RAND Journal of Economics
Vol. 00, No. 0, xxxx 2019
pp. 1–35

The roles of energy markets and environmental regulation in reducing coal-fired plant profits and electricity sector emissions

Joshua Linn*

and

Kristen McCormack**

Between 2005 and 2015, US electricity sector emissions of nitrogen oxides and sulfur dioxide, which harm human health and the environment, declined by two thirds, and many coal-fired power plants became unprofitable and retired. Intense public controversy has focused on these changes, but the literature has not identified their underlying causes. Using a new electricity sector model of the US eastern interconnection that accurately reproduces unit operation, emissions, and retirement, we find that electricity consumption and natural gas prices account for nearly all the coal plant profitability declines and resulting retirements. Environmental regulations had little effect on these outcomes.

1. Introduction

■ Electricity sector emissions of nitrogen oxides (NO_x) harm human health and the environment by raising ambient concentrations of ozone and particulates. The United States began regulating electricity sector NO_x emissions in the 1970s, and emissions declined gradually and steadily from then until around 2000, after which emissions declined sharply. Between 2000 and 2015, emissions declined at a rate four times greater than between 1990 and 2000, and emissions in 2015 were just one fifth of 1990 emissions. Coinciding with these changes are the tightening stringency and broadening scope of NO_x emissions caps that the US Environmental Protection Agency (EPA) administers. Likewise, emissions of other pollutants, such as sulfur dioxide, have declined dramatically since 2000. During the same period, many coal-fired plants became unprofitable and about one third of coal-fired plants prepared to retire.

In the political debate over electricity sector policy, two views have emerged about the cause of the decline in electricity sector emissions and the retirement of coal-fired plants. The first

* University of Maryland and Resources for the Future; linn@rff.org.
** Harvard University; kmccormack@g.harvard.edu.

view credits technological innovation and pro-renewables policies for reducing costs of natural gas-fired plants and renewables and causing a shift from coal to lower-emitting sources. Many adherents of this view favor tightening emissions caps and other regulations in light of their benefits and lower-than-expected compliance costs. The second view is that by raising the costs of coal-fired power plants relative to other technologies, emissions regulations have excessively harmed coal-fired plant profits, jobs, local communities, and the reliability of electricity supply. Some adherents of this view call for weakening regulations to end the "war on coal." These two views crystalized during the 2016 presidential election, and favoring the second view, the Trump administration has begun to weaken regulations on the coal sector.

The economics literature suggests that on the margin, low natural gas prices reduce coal-fired generation, but it does not provide direct evidence for the aggregate effects of natural gas prices and other market forces and does not identify the causes of the retirements. Several recent articles examine the statistical relationships among natural gas prices, wind generation, fossil fuel-fired generation, and emissions (e.g., Cullen and Mansur, 2017; Holladay and LaRiviere, 2017; Fell and Kaffine, 2018; Linn and Muehlenbachs, 2018; Johnsen, LaRiviere, and Wolff, forthcoming). However, because of their reduced-form approach, these articles focus only on the short-run and marginal effects of natural gas prices and wind generation; long-run responses may differ. The long run includes entry and exit decisions that depend on fixed costs, whereas short-run responses depend only on variable costs. For example, low gas prices may lead to an increase in natural gas plant investment, potentially compounding the short-run effects on coal plant profits and retirements. In addition, when emissions caps are binding, low natural gas prices may reduce emissions credit prices, lowering costs for coal plants and opposing the short-run effects of gas prices on coal plant profits. Thus, the long-run effects of market shocks may differ positively or negatively from the short-run effects that have been the focus of the literature, leading to different conclusions about the historical effects of these shocks.[1]

Moreover, although several studies have compared expected and realized costs of sulfur dioxide emissions reductions under the Acid Rain Program (ARP) (e.g., Carlson et al., 2000; Ellerman et al., 2000), the literature has not compared the effects of market forces with the effects of regulation. Most electricity sector NO_x emissions in the East are covered by EPA emissions caps. When they were established, the EPA expected them to cost the sector at least \$3 billion per year (2005 dollars).[2] These costs fall largely on the coal-fired fleet (Linn, 2010), implying roughly a 10% cost increase for those units, with most of the costs incurred by the older and higher-emitting coal-fired units. The Mercury and Air Toxics Standards (MATS) were expected to cost substantially more than the NO_x caps. Thus, environmental regulation may have substantially reduced profitability of coal-fired units, leading to their retirement. Fowlie and Muller (2013) analyze the costs of achieving NO_x emissions caps, but they consider only the early portion of the program and do not evaluate whether either view of the electricity sector trends is correct. We are not aware of any *ex post* analysis of MATS.

In this article, we use a new computational operational and investment model of the electricity system and quantify the effects of market shocks and emissions regulations on emissions, profits, and retirements of coal-fired plants. We focus on the eastern United States, which accounts for about 90% of electricity sector NO_x emissions.[3]

[1] Houser, Bordoff, and Marsters (2017) compare the effects of electricity consumption, natural gas prices, and renewables on coal consumption between 2006 and 2016 and conclude that natural gas prices were the most important factor, followed by electricity consumption. They do not analyze the effects of these factors on emissions or coal plant profits. The Department of Energy (DOE) (2017) argues that natural gas prices are the most important factor explaining coal plant retirements but provides little evidence.

