.902	0.738	0.626	0.706
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
Number of Counties	166	166	166	166	166
Observations	854	854	3,355	10,010	10,010

Standard errors clustered by state in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Panel B: Monitor Level (70 Mile Radius)

	(1)	(2)	(3)	(4)	(5)
Measure of Plant Activity					
B-1: Avg Daily Gross Load (Twh)	17.79***	16.27***	16.11***	18.10***	17.39***
	(3.560)	(3.352)	(3.068)	(2.888)	(2.489)
R-squared	0.853	0.862	0.693	0.591	0.696
Adjusted R-squared	0.832	0.842	0.682	0.586	0.689
B-2: Avg Daily SO2 (Million Tons)	0.481***	0.420***	0.498***	0.584***	0.549***
	(0.0798)	(0.0762)	(0.0761)	(0.0735)	(0.0821)
R-squared	0.848	0.858	0.687	0.585	0.690
Adjusted R-squared	0.826	0.837	0.676	0.579	0.684
B-3: Avg Daily NOx (Million Tons)	1.704***	1.636***	-0.78***	-0.514**	-0.376
	(0.392)	(0.405)	(0.256)	(0.239)	(0.401)
R-squared	0.851	0.860	0.691	0.588	0.693
Adjusted R-squared	0.829	0.839	0.680	0.583	0.687
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
Number of monitors	387	387	387	387	387
Observations	2,316	2,316	9,255	27,763	27,763

Standard errors clustered by state in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Note: Table 1 reports the results of 30 different linear regressions, each of which seeks to quantify the relationship between power plant activity and ambient $PM_{2.5}$. Within each panel, 15 regression models are specified, five for each of three measures of plant activity: gross electrical load, emitted SO_2, and emitted NOx. Results are presented in two panels: Panel A reports results from the county-level unit of observation, and Panel B contains those results obtained by defining the unit of observation as the 70-mile radius circle around a pollution monitor. Each regression in Panel A is based on a sample of 166 counties, each of which contain both a $PM_{2.5}$ monitor and a coal-fired power plant. Reported electrical load and emitted pollution are the sums of the respective quantities from all coal-fired power plants within a given county, and measured ambient $PM_{2.5}$ is the simple mean of all monitors with that county. Each regression in Panel B is based on 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 columns (1) – (2) we employ the annual time unit of observation, progressively adding weather controls and county fixed effects and year fixed effects. Columns (3) reports results of our quarterly aggregation, and columns (4) and (5) report results of monthly aggregation, while (5) also includes county as well as US Census Region by year by month fixed effects.

Table 2: IV regression of average $PM_{2.5}$ on instrumented average daily coal generation

	(1)	(2)	(3)	(4)	(5)
VARIABLES	70 Mile	70 Mile	70 Mile	70 Mile	70 Mile
1: Avg Daily Gross Load (TWhs)	17.61***	15.06***	27.64***	29.43***	20.76***
	(6.140)	(5.001)	(7.112)	(6.499)	(4.614)
R-squared	0.853	0.862	0.694	0.591	0.696
Adjusted R-squared	0.832	0.842	0.683	0.586	0.689
2: Avg Daily SO2 (Million Tons)	1.396***	1.157***	2.303***	2.715***	1.848***
	(0.362)	(0.305)	(0.487)	(0.494)	(0.471)
R-squared	0.832	0.849	0.642	0.542	0.676
Adjusted R-squared	0.808	0.826	0.629	0.536	0.670
3: Avg Daily NOx (Million Tons)	3.969**	3.369***	6.422***	7.805***	5.735***
	(1.542)	(1.113)	(1.970)	(2.171)	(1.715)
R-squared	0.843	0.856	0.639	0.536	0.664
Adjusted R-squared	0.821	0.834	0.626	0.530	0.657
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
Observations	2,316	2,316	9,255	27,763	27,763

Standard errors clustered by state in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Note: Table 2 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. 15 regression models are specified, five for each of three measures of plant activity: gross electrical load, emitted SO_2, and emitted NOx. In columns (1) – (2) we employ the annual time unit of observation, progressively adding weather controls and county fixed effects and year fixed effects. Columns (3) reports results of our quarterly aggregation, and columns (4) and (5) report results of monthly aggregation, while (5) also includes county as well as US Census Region by year by month fixed effects.

