

The Effects of Fuel Prices, Environmental Regulations, and Other Factors on U.S. Coal Production, 2008-2016

John Coglianesse*

Economist,

Research & Statistics Division, Board of Governors of the Federal Reserve System

Todd D. Gerarden

Assistant Professor,

Charles H. Dyson School of Applied Economics and Management, Cornell University

James H. Stock

Harold Hinchings Burbank Professor of Political Economy,

Department of Economics and Harvard Kennedy School, Harvard University, and

National Bureau of Economic Research

March 30, 2018

This revision: December 2018

Abstract

We decompose the decline in coal production from 2008 to 2016 into the contributions of several sources. In particular, we estimate the effects of declining natural gas prices and the introduction of new environmental regulations along with several other factors, using both monthly state-level data and annual information on coal plant closings. We estimate that falling price of natural gas relative to coal is responsible for 92% of the total decline in coal production over this period and that environmental regulations account for an additional 6%, with other factors making small and offsetting contributions.

Keywords: energy transition; electricity; environmental regulation.

JEL Codes: Q40, Q58

* The authors thank Ken Gillingham and Chris Knittel for helpful discussions. Corresponding author: Todd Gerarden, Charles H. Dyson School of Applied Economics and Management, Cornell University, Ithaca, NY 14853, USA. Email: gerarden@cornell.edu. The views expressed in this paper are those of the authors and do not necessarily represent the views or policies of the Board of Governors of the Federal Reserve System or its staff.

1. Introduction

From 2008 to 2016, U.S. coal production fell 37%, from 1,172 million short tons to 739 million short tons (EIA, 2017). During this period, carbon dioxide emissions from coal-fired power plants fell by 718 million metric tons, although some of those reductions were offset by increased CO₂ emissions from natural gas for power generation. The decline in coal CO₂ emissions drove a 23% reduction in total power sector CO₂ emissions 2008 to 2016. The decline in the burning of coal to generate electricity also provided significant health benefits arising from reductions in emissions of fine particulates and other pollutants. At the same time, this decline in coal consumption has had a substantial impact on the economic prospects of coal mining communities, with employment in coal mining falling from 87,000 in 2008 to 52,000 in 2016.

Public discussions of the reasons for this decline focus on three explanations: the decline in natural gas prices as a result of fracking, environmental regulations affecting coal-fired power plants, and the role of state-level renewable mandates in increasing wind and solar generation. Understanding the reasons for the decline in coal is important. If the primary reason is the decline in the price of natural gas, then policies that spur additional production of natural gas will serve to keep prices low and hasten the transition away from coal. If, on the other hand, the main reason for the decline was environmental rules regulating coal emissions, then softening or reversing environmental regulations could spur a revival of coal. Thus, quantifying these and other factors informs the likely effect of environmental and climate policies in the power sector going forward.


We decompose the decline in coal production from 2008 to 2016 into nine factors. Six pertain to domestic steam coal use for electricity: changes in the price of natural gas relative to coal, environmental regulations affecting coal-fired power plants, Renewable Portfolio Standards

(RPSs), heat rates (i.e., thermal efficiency), overall electricity demand, and a small unexplained residual. The remaining three factors relate to the other uses of U.S. coal: changes in industrial use, net exports, and metallurgical coal.

Our decomposition combines three methodological approaches. First, we use monthly state-level data from 2001 to 2016 to decompose econometrically the decline in the coal share of generation into the effect of changing relative prices, air regulations other than the Mercury and Air Toxics Standards (MATS), and RPSs on coal's share of electricity generation, allowing for seasonal patterns that differ across states. We use state-level data (where prices and quantities are aggregated from plant-level data) because the RPSs and many of the air regulations are implemented at the state level. Second, the one air regulation that took effect over this period that does not have regional variation – and thus the effect of which is not identified by variation across states – is the MATS rule. To estimate its effect on steam coal demand, we perform a differences-in-differences analysis of the effect of the rule on planned plant closures and add the results to the econometric decomposition. Third, to complete the decomposition, we use an accounting approach to add in the relatively small amount of production arising from changes in electricity consumption, steam coal exports, and metallurgical coal demand.

Figure 1 presents the resulting decomposition of the decline in coal production tonnage from 2008 to 2016 into its nine components. According to our estimates, the declining price of natural gas relative to coal, on an energy-adjusted basis, explains 92% (SE = 2.5 percentage points) of the decline in coal production, or approximately 397 of the 433 million ton decline. An additional 6% (SE = 2.2 percentage points) of the decline is explained by environmental regulations, primarily the Cross-State Air Pollution Rule (CSAPR) and the MATS rule. The remaining seven factors contribute small, largely offsetting amounts to the change in coal

production. We attribute a small amount of the decline, 2 percent (9 million tons) to the adoption of Renewable Portfolios Standards. For the middle part of this period the growth in coal exports contributed positively, although the export market has declined subsequently and changes in coal exports over the full 2008-2016 period are negligible. There is also a small unexplained component (20 million tons), which arises from measurement error from combining data from different surveys and the regression residual in the state-level econometric model of shares. The implication of this decomposition is that the decline in coal consumption in the power sector, and thus of CO₂ emissions in the power sector, over this period was predominately driven by the fracking revolution, not by public policy.



Sources: EIA Form 923, Electric Power Monthly, EIA API, FRED.

Notes: Shaded areas denote the contributions, in millions of short tons, to the change in coal production (solid marked line) from 2008 to the indicated year. Including the residual, the components sum to the change in coal production from 2008 to that year. *Source:* Energy Information Administration data and authors' calculations.

Figure 1. Decomposition of changes in annual coal production, 2008-2016

These results inform the prospects for either a rebound, or a continuation of the decline, in coal production. Because the major driver in the decline is the relative price of natural gas to coal, prospects for a rebound in coal production largely hinge on the path of this relative price. The U.S. Energy Information Administration's Annual Energy Outlook projects natural gas prices to rise gradually to approximately \$5/MMBtu in 2040, still well below 2008 prices, and for national electricity demand to be flat (EIA, 2018a). Policy actions could affect these prices. Greater access to gas and coal deposits on federal lands would tend to lower both coal and gas prices; however, the time required to open a new mine is far longer than to drill a new gas well and the price of coal from public lands is already very low, so opening public lands further would tend to make gas more, not less, competitive. Other deregulatory actions, such as the suspension of the Bureau of Land Management's waste methane rule (BLM 2017), could also push down the gas-coal relative price. Although the price of renewables is projected to decline further (e.g., Shankleman and Warren, 2017), the full cost of building new renewables must compete with the marginal cost of burning coal in existing facilities. Moreover, federal production and investment tax credits for renewables expire in 2022, and there may be additional economic challenges associated with integrating large amounts of renewable energy given its intermittent nature.

Taken together, our results suggest that in the projected environment of stable relative prices of gas to coal, existing renewables penetration, and stable total electricity demand, the prospects for a rebound in coal production are slim. At the same time, barring major price decreases for gas and/or renewables or major climate regulations or legislation, our results are consistent with a plateau in coal use, not a continuing decline.

Relation to the Literature Several recent studies seek to understand the effects of low natural gas prices on electricity generation and emissions (Knittel, Metaxoglou, and Trindade 2015; Johnsen, LaRiviere, and Wolff 2016; Cullen and Mansur 2017; Holladay and LaRiviere 2017; Linn and Muehlenbachs 2018). A few studies seek to understand the effects of a broader set of factors – including natural gas prices, environmental regulations, and renewables, among others – on a variety of coal and electricity market outcomes (Houser, Bordoff, and Marsten 2017; Linn and McCormack 2017; U.S. DOE 2017; Fell and Kaffine 2018; Jordan, Lange, and Linn 2018). In this section we provide a brief discussion of the papers most closely related to this one.

The U.S. Department of Energy (2017) finds that the biggest driver of recent coal plant retirements has been low relative prices for natural gas, consistent with our results. It also cites low electricity demand, generation by renewable technologies, and environmental regulations as factors that “negatively impacted the economics of baseload plants,” but the report does not provide quantitative estimates of their importance.¹ In contrast, our paper provides quantitative estimates of the importance of these factors on coal *consumption* by all coal-fired electricity generators, rather than focusing more narrowly on coal-fired electricity generation *capacity*.

Like us, Houser, Bordoff, and Marsten (2017) decompose the decline in aggregate coal production into the contributions of environmental regulations, coal-to-gas substitution, renewables penetration, international demand, and other effects. There are several important differences between our approach and theirs. First, they use *ex-ante* estimates of the impacts of environmental regulations, whereas we use econometric methods to estimate these effects *ex-*

¹ The study cites anecdotes about the negative effect of renewables generation driven by state and federal policy on the profitability of traditional baseload capacity, but finds that “To date, however, the data do not show a widespread relationship between VRE [variable renewable energy – wind and solar] penetration and baseload retirements” (DOE 2017, p. 50).

post. Second, they use an accounting identity approach based on *ex-ante* electricity generation forecasts to quantify the role of natural gas prices that may conflate multiple factors affecting changes in generation over time, whereas we explicitly estimate this relationship using variation over both time and space. Third, our use of econometric methods allows us to estimate the uncertainty associated with the decomposition. Fourth, their decomposition of the decline in coal is relative to 2006 projections by the EIA, whereas ours is relative to the actual value of production in 2008, which makes it difficult to directly compare results.² That said, they attribute 49 percent of the decline of domestic coal consumption to cheap natural gas, 26 percent to lower-than-expected demand, 18 percent to growth in renewable energy, and a significantly smaller amount to environmental regulations.