[2] These costs represent a large share of overall estimated costs of federal environmental regulations of the electricity sector. Between 2003 and 2015, the EPA implemented the emissions caps in three phases (see Section 2). The agency reports costs of complying with each phase (EPA, 1998, 2005, 2011). For the latter two phases, the costs are combined with the costs of achieving the sulfur dioxide caps. In the main text, we use only the cost estimate from EPA (1998).

[3] The United States has three major interconnections, across which there is little available transmission. Throughout the article, East refers to the eastern interconnection, which spans the Great Plains to the East Coast (see Figure A1).

The model includes 3500 generation units in the eastern United States and characterizes unit construction, retirement, emissions abatement, and hourly operation. We approximate uncertainty in consumption, uncertainty in unit availability, and constraints on unit operation by extending the approach of Davis and Hausman (2016). The model accurately predicts observed hourly operation and emissions and coal plant retirements. We show that a conventional economic dispatch model, which is constructed using the same underlying data but omits these features, would overpredict the effects of changing natural gas prices on coal-fired plants.

We model three market shocks: natural gas prices, renewables generation, and electricity consumption. Largely because of the rise in production from shale formations, natural gas prices in 2015 were 30% lower than the level projected in 2005. Improved wind generator performance and subsidies caused wind generation in 2015 to be 10 times higher than had been expected. Because of the 2008–2009 economic recession and other factors, 2015 electricity consumption was 20% below 2005 expectations. For convenience, we refer to differences between the 2005 projections and the 2015 realized outcomes as energy market shocks, noting that policies have contributed to them.

We model two environmental regulations: NO_x emissions caps that were adopted between 2005 and 2015, and MATS.[4] The emissions caps require that 2015 emissions be about half of 2005 levels. MATS requires plants to reach specific emissions standards for mercury and other pollutants. Using 2005 projections for electricity consumption, wind generation, and fuel prices, we estimate that without shocks, NO_x abatement costs would have been about $2.9 billion per year, which roughly agrees with *ex ante* EPA assessments.[5] Note that we model explicitly the compliance decisions for the NO_x abatement caps, whereas we estimate MATS costs based on observed decisions. We make this distinction because the caps applied throughout the period of analysis, 2005 through 2015, whereas the initial compliance period of MATS falls near the end of the period, after the market shocks occurred.

The market shocks explain 80% of the coal-fired plant retirements observed between 2005 and 2015.[6] After accounting for these shocks, the emissions caps had a small effect on coal-fired plant profits and retirements. The three shocks collectively reduced regulatory costs from $2.9 billion to $0.4 billion per year (86%) and reduced coal-fired plant profits by 89%. We find that, after accounting for the shocks, MATS had a small effect on retirements and profits. These results confirm the first of the two views, that factors other than environmental regulation explain most of the decline in the profits of coal-fired plants and the resulting retirements.

[4] We do not model sulfur dioxide emissions caps because emissions credit prices were close to zero in 2015, indicating that the caps were not binding. The EPA's Clean Power Plan, finalized in 2015, established carbon dioxide emissions standards for fossil fuel-fired generators. Linn, Burtraw, and McCormack. (2016) see two reasons why it is unlikely to have caused any coal plant retirements in 2015: (i) it would not have taken effect until 2022, and therefore would not have affected power plant operational decisions in 2015, and (ii) compliance decisions for MATS were made prior to 2015. The only plants that would retire in 2015 because of the Clean Power Plan are those for which the firm needed to make a life-extending investment in 2015, and which would have been unprofitable under the Clean Power Plan. Our data do not appear to contain any such plants.

[5] We cannot compare our estimated costs directly with EPA estimates. Although the agency reports costs of complying with each of three regulatory phases, the costs are estimated relative to different baselines, making it inappropriate to add the three cost estimates.

[6] There is some disagreement about retirements in the Energy Information Administration (EIA) and EPA data. According to the Energy Information Administration (EIA), between 2005 and 2015, firms announced the retirements of about 88 gigawatts (GW) of coal-fired plant capacity. Of this amount, about 20 GW did not operate in 2005, according to EPA data. According to EPA data, about 41 gigawatts stopped operating by 2015. Much of the remaining 27 GW of retirements in the EIA data appears to have stopped operating since 2015. Because we model market shocks and regulation through 2015, we focus on units that stopped operating by 2015 and the percentage of retirements cited in the text includes only those plants that stopped operating between 2005 and 2015, according to EPA data. The model predicts low profits for the units that stop operating after 2015, indicating that many of those units would be predicted to retire after 2015 if we model market shocks after 2015.

The analysis implies that reducing the stringency of emissions caps would have little effect on the profitability of existing coal-fired plants. Reducing stringency would affect emissions directly if the emissions caps continue to bind.