Table 3: IV regression of average $PM_{2.5}$ on instrumented average daily coal generation

	(1)	(2)	(3)	(4)	(5)
Panel A	40 Mile	40 Mile	40 Mile	40 Mile	40 Mile
Ave Daily Coal (TWhs)	8.821	6.935	14.82	18.26*	12.20
	(11.32)	(9.938)	(11.20)	(10.40)	(9.253)
Constant	12.51***	-35.72	-10.08	96.80***	79.16**
	(0.270)	(38.14)	(40.74)	(34.98)	(30.89)
Observations	1,980	1,980	7,899	23,695	23,695
R-squared	0.842	0.852	0.697	0.602	0.694
Adjusted R-squared	0.820	0.830	0.687	0.596	0.687
Panel B	100 Mile	100 Mile	100 Mile	100 Mile	100 Mile
Ave Daily Coal (TWhs)	18.50	16.12	26.63***	26.64***	19.70***
	(11.72)	(10.92)	(9.830)	(7.272)	(6.397)
Constant	10.22***	39.63	61.97	139.4***	121.7***
	(1.530)	(48.47)	(46.10)	(32.94)	(27.60)
Observations	2,514	2,514	10,056	23,695	23,695
R-squared	0.848	0.860	0.683	0.602	0.694
Adjusted R-squared	0.826	0.840	0.672	0.596	0.687
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

Standard errors clustered by state in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Note: Table 3 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 the three of the five specifications and similar in magnitude to the 70-mile specifications shown in table 2.

Table 4: IV Robustness regression results for PM2.5 on coal and natural gas power plant generation

	--------OLS--------					IV
	(1)	(2)	(3)	(4)	(5)	(6)
	county	county	county	county	county	70 miles
Avg Daily Load Coal (Twh)	57.85***	44.11**	50.77***	41.79***	25.86**	20.54***
	(18.32)	(19.20)	(14.49)	(9.941)	(10.13)	(5.395)
Avg Daily Load of Natural Gas (Twh)	27.97	19.04	34.39	39.62	14.46	-101.3
	(20.14)	(16.87)	(26.47)	(23.98)	(15.19)	(110.1)
Constant	11.30***	-175.7*	20.41	84.58**	31.91	79.57***
	(0.369)	(100.7)	(45.24)	(36.90)	(40.44)	(25.16)
Observations	1,608	1,608	6,345	18,942	18,942	27,772
R-squared	0.890	0.902	0.646	0.516	0.638	0.693
Adjusted R-squared	0.865	0.879	0.628	0.506	0.626	0.687

Standard errors clustered by state in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Note: Table 4 displays the IV regression results when in addition to coal generation, natural gas generation is added as an independent right hand side variable of interest. The results show that coal generation has a qualitatively similar effect on $PM_{2.5}$ as in our main regression Table 2, in particular for our preferred IV regression specification (6). In contrast, the generation of electricity by natural gas has not statistical significant impact on air quality.

Table 5: Gross externality benefits measured in USD from improved ambient air quality due do the displacement of coal from the fracking revolution

		Scenario of Natural Gas Price Reduction due to Fracking		
		Lower	Medium	Upper
Health Cost Scenarios	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" scenario 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 un-adjusted (uniform) VSL of 7.7 million, respectively, as described in Muller, Mendelsohn, Nordhaus (2011). All prices are in 2012 USD.

