

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

John Coglianesi

Department of Economics and Harvard Kennedy School, Harvard University

Todd Gerarden

Harvard Kennedy School, Harvard University

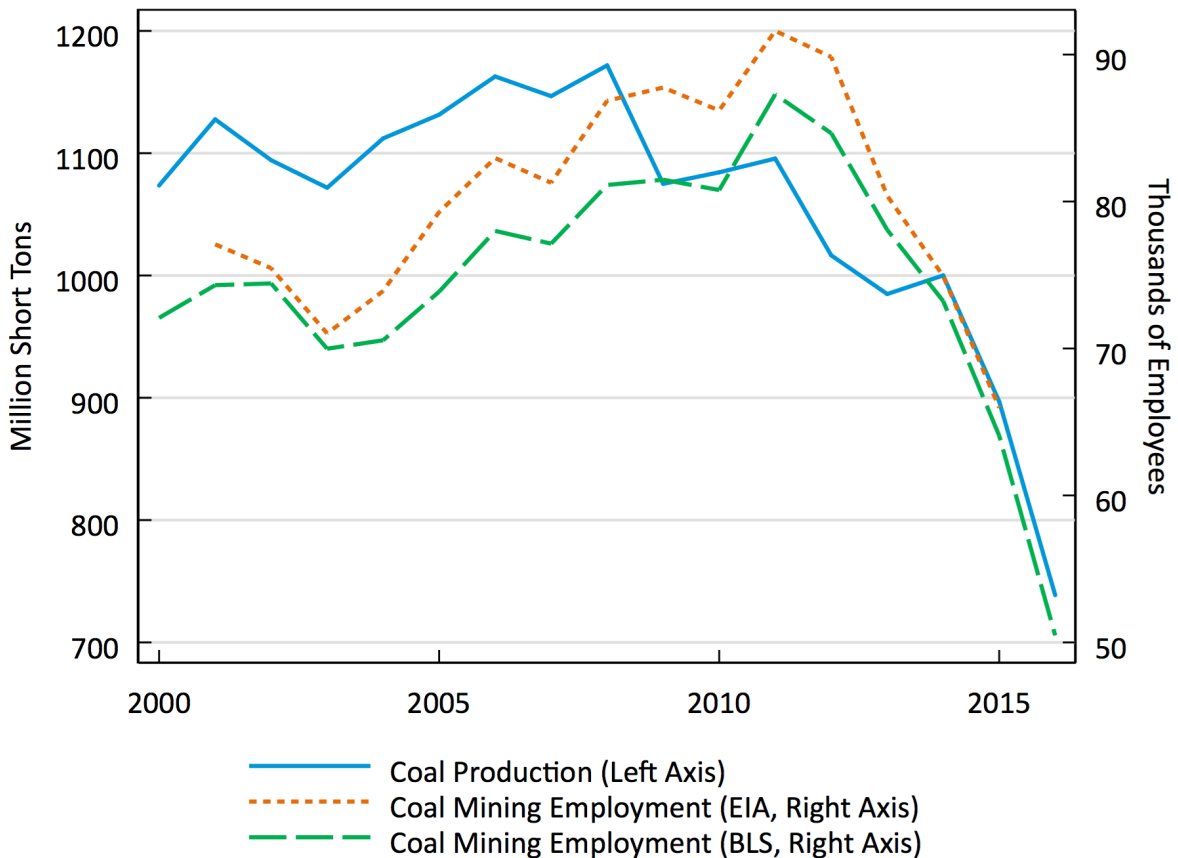
James H. Stock

Department of Economics and Harvard Kennedy School, Harvard University

June 15, 2017

1. Introduction and Summary

As is shown in Figure 1, between 2008 and 2015, U.S. coal production fell from 1,172 million tons to 897 million tons and coal employment fell from 87,000 to 66,000. In 2016, coal production declined further, to 739 million tons, 37% below its 2008 level. It is widely understood that a primary factor in this decline has been the sharp decline in natural gas prices, which has led to the substitution of natural gas for coal in electricity generation. In 2008, the national average price of natural gas delivered to an electricity generator nationally was 4.3 times the price of coal, on a Btu basis; by 2016, this relative price had fallen to 1.4 as a result of the development and spread of fracking. This national decline masks regional variation, with natural gas prices being even more competitive in some regions. For the first time, in 2016 electricity generated from gas overtook generation from coal.



Source: Energy Information Administration and Bureau of Labor Statistics. EIA mining employment data go through 2015.

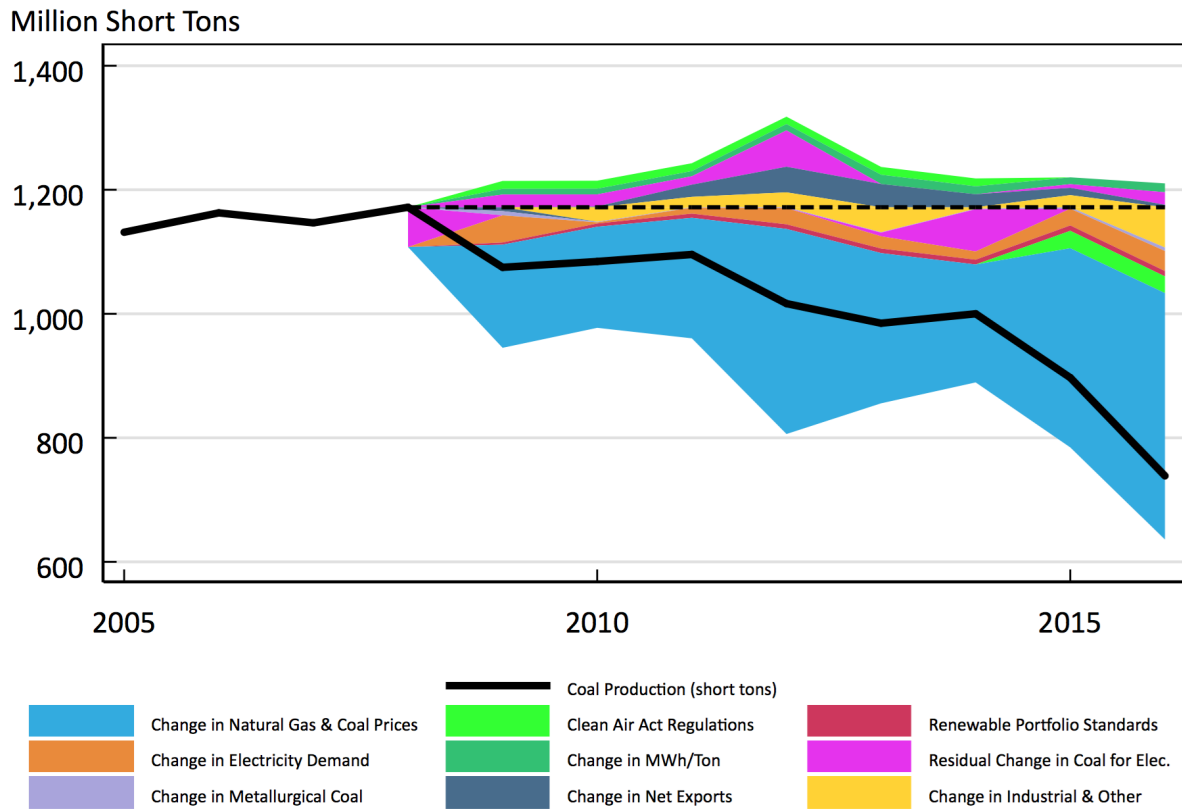
Figure 1. U.S. coal production and coal mining employment, 2002 – 2016.