Like the findings in the recent literature (e.g., Fell and Kaffine, 2018), our results confirm the effects of natural gas prices on coal- and gas-fired generation, but in contrast to the empirical literature that has examined natural gas prices and local air pollution (Linn and Muehlenbachs, 2018; Johnsen, LaRiviere, and Wolff, forthcoming), we show that the natural gas price shock had little effect on NO_x emissions. This demonstrates the importance of accounting for long-run interactions between emissions caps and market shocks and offers evidence that claims about the environmental benefits of low natural gas prices may be overstated. Our results differ from those in the literature in that we find that the consumption shock affected the profitability of coal-fired power plants as much as the natural gas price shock and substantially more than the renewables generation shock. The previous literature has not considered the quantitative effects of the consumption shock (DOE, 2017).

This article builds on the extensive literature that has used computational or structural models of the electricity sector to address economic and environmental questions (e.g., Borenstein, Bushnell, and Wolak, 2002; Reguant, 2014; Cullen and Reynolds, 2016). We demonstrate that expanding the model beyond a standard economic dispatch model substantially improves model performance, particularly regarding the substitution between coal- and natural gas-fired generation and changes in both the intensive and extensive generation margins. Moreover, the model relies entirely on publicly available data and is relatively simple to operate. The structure approximates dynamics without explicitly modelling those dynamics, which allows us to model thousands of units at an hourly time-step and to find the equilibrium in a reasonable amount of time. These features enable a transparent analysis of the factors affecting entry and exit and relate to the broader literature that assesses the role of environmental regulation, international trade, and other factors on plant entry and exit and on industry dynamics (e.g., Ryan, 2012; Curtis, 2018; Shapiro and Walker, forthcoming). As in that literature, we use the model to estimate entry costs and compare counterfactuals to disentangle the long-run roles of regulation and market forces in explaining observed entry and exit decisions.

Section 2 summarizes the history of regulation, describes the data, and presents summary statistics. Section 3 outlines the computational model, section 4 describes the scenarios, section 5 presents the results, and section 6 concludes.

2. Background

■ This section provides a brief history of NO_x regulation and MATS, describes the data sources, and summarizes recent electricity sector trends. We end the section by reporting several stylized patterns of unit-level generator operation that we aim to reproduce with our model.

□ **Overview of electricity sector NO_x regulations and MATS.** Stationary and mobile sources emit NO_x when they burn fuel at high temperatures. Emissions of NO_x adversely affect health and the environment by contributing to the formation of ground-level ozone, particulate matter, and acid deposition, among other effects. These environmental and health effects create a role for government regulation, because otherwise, electricity generators and consumers would not account for them when making decisions about generation and consumption. Under the 1970 Clean Air Act (CAA), the EPA established air quality standards for NO_x and ground-level ozone that protect human health and welfare. States submit plans to demonstrate their strategies for meeting the standards. The CAA also authorizes the EPA to create emissions standards for certain sources.

Between passage of the CAA in 1970 and the late 1980s, these regulations and state plans proved ineffective at reducing NO_x emissions. Burtraw et al. (2005) suggest that because the regulations did not apply to most existing sources, they raised the costs of generating electricity from new plants relative to existing plants, causing older plants to retire more slowly than expected—a manifestation of vintage-differentiated regulation (Stavins, 2005). In addition, laws and regulations did little to address the problem of air transport. Indeed, the CAA created an incentive for power generators to construct tall smokestacks, which improved local air quality but exacerbated downwind pollution (Burtraw and Palmer, 2003).

In the 1980s, policy makers became increasingly aware of the contributions of NO_x and sulfur dioxide to acid rain. The shortfalls of initial regulations and the new information contributed to the CAA Amendments (CAAA) in 1990. The law capped national sulfur dioxide emissions, set maximum NO_x emissions rates for most existing coal-fired boilers, and required the installation of NO_x abatement equipment at boilers in regions that did not attain the air quality standards.

The CAAA also meaningfully addressed, for the first time, the long-distance transport of NO_x emissions. Because the Northeast had some of the most severe ozone problems in the country, the CAAA created the Ozone Transport Commission, which led to a NO_x cap-and-trade program covering large electricity and industrial sector boilers in the Northeast. Analysis conducted in the mid-1990s, however, suggested that NO_x emissions outside the region would cause many areas in the Northeast to exceed the ozone air quality standards even after the emissions cap was fully implemented. Based on these conclusions, the EPA created the NO_x Budget Trading Program. The program, which included 19 states and the District of Columbia, began in 2003 and capped NO_x emissions occurring each year between May and September, when ozone levels tend to be highest. The program reduced emissions by more than half from 1990 levels.

Because of continuing concerns about achieving air quality standards, the EPA ultimately created the Cross State Air Pollution Rule (CSAPR).[7] The CSAPR program caps sulfur dioxide and NO_x emissions, with summer and annual caps on NO_x. The program restricts cross-state credit trading, and it began capping NO_x emissions from 27 states and the District of Columbia in 2015. Figure A1 shows that CSAPR covers most of the eastern interconnection. Thus, over time, the NO_x emissions caps have expanded geographically and increased in stringency.