Fell and Kaffine (2018) estimate the dependence of coal-fired electricity generation on the amount of electricity generated from wind and the relative price of coal to natural gas. Their use of daily price and generation data means that their estimates focus on dispatching decisions (the intensive margin), whereas our analysis using monthly data captures plant closures as well (the extensive margin). They find that changes in the price of natural gas relative to coal and increased wind generation contributed significantly to the decline in coal generation between 2008 and 2013, and they are unable to reject the null hypothesis that these two factors explained

² The main distinction here is between projections and actuals. The 2006 EIA total demand projections did not anticipate the Great Recession and thus showed steadily growing electricity demand, whereas in actuality the recession led to a decline in demand which then plateaued. Relative to the EIA projections, Houser, Bordoff, and Marsten (2017) attribute a substantial amount of the decline in coal production to lower-than-projected electricity demand. In contrast, our baseline is the amount of coal consumed in 2008, and we decompose the change in actual consumption from 2008 to 2016. Because electricity demand was roughly flat from 2008 to 2016, we attribute little of the change since 2008 to a change in total electricity demand. Said differently, Houser, Bordoff, and Marsten (2017) answer the question, “What accounts for coal production being less than EIA thought it would be circa 2006?” whereas we answer the question, “What accounts for the decline of coal production from 2008 to 2016?”

all of the observed reduction of capacity factors for coal generators across the four electricity markets they study.³ In contrast, we focus more broadly on coal production rather than coal-fired electricity generation, and we explicitly estimate the impact of environmental regulations on coal-fired electricity generation in addition to the impacts of natural gas prices and renewables. Their results are broadly similar to ours, although they find that wind power plays a larger role, and natural gas prices a somewhat smaller role, than our estimates imply.

Jordan, Lange, and Linn (2018) also study the determinants of trends in the U.S. coal industry using a model of coal mine profitability. They find that natural gas prices and electricity demand explain a significant number of Appalachian coal mine closures, but that declining labor productivity had an even larger effect. In the context of our analysis, this declining labor productivity acts as a disturbance to the coal supply curve that induces an increase in price of coal, all else equal, contributing to the decline in the relative price of gas to coal. At the national level, 89% of the decline in the log relative price is due to the fall in gas prices, so that is how we describe our results; however the Jordan et. al. (2018) mine closure effect could further have exacerbated the effect of fracking for states that use Appalachian coal.

Two other papers provide context for our results. First, Johnsen, LaRiviere, and Wolff (2016) use a least cost dispatch model to find that low natural gas prices due to fracking displaced 28 percent of coal-fired generation over the period 2007 to 2012. Our results suggest

³ Fell and Kaffine (2018) allow for natural gas prices and wind generation to have nonlinear effects on coal generation. Considered separately, natural gas prices play a larger role than wind generation, with relative prices explaining between 25 and 100 percent of the decline in electricity generation from coal across four regions (ERCOT, MISO, PJM, and SPP). Considered together, these two factors explain between 68 and 135 percent of the decline across the four regions. In the case of ERCOT, the combined effects of relative prices and wind generation explain more than the full decline in electricity generation from coal, which could be due to uncertainty not reflected in point estimates or changes in electricity demand, among other factors.

that low natural gas prices caused a 38 percent decline in coal consumption for electricity egeneration in 2016 relative to 2008. Second, using a detailed model of operations and investment in the eastern interconnection, Linn and McCormack (2017) find that natural gas prices, renewables generation, and lower than expected electricity demand explain a 41 percent reduction in coal use by electricity generators between 2005 and 2015. Their quantitative results contrast with ours, as they find electricity demand was the most important of these three factors, followed closely by natural gas prices.⁴

A methodologically related paper is Knittel, Metaxoglou, and Trindade (2015). They use plant-level panel data to estimate econometrically the response of electricity generation from coal to changes in the relative price of coal to natural gas. They focus on how this response varies by ownership and market type, and do not provide aggregate estimates of the effects of natural gas price declines, environmental regulations, or flattening demand.

Road Map The remainder of this paper is organized as follows: Section 2 describes the data and Section 3 presents the methods and results for the state panel data econometric analysis. Section 4 presents the MATS event study, and Section 5 explains the overall decomposition. Section 6 concludes.

⁴ A number of factors make it difficult to directly compare Linn and McCormack's (2017) results to ours. First, they do not provide a detailed breakdown of the impacts of these three factors on coal use (they report results on electricity generation shares, mean capacity factors, profits, and emissions, but not coal consumption). They find that the impacts of the demand shock and relative fuel prices on the mean capacity factor of coal plants are approximately equal, so the impacts of these two factors on coal use are likely to be of similar magnitudes. Furthermore, differences in the time periods and approaches of these two studies make it difficult to directly compare results. For example, Linn and McCormack (2017) compare 2005 projections of 2015 outcomes to actual 2015 outcomes to define the energy market shocks they study, whereas our decomposition is based on comparing actual 2008 outcomes to actual 2016 outcomes (as well as outcomes in the intervening years).

2. Data

2.1 Data sources

Data on annual coal production come from EIA Office of Oil, Gas, and Coal Supply Statistics, which compiles data from Form EIA-7A and Form 7000-2.⁵ Coal consumption for electricity generation, average coal prices for electricity, and average natural gas prices for electricity are compiled by the EIA Office of Energy Statistics using Forms EIA-826, EIA-923, EIA-860, and EIA-861.⁶ We use quarterly domestic consumption of metallurgical coal compiled by EIA from Forms EIA-3 and EIA-5 and aggregate it to an annual frequency.⁷ We use data on quarterly exports and imports of coal by origin and destination country compiled by EIA from US Census Bureau Monthly Report 545 and Monthly Report IM 145 and aggregate to a national annual time series of net exports for the US.⁸

We also construct panel data on electricity production and the coal consumption for electricity at the state level. We obtain data on electricity generation both from coal and in total across all fuels (both in MWh) at the state level from the EIA's API covering the period 2001-2016. Data on consumption of coal by electricity generating units (EGUs), along with the delivered prices of coal and natural gas used in electricity generation are taken from the EIA publication Electric Power Monthly over the 2003-2016 period. EIA constructs the delivered prices of coal and natural gas as a weighted average over EGUs, weighted by the quantity of

⁵ These data were accessed through the EIA's Coal Data Browser (EIA 2017).

⁶ These data were accessed through the EIA's Electricity Data Browser.

⁷ We downloaded this data from the EIA's Application Programming Interface (API).

⁸ These data were downloaded from the EIA's API.

each fuel used by EGU. We estimate the heat rate using a 12-month moving average of state coal consumption divided by a 12-month moving average of state coal-fired generation.

The data contain some missing values on delivered fuel prices.⁹ We handle these missing values two ways. For our primary results, we omit observations with missing prices when estimating the regressions. As a sensitivity check, we re-estimate the panel data regressions using imputed prices, computed as the predicted value from a regression of log prices on state and time fixed effects (so the imputed log price is the national log price for that month, adjusted for a constant state departure from the national log price). Given the regression coefficients, the coal decline decomposition is computed using the full set of prices, observed and imputed.

We also create panel data on environmental regulations at the state level. We construct indicators to capture whether individual Clean Air Act regulations affect coal generators in a given state based on the date each regulation went into effect and the states each regulation covered. This information was gathered from the Environmental Protection Agency (EPA) website and regulatory filings. The Clean Air Act regulations included are:

- **Ozone Transport Commission (OTC) NO_x Budget Program, 1999-2002:** allowance trading program to reduce summertime NO_x emissions in order to reduce ground-level ozone concentrations. Applied to EGUs and large industrial boilers in Northeastern states.
- **NO_x Budget Trading Program, 2003-2008:** cap and trade program to reduce summertime NO_x emissions from EGUs and industrial sources in eastern states. The program was a component of the NO_x State Implementation Plan (SIP) Call. It replaced

⁹ 18% of the prices were missing, accounting for 14% of coal generation. EIA only reports state average delivered prices when sufficiently many EGUs reporting delivered prices that averages do not disclose confidential business information.

the OTC NO_x Budget Program. The initial compliance deadline varied by state (see Table A.2).

- **Clean Air Interstate Rule (CAIR):** three separate cap and trade programs to reduce NO_x and SO₂ emissions from EGUs and industrial sources in eastern states.
 - **CAIR Annual NO_x Trading Program, 2009-2014:** limited total annual emissions of NO_x.
 - **CAIR Ozone Season NO_x Trading Program, 2009-2014:** limited summertime NO_x emissions. Effectively replaced the NO_x Budget Trading Program due to its greater stringency.
 - **CAIR SO₂ Trading Program, 2010-2014:** limited total annual SO₂ emissions.
- **Cross-State Air Pollution Rule (CSAPR), 2015-Present:** implemented in 2015 as a replacement for the Clean Air Interstate Rule. Like CAIR, it implemented cap and trade programs designed to reduce NO_x and SO₂ emissions:
 - **CSAPR Annual NO_x Trading Program, 2015-Present**
 - **CSAPR Ozone Season NO_x Trading Program, 2015-Present**
 - **CSAPR SO₂ Groups 1 and 2, 2015-Present:** two separate programs to regulate different entities.
- **Mercury and Air Toxics Standards (MATS), 2015-Present:** regulation that imposed emissions standards on coal- and oil-fired EGUs to reduce emissions of toxic air pollutants including mercury and arsenic. The initial compliance date was April 2015, but regulated entities could file for extensions to comply by April 2016. We used two different indicator variables to capture the impact of these two compliance deadlines.

This process resulted in a binary variable for each regulation, with a one indicating that one or more (typically, all) EGUs in that state/month were subject to the regulation. Preliminary analysis showed that there is insufficient variation in the data to estimate separate coefficients for each of these rules. We therefore merged all the CAIR rules into a single CAIR dummy variable, and we also merged the CSAPR rules into a single CSAPR dummy variable. For all regulations except for MATS, this resulted in regulations that have variation across states in multiple months, so that for those regulations the effect of the regulations is identified from state variation. This leaves us with four regulatory dummy variables with state-level variation over this period: OTC NBP Seasonal NO_x, NO_x NBTP, CAIR, and CSAPR.