The drop in gas prices was not the only relevant change in the market for coal in this period, however. These other factors include the proliferation and expansion of renewable portfolio standards (RPSs), new environmental rules, changes in overall electricity demand, and changes in demand for steam coal exports. In addition, there were fluctuations in demand for metallurgical coal, spurred by changing global economic conditions. Although all these factors are qualitatively relevant, surprisingly, it remains an open question how much of a quantitative role these factors played in the decline in demand for coal. Understanding why this decline occurred, in a quantitative sense, is important because it sheds light on the likely path for coal production going forward.

This paper undertakes a quantitative decomposition of the reasons for the decline in coal production from 2008 to 2016 into nine factors. Six of these factors explain the decline in domestic steam coal use for electricity, which accounts for the bulk of U.S. coal production: changes in the price of coal relative to gas, environmental regulations, renewable portfolio standards, heat rates, overall electricity demand, and an unexplained residual. The remaining three factors relate to the other uses of U.S. coal: changes in industrial use, net exports, and metallurgical coal. The decomposition combines three methodological approaches. First, we decompose coal's share of electricity generation into the effect of changing relative prices, air regulations (other than MATS), and renewable portfolio standards on coal's share of electricity generation. This decomposition is done econometrically using monthly data on generation, delivered fuel prices, regulations, and RPSs by state, allowing for different seasonal patterns across states. Second, the one air regulation that took effect over this period that does not have regional variation – and thus is not identified by variation across states – is the MATS rule. Here, we undertake a novel “event study” analysis of the electricity generating units (EGUs) that likely were closed as a result of this rule, and add the results in to the econometric decomposition. Because steam coal for electricity generation is the dominant use of coal in the United States and because total electricity generation was roughly constant over this period, the econometric decomposition plus the MATS event study accounts for most of the changes in coal demand. To complete the decomposition, we use an accounting approach to add in the relatively small amount (in tons) of production arising from changes in electricity consumption, in exports of steam coal, and in metallurgical coal demand.

Figure 2 presents the resulting decomposition of the decline in coal production tonnage from 2008 to 2016 (numerical values are presented in Section 5). According to our estimate, the declining price of natural gas relative to coal, on an energy-adjusted basis, explains 92% of the decline in coal production. An additional 10% of the decline is explained by environmental regulations, primarily the Cross-State Air Pollution Rule (CSAPR) and the Mercury and Air Toxics

Standard (MATS), which took effect in 2015 through 2016. The remaining seven factors contribute small, largely offsetting amounts to the change in coal prices. We attribute a small amount of the decline, 9 million tons, to the adoption of RPSs, while for the middle part of this period the growth in coal exports contributed positively. There is also a small unexplained component, which arises from measurement error from combining data from different surveys and the regression residual in the state-level econometric model of shares.



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

Source: EIA data and authors' calculations

Figure 2. Decomposition of changes in coal production, 2008-2016

The dominant role of the relative price of coal to natural gas in this decomposition informs the prospects for a rebound in the coal market. Because the major driver in the decline is the relative price of coal to natural gas, prospects for a rebound largely hinge on the path of this relative price. EIA (2017) projects natural gas prices rising slightly over the next decade in their reference scenario, rising to around \$5/MMBtu in 2040, well below 2008 prices. Some market developments and policy actions could affect these paths. Greater access to gas and coal

deposits on federal lands would tend to lower both coal and gas prices and the effect on relative prices is in principle ambiguous. However, because some of the reduction in coal consumption in favor of gas occurred through early retirements of coal units, we would not expect a reversal in the coal-gas price trend to result in a full reversal in coal demand. Similarly, because the air regulations were met either by installation of emissions controls or by closing plants, possible weakening of those rules in the future are unlikely to cause a reversal of their effects for generators now in compliance. EIA (2017) also projects total electricity demand to be flat over the next two decades under all their scenarios. Even if states scale back RPS goals, past RPSs have incentivized the construction of renewable generation and those facilities will remain. Taken together, our model and decomposition suggests that in the projected environment of stable relative prices of coal to gas, existing renewables penetration, and stable total electricity demand, the prospects for a rebound in coal production are slim. At the same time, barring price decreases of competitors – gas or renewables – our results are consistent with a plateau, not a continuing decline, in coal use, under current regulations. The two market threats to this plateau going forward would be further declines in the price of natural gas and potential declines in the price of renewables. Concerning renewables, Shankleman and Warren (2017) project substantial renewable price declines, but we note that the full cost for building new renewables must be able to compete with the marginal cost of coal in existing facilities for renewables to produce substantial penetration of the coal market without policy support. At the same time, we find evidence of high price sensitivity of the share as the prices approach energy parity, so further drops in the price of natural gas relative to coal could substantially reduce coal demand even further.

This paper is most closely related to three other papers in the literature. Knittel, Metaxoglou, and Trindade (2015) use plant-level panel data to identify the response of electricity generation from coal to changes in the relative price of coal to natural gas and then examine how this response varies by ownership and market type. Our analysis uses a similar methodology to measure responses to relative prices, but focuses on using these estimates to decompose changes in overall coal demand. Hauser, Bordoff, and Marsten (2017) decompose the decline in coal demand into the contributions of several factors including environmental regulations, measuring the effects of each regulation from EPA's ex ante estimate in the Regulatory Impact Analysis, and measuring gas-coal substitution by departures of those fuels from 2006 EIA projections; their analysis does not use coal or electricity prices. Our analysis uses ex post data to measure the effects of environmental regulations. Lastly, Culver and Hong (2016) argue that the lower fuel costs of natural gas and lower volatility of these prices mean that new power plants should be powered by natural gas instead of coal; they also discuss how regulations and prices have affected coal-fired generation from existing plants but do not provide any formal

empirical model or numerical decompositions. Our analysis examines changes in fuel usage across all EGUs, both new and existing.