Under a process established in the 1990 CAAA, the EPA regulates emissions of air toxics. In 2000, the EPA determined that mercury and other air toxin emissions from coal- and oil-fired electricity generators should be regulated. The Clean Air Mercury Rule was finalized in 2005 and would have capped mercury emissions, but in 2008, the law was vacated by the courts. In response, the EPA proposed MATS in 2011 and finalized the regulation in 2012. MATS creates emissions limits for mercury and other air toxins from coal- and oil-fired power plants larger than 25 megawatts. Units had to meet these limits by 2015. The rule has been highly controversial and heavily litigated, but the program has survived legal challenges.

□ **Data.** Our main source of data is the EPA Continuous Emissions Monitoring System (CEMS). The data set comprises nearly all emissions from fossil fuel-fired units that operate in the eastern interconnection. Using 2005–2015 CEMS data, for each fossil fuel-fired unit, we compile hourly fuel consumption and generation; hourly emissions of NO_x, sulfur dioxide, and carbon dioxide; and unit characteristics. Unit characteristics include the state in which the unit is located, whether the unit has specific emissions abatement equipment, and rated capacity and fuel type.

We complement the CEMS data with Energy Information Administration (EIA) data from 2000 through 2015. The EIA data include information about generators that collectively account for nearly all generation from large plants. We use these data to create some of the summary statistics reported in the next subsection. We also use the data to compute fuel prices and construct the set of potential entering plants in the model.

[7] EPA developed CSAPR after legal challenges to its previous regulation, the Clean Air Interstate Rule. This program included three separate emissions caps: May through September NO_x emissions, annual NO_x emissions, and annual sulfur dioxide emissions. Twenty-seven states and the District of Columbia participated in at least one of the three caps.

FIGURE 1

NATIONAL NITROGEN OXIDES AND SULFUR DIOXIDE EMISSIONS

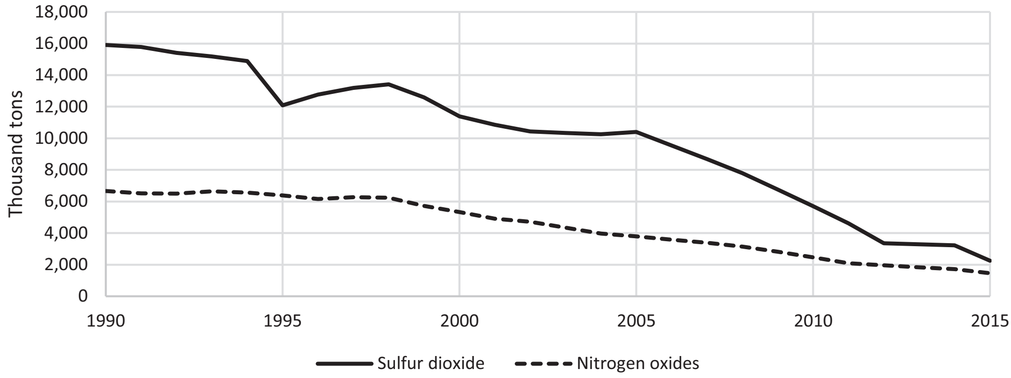

Note: Data are from the EPA National Emissions Inventory and include emissions from fossil fuel combustion in the electricity sector.

□ **Electricity sector trends.** In this subsection, we document declining NO_x emissions and changes in the electricity sector that have contributed to this decline. Figure 1 shows that national NO_x emissions declined by 75% between 1990 and 2015, with most of the decline occurring after 2000. For comparison, sulfur dioxide emissions declined by 85% during this period.

We focus on the eastern interconnection, which accounts for about 90% of national electricity sector NO_x emissions. Between 2000 and 2015, NO_x emissions in the East declined by nearly 75%, mirroring the national trend. Total NO_x emissions are the product of total generation and the average rate of emissions. Therefore, reductions in total generation or average emissions rates could explain the declining emissions. Total generation in the East increased steadily between 2001 and 2007, at about 2% per year, then declined between 2007 and 2009, and remained roughly flat from 2009 through 2015. The 2007–2009 decline coincides with the macroeconomic recession, but partly because of gains in energy efficiency, electricity generation in 2015 was slightly lower than generation in 2009. Growth in fossil fuel-fired generation, which accounts for nearly all NO_x emissions from the electricity sector, experienced a similar leveling off after 2007. The fact that fossil fuel-fired generation was the same in 2001 and 2015 implies that declining average emissions rates, and not total fossil generation, explain the emissions reduction.

The decline in average emissions rate appears to be due both to a reduction in emissions rates at individual units and to a shift to lower-emitting fuels. Coal-fired units have steadily adopted emissions-reducing technology, such as selective catalytic reduction (SCR), which reduces emissions rates by roughly 90%. As Figure 2 illustrates, generation shifts have also contributed to the decline in emissions rates. The shift from coal- to gas-fired generation between 2000 and 2015 reduced the average emissions rate across fossil generation units. The increase in the wind generation share, from close to zero in 2000 to 4% in 2015, further reduced emissions. In short, changes in unit emissions rates and a shift from coal to cleaner fuels contributed to the reduction in emissions rates—as it turns out, about equally (not shown).