Table 6: Year by year IV regression of average $PM_{2.5}$ on instrumented average daily coal generation

YEAR	(2007)	(2008)	(2009)	(2010)	(2011)	(2012)
VARIABLES	70 Mile	70 Mile	70 Mile	70 Mile	70 Mile	70 Mile
Avg Daily Gross Load (TWhs)	32.10***	14.16**	11.88	24.85***	28.23***	31.78***
	(10.87)	(6.471)	(7.667)	(6.884)	(6.604)	(4.834)
R-squared	0.722	0.722	0.709	0.709	0.686	0.672
Adjusted R-squared	0.700	0.700	0.686	0.686	0.662	0.646
Time unit of observation	Monthly	Monthly	Monthly	Monthly	Monthly	Monthly
Weather controls	yes	yes	yes	yes	yes	yes
Fixed Effects:						
County	yes	yes	yes	yes	yes	yes
Region*Month	yes	yes	yes	yes	yes	yes
Observations	4,632	4,632	4,632	4,631	4,631	4,605

Standard errors clustered by state in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Note: Table 5 reports the results of single year IV regressions of average $PM_{2.5}$ on average daily load for the years 2007 through 2012. The regression specification is identical across all 6 cases.

[Web Supplemental Online Material]

Appendix I: Ambient SO_2 as the Dependent Variable

Our model for estimating the effect of plant activity on ambient SO_2 is identical to that of plant activity on ambient $PM_{2.5}$, which we describe in section 5. The two stage least squares model takes the following form:

$$Generation_{it} = \alpha + Z_{it}\theta + Inst.Generation_{it}\delta + \varepsilon_{it}$$

$$SO_{2jt} = \alpha + X_{it}\varphi + \sum_j \widehat{Generation}_{jit}\,\beta + \varepsilon_{it}$$

Fewer SO_2 monitors meet the selection criteria of section 3.1 than do $PM_{2.5}$ monitors, so our sample is necessarily smaller. The county-level regressions (Panel A) consist of 66 counties that contain at least one monitor and at least one coal-fired power plant, while the monitor-level regressions consist of the 73 monitors for which there is at least one coal-fired power plant located within 70 miles. Figures A1 and A2 depict our samples for the county-level and 70-mile radius specifications respectively.

Table A2 reports the results the results for 32 regressions in which ambient SO_2 is taken as the dependent variable. We find that average daily gross electrical load is not a statistically significant determinant of ambient SO_2 in any of our model specifications. When emitted SO_2 is taken as the independent variable and ambient SO_2 is taken as the dependent variable (shown in Table 1 A-2 and B-2), we find a statistically significant relationship at the county-level unit of observation only. One possible explanation for this finding is that SO_2 dissipates at some distance less than 70 miles from the point at which it is emitted. The counties that are included in the county-level sample are substantially smaller than a 70-mile radius circle, and thus might more accurately capture the effect of emitted SO_2 on ambient SO_2.

Appendix 2: Log-linear IV specification and results

In section 5 we define a linear-linear model of ambient air pollution on power plant activity. In this appendix, we examine a log-linear IV specification, in which the log of ambient air pollution is regressed against instrumented plant activity. The second stage takes the following form:

$$\ln(Ambient_{jt}) = \alpha + X_{it}\varphi + \sum_{j} \widehat{Generation}_{jit}\,\beta + \varepsilon_{it}$$

in which $\ln(Ambient_{it})$ denotes the natural log of ambient $PM_{2.5}$ measurements. All other variables are defined as in section 5.

Table A2 summarizes the results obtained by applying the log-linear model. In our monthly specifications of the 70-mile radius unit of observation (column 5 of Panel B), we find that an increase of 1 Twh (or 1,000 Gwh) of electrical load from coal-fired power plants corresponds to an 129% increase in ambient PM2.5 as measured in nearby populated areas. At the county level (Panel A), however, we find statistically significant results in only two of five model specifications. To put these results into context, note that these estimates are a bit more noisy and smaller in size compared to the linear model. For example, shutting down the largest Texas power plant leads to an 21% reduction in $PM_{2.5}$ in our main specification, but only to a reduction of 13% (column 5) to 19% (column 4) in the log-linear specification.

Rows A-2 and B-2 give the results obtained when emitted SO_2 is used as the independent variable of interest. We find small but statistically significant results at the 70-mile circle levels. Under the 70-mile circle specification, our preferred model indicates that an increase of 1 million tons of emitted SO_2 is associated with a 13% increase in ambient $PM_{2.5}$, while other time unit specifications estimate an effect between 8% and 20%. At the county level, estimates of the effect are again more noisy.