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.

2. Data

2.1 Data sources

Parts of this analysis use aggregate time series data on the coal market. Data on annual coal production come from the Energy Information Administration (EIA) Coal Data Browser. Coal consumption for electricity generation, average coal prices for electricity, and average natural gas prices for electricity are taken from the EIA Electricity Data Browser. Quarterly domestic consumption of metallurgical coal was downloaded from the EIA Open Data API and aggregated to an annual frequency. We used data on quarterly exports and imports of coal by origin and destination country from the EIA Open Data API and aggregated to a national annual time series of net exports for the US.

We also construct panel data on electricity production and the coal market 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 Open Data API covering the period 2001-2016. Data on consumption of coal by 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 each fuel used by EGU. The heat rate is estimated as the ratio of a 12-month moving average of state coal consumption to a 12-month moving average of state generation from coal power.

The data contain some missing values on delivered fuel prices.¹ Although 18% of the prices were missing, those observations account for only 14% of coal generation. 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. The imputed price is the predicted value from a regression of

¹ EIA only reports state average delivered prices when sufficiently many EGUs reporting delivered prices that averages do not disclose confidential business information.

log prices on state and time fixed effects, so that 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 created panel data on environmental regulations at the state level. We construct indicators to capture whether individual Clean Air Act regulations affect a given coal generator 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:

- OTC NBP Seasonal NO_x
- NO_x Sip Call NBTP Seasonal NO_x
- CAIR Annual NO_x
- CAIR Seasonal Ozone
- CAIR Annual SO₂
- CSAPR Annual SO₂/NO_x
- CSAPR Seasonal Ozone
- MATS

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 combined all the CAIR rules into a single CAIR dummy variable, and we also combined 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.

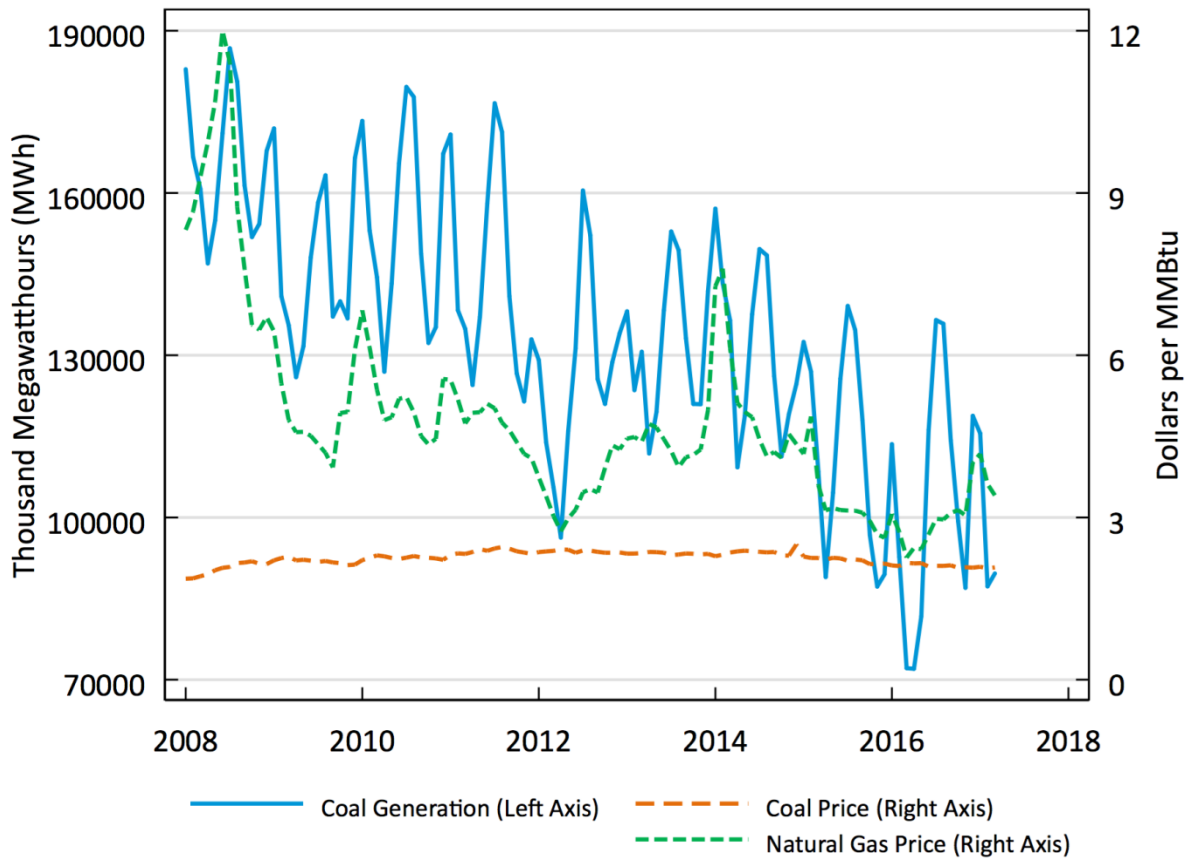
We incorporate data on Renewable Portfolio Standards (RPS) from the Database of State Incentives for Renewables and Efficiency (DSIRE) maintained by the NC Clean Energy

Technology Center. Ideally, we would like to have total generation in-state mandated by the RPSs; however in practice 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, which is a binary indicator for the presence of an RPS based on the date a state first established its RPS.

Finally, of a total of 8,429 observations, 89 observations have no coal-fired electricity generation. Most of these observations are for states (like Oregon, Washington, and California) in which the final coal-fired EGE was retired during the sample period. These 89 observations were dropped from the analysis. We also dropped states with no coal-fired generation over the entire period (Idaho, Rhode Island, Vermont, and D.C.).

2.2 Data description

Figure 3 shows total electricity generation from coal power plants in the US since 2008 (left axis) as well as the national average delivered prices of coal and natural gas used for electric power (right axis). Coal generation fell nearly 40% over this time period. It is also highly seasonal, with seasonal peaks for summer cooling and winter heating. Although coal prices were fairly stable, natural gas prices fell by almost two-thirds between 2008 and 2016.