The change in the capital equipment used to generate electricity is consistent with these changes. Figure 3 shows that about 90 gigawatts (GW) of coal-fired capacity, almost one third of the initial capacity, began retirement preparations between 2005 and 2015. Table 1 compares the attributes of coal-fired units that retired with those that continued operating. The retiring units tended to be smaller, older, less efficient, and less heavily utilized than the continuing units. Figure 4 shows a large decrease in natural gas prices after 2008, relative to previous prices and to coal and oil prices.

FIGURE 2

GENERATION SHARE BY ENERGY SOURCE FOR EASTERN INTERCONNECTION, 2000–2015

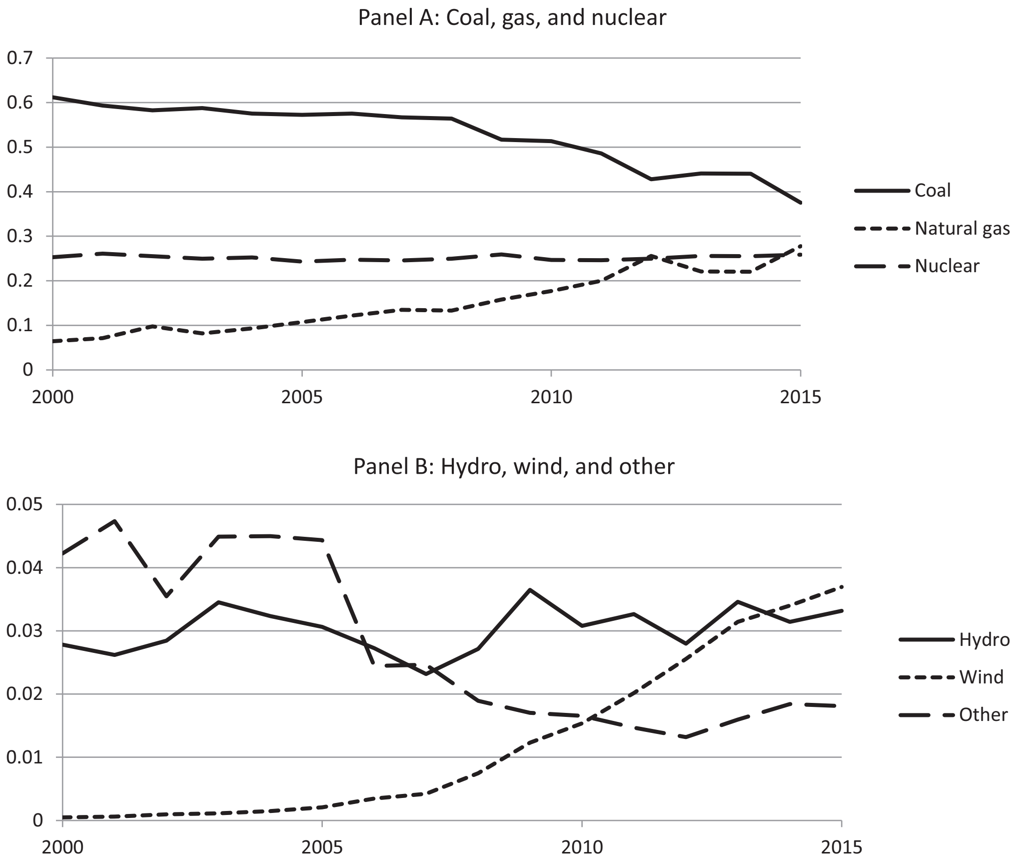

Notes: The figure plots the share of generation in total generation in the eastern interconnection, by technology. The other category includes solar, geothermal, and a few additional technologies. Data are from EIA Form 920.

We summarize these recent developments by comparing projections of the electricity system made in 2005 and 2015. In Figure 5, we compare projections that the EIA made in the 2005 Annual Energy Outlook with outcomes between 2005 and 2015. Compared with the projections, 2015 natural gas prices were 25% lower, generation from renewables was 2.5 times higher, and total electricity sector generation was 15% lower. These differences represent the unanticipated changes in fuel prices, renewables generation, and aggregate electricity consumption that occurred during this 10-year period. We use the term *shocks* to describe the difference between the observed 2015 outcomes and the 2005 projections of 2015 outcomes. Projections through 2025 suggest that these shocks were permanent rather than temporary.