For MATS, the regulation affected all coal-fired EGUs so the MATS binary indicator is in effect a time dummy variable that takes the value of one after the compliance date. We do not use time effects in our regressions because doing so would attribute the decline to time effects without energy-economic substantive content. Instead of identifying MATS from time series variation, we undertake a plant closing event study, which is reported in Section 4. For that study, we used data on planned plant closures from EIA Form 860; the details are discussed in Section 4.

We incorporate data on Renewable Portfolio Standards (RPSs) from the Database of State Incentives for Renewables and Efficiency (DSIRE) maintained by the NC Clean Energy Technology Center at North Carolina State University. Ideally, we would like to have total in-state generation mandated by the RPSs; however it is not feasible to construct that series, both because of lack of renewable generation quantity target data for some states and because RPS requirements typically can be met, to varying degrees, by interstate trading of renewable energy credits. We therefore used a simpler RPS measure: a binary indicator for the presence of an RPS


based on the date a state first established its RPS. Table A.2 contains a detailed summary of when each RPS and environmental regulation was in effect in each state.

Finally, of a total of 9,231 state-month observations, 813 observations have no coal-fired electricity generation because there are no operating coal plants in the state. Of these, 724 observations are for states with no coal-fired generation over the entire period (Idaho, Rhode Island, Vermont, and D.C.), and the remaining 89 are for states (like Oregon, Washington, and California) in which the last coal-fired EGU was retired during the sample period. These observations are excluded from our econometric analysis and decomposition.

2.2 Data description


Figure 2 shows coal production and employment from 2000 to 2016. Coal production grew slightly from 2000 to 2008, then declined, with the sharpest declines occurring in 2015 and 2016. Coal mining employment, both as measured by the Bureau of Labor Statistics and by the EIA, grew from 2002 to 2011, then also fell sharply through 2016.

Figure 3 highlights the three main trends in electricity generation from 2001-2016: the decline in coal-fired generation, the rise in generation by natural gas, and the increasing (but still small) role of renewables.



Source: Energy Information Administration and Bureau of Labor Statistics.


Figure 2. Annual U.S. coal production and coal mining employment, 2000 – 2016.



Source: Energy Information Administration and Bureau of Labor Statistics.

Figure 3. Annual U.S. electricity generation by source

Figure 4 shows total electricity generation by coal- and gas-fired power plants at the monthly level since 2008 (top panel) as well as the national average delivered prices of coal and natural gas used for electric power (bottom panel). Coal generation fell nearly 40% over this period, which corresponds to a fall in coal consumption absent a large decline in efficiency. Both coal- and gas-fired generation are highly seasonal, with seasonal peaks for summer cooling and winter heating. Although coal prices were fairly stable, average natural gas prices to the power sector fell by almost two-thirds between 2008 and 2016, from \$9.30/MMBtu to \$2.99/MMBtu.




Source: Energy Information Administration

Figure 4. Monthly U.S. electricity generation from coal and natural gas, and the delivered prices of coal and natural gas.

The strong seasonality in coal generation obscures the relationship between coal generation and gas prices, or relative prices of coal to gas. Figure 5 plots U.S. coal use for electricity generation and the ratio of the national-average delivered price of natural gas to coal,

both seasonally adjusted in logarithms.¹⁰ After seasonal adjustment, a strong relationship between coal consumption and natural gas prices is apparent. Coal consumption moves closely with the relative price. This co-movement is not just a consequence of both series containing a downward trend, rather the increase of the relative price in 2012 through 2013, and again in the second half of 2016, was matched by an increase in coal generation.



Source: Energy Information Administration and authors' calculations

Figure 5. Seasonally adjusted log monthly U.S. coal use at EGUs and log relative price of natural gas to coal, 2009-2016

3. State Panel Data Analysis of the Coal Share in Electricity Generation

3.1 Methods

We use panel data regressions to estimate the response of the share of electricity generation by coal to the relative price of coal to gas (on an energy content basis), environmental

¹⁰ The seasonals were estimated separately for each state in a regression of the logarithm of the series on monthly dummy variables. This mirrors the treatment of seasonals in the panel data regressions, which include state-by-month effects. An alternative would be to estimate the seasonals using a method that allows the seasonals to evolve, such as Census X-11; however given the short data set here the end-point problems of seasonal adjustment with changing seasonals becomes important so we elected to fix the seasonals.

regulations, and the presence of RPS requirements. All regressions are on monthly data by state and include state effects and a full set of monthly seasonals for each state; that is, all regressions include state effects and interactions of the state effects with 11 monthly binary indicators.

Because the shares are constrained to be between zero and one, we use a logistic transform of the shares. We allow for the possibility of adjustment lags, so that the effect of a change in relative fuel prices might not take full effect for several months, perhaps due to physical or contractual adjustment constraints. We allow for these lags in two ways: first by using the log relative price of multi-month moving averages of the fuel prices (our baseline specification), and second by using a distributed lag of the monthly fuel prices. We refer to these respectively as the static and the dynamic specifications. For the static specification, we assume that prices are exogenous, conditional on the state, year, and state-by-month fixed effects, an assumption we relax in Section 3.3 when we use an instrumental variable (proximity to fracking reserves) for relative prices. For the dynamic specification, we assume that relative prices are conditionally weakly exogenous, that is, the demand error is uncorrelated with current and past relative prices, conditional on the state-by-month fixed effects.

Specifically, let g_{it} denote the fraction of electricity generated by burning coal in state i and time t . The static specification is,

$$\text{logit}(g_{it}) = \alpha_{im(t)} + \beta_1 p_{it}^{MA} + \beta_2 (p_{it}^{MA})^2 + \gamma x_{it} + \theta r_{it} + u_{it}, \quad (1)$$

where $\alpha_{im(t)}$ are state-by-calendar-month fixed effects (including the main effects for state and calendar month), p_{it}^{MA} is a moving average of current and past log relative prices, x_{it} is a vector of regulatory dummy variables, and r_{it} is a dummy variable for whether RPSs were in effect. In the base specification, p_{it}^{MA} is an equal-weighted moving average of current and 5 lags of log relative prices. The vector of regulatory variables includes: 1) a dummy for whether any of the

Clean Air Interstate Rule (CAIR) regulations governing NO_x, ozone, or SO₂ were in effect; 2) a dummy for whether either the year-round NO_x/SO₂ rule or the seasonal ozone rule from the Cross-State Air Pollution Rule (CSAPR) were in effect; 3) a dummy for whether the NO_x Sip Call NBTP Seasonal NO_x rule was in effect; and 4) a dummy for whether the OTC NBP Seasonal NO_x rule was in effect.

The dynamic specification allows for richer dynamics, but does not parametrically incorporate nonlinearities. Let L be the lag operator and let $\beta(L)$ be a lag polynomial. The dynamic specification is,

$$\text{logit}(g_{it}) = \alpha_{im(t)} + D_{it} \times \beta(L)p_{it} + (1 - D_{it}) \times \beta(L)p_{it} + \gamma x_{it} + \theta r_{it} + u_{it}, \quad (2)$$

where $p_{it} = \ln(P_{it}^{\text{coal}} / P_{it}^{\text{gas}})$ and D_{it} is a binary indicator that equals one if p_{it} is above the median relative price. The term $D_{it} \times \beta(L)p_{it}$ denotes the current value and nine lags of log relative prices, where each of the regressors is interacted with D_{it} (so the regressor for the k^{th} lag would be $D_{it}p_{it-k}$). The dummy variable interaction specification (2) allows for different price dynamics, including different cumulative elasticities, depending on whether the relative price is high or low, while imposing that the coefficients on the remaining regressors do not depend on the relative price.

We will interpret the responses to a small change in log prices, as estimated by these regressions, as elasticities or semi-elasticities of demand, depending on the setting. If prices were exogenous then we could associate these regression estimates with estimates of demand elasticities. Over this period, a large portion of the decline in the gas price stems from the development, improvement, and deployment of fracking technology, which from the perspective of demand estimation constitutes an exogenous shift in the supply of gas. The availability of shale gas depends on local pipeline infrastructure and the changing location of fracking fields

which provides additional state-level exogenous variation in prices. Although there are seasonal swings in gas prices as a result of seasonal changes in demand, all the specifications include state-level seasonals which absorb this source of potential endogeneity. To us, these features suggest that treating relative prices as exogenous is a plausible approximation.

3.2 Results

Table 1 presents the estimated coefficients in the baseline regressions. We highlight four features of these results.

First, the effects of air regulations and RPS requirements are estimated to reduce the coal share, although most coefficients are not statistically significant. In the baseline regressions (1), (2), and (3), the coefficients for RPS, CSAPR, and NBTP NO_x are all estimated to reduce the coal share. Although the baseline regressions estimated that CAIR could have increased or decreased the coal share, these coefficients are both small and statistically indistinguishable from zero. We view this as consistent with CAIR rules having a small (possibly negative) effect on the coal share, which is difficult to identify precisely from state-level variation.¹¹

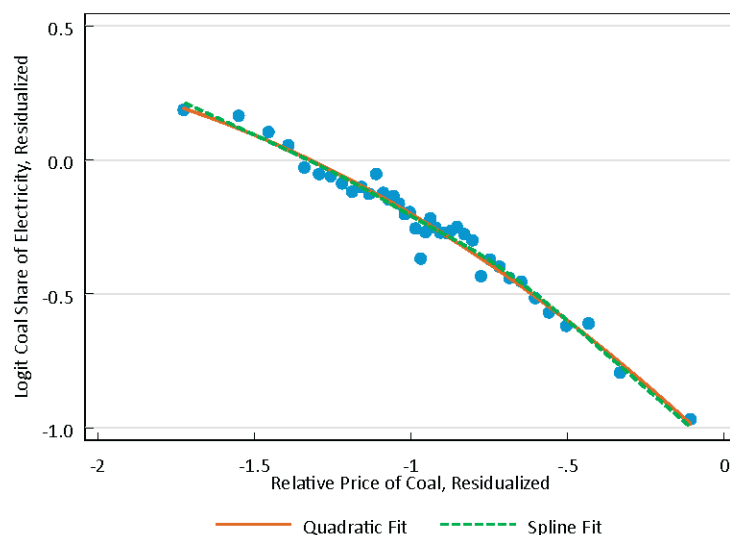
¹¹ Our list of air regulations does not include the Clean Power Plan (CPP), which was proposed in June, 2014 and finalized in August, 2015 but the implementation of which was stayed by the Supreme Court in February, 2016. The CPP would have had the effect of adding a carbon price that would differentially benefit natural gas over coal, beyond the delivered relative price which we use as data. However, because the CPP never took effect, the measured delivered prices reflect the marginal relative price of coal to gas. The CPP would plausibly have led to coal plant retirements, however any such retirements would have occurred after the period covered by our data and in any event would only have occurred had the CPP gone into force. To the extent that there is an anticipatory effect of the CPP in reducing coal-fired generation in 2015, beyond that accounted for by relative prices, the other air regulations, and RPSs, that decline would appear as a residual in our decomposition.