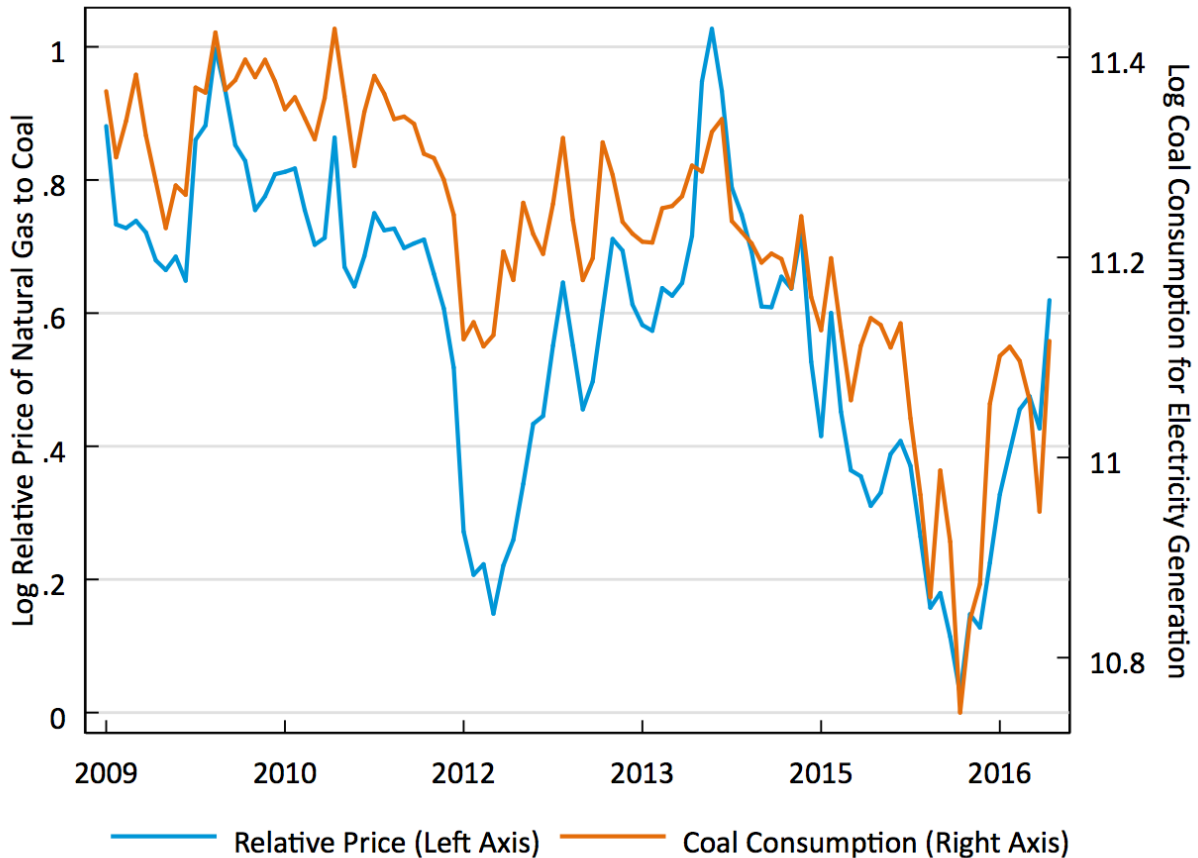
Source: EIA

Figure 3. Electricity generation from coal and the delivered prices of coal and natural gas (monthly for the U.S.)

The strong seasonality in coal generation obscures the relation between coal generation and gas prices, or relative prices of coal to gas. Figure 4 therefore 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.² 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

² 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.

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: EIA and authors' calculations

Figure 4. 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

The panel data regressions estimate the response of the share of electricity generation by coal to the relative price of coal to gas (on an energy content basis), to environmental regulations, and to 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. This does not ensure linearity of the resulting specification and we therefore examine possible nonlinearities, in particular nonlinearities in the relationship to relative prices. In addition, we want to allow for the possibility of adjustment lags, so that the effect of a change in 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 (in our baseline specification) by using the log relative price of multi-month moving averages of the fuel prices, and second by using a distributed lag of the monthly fuel prices. We refer to these respectively as the static and the dynamic specifications.

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

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

where $\alpha_{im(t)}$ are state-by-calendar-month fixed effects, 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 Renewable Portfolio Standards were in effect. In the base specification, p_{it}^{MA} is a 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 $p_{it} = \ln(P_{it}^{coal} / P_{it}^{gas})$ 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. This is a parametric (interaction) version of the partially linear regression model.

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 fracked 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. That said, transient variations in regional demand (e.g. a particularly cold winter) would plausibly lead to regional price variations so coal, or gas, or relative prices are, strictly speaking, not exogenous. Resolving this challenge would require a plausibly exogenous instrument. There are credible sources of exogenous shifts in supply, for example related to fracking technology development, rail transport price spikes for coal (as there were in 2012-2014 as rail transport of oil increased), or regional variation in successful gas exploration wells. For our purposes, however, the demands for such instruments are high because they should map out regional variation, which would allow us to identify the effects of environmental regulations and RPS. But developing such instruments is beyond the scope of this paper.