□ **A few stylized facts and the importance of the intensive generation margin.** We observe several patterns in unit-level emissions and generation in the CEMS data; an objective of our model is to reproduce these patterns. If emissions rates varied greatly across fuel types but little within fuel types, accurately predicting emissions would require only an accurate prediction of generation shares by fuel type. However, Figure A2 illustrates substantial within-fuel-type variation in emissions rates. For each unit in the sample, we compute the average emissions rates of NO_x, sulfur dioxide, and carbon dioxide, using hourly emissions and generation data from

FIGURE 3

CUMULATIVE RETIREMENTS AND CAPACITY ADDITIONS IN THE EAST, 2005–2015

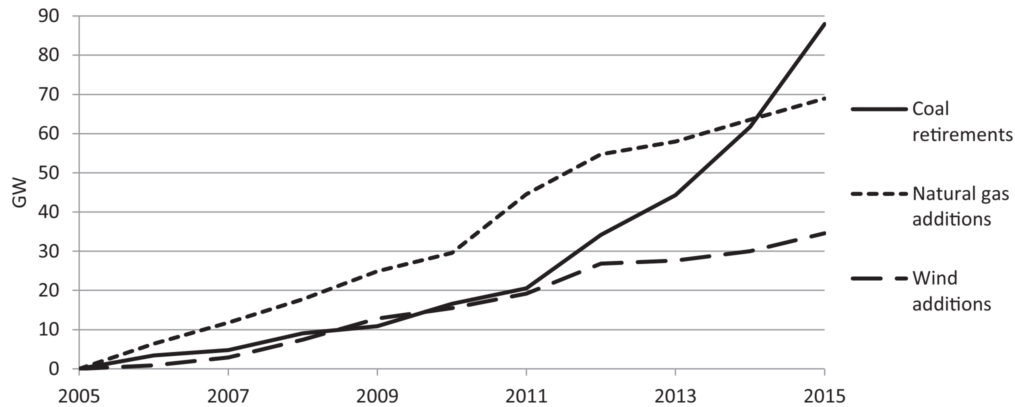

Note: Cumulative retirements and capacity additions by fuel type, in gigawatts (GW), are computed from EIA Form 860 for the years 2005 through 2015. Both retirements and additions are gross amounts and not net. Retirements include plants that announced their retirement but which may not shut down until after 2015.

TABLE 1 Summary Statistics of Eastern Coal-Fired Units

	Units that Retire between 2005 and 2015	Units that Continue Operating	Difference (continue - retire)	t-statistic
Number of units	148	667		
Capacity (MW)	164	372	208	15.28
	(111)	(262)		
Vintage (year)	1956	1967	12	15.94
	(7.27)	(11.23)		
Heat rate (mmBtu/MWh)	10.71	10.07	−0.64	−5.21
	(1.37)	(1.31)		
Capacity factor	0.44	0.65	0.22	12.31
	(0.20)	(0.17)		

Notes: Coal-fired units operating in 2005 are separated into two sets: units that retire by 2015 and units that continue operating through 2015. Capacity (in megawatts, MW) and vintage (initial operating year) are obtained from the CEMS unit characteristics. Heat rate and capacity factor are computed from hourly fuel input and generation from CEMS. The right-most column reports the t-statistic from a test on the equality of the means of the variables across the two samples.

2005. The observed variation in NO_x emissions rates within fuel types suggests that predicting emissions requires an accurate prediction of unit-level generation.

A generation unit's extensive margin refers to whether the unit operates at all, and the intensive margin refers to how much the unit operates, conditional on its operating. Table A1 demonstrates that generation units experienced changes along both the extensive and the intensive margins between 2005 and 2015. Between 2005 and 2008, coal units operated an average of 85% of all hours. This fell to 73% between 2009 and 2015, representing a 12-percentage point decline along the extensive margin. Across the same two periods, natural gas-fired unit operation increased by seven percentage points along the extensive margin. We also observe changes along the intensive margin between 2009 and 2015. Capacity factors conditional on operation declined by six percentage points for coal. For natural gas, conditional capacity factors increased by five percentage points.

In short, the data indicate substantial variation over time in both the probability that units operate and their capacity factors conditional on operating. Because a unit's profits depend on the correlation across hours between its generation and the equilibrium electricity price, accurately

FIGURE 4

FUEL PRICES FOR EASTERN INTERCONNECTION, 2000–2015

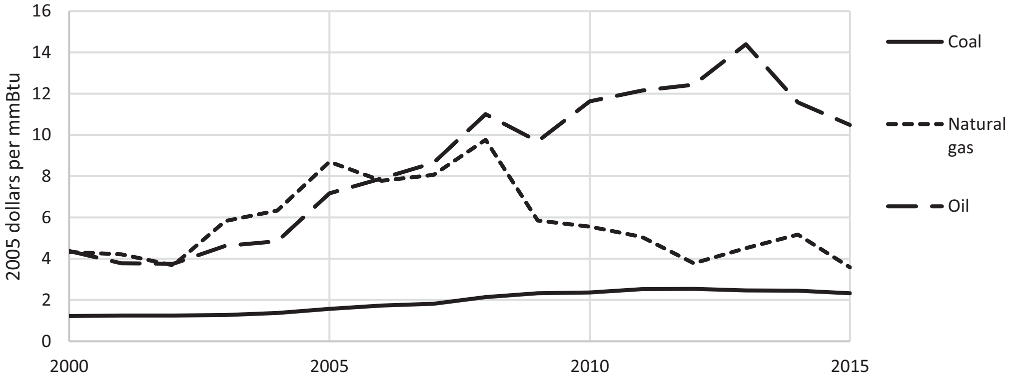

Note: Fuel prices are measured in dollars per million British thermal unit (mmBtu) and are Btu-weighted means across plants reporting to EIA Forms 423 and 923.

modelling a unit's hourly operation, including the extensive and intensive margins, is essential for estimating its annual profits and emissions.