	(1)	(2)	(3a)	(3b)
	Logit(Coal Share of Elec.)			
p_{it}^{MA}	-1.327*** (0.181)	--	--	--
$p_{it}^{MA,2}$	-0.303*** (0.0770)	--	--	--
Dynamic Coefficients	--	All	Below Median	Above Median
Δp_{it}	--	-0.444*** (0.0807)	-0.356*** (0.0528)	-0.935*** (0.163)
Δp_{it-1}	--	-0.549*** (0.0900)	-0.440*** (0.0580)	-1.086*** (0.172)
Δp_{it-2}	--	-0.580*** (0.102)	-0.504*** (0.0624)	-1.023*** (0.183)
Δp_{it-3}	--	-0.510*** (0.0938)	-0.512*** (0.0709)	-0.861*** (0.142)
Δp_{it-4}	--	-0.542*** (0.0933)	-0.544*** (0.0731)	-0.893*** (0.137)
Δp_{it-5}	--	-0.591*** (0.0906)	-0.639*** (0.0825)	-0.819*** (0.124)
Δp_{it-6}	--	-0.600*** (0.0880)	-0.592*** (0.0805)	-0.905*** (0.129)
Δp_{it-7}	--	-0.679*** (0.0941)	-0.674*** (0.0933)	-0.977*** (0.141)
Δp_{it-8}	--	-0.562*** (0.0921)	-0.582*** (0.0927)	-0.879*** (0.145)
p_{it-9}	--	-0.761*** (0.0963)	-0.799*** (0.0962)	-1.078*** (0.149)
CAIR Dummy	0.0312 (0.0954)	0.0583 (0.0930)	-0.00169 (0.0984)	
CSAPR Dummy	-0.0941 (0.128)	-0.175 (0.116)	-0.172 (0.115)	
OTC NOx Dummy	--	--	--	
NBTP NOx Dummy	-0.0544 (0.0885)	-0.0244 (0.0916)	-0.0614 (0.0758)	
RPS Dummy	-0.180 (0.0955)	-0.0922 (0.109)	-0.115 (0.105)	
R-squared	0.936	0.939	0.941	
N	6325	6037	6037	

Notes: Dependent variable is the logistic transform of the state-month coal share in generation. Regressors are described in Section 3.1. p_{it}^{MA} is computed using a six-month moving average of the relative price. Standard errors are two-way clustered by state (47 states) and time period (180 time periods). All regressions are estimated on 2002m1 – 2016m12. Observations in which the moving averages or distributed lags include imputed coal or natural gas prices are excluded. Significant at the *5%, **1%, ***0.1% level.

Table 1. Panel regression results, shares regressions

Second, as illustrated in Figure 6, the nonlinear term in the relative prices is statistically significant and consistent with the demand being more elastic at higher relative prices. Consistent with the good fit of the quadratic specification in the figure, a cubic spline (shown in the figure) delivers essentially the same fit as the quadratic specification. Restricting to a linear specification results in a worse fit particularly during the later years of our sample period when falling natural gas prices raised the relative price of coal.

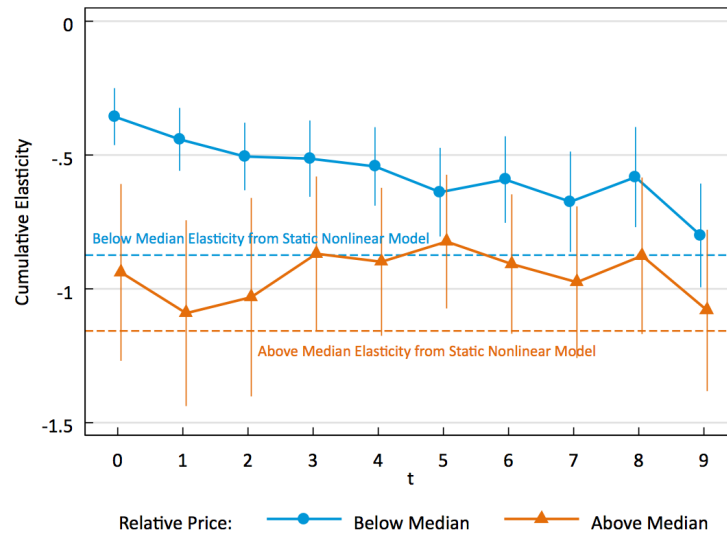


Notes: Both variables are regressed on month-by-state fixed effects and regulatory dummies, then residuals from this regression are divided into 20 equal-sized bins. The means of each variable within each bin are plotted as a scatterplot. The two lines are the regression line from regressing the residualized logit share on the residualized price and the residualized price squared, and a cubic spline version of that regression (5 knots). *Source:* Authors' calculations.

Figure 6. Binned scatterplot of shares (logit transform) v. log relative price, controlling for non-price variables in the static regression of Table 1, column (1).

Third, the estimated distributed lag coefficients are consistent with most of the effect of a relative price change occurring quickly, within the first few months. Figure 7 shows the cumulative dynamic effect of a one percentage point change in the relative price using various specifications in Table 1. When the relative price is high (approaching one), the effect of a one percent change in the relative price on the logit transform of the share is roughly twice what it is

when the relative price is low. Thus the figure provides additional evidence that the static specification provides an accurate approximation to the more complicated dynamics of the dynamic specification. The straight dashed lines are the values of this effect, estimated using the static model. The static estimates closely approximate, and are within one standard error of, the cumulative effects in the dynamic specification. We interpret this figure as providing support for using the static specification for the decomposition.



Notes: The dynamic responses correspond to columns (3a) and (3b) of Table 1. The constant dashed lines are the implied responses using the static specification in Table 1, column (1), evaluated at the mean relative price in the subsamples used to estimate columns (3a) and (3b). Vertical bars indicate 95% confidence intervals.

Figure 7. Cumulative dynamic response of coal share (logistic transform) to a 1 percentage point change in the coal-gas relative price.

Fourth, the static and dynamic specifications give similar predictions for changes in coal demand, and both provide good fits to the state-level share data. Figure 8 shows the predicted values from the static specification and from the dynamic specification; both are fit over the full sample for four representative states. Pennsylvania has a moderate use of coal with little seasonality and both regressions are similarly close to the observed data. In contrast, Montana has strong seasonal patterns, but both regressions are able to fit this data owing to the state by

calendar month fixed effects. We note that comparable figures (not shown) without the logistic transform provides poor fits for states with shares near zero or one as are New York and West Virginia, and occasionally produce predicted shares outside their 0-1 range. The fit of the dynamic specification is marginally worse in some cases than the static linear model, presumably because the nonlinearity is approximated in a way that is not smooth. This plot also supports the use of the static specification.

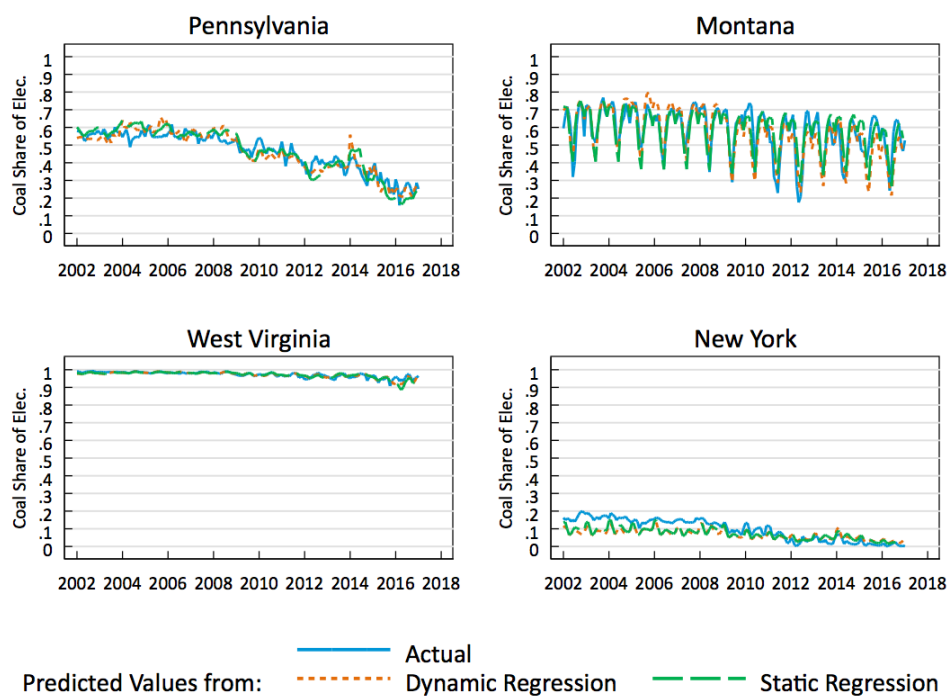


Figure 8. Actual and predicted shares for four representative states: static and dynamic specifications (full-sample estimates).