3.2 Results

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

	(1)	(2a)	(2b)	(3)	(4a)	(4b)	(5)	(6)
	Logit(Coal Share of Elec.)							
p_{it}^{MA}	-1.329***	--	--	-1.425***	--	--	-0.724***	-1.405***
	(0.181)	--	--	(0.283)	--	--	(0.0954)	(0.188)
$p_{it}^{MA,2}$	-0.304***	--	--	-0.314**	--	--	--	-0.334***
	(0.0769)	--	--	(0.108)	--	--	--	(0.0752)
Dynamic Coefficients for Above/Below Median Relative Price	--	Below Median	Above Median	--	Below Median	Above Median	--	--
Δp_{it}	--	-0.357***	-0.939***	--	-0.371***	-0.818***	--	--
	--	(0.0528)	(0.164)	--	(0.0570)	(0.204)	--	--
Δp_{it-1}	--	-0.441***	-1.091***	--	-0.465***	-1.035***	--	--
	--	(0.0583)	(0.172)	--	(0.0699)	(0.219)	--	--
Δp_{it-2}	--	-0.505***	-1.031***	--	-0.523***	-0.988***	--	--
	--	(0.0627)	(0.184)	--	(0.0773)	(0.229)	--	--
Δp_{it-3}	--	-0.514***	-0.869***	--	-0.582***	-0.847***	--	--
	--	(0.0708)	(0.143)	--	(0.0879)	(0.196)	--	--
Δp_{it-4}	--	-0.543***	-0.899***	--	-0.609***	-0.867***	--	--
	--	(0.0728)	(0.137)	--	(0.0953)	(0.183)	--	--
Δp_{it-5}	--	-0.639***	-0.824***	--	-0.685***	-0.808***	--	--
	--	(0.0822)	(0.124)	--	(0.0996)	(0.188)	--	--
Δp_{it-6}	--	-0.592***	-0.907***	--	-0.629***	-0.931***	--	--
	--	(0.0803)	(0.129)	--	(0.102)	(0.203)	--	--
Δp_{it-7}	--	-0.675***	-0.975***	--	-0.703***	-0.988***	--	--
	--	(0.0931)	(0.141)	--	(0.111)	(0.208)	--	--
Δp_{it-8}	--	-0.583***	-0.877***	--	-0.688***	-0.909***	--	--
	--	(0.0929)	(0.145)	--	(0.107)	(0.218)	--	--
p_{it-9}	--	-0.800***	-1.081***	--	-0.896***	-1.082***	--	--
	--	(0.0962)	(0.150)	--	(0.135)	(0.226)	--	--
CAIR Dummy	0.0245	-0.00991		0.0986	0.0998		0.0492	0.0261
	(0.0954)	(0.0983)		(0.103)	(0.105)		(0.0983)	(0.0981)
CSAPR Dummy	-0.0992	-0.179		0.00296	-0.0544		-0.190	-0.0771
	(0.127)	(0.114)		(0.173)	(0.160)		(0.126)	(0.135)
OTC NOx Dummy	--	--		0.395*	0.446*		--	--
	--	--		(0.180)	(0.174)		--	--
NBTP NOx Dummy	-0.0786	-0.0930		0.0237	0.0391		-0.00244	-0.145
	(0.0898)	(0.0777)		(0.126)	(0.127)		(0.104)	(0.0840)
RPS Dummy	-0.183	-0.117		-0.278**	-0.235**		-0.129	-0.188*
	(0.0956)	(0.105)		(0.0802)	(0.0828)		(0.102)	(0.0920)
Includes Imputed Prices	No	No	No	Yes	Yes	Yes	No	No
Moving Average Length	6 Months	--	--	6 Month	--	--	6 Months	3 Months
N	6325		6037	8340		8340	6325	6585

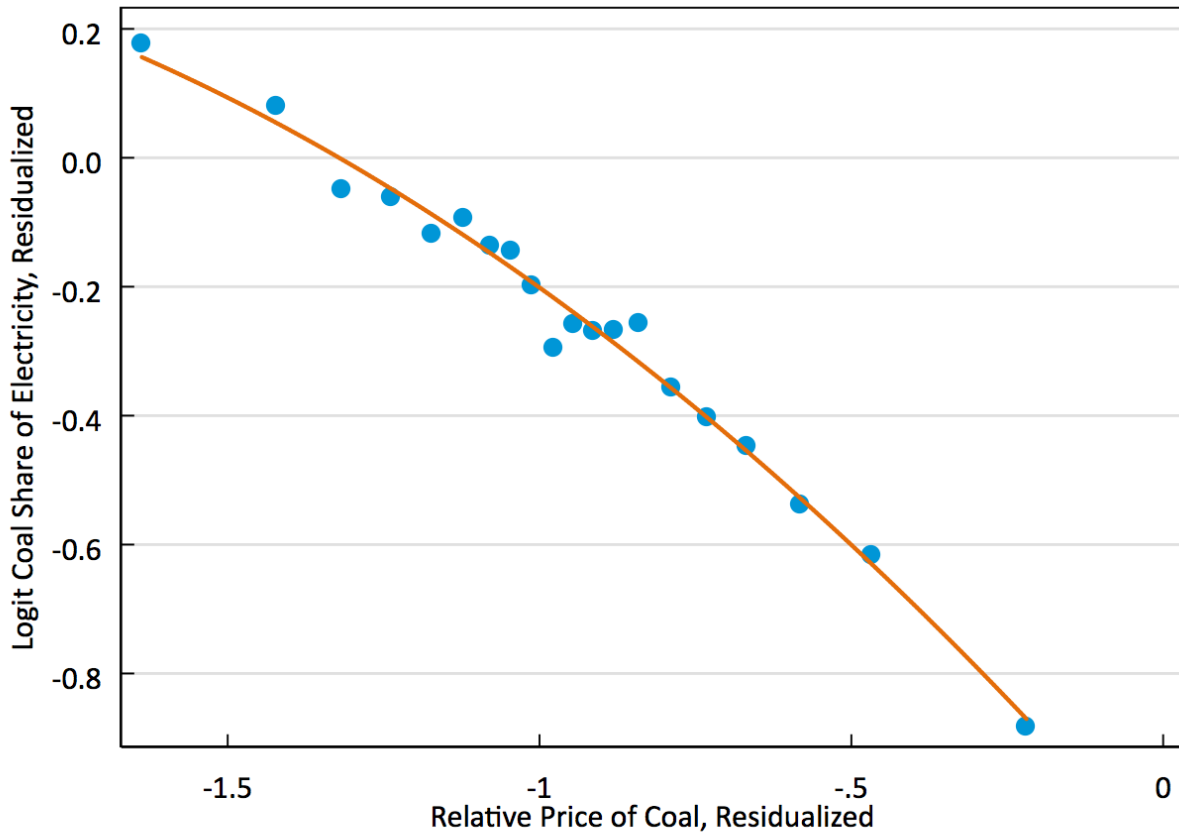
Notes: Dependent variable is the logistic transform of the state-month coal share in generation. Regressors are described in Section 3.1. Standard errors are two-way clustered by state and time period. All regressions are estimated on 2002m1 – 2016m12. Significant at the *5%, **1%, ***0.1% level.

Table 1. Panel regression results, shares regressions

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) and (2), the coefficients for RPS, CSAPR, and NBTP NOx are all estimated to reduce the coal share. Although the baseline regressions estimated that CAIR increased 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.

Second, as illustrated in Figure 5, 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, higher order polynomial

terms do not substantially change the estimated fit. Restricting to a linear specification, as in regression (5), 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.