3. Computational model

■ We develop a computational model that combines attributes of structural or computational models of the electricity system, such as Bushnell, Chen, and Zaragoza-Watkins (2014) and Reguant (2014), with attributes of the reduced-form model in Davis and Hausman (2016). The model has limited data requirements, all data are publicly available, and simulations are relatively easy because of the reduced-form treatment of dynamics. The model reproduces observed unit-level operation and emissions, as well as abatement and plant retirement decisions.

□ **Overview of model structure.** The model consists of three phases: retirements and new construction, pollution abatement investment, and hourly operation. The hourly operation phase is simulated over a single representative year (i.e., 8760 hours) to estimate revenues and operating costs for each unit in the model. Next, we simulate investments in pollution abatement equipment. Finally, we simulate entry and exit decisions (annualizing all capital costs), iterating until we achieve convergence to ensure consistency across the phases.[8] Intuitively, an increase in natural gas prices reduces profits of coal-fired units in the operation phase, particularly for inefficient units, and potentially causing their retirement. Because pollution abatement has high fixed costs and low marginal operating costs, abatement equipment is installed typically at large units with high emissions and capacity factors, making it more likely for other units to exit because of regulation.

The structure is similar to that of planning models used in the power sector and in the economics literature (e.g., Borenstein and Holland, 2005; Fell and Linn, 2013). This type of model is particularly useful for comparing long-run steady states rather than transitional dynamics.

Two considerations motivate the use of this class of model. First, our main interest lies in the long-run effects of market shocks and emissions regulation, not the transitional dynamics. Second, although in principle we could use equilibrium electricity prices to estimate a dynamic model in which generation units make investment and operational decisions based on current

[8] In principle, we could specify the model as a constrained optimization problem, where the objective would be to minimize total costs subject to emissions and operating constraints. However, because of the discrete nature of many decisions in the model and the large number of units in the system, this approach is computationally infeasible.