3.3 Additional Results

We undertook a number of sensitivity checks, some of which are reported in Table 2.

	(1) IV Logit(Coal Share of Elec.)	(2) IV Logit(Coal Share of Elec.)	(3) OLS Logit(Coal Share of Elec.)	(4) OLS Logit(Coal Share of Fossil Fuel Elec.)	(5) OLS Logit(Coal Share of Elec.)	(6a) OLS Logit(Coal Share of Elec.)	(6b) OLS Logit(Coal Share of Elec.)
p_{it}^{MA}	-0.904*** (0.143)	-1.847*** (0.509)	-1.328*** (0.184)	-1.978*** (0.318)	-1.418*** (0.281)	--	--
$p_{it}^{MA,2}$	--	-0.464 (0.242)	-0.296*** (0.0768)	-0.411 (0.232)	-0.312** (0.107)	--	--
Dynamic Coefficients for Above/Below Median Relative Price	--	--	--	--	--	Below Median	Above Median
Δp_{it}	--	--	--	--	--	-0.371*** (0.0573)	-0.814*** (0.202)
Δp_{it-1}	--	--	--	--	--	-0.463*** (0.0696)	-1.029*** (0.217)
Δp_{it-2}	--	--	--	--	--	-0.520*** (0.0765)	-0.980*** (0.225)
Δp_{it-3}	--	--	--	--	--	-0.580*** (0.0871)	-0.838*** (0.192)
Δp_{it-4}	--	--	--	--	--	-0.608*** (0.0946)	-0.860*** (0.179)
Δp_{it-5}	--	--	--	--	--	-0.683*** (0.0988)	-0.802*** (0.186)
Δp_{it-6}	--	--	--	--	--	-0.626*** (0.100)	-0.928*** (0.201)
Δp_{it-7}	--	--	--	--	--	-0.699*** (0.110)	-0.987*** (0.208)
Δp_{it-8}	--	--	--	--	--	-0.684*** (0.105)	-0.909*** (0.218)
p_{it-9}	--	--	--	--	--	-0.893*** (0.134)	-1.077*** (0.224)
CAIR	0.187 (0.121)	0.160 (0.122)	0.206 (0.120)	0.0579 (0.248)	0.105 (0.102)	0.103 (0.104)	
CSAPR	0.00883 (0.167)	0.165 (0.183)	0.0923 (0.173)	-0.245 (0.287)	0.00626 (0.171)	-0.0536 (0.158)	
OTC NOx	--	--	--	--	0.438* (0.191)	0.496** (0.177)	
NBTP NOx	0.00987 (0.105)	-0.120 (0.106)	0.183 (0.127)	-0.00660 (0.161)	0.0712 (0.124)	0.0737 (0.126)	
RPS	-0.0582 (0.114)	-0.131 (0.123)	-0.158 (0.0934)	--	-0.280** (0.0798)	-0.237** (0.0820)	
Cragg-Donald Statistic	4746.1	129.7	--	--	--	--	--
Regulation Variables Includes Imputed Prices	Binary No	Binary No	Continuous No	Binary No	Binary Yes	Binary Yes	
Moving Average Length	6 Month	6 Month	6 Month	6 Month	6 Month	--	
R-squared	0.416	0.396	0.928	0.901	0.925	0.924	
N	6271	6271	6159	6283	8340	8340	

Notes: Dependent variable is the logistic transform of the state-month coal share of either total generation (columns 1, 2, 3, 5, & 6) or fossil-fuel generation (column 4). Instruments and regressors are described in Section 3.3. Standard errors are two-way clustered by state and time period. In the scalar case, the Cragg-Donald statistic is the first-stage F statistic. All regressions are estimated on 2002m1 – 2016m12. Significant at the *5%, **1%, ***0.1% level.

Table 2. Additional regression results, shares regressions

First, we instrument for prices with the availability of shale gas production to verify that our results are not driven by endogeneity in prices. In our baseline specifications, we assumed that all variation in the relative price of coal aside from seasonal variation is due to the growth of fracking technology. To the extent that this is not the case, prices may be somewhat endogenous. We test this assumption by instrumenting for the relative price of coal with a measure of the availability of fracking in order to identify the demand elasticity using only this exogenous source of variation.

We construct our measure of fracking availability by interacting the total amount of available shale gas locally with a time trend. Specifically, we instrument for the relative price p_{it}^{MA} and the quadratic term $(p_{it}^{MA})^2$ with the instruments,

$$z_{it}^{(1)} = \log \left(s_i + \sum_{j \neq i} \frac{s_j}{d_{ij}} \right) * t, \quad z_{it}^{(2)} = \log \left(s_i + \sum_{j \neq i} \frac{s_j}{d_{ij}} \right) * t^2$$

where s_i is cumulative shale gas production for state i over the 2007-2016 period reported by EIA Office of Oil and Gas (based on Form EIA-23 data) and d_{ij} is the distance in miles between the centroids of states i and j . We compute a weighted average of total shale gas production for each state, weighting by inverse distance to approximate the costs of transportation, taking the log of the resulting average. In this way, we measure the availability of shale gas for each state, including both its own reserves as well as the size of reserves in nearby states. By interacting this measure with a time trend, we capture that states with larger availability of shale gas experienced larger declines in the price of natural gas, and thus larger increases in the relative price of coal, over this time period. We also include as an instrument the same fracking availability measure interacted with a squared time trend to instrument for the quadratic term in prices.

Instrumenting for prices to isolate variation due to fracking technology produces similar estimates to our baseline specification, as can be seen in columns (1) and (2) of Table 2. Column (1) uses only the linear term and delivers a similar result as our baseline, although this does not capture nonlinearities.¹² Column (2) also includes the quadratic term and delivers close results to our baseline specification. The coefficients on the linear and quadratic price terms are about -1.8 and -0.5, respectively, which indicates slightly more elastic estimates of demand for coal generation. Both IV regressions are easily able to reject that the instruments are weak, as evidenced by the large first stage F statistics.

Our second set of additional regression results relax the assumption that regulations are either completely or not at all binding for a given state-month observation. In fact, Clean Air Act regulations are binding at the plant-level, since only some types of generators are subject to these regulations. We aggregate the EPA's Air Markets Program Data to the facility level and merge this with the average electricity generating capacity at each facility over the years 2007-2014 from the EIA 860 data. We use this to compute the capacity-weighted share of plants affected by each regulation for each state-month observation, giving us a continuous measure between 0 and 1 of the extent to which each regulation was binding. Column (3) of Table 2 indicates that using these continuous regulation variables produces coefficient estimates that are not statistically or economically significant, as in our baseline specification.

¹² We have additionally explored instrumenting with a quantity that interacts our measure of cumulative shale gas production with a dummy variable for whether production exceeded a minimum threshold in a given time period. This allows for variation along the time dimension not due to a time trend, but instead coming from the introduction of extraction in different locations. As fracking technology developed, extraction began in some areas before others, presenting useful identifying variation. This instrument yields very similar results when using the linear term only as in column (1), but does not sufficiently pass a weak instruments test when the quadratic terms are included.

Third, we show that our results are robust to alternative assumptions about the use of renewable fuels. In our baseline specification, we controlled for the presence of RPS address this issue. Alternatively, we can allow for any exogenous change in the use of renewable fuels by modeling electricity generation from coal as a share of electricity generation from only fossil fuel sources (coal + natural gas). This allows for the impact of RPS and other factors driving the rise of renewable fuels to potentially vary by state and over time. Column (4) of Table 2 repeats our baseline specification using the coal share of fossil fuel generation as the dependent variable without the RPS covariate. The resulting coefficient estimates are similar but indicate slightly more elastic demand, indicating that substitution into renewable fuels was not a confounding factor in our baseline specification.

Fourth, we repeat our baseline regressions adding in observations for which prices have been imputed. Columns (5), (6a), and (6b) indicate that including these observations makes the resulting demand elasticity estimates more elastic.

Fifth, we examined the robustness of the regressions to using fewer or more lags. For the static specification, this amounts to using moving average windows of 1, 2, ..., 18 months. The results are presented in Figure A.1 for the linear static specification. Evidently, the estimated elasticity is insensitive to the moving average window length, at least for windows exceeding three months. In addition, Table A.3 reports dynamic specifications with 6, 9, 12, 15, and 18 lags. Although the dynamics differ somewhat across lag lengths, the long-run (cumulative) elasticity, which is the key elasticity for the decomposition, is very stable across specifications, increasing from -0.759 (SE = 0.097) for six lags to -0.7921 (SE = 0.101) for 18 lags. For all dynamic specifications, the elasticities are close to the elasticities in the linear static model shown in Figure A.1.

Sixth, in unreported results, we estimated specifications in which the price of coal and natural gas entered separately (in logarithms); the hypothesis that the coefficients have equal and opposite signs is not rejected.

4. MATS Event Study

4.1 Methods

The MATS rule regulates emissions of toxic air pollutants including mercury, arsenic, and heavy metals from coal- and oil-fired power plants. The Clean Air Act Amendments of 1990 required EPA to prepare a study on the health effects of hazardous pollution from power plants, and EPA submitted the study in 1998. In 2000, EPA determined that regulating those pollutants was appropriate and necessary. After litigation and court delays, the MATS rule was proposed on March 16, 2011 and was finalized on December 21, 2011.

The MATS rule sets out technology-based standards, with a compliance deadline of March 2015.¹³ EPA expected that it would be economically more cost-effective to retire some plants than to retrofit them. In its Regulatory Impact Analysis, EPA estimated that 4.7 GW of capacity would be retired for MATS compliance and that power sector coal consumption would fall by one percent as a result of MATS (EPA, 2011).