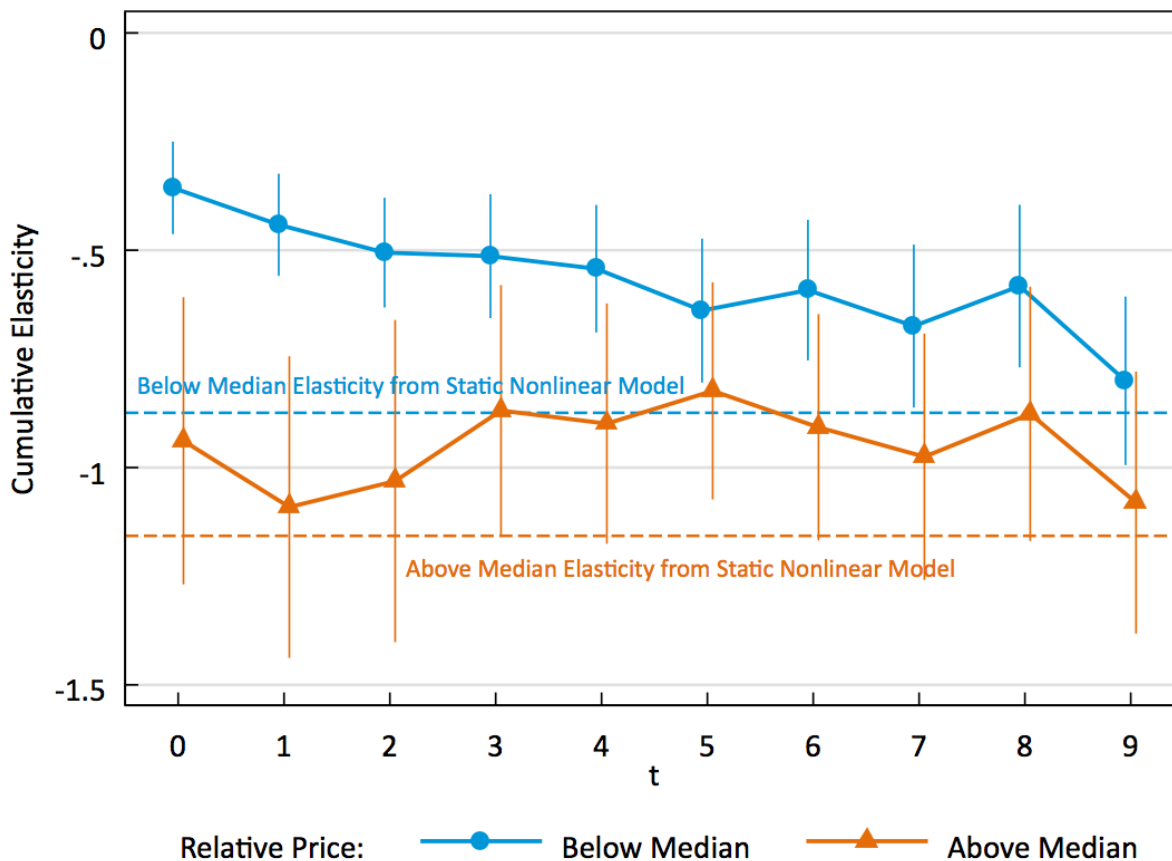


Notes: Both variables are residualized on month-by-state fixed effects and regulatory dummies before plotting. Source: authors' calculations.

Figure 5. 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 6 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 two-regime dynamic specification. We interpret this figure as providing support for using the static specification for the decomposition.



Notes: The dynamic responses correspond to columns (2a) and (2b) 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 (2a) and (2b).

Figure 6. Cumulative dynamic response of coal share (logistic transform) to a 1 percentage point change in the coal-gas relative price.

Fourth, the dynamic and static regressions give similar predictions for changes in coal demand, and both provide good fits to the state-level share data. Figure 7 shows the predicted values from the static regression and from the dynamic regression; both are fit over the full sample (not split-sample for the linear specification) 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, respectively, New York and West Virginia, and in particular occasionally produce predicted shares outside their 0-1 range. The fit of the dynamic two-regime model 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 too supports the use of the static specification.