Rows A-3 and B-3 summarize the IV results obtained when emitted NOx is used as the independent variable of interest. As in the linear specification, while the OLS results are unstable likely due to the NOx Budget Program's seasonal cap and trade system, our 70-buffer IV estimates on NOx are all significant and have the correct sign.

Appendix 3

Figure A3 of Appendix 3 provides background data for the nonlinear supply curve of natural gas electricity generation. Panel (a) shows the distribution of observed heat rates for natural gas boilers in SERC in 2012 and Panel (b) shows the dispatch model's results for the lowest possible counterfactual natural gas price without fracking (ΔP^{Upper} scenario) versus the highest possible natural gas price (ΔP^{Lower} scenario) without fracking relative to observed natural gas prices in the 2012 EIA SEDS database. Panel (a) shows the well-known bimodal distribution of heat rates for NG boilers by combined cycle versus gas turbine technologies. Panel (b) shows that the increase in generating hours for natural gas units during 2012 across low versus high natural gas prices would only be met by combined cycle generators. Observed large differences are only present for a couple of individual boilers. As a result, the medium run impacts we estimate are limited by installed combined cycle natural gas capacity.

Figure A1: Counties that contain at least one SO_2 monitor and at least one coal-fired power plant

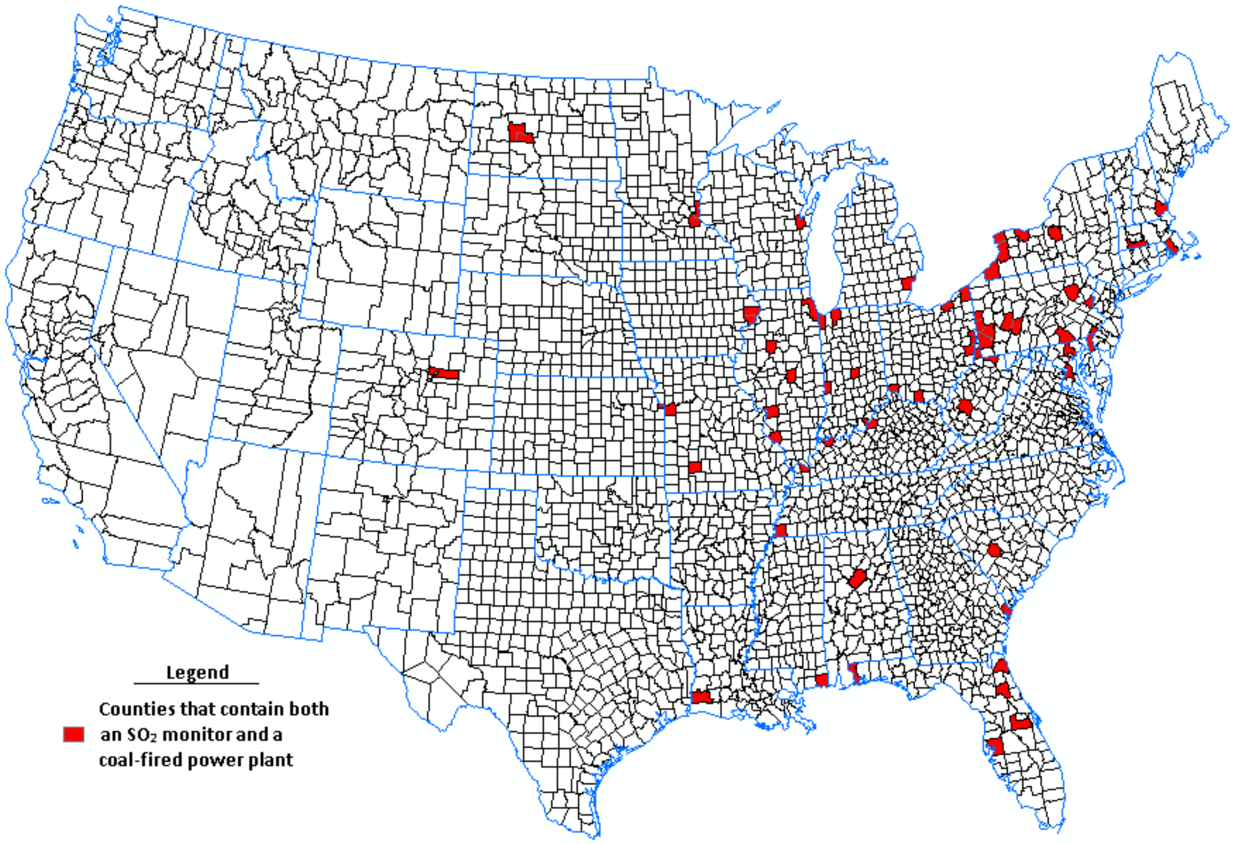

For regressions using county as the unit of observation, only those counties that contain both an SO_2 monitor and a coal-fired power plant are included.

Figure A2: 70-Mile Buffers around SO_2 Monitors

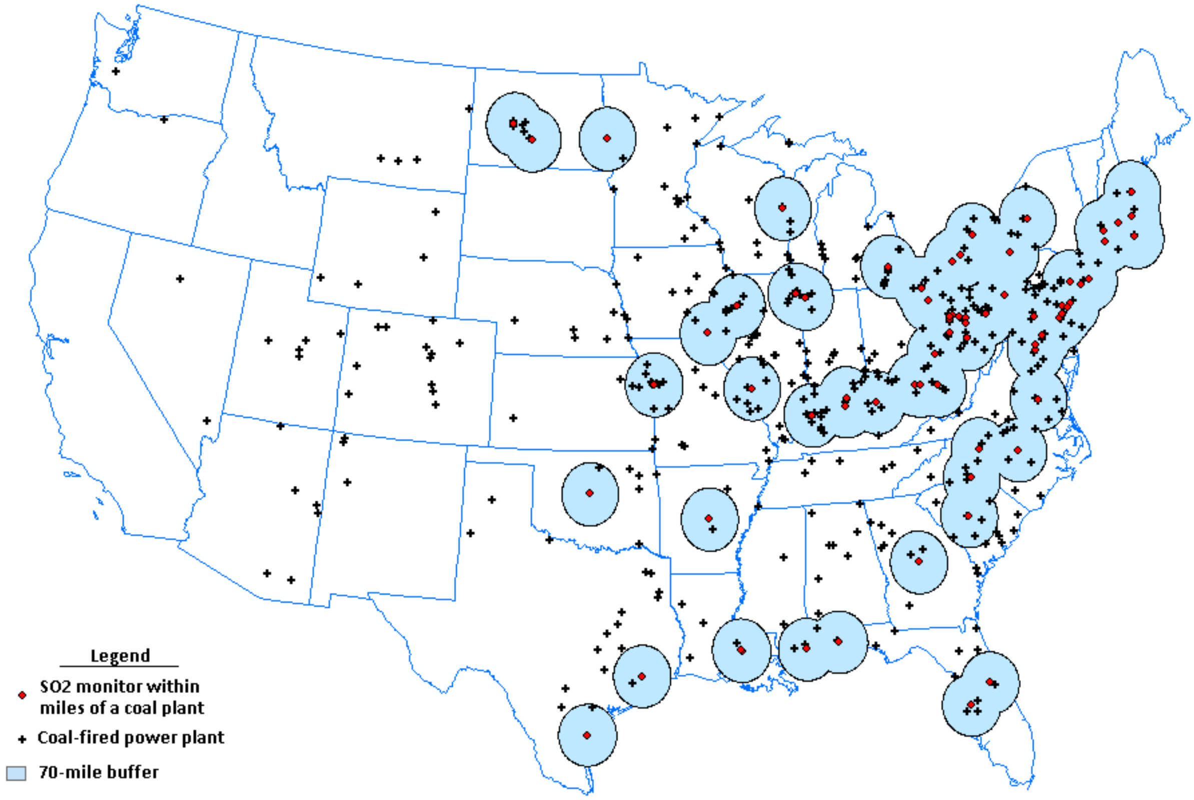

In those regressions using the 70 mile radius circle as the unit of observation, a circle is drawn around each SO_2 monitor. Within each circle, the total amount of electricity generated and pollution emitted is summed. Each circle is then treated as a single observation.