Because the MATS rule applied nationally, a dummy variable indicating the MATS compliance date of 2015 is not separately identified from time effects including national mean changes in prices. Thus estimating the effect of MATS on coal generation share is not amenable

¹³ MATS did allow coal plants to apply for a one-year compliance date extension on a case-by-case basis, but the small change between 2016 planned retirements in the 2010 and 2011 data suggests this did not influence retirement plans made upon the release of the MATS formal rule in 2011.

to the regression methods of the previous section. However, a convenient institutional feature of the MATS rulemaking makes it possible to estimate its effect on newly planned plant closings resulting from the MATS rule. EIA collects data (EIA Form 860) on whether an EGU is planned to be retired and, if so, when the retirement is planned to occur. The deadline for EGU owners to submit these data to EIA is the end of February, two months after the relevant reporting year. Thus forms filled out in early 2011 would not have taken into account the MATS rule because it had not yet been proposed, while forms filled out in early 2012 would take into account the finalized MATS rule. Because the final rule included the compliance schedule, with final compliance in March 2015, some of the changes in EIA Form 860 data between early 2011 and early 2012 for retirements planned to occur in 2015 can be associated with the MATS rule.

There was in fact a 7.2 GW spike in retirements in 2015 which were not planned as of February 2011 but were planned as of February 2012 (Figure 9). There was also an increase in 2014 planned retirements, and a smaller increase in 2016 planned retirements. Neither the 2014 nor 2016 increase can plausibly be attributed to MATS: there is no reason to retire a unit early for MATS compliance if it is economical without the MATS compliance upgrades, while retiring it later than 2015 would place it out of compliance. Rather, these newly announced retirements for 2014 and 2016 likely reflect other confounding factors that changed between early 2011 and early 2012. This suggests that not all of the 7.2 GW increase in 2015 planned retirements is attributable to the MATS rule, so estimating the MATS effect needs to adjust for these confounding factors.

The analysis of the preceding section suggests that the leading confounding factor is the relative price of natural gas to coal. During 2011 fracking was having its first major impact on gas prices, and gas price projections were being adjusted to take fracking into account.

Moreover, because of pipeline infrastructure and the location of shale gas, both gas prices and projected gas price paths during this period varied regionally.

We therefore use data on price projections for 9 U.S. regions, published by the EIA in early 2011 and again in early 2012 as the reference case in their *Annual Energy Outlooks*, to control for the effect of changing natural gas price forecasts over this period. These regional projections were distributed to the state level so that the final year of actual prices in the AEO regional forecast aligns at the state level with actual prices in the final quarter of that year, resulting in state-level projections of gas-coal relative prices. With these data in hand, it is possible to estimate the effect of the MATS rule as the excess retirements in 2015, above and beyond those arising from a change in gas prices.


4.2 Results

The regression results are given in the first column of Table 3 and are summarized in Figure 9. The negative coefficient on the log relative gas-coal price indicates that lower gas prices would have increased planned retirements, even absent the MATS rule. The total increase in planned retirements in 2015 is 7.2 GW. Based on the estimated regression coefficients in Table 3, of this 7.2 GW, 2.0 GW would have occurred without the MATS rule as a result of changes in projected relative prices. The remaining 5.2 GW of retirements are attributed to the MATS rule.

Dependent variable		Δ Capacity	Δ Coal consumption
Retirement year span		2011-2016	2011-2016
Regressors			
	<i>2015 binary indicator</i>	0.20 (0.13)	0.25 (0.19)
	<i>State relative price forecasts (logs)</i>	-0.47* (0.20)	-0.61 (0.33)
	<i>Intercept</i>	-0.05 (0.05)	-0.05 (0.08)
R-squared		0.066	0.041
N		156	156
Estimated MATS effect (national)		5.2 GW (3.4)	6.4 MT (4.9)

Notes: Regressions are estimated using state-level data on projected capacity retirements, by retirement year, and log projected relative prices. Capacity and price projections are specified in changes in projections between early 2011 and early 2012. The regressions were estimated on the 26 states that reported changes in planned retirements, over the retirement years indicated in the first numerical row. The first three rows provide regression coefficients and, in parentheses, heteroskedasticity-robust standard errors. *Data sources:* EIA Forms 860 and 923, EIA *Annual Energy Outlook* 2011 and 2012, and authors' calculations. Significant at the *5% **1% ***0.1% level.

Table 3. Regression estimate of MATS retirement effect on nameplate capacity and direct coal consumption.



Notes: To derive the econometrically adjusted changes, we use estimates from the first model in Table 3 to remove the effects of the constant and the change in relative fuel prices from the raw data. The result is the composite effect of the MATS rule and residual variation unexplained by the model, which we plot side-by-side with the raw data. Source: Authors' calculations.

Figure 9. Change in planned coal-fired power plant retirements by year of retirement, between February 2011 and February 2012: raw data and econometrically adjusted for changes in forecasted gas and coal prices.

The effect of the MATS rule on coal consumption can also be estimated by incorporating data from EIA Form 923. For each unit planned for retirement according to Form 860, we incorporate data from Form 923 on coal receipts during that reporting year and distribute the total quantity for the plant to each coal-fired unit within the plant proportionally based on the nameplate capacity of each unit. We then aggregate these data for all planned retirements in a given 860 reporting year to the state level. As with capacity, some of the reduction in coal consumption at these newly planned retirements is likely due to the decline in gas price forecasts, and the net effect of the MATS rule can be estimated in the same way as its net effect on retirements. The results are reported in the second column of Table 3. The direct effect of the MATS rule on coal consumption is estimated to be a reduction of 6.4 million short tons, or 0.6% of total 2014 coal production.

5. Decomposition

5.1 Methods

We start by decomposing the production of coal in year y into coal used for domestic electricity generation, net steam coal exports, metallurgical coal, domestic industrial steam coal use, and all other sources. Thus the change in coal production from 2008 to year y is the sum of the changes of these components:

$$\Delta C^{Total} = \Delta C^{elec} + \Delta C^{industrial} + \Delta C^{netExports} + \Delta C^{Met}, \quad (3)$$

where ΔC^{Total} is the change in total coal production from 2008 to year y , that is, $\Delta C^{Total} = C_y^{Total} - C_{2008}^{Total}$, and so forth for the other terms. We include changes in coal stocks, as well as any measurement error arising from differences in the data source, in the term containing domestic industrial steam coal. The units are millions of tons of coal.

The analysis of Sections 3 and 4 allows us further to decompose coal used for electricity generation into changes in relative prices (p), environmental regulations (x), RPS (r), heat rates (h), and electricity demand (E) as well as an unexplained component. Write coal consumed for electricity in year y as the sum of coal for electricity in each of the 50 states plus Washington D.C., that is, $C_y^{elec} = \sum_{i=1}^{51} C_{iy}^{elec}$, where coal tonnage consumed in state i in year y is the product of its share in generation, the heat rate in that state-year, and total generation in that state-year:

$C_{iy}^{elec} = g_{iy} h_{iy} E_{iy}$. The econometric model of Section 3 represents g_{iy} in terms of p_{iy} , x_{iy} , and r_{iy} , so that the estimated contributions of relative prices, environmental regulations, and RPSs enter through g_{iy} .

Because coal for electricity is a nonlinear function of the prices and the other determinants, we use a decomposition based on repeated conditional expectations; this approach specializes to the familiar linear decomposition in the case that the determinants enter additively. Let p_y denote the set of state prices $\{p_{iy}\}$ in year i and so forth. Now define v_y to be the unexpected component of coal electricity, given these determinants, that is, $v_y =$

$$C_y^{elec} - E(C_y^{elec} | p_y, x_y, r_y, h_y, E_y), \text{ so that (identically) } C_y^{elec} = E(C_y^{elec} | p_y, x_y, r_y, h_y, E_y) + v_y.$$

Thus the change in coal burned for electricity between 2008 and year y can be written as,

$$\Delta C^{elec} = \left[E(C_y^{elec} | p_y, x_y, r_y, h_y, E_y) + v_y \right] - \left[E(C_{2008}^{elec} | p_{2008}, x_{2008}, r_{2008}, h_{2008}, E_{2008}) + v_{2008} \right]. \quad (4)$$

This expression in turn can be expanded as the sum of differences of conditional expectations, changing one conditioning variable at a time:¹⁴

¹⁴ We make two technical notes concerning the decomposition (5). First, because of the nonlinearity, the ordering of the variables matters. Numerically, however, it turns out that changing the order of the variables makes a negligible difference to the decomposition (linearization provides a good approximation). Second, because of the logistic transformation in the shares model, the shares are not linear functions of the regression error, so in principle the

$$\begin{aligned}
\Delta C^{elec} = & \left[E\left(C_y^{elec} \mid p_y, x_y, r_y, h_y, E_y\right) - E\left(C_y^{elec} \mid p_{2008}, x_y, r_y, h_y, E_y\right) \right] \\
& + \left[E\left(C_y^{elec} \mid p_{2008}, x_y, r_y, h_y, E_y\right) - E\left(C_y^{elec} \mid p_{2008}, x_{2008}, r_y, h_y, E_y\right) \right] \\
& + \left[E\left(C_y^{elec} \mid p_{2008}, x_{2008}, r_y, h_y, E_y\right) - E\left(C_y^{elec} \mid p_{2008}, x_{2008}, r_{2008}, h_y, E_y\right) \right] \\
& + \left[E\left(C_y^{elec} \mid p_{2008}, x_{2008}, r_{2008}, h_y, E_y\right) - E\left(C_y^{elec} \mid p_{2008}, x_{2008}, r_{2008}, h_{2008}, E_y\right) \right] \\
& + \left[E\left(C_y^{elec} \mid p_{2008}, x_{2008}, r_{2008}, h_{2008}, E_y\right) - E\left(C_y^{elec} \mid p_{2008}, x_{2008}, r_{2008}, h_{2008}, E_{2008}\right) \right] \\
& + \left[v_y - v_{2008} \right].
\end{aligned} \tag{5}$$

The six terms in (5) respectively are the contributions to the change in coal for electricity of the change in prices, environmental regulations, RPSs, heat rates, electricity demand, and an unexplained component.¹⁵ The unexplained component encompasses residual modeling error and discrepancies in the heat rate identity because data come from different sources. These six terms, plus the three final terms in (3) – that is, industrial use, steam coal net exports, and metallurgical coal – comprise our nine-fold decomposition.