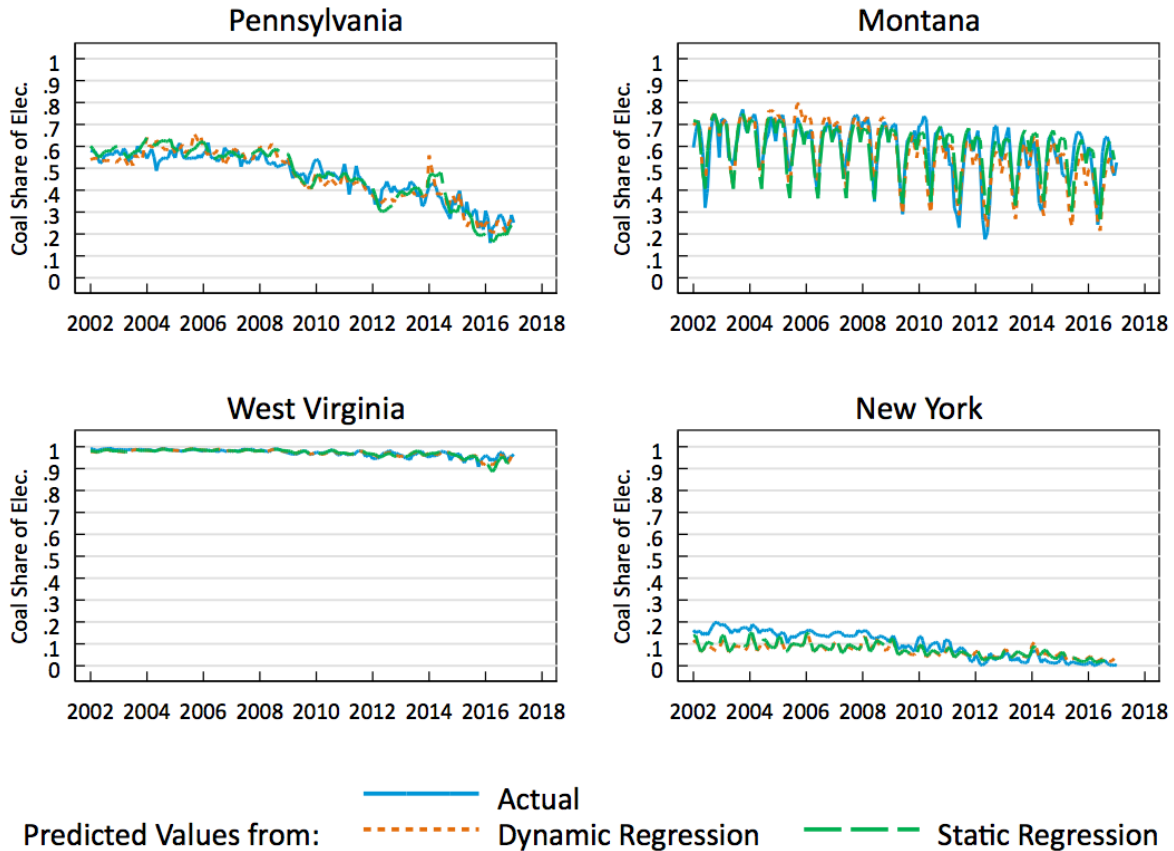


Figure 7. Actual and predicted shares for four representative states: static and dynamic specifications (full-sample estimates).

4. MATS Event Study

The MATS rule regulated emissions of toxic air pollutants including mercury, arsenic, and heavy metals from coal- and oil-fired power plants. The rules set out technology-based standards. EPA expected that it would be economically more cost-effective to retire some plants than to implement new technology. In its Regulatory Impact Analysis, EPA estimated that 4.7 GW of

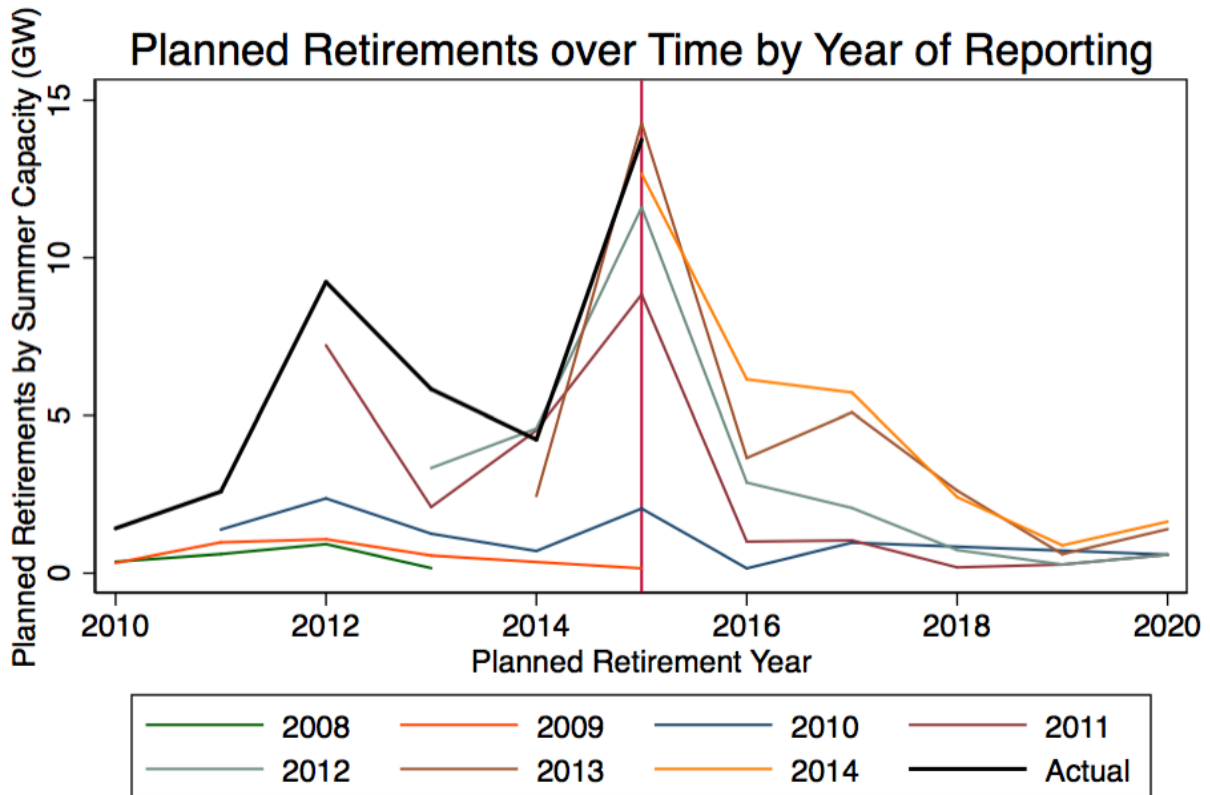
capacity would be retired for MATS compliance and that power sector coal consumption in 2015 would fall by one percent as a result of MATS (EPA, 2011).

The MATS rule took effect in 2015. Because the MATS rule applied nationally, a dummy variable indicating the MATS rule is not separately identified from time effects including national mean changes in prices, so estimating the effect of MATS on coal generation share is not amenable to the regression methods of the previous section. In this section, we therefore take advantage of a convenient institutional feature of the MATS rulemaking to undertake an event study that estimates plant closings as a result of MATS. Given this plant closings estimate, we estimate the subsequent effect of these closings on coal consumption.

The MATS rule had been in train for nearly two decades. 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.

Using data from the EIA, we are able to estimate the impact of the MATS rule promulgation on planned plant closings. Conveniently, the EIA data (EIA Form 860) are collected by a survey of EGU owners, with a deadline of 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, which had not yet been proposed, while forms filled out in early 2012 would take into account the finalized MATS rule. Among other things, EIA Form 860 asks for whether a unit is planned to be retired and, if so, when the retirement is planned to occur. Because the final rule included a compliance schedule, with final compliance in March 2015, changes in EIA Form 860 between 2011 and 2012 for retirements planned to occur in 2015 can plausibly be associated with the MATS rule. Some of these retirements might be expedited planned retirements, while others might be newly planned retirements.

Figure 8 presents aggregate planned retirements of coal capacity over time, by year of reporting. As the figure shows, there is a spike in retirements in 2015, the reporting year. Much of that spike is associated with retirements that were not planned in the 2010 data, but were planned in the 2011 data; that is, retirement decisions that occurred between February 2011 and February 2012.



Sources: EIA-860 and EIA-923. EIA-860 reporting deadline is last day of February in year following data year (e.g., February 2012 for 2011 data).