Figure A3: Sensitivity Analysis around Different Fracking Impacts

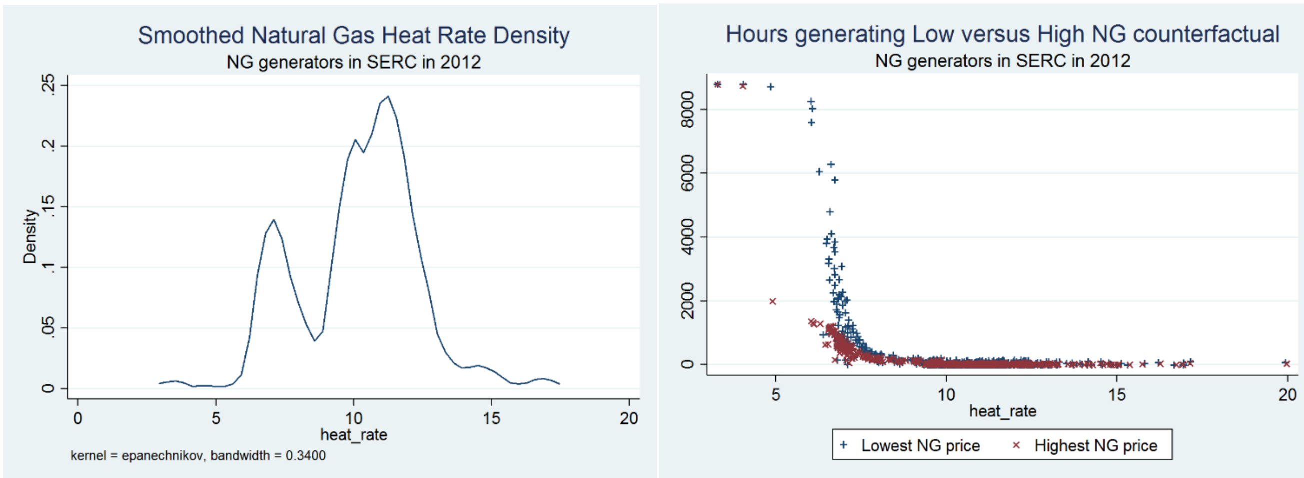

Panel (a) *Panel (b)*

Panel (a) shows the distribution of observed heat rates for natural gas boilers in SERC in 2012 and Panel (b) shows the dispatch model's results for the lowest possible counterfactual natural gas price without fracking (ΔP^{Upper} scenario) *versus the highest possible natural gas price* (ΔP^{Lower} scenario) *without fracking relative to observed natural gas prices in the 2012 EIA SEDS database. Panel (a) shows the well-known bimodal distribution of heat rates for NG boilers by combined cycle versus gas turbine technologies. Panel (b) shows that the increase in generating hours for natural gas units during 2012 across low versus high natural gas prices would only be met by combined cycle generators. Observed large differences are only present for a couple of individual boilers. As a result, the medium run impacts we estimate are limited by installed combined cycle natural gas capacity.*

Table A1: IV Regression results for ambient SO_2 on various specifications of plant activity.