Standard errors for the estimated contributions are based on the sampling uncertainty of the predicted effects in the state-level shares regressions and are estimated using the block bootstrap resampling states.¹⁶

conditional expectation includes an adjustment for this nonlinearity. However, we found that this adjustment (the second order term in the Taylor series expansion of the conditional expectation) is numerically negligible, so the results here do not include that adjustment and are based on the leading term in the Taylor series expansion of the conditional expectation of the shares.

¹⁵ We add 6.4 million short tons to the environmental regulations component in 2015 and 2016, representing the contribution of MATS estimated in Section 4. We accordingly subtract 6.4 million short tons in 2015 and 2016 from the unexplained component to preserve the additive decomposition.

¹⁶ To implement the block bootstrap, we draw with replacement a random sample of 47 states from our baseline dataset and replicate our main analysis on this bootstrapped sample. We repeat this procedure 100 times, storing the estimated decomposition of coal production into components from each bootstrapped sample. We compute standard errors for each component as the standard deviation of the estimated contribution of the component across the 100 bootstrapped draws.

5.2 Results

The decomposition results are shown graphically in Figure 1, and numerical values are given in Table 4. Entries in Table 4 are the change in coal consumption from 2008 to the column year attributed to the source in the first column of each row. The most striking feature of the decomposition is the role played by the declining price of natural gas. Over the full period, of the 433 million ton decline in production, 397 million tons is attributable to a decrease in the price of gas relative to coal. This amounts to 92% (SE = 2.5 percentage points) of the reduction in coal production. In contrast, environmental regulations explain 6% (SE = 2.2 percentage points) of the decline in coal production, primarily driven by the effects of the CSAPR and MATS rules. The remaining seven factors contribute small, largely offsetting amounts to the change in coal production.

The main driver of the contribution of the air regulations is the MATS rule; absent the MATS rule, the point estimate of the effect of air regulations is counterintuitively positive. This positive effect is driven by the positive but statistically insignificant coefficient on the Clean Air Interstate Rule in Table 1, column (1). The near-zero coefficient on CAIR is consistent with the small effects that Glasgow and Zhao (2017) found on the reduction of pollution by the CAIR. If one sets the CAIR coefficient in Table 1 column (1) to zero, the contribution of air regulations to the coal decline from 2008 to 2017 is slightly larger, -34 million tons (8%) compared to -28 million tons (6%) in Table 4, and the relative price contribution drops to 391 million tons, while the other contributions remain unchanged (both in tons and percentage points).

	2008	2009	2010	2011	2012	2013	2014	2015	2016
Total Coal Production	1172	1075	1084	1096	1016	985	1000	897	739
<u>Change from 2008</u>									
Total	--	-97	-87	-76	-155	-187	-172	-275	-433
Electricity	--	-107	-64	-112	-219	-184	-190	-304	-367
Relative Prices	--	-166	-163	-194	-331	-242	-190	-321	-397
CAA Regulations	--	13	13	13	12	13	13	-29	-28
RPS	--	-3	-5	-7	-7	-7	-8	-9	-9
Electricity Demand	--	-44	-2	-10	-26	-20	-13	-27	-32
MWh/Ton	--	9	9	8	10	15	13	11	14
Other/unexplained	--	84	83	77	122	58	-4	70	85
Metallurgical Coal	--	-7	-1	-1	-1	-1	-1	-2	-6
Net Exports	--	-6	1	20	41	38	21	12	4
Industrial and Other	--	23	-23	17	24	-40	-2	20	-65

Notes: Estimates from column (1) of Table 1 are used to compute the six electricity components. The residual, unexplained coal demand for electricity production, is normalized to have mean zero over the 2002-2016 time period, leading to a non-zero average over the 2008-2016 period. *Source:* Authors' calculations

Table 4. Decomposition of changes in U.S. coal production, 2008-2016 (million short tons)

6. Conclusion

By using econometric methods to estimate the *ex-post* effects of changes in natural gas prices, environmental regulations, and several other factors, we are able to provide a detailed quantitative decomposition of the decline in coal production. We find that the decline in the price of natural gas relative to coal explains the majority of the decline in coal production, and that environmental regulations under the Clean Air Act had a modest impact on coal production, with other factors making small and offsetting contributions.

Looking forward, our results are consistent with stable coal production, assuming that the relative price of natural gas to coal remains low and stable and absent new regulations. This assessment is consistent with large-scale model projections by the EIA, which projects flat coal demand through mid-century (EIA, 2018b). The main caveat to this assessment is the possibility

of a decline in the prices of wind and solar generation, spurred by increasingly aggressive state-level RPSs, that is sufficient to make up for the anticipated sunset of wind and solar tax credits. Even though the growth of wind and solar were a modest factor over the period covered by this study, if they can be combined with inexpensive grid-scale storage they could turn out to be a key part of the ongoing transition in the power sector.

References

- Cullen, Joseph A., and Erin T. Mansur. 2017. “Inferring Carbon Abatement Costs in Electricity Markets: A Revealed Preference Approach Using the Shale Revolution.” *American Economic Journal: Economic Policy* 9 (3): 106–33.
- Energy Information Administration (EIA). 2017. Coal Data Browser. U.S. Department of Energy. Accessed March 9, 2018. <https://www.eia.gov/beta/coal/data/browser/>
- Energy Information Administration (EIA). 2018a. Annual Energy Outlook (AEO). U.S. Department of Energy.
- Energy Information Administration (EIA). 2018b. “EIA projects that U.S. coal demand will remain flat for several decades,” *Today in Energy*, March 30, 2018 at <https://www.eia.gov/todayinenergy/detail.php?id=35572>.
- Environmental Protection Agency (EPA). 2011. Regulatory Impact Analysis for the Final Mercury and Air Toxics Standards.
- Fell, Harrison, and Daniel T. Kaffine. 2018. “The Fall of Coal: Joint Impacts of Fuel Prices and Renewables on Generation and Emissions.” *American Economic Journal: Economic Policy* 10 (2): 90–116.
- Glasgow, Derek and Shuang Zhao. 2017. “Has the Clean Air Interstate Rule Fulfilled its Mission? An Assessment of Federal Rule-Making in Preventing Regional Spillover Pollution.” *Review of Policy Research* 34 (March): 186-207.
- Holladay, J. Scott, and Jacob LaRiviere. 2016. “The Impact of Cheap Natural Gas on Marginal Emissions from Electricity Generation and Implications for Energy Policy.” *Journal of Environmental Economics and Management* 85 (September): 205–27.

- Houser, Trevor, Jason Bordoff, and Peter Marsters. 2017. "Can Coal Make a Comeback?" Columbia Center on Global Energy Policy.
- Johnsen, Reid, Jacob LaRiviere, and Hendrik Wolff. 2016. "Estimating Indirect Benefits: Fracking, Coal and Air Pollution." Discussion Paper 10170. Institute for the Study of Labor (IZA).
- Jordan, Brett, Ian Lange, and Joshua Linn. 2018. "Coal Demand, Market Forces, and US Coal Mine Closures." Resources for the Future Working Paper 18-13. Washington, D.C.
- Knittel, Christopher R., Konstantinos Metaxoglou, and Andre Trindade. 2015. "Natural Gas Prices and Coal Displacement: Evidence from Electricity Markets." National Bureau of Economic Research Working Paper 21627.
- Linn, Joshua, and Kristen McCormack. 2017. "The Roles of Energy Markets and Environmental Regulation in Reducing Coal-Fired Plant Profits and Electricity Sector Emissions." Resources for the Future Report.
- Linn, Joshua, and Lucija Muehlenbachs. 2018. "The Heterogeneous Impacts of Low Natural Gas Prices on Consumers and the Environment." *Journal of Environmental Economics and Management* 89 (May): 1–28.
- Shankleman, Jess, and Hayley Warren. 2017. "Solar Power Will Kill Coal Faster Than You Think." Accessed June 15, 2017. <https://www.bloomberg.com/news/articles/2017-06-15/solar-power-will-kill-coal-sooner-than-you-think>
- U.S. Bureau of Land Management (BLM). 2017. "Waste Prevention, Production Subject to Royalties, and Resource Conservation; Delay and Suspension of Certain Requirements," 82 FR 46458-46475, Oct. 5, 2017.

U.S. Department of Energy (DOE). 2017. Staff Report to the Secretary on Electricity Markets and Reliability. Accessed March 9, 2018. <https://www.energy.gov/downloads/download-staff-report-secretary-electricity-markets-and-reliability>.