Figure 8. Planned coal plant retirements by year of retirement, for different EIA 860 reporting periods

There was also an increase in 2014 planned retirements between the 2010 and 2011 data, 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.³ Rather, a plausible explanation for these newly announced retirements for 2014 and 2016 is the decline in gas prices that occurred during 2011, the period of the first large drop in gas prices (Figure 3).

³ 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 immediately following the release of the MATS formal rule.

We therefore estimate the additional MATS-related retirements as the difference between the increase in the 2015 planned retirements, minus the average increase in 2013, 2014, and 2016 planned retirements. This estimated decline in nameplate capacity is shown in Table 2. Because the individual plants are identified in the Form 860 data, along with their coal use in the reporting year, we can compute the decline in coal demand directly associated with these plant closures. Note that this decline, also shown in Table 2, is arguably an overestimate of the reduction in coal demand, because it does not account for a potential increase in coal use in non-retired units to make up for the lost generation at the retired units.

A. Nameplate Capacity (GW)						
Survey year	Retirement year					
	2012	2013	2014	2015	2016	
2010	2.7	1.4	0.8	2.4	0.2	
2011	7.5	2.4	5.0	9.6	1.1	
Difference, 2011-2010	4.8	1.0	4.3	7.2	1.0	
non-MATS estimate				2.1		
MATS estimate				5.1		

B. Coal Consumption (million short tons)						
Survey year	Retirement year					
	2012	2013	2014	2015	2016	
2010	4.6	2.3	1.2	4.6	0.3	
2011	11.2	3.0	8.0	14.1	2.6	
Difference, 2011-2010	6.6	0.7	6.8	9.4	2.3	
non-MATS estimate				3.3		
MATS estimate				6.2		

Source: EIA Form 860 and authors' calculations.

Table 2. Event-study estimate of MATS retirement effect on nameplate capacity and direct coal consumption.

As can be seen in Table 2, the event-study estimate of the effect of the MATS rule is a reduction in nameplate capacity of 5.1 GW, very slightly greater than the EPA's estimate, and an associated reduction in coal demand of 6.2 million tons.

5. Decomposition

In this section, we decompose the decline of coal production in the US since 2008 into several components, building on results from the previous two sections.

5.1 Methods

Our decomposition starts by decomposing the consumption of coal in year y into coal used for domestic electricity generation, domestic industrial steam coal use, net steam coal exports, and metallurgical coal. Thus the change in coal consumption 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 consumption from 2008 to year y , that is, $\Delta C^{Total} = C_y^{Total} - C_{2008}^{Total}$, and so forth for the other terms. We ignore changes in coal stocks and thus equate annual consumption and annual production. 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 further represents g_{iy} in terms of p_{iy} , x_{iy} , and r_{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, $v_y = C_y^{elec} - E(C_y^{elec} | p_y, x_y, r_y, h_y, E_y)$, 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:⁴

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

The six terms in (5) respectively are the contributions to the change in coal for electricity of the change in prices, environmental regulations, RFSs, heat rates, electricity demand, and an unexplained component.⁵ 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.

5.2 Results

The decomposition results are shown graphically in Figure 2, and numerical values are given in Table 3.

⁴ We make two technical notes concerning the decomposition (5). First, because of the nonlinearity, the ordering of the variables matters in theory. Numerically, however, changing the order of the variables makes a negligible change in the quantitative decomposition. Second, because of the logistic transformation in the shares model, the shares are not linear functions of the regression error, so in principle the 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.

⁵ We add 6.2 million short tons to the environmental regulations component in 2015 and 2016, representing the contribution of MATS estimated in Section 4. We accordingly subtract 6.2 million short tons in 2015 and 2016 from the unexplained component to preserve the additive decomposition.

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 – 92% of the reduction – is attributable to cheaper gas relative to coal. In many years, the amount of the decline attributable to cheap gas more than explains the overall decline, because of offsetting factors. Most notably, exports grew in 2011-2013, partially offsetting the effect of declining gas prices during those years. A decline in overall electricity demand makes a negative contribution to the change, although inspection of Table 3 indicates that the main decline in overall electricity demand occurred during the recession (2009), and it has stayed below its 2008 value for remaining period. RPSs made a small negative contribution.

The econometric model estimates a small and statistically insignificant coefficient on the CAIR regulation, which took effect in 2009; counterintuitively, this insignificant coefficient is positive, indicating a positive contribution, relative to 2008. We view this as a consequence of the difficulty of estimating the regulatory impacts using state variation, combined with at least the CAIR regulation having a relatively small effect.

From 2014 to 2016, the contribution of the air regulations is estimated to have contributed a decline in coal demand of 40 million tons; this is the combined effect of CSAPR and the MATS regulations. We view the statistically insignificant positive coefficient on CAIR as anomalous, and treat this 40 million ton decline as the effect of air regulations over this period. This 40 million ton decline constitutes 3.4% of total coal production in 2008, or 9.2% of the decline from 2008 to 2016.

Million Short Tons

	2008	2009	2010	2011	2012	2013	2014	2015	2016
Total Coal Production	1172	1075	1084	1096	1016	985	1000	897	739
Total Change from 2008	--	-97	-87	-76	-155	-187	-172	-275	-433
Met Coal	--	-7	-1	-1	-1	-1	-1	-2	-6
Net Exports	--	-6	1	20	41	38	21	12	4
Industrial and Other	--	0	-23	17	24	-40	-2	20	-65
Electricity	--	-171	-128	-176	-284	-248	-254	-368	-431
Relative Prices	--	-166	-163	-194	-331	-242	-190	-321	-397
CAA Regulations	--	13	13	13	12	13	13	-28	-27
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	--	21	19	14	59	-6	-68	6	20

Source: Authors' calculations

Table 3. Decomposition of changes in coal production since 2008

References

- Culver, Walter J., and Mingguo Hong. 2016 "Coal's decline: Driven by policy of technology?" *The Electricity Journal* 29(7) 50-61.
- Energy Information Administration (EIA). 2017. Annual Energy Outlook (AEO). U.S. Department of Energy.
- Environmental Protection Agency (EPA). 2011. Regulatory Impact Analysis for the Final Mercury and Air Toxics Standards.
- Houser, Trevor, Jason Bordoff, and Peter Marsters. 2017. "Can Coal Make a Comeback?" Columbia Center on Global Energy Policy.
- Knittel, Christopher R., Konstantinos Metaxoglou, and Andre Trindade. 2015. "Natural Gas Prices and Coal Displacement: Evidence from Electricity Markets". *NBER Working Paper*.
- 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>