Panel A: County Level

	(1)	(2)	(3)	(4)	(5)
Measures of Plant Activity					
A-1: Avg Daily Gross Load (Twh)	-94.51	-101.3	-44.76	-25.79	-57.94
	(114.5)	(120.2)	(96.86)	(70.00)	(76.51)
R-squared	0.700	0.712	0.642	0.583	0.608
Adjusted R-squared	0.632	0.639	0.617	0.569	0.573
A-2: Avg Daily Emitted SO2 (tons)	-11.45	-13.23	-6.175	-4.494	-8.344
	(23.60)	(27.17)	(13.04)	(10.80)	(12.73)
R-squared	0.468	0.423	0.551	0.524	0.460
Adjusted R-squared	0.348	0.275	0.520	0.508	0.412
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
Number of Counties	66	66	66	66	66
Observations	365	365	1,447	4,315	4,315

Standard errors clustered by state in parentheses

*** $p<0.01$, ** $p<0.05$, * $p<0.1$

Panel B: Monitor Level

	(1)	(2)	(3)	(4)	(5)
Measures of Plant Activity					
B-1: Avg Daily Gross Load (Twh)	2.734	-1.407	1.356	1.890	3.643
	(12.77)	(14.07)	(8.729)	(5.975)	(6.704)
R-squared	0.691	0.701	0.616	0.566	0.602
Adjusted R-squared	0.627	0.632	0.593	0.553	0.578
B-2: Avg Daily Emitted SO2 (tons)	0.237	-0.120	0.117	0.193	0.381
	(1.115)	(1.204)	(0.753)	(0.609)	(0.702)
R-squared	0.693	0.697	0.618	0.568	0.600
Adjusted R-squared	0.629	0.627	0.595	0.555	0.577
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
Number of Monitors	73	73	73	73	73
Observations	438	438	1,750	5,246	5,246

Standard errors clustered by state in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table A2: IV Regression results for log (ambient $PM_{2.5}$) on various specifications of plant activity

Panel A: County Level

	(1)	(2)	(3)	(4)	(5)
Measure of Plant Activity					
A-1: Avg Daily Gross Load (Twh)	7.464	4.487	8.842*	9.625**	4.550
	(5.474)	(5.310)	(5.065)	(4.309)	(3.880)
R-squared	0.934	0.939	0.789	0.673	0.753
Adjusted R-squared	0.918	0.924	0.777	0.665	0.741
A-2: Avg Daily SO2 (Million Tons)	0.663	0.409	0.927*	1.098*	0.511
	(0.447)	(0.455)	(0.552)	(0.590)	(0.471)
R-squared	0.922	0.936	0.769	0.651	0.749
Adjusted R-squared	0.904	0.920	0.755	0.643	0.737
A-3: Avg Daily NOx (Million Tons)	3.068	1.930	3.597	4.129*	2.034
	(2.500)	(2.408)	(2.353)	(2.151)	(1.892)
R-squared	0.912	0.931	0.743	0.625	0.739
Adjusted R-squared	0.892	0.914	0.728	0.616	0.727
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
Number of Counties	166	166	166	166	166
Observations	854	854	3,355	10,010	10,010

Standard errors clustered by state in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Panel B: Monitor Level

	(1)	(2)	(3)	(4)	(5)
Measure of Plant Activity					
B-1: Avg Daily Gross Load (Twh)	1.075***	0.854***	1.965***	1.908***	1.287**
	(0.353)	(0.304)	(0.711)	(0.619)	(0.557)
R-squared	0.873	0.882	0.673	0.598	0.694
Adjusted R-squared	0.855	0.865	0.661	0.592	0.688
B-2: Avg Daily SO2 (Million Tons)	0.0967***	0.0799**	0.199***	0.205***	0.131**
	(0.0372)	(0.0377)	(0.0715)	(0.0620)	(0.0602)
R-squared	0.860	0.873	0.626	0.564	0.682
Adjusted R-squared	0.840	0.855	0.613	0.559	0.675
B-3: Avg Daily NOx (Million Tons)	0.284**	0.228**	0.555**	0.594***	0.413**
	(0.132)	(0.0947)	(0.229)	(0.228)	(0.203)
R-squared	0.865	0.877	0.635	0.565	0.677
Adjusted R-squared	0.846	0.859	0.622	0.559	0.670
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
Number of monitors	387	387	387	387	387
Observations	2,316	2,316	9,255	27,763	27,763

Standard errors clustered by state in parentheses

*** p<0.01, ** p<0.05, * p<0.1