Appendix

I. Data Sources

Data	Period	Frequency	Geography	Source	Accessed By
Coal Production	2000-2016	Annual	USA	EIA, Form EIA-7A & 7000-2	EIA Coal Data Browser
Coal Consumption for Electricity Generation	2008-2016	Monthly	USA	EIA, Forms EIA-826, EIA-923, EIA-860, & EIA-861	EIA Electricity Data Browser
Coal Price for Electricity	2008-2016	Monthly	USA		
Natural Gas Price for Electricity	2008-2016	Monthly	USA		
Coal Imports and Exports	2008-2016	Annual	Country	EIA, US Census Bureau Monthly Report 545 & IM 145	EIA API
Electricity Generation: Total, Coal, & Natural Gas	2001-2016	Monthly	State	EIA	EIA API
Delivered Price: Coal & Natural Gas	2003-2016	Monthly	State	EIA, Electric Power Monthly	Electric Power Monthly
Environmental Regulations (binary)	2000-2016	Monthly	State	EPA Clean Air Markets Progress Reports	EPA Website
Environmental Reg. (continuous)	2000-2016	Monthly	State	EPA Air Markets Program Data	EPA AMPD tool
Electricity Capacity	2000-2016	Annual	Plant	Form EIA-860	EIA-860 Website
Planned Coal Plan Retirements	2010-2016	Annual	Plant	Form EIA-860	EIA-860 Website

Table A.1. Data Sources used in Econometric Analysis and Decomposition


II. Environmental Regulations

State	OTC NOx (S)	NBP (S)	CAIR Ozone (S)	CAIR NOx (A)	CAIR SO2 (A)	CSAPR Ozone (S)	CSAPR SO2/NOx (A)	RPS
Alabama		2004-2008	2009-2014	2009-2014	2010-2014	2015-2016	2015-2016	
Alaska								
Arizona								2000-2016
Arkansas			2009-2014			2015-2016		
California								2002-2016
Colorado								2004-2016
Connecticut	1999-2002	2003-2008	2009-2014					2000-2016
Delaware	1999-2002	2003-2008	2009-2014	2009-2014	2010-2014			2005-2016
Florida			2009-2014	2009-2014	2010-2014	2015-2016		
Georgia				2009-2014	2010-2014	2015-2016	2015-2016	
Hawaii								2001-2016
Idaho								
Illinois		2004-2008	2009-2014	2009-2014	2010-2014	2015-2016	2015-2016	2001-2016
Indiana		2004-2008	2009-2014	2009-2014	2010-2014	2015-2016	2015-2016	2011-2016
Iowa			2009-2014	2009-2014	2010-2014	2015-2016	2015-2016	2000-2016
Kansas							2015-2016	2009-2016
Kentucky		2004-2008	2009-2014	2009-2014	2010-2014	2015-2016	2015-2016	
Louisiana			2009-2014	2009-2014	2010-2014	2015-2016		
Maine								2000-2016
Maryland	1999-2002	2003-2008	2009-2014	2009-2014	2010-2014	2015-2016	2015-2016	2004-2016
Massachusetts	1999-2002	2003-2008	2009-2014					2000-2016
Michigan		2004-2008	2009-2014	2009-2014	2010-2014	2015-2016	2015-2016	2008-2016
Minnesota							2015-2016	2007-2016
Mississippi			2009-2014	2009-2014	2010-2014	2015-2016		
Missouri		2007-2008	2009-2014	2009-2014	2010-2014	2015-2016	2015-2016	2007-2016
Montana								2005-2016
Nebraska							2015-2016	
Nevada								2000-2016
New Hampshire	1999-2002							2007-2016
New Jersey	1999-2002	2003-2008	2009-2014	2009-2014	2010-2014	2015-2016	2015-2016	2001-2016
New Mexico								2002-2016
New York	1999-2002	2003-2008	2009-2014	2009-2014	2010-2014	2015-2016	2015-2016	2004-2016
North Carolina		2004-2008	2009-2014	2009-2014	2010-2014	2015-2016	2015-2016	2007-2016
North Dakota								2007-2016
Ohio		2004-2008	2009-2014	2009-2014	2010-2014	2015-2016	2015-2016	2008-2016
Oklahoma						2015-2016		2010-2016
Oregon								2007-2016
Pennsylvania	1999-2002	2003-2008	2009-2014	2009-2014	2010-2014	2015-2016	2015-2016	2004-2016
Rhode Island	1999-2002	2003-2014						2004-2016
South Carolina		2004-2008	2009-2014	2009-2014	2010-2014	2015-2016	2015-2016	2014-2016
South Dakota								2008-2016
Tennessee		2004-2008	2009-2014	2009-2014	2010-2014	2015-2016	2015-2016	
Texas				2009-2014	2010-2014	2015-2016	2015-2016	2000-2016
Utah								2008-2016
Vermont								2015-2016
Virginia		2004-2008	2009-2014	2009-2014	2010-2014	2015-2016	2015-2016	2007-2016
Washington								2006-2016
West Virginia		2004-2008	2009-2014	2009-2014	2010-2014	2015-2016	2015-2016	
Wisconsin			2009-2014	2009-2014	2010-2014	2015-2016	2015-2016	2000-2016
Wyoming								

Notes: This table summarizes the years during which each environmental regulation and Renewable Portfolio Standard (RPS) was in effect in each state. The Mercury and Air Toxics Standards (MATS) are omitted because they apply uniformly to all states starting in April 2015 (with discretionary compliance extensions granted until April 2016). For other regulations, (A) denotes annual and (S) denotes seasonal compliance requirements. Seasonal regulations are in force from May to September with two exceptions: (1) the 11 states that were covered by the NOx SIP Call's NOx Budget Trading Program (NBP) beginning in 2004 were given a compliance deadline of May 31, 2004 rather than May 1 due to litigation that delayed implementation; and (2) Rhode Island was covered under the NBP but not CAIR or CSAPR, and so the NBP was in effect in Rhode Island until it was repealed by the state effective July 29, 2014.

Table A.2. Coverage Period for Environmental Regulations and Renewable Portfolio Standards

III. Additional Specifications



Notes: This figure shows the coefficient on the relative price in our linear static specification, repeated for different choices of the window used to create a moving average of the relative price. Our baseline specification uses a six-month moving average. Bars denote 95% confidence intervals.

Figure A.1. Linear Coefficient on Relative Price for Different Moving Average Windows

	(1)	(2)	(3) Logit(Coal Share of Elec.)	(4)	(5)
Δp_{it}	-0.446*** (0.0861)	-0.450*** (0.0778)	-0.416*** (0.0738)	-0.403*** (0.0753)	-0.410*** (0.0705)
Δp_{it-1}	-0.583*** (0.105)	-0.557*** (0.0908)	-0.519*** (0.0857)	-0.501*** (0.0893)	-0.480*** (0.0829)
Δp_{it-2}	-0.581*** (0.103)	-0.566*** (0.0967)	-0.544*** (0.0953)	-0.520*** (0.0963)	-0.489*** (0.0871)
Δp_{it-3}	-0.531*** (0.0972)	-0.509*** (0.0947)	-0.493*** (0.0930)	-0.466*** (0.0910)	-0.432*** (0.0845)
Δp_{it-4}	-0.574*** (0.0959)	-0.553*** (0.0943)	-0.508*** (0.0897)	-0.486*** (0.0886)	-0.461*** (0.0853)
Δp_{it-5}	-0.612*** (0.0907)	-0.598*** (0.0896)	-0.563*** (0.0863)	-0.536*** (0.0852)	-0.500*** (0.0816)
Δp_{it-6}	--	-0.609*** (0.0874)	-0.580*** (0.0846)	-0.577*** (0.0875)	-0.545*** (0.0827)
Δp_{it-7}	--	-0.676*** (0.0936)	-0.640*** (0.0902)	-0.632*** (0.0917)	-0.589*** (0.0874)
Δp_{it-8}	--	-0.555*** (0.0906)	-0.477*** (0.0907)	-0.466*** (0.0901)	-0.441*** (0.0815)
Δp_{it-9}			-0.491*** (0.0874)	-0.456*** (0.0866)	-0.448*** (0.0861)
Δp_{it-10}			-0.551*** (0.0934)	-0.516*** (0.0923)	-0.482*** (0.0899)
Δp_{it-11}			-0.621*** (0.0944)	-0.585*** (0.0923)	-0.556*** (0.0896)
Δp_{it-12}				-0.669*** (0.101)	-0.639*** (0.100)
Δp_{it-13}				-0.687*** (0.108)	-0.660*** (0.108)
Δp_{it-14}				-0.667*** (0.109)	-0.611*** (0.103)
Δp_{it-15}					-0.571*** (0.0929)
Δp_{it-16}					-0.586*** (0.104)
Δp_{it-17}					-0.649*** (0.0992)
p_{it-6}	-0.759*** (0.0974)				
p_{it-9}		-0.767*** (0.0972)			
p_{it-12}			-0.775*** (0.0981)		
p_{it-15}				-0.780*** (0.0992)	
p_{it-18}					-0.792*** (0.101)
CAIR Dummy	0.0770 (0.0979)	0.0598 (0.0939)	0.0642 (0.0919)	0.0689 (0.0911)	0.0645 (0.0899)
CSAPR Dummy	-0.154 (0.127)	-0.171 (0.118)	-0.191 (0.114)	-0.196 (0.113)	-0.207 (0.111)
OTC NOx Dummy	--	--	--	--	--
NBTP NOx Dummy	0.0142 (0.0961)	-0.0287 (0.0910)	-0.00208 (0.0913)	0.0259 (0.0914)	0.0323 (0.0915)
RPS Dummy	-0.107 (0.103)	-0.0935 (0.108)	-0.0724 (0.111)	-0.0475 (0.117)	-0.0278 (0.125)
R-squared	0.935	0.938	0.940	0.940	0.942
N	6325	6105	5909	5727	5553

Notes: Dependent variable is the logistic transform of the state-month coal share in generation. Regressors are described in Section 3.1. Standard errors are two-way clustered by state (47 states) and time period (180 time periods). All regressions are estimated on 2002m1 – 2016m12. Observations in which the distributed lags include imputed coal or natural gas prices are excluded. For each specification, a joint F-test of the last three lags is conducted and the F-statistic and p-value are reported. Significant at the *5%, **1%, ***0.1% level.

Table A.3. Dynamic specification results, varying distributed lag length