FIGURE 5

COMPARISON OF 2005 AND 2015 EIA PROJECTIONS

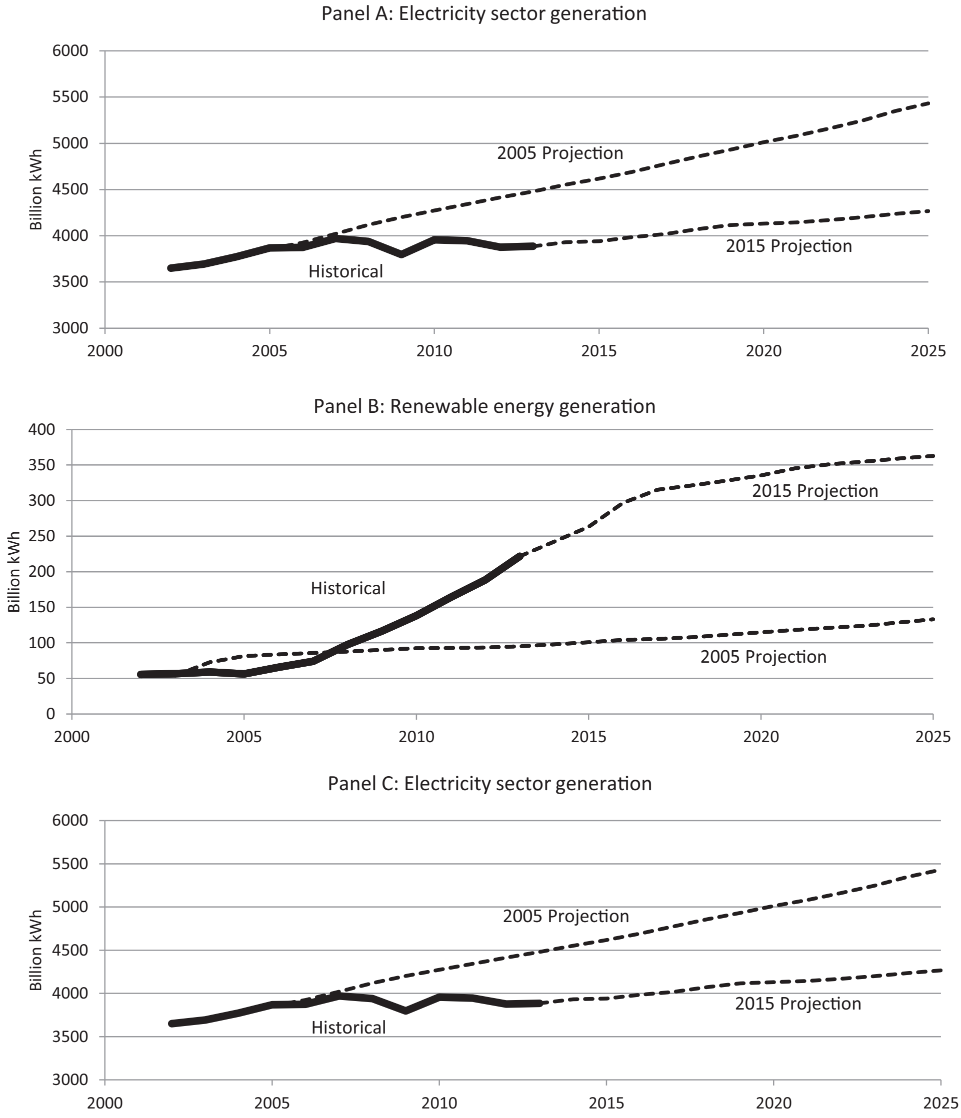

Notes: The dashed lines show projections from the Annual Energy Outlook (AEO) for the indicated years. Solid lines show the estimated historical price or generation level using data reported in the AEOs between 2005 and 2015. For example, historical electricity sector generation for the year 2002 is from the 2005 AEO. Data for the projections begin two years prior to the projection year.

and expected future state variables (e.g., Mansur, 2007), prices are not available in much of the eastern United States, particularly in the Southeast.[9] Consequently, a dynamic model would omit much of the eastern emissions.

□ **Phase 1: Retirements and new construction.** At the outset of the first stage, there exists a set of generation units that have already been constructed. Each firm owns one unit. The owner

[9] Cost metrics are available for the regions that do not have active markets, but these metrics are not directly comparable with the prices that are observed in other regions (Linn and Muehlenbachs, 2018).

of an existing unit decides whether to retire the unit or continue operating. The owner retires the unit if expected profits in the subsequent abatement and operational stages would be negative, where profits equal the difference between discounted revenues and costs. Costs include pollution abatement costs, variable costs, and fixed costs. We assume perfect foresight over revenues and costs. For simplicity, there are no retirement costs and the unit does not have any scrappage value. Consequently, retiring is equivalent to not operating. We refer to units that are not retired as continuing units.

There also exists a set of firms that each decide whether to construct a single new generation unit. Each of these firms has an exogenous fuel type, heat rate, and generation capacity. The heat rate, a standard metric in the electric power sector, is inversely related to the unit's efficiency: more efficient units have lower heat rates and lower fuel costs. Fixed costs are associated with unit permitting and construction. The potential entrant decides to construct the new unit if the expected profits are positive, where profits take into account permitting and construction costs, as well as costs and revenue from the abatement and operation phases.

□ **Phase 2: Pollution abatement investments.** After each firm has decided whether to continue operating its existing unit or to construct a new one, firms with continuing or entering units must decide whether to invest in pollution abatement equipment. As in CSAPR, the NO_x regulation in the model includes both annual and summer emissions caps that cover most units in the East. For each state with an annual or summer cap, the cap is denominated in tons of NO_x. All states implement the caps by allocating emissions credits to each firm, and the total number of credits allocated equals the cap. Allocation to each firm depends on its unit's historical generation, with a certain fraction of credits set aside for entering units. Firms can trade credits with other firms in the same state. At the end of the year, each unit's emissions cannot exceed the number of credits its owner holds. Units in some states face both annual and summer caps; in the model, it is endogenous whether either or both caps are binding.[10]

Each generation unit can abate its emissions by reducing its generation or by installing pollution abatement equipment. In this subsection, we focus on the decision to install SCR.[11] Here, we discuss the annual caps, assuming the summer cap is not binding (we relax this assumption in the solution algorithm, as explained later).

Installing abatement equipment involves a fixed cost as well as an operational cost that scales linearly with generation. For a firm that installs abatement equipment, we define the abatement cost as K_i^a, where K_i^a is the annualized capital cost of the abatement equipment for unit i. The capital costs depend on the unit's size, age, and other attributes.[12]

The abatement equipment reduces the unit's annual emissions by $(e_i - e')g_i$, where e_i is the emissions rate (tons of NO_x per megawatt hour of generation) in the absence of the abatement equipment, e' is the emissions rate with the abatement equipment, and g_i is the unit's annual generation. Each unit has the same emissions rate with abatement equipment installed, and abatement increases with the unit's uncontrolled emissions rate (e_i) and generation (g_i).

[10] The modelled emissions program has simplifications to reduce computational burden. First, whereas CSAPR allows some emissions credit trading across states, we assume emissions credit trading within but not across states. Second, we assume that each state's credit market is perfectly competitive, consistent with EPA anlysis of CSAPR. In practice, firms may have market power in credit markets, but a considerable number of cross-state trades are observed and there is a single market-clearing emissions credit price. This implies that even though some states contain few firms subject to the caps, these firms are subject to competition with firms located in other states. Finally, we assume no banking of compliance credits. These assumptions may cause us to overestimate compliance costs and emissions credit prices (see Section 6).

[11] For simplicity, we model SCR but not other abatement technologies, such as low-NO_x burners. This assumption may cause us to overestimate the costs of complying with CSAPR to the extent that including other technologies would reduce retirements or abatement costs (see Section 6).

[12] Age can affect annualized abatement costs because installation costs may be higher at older units and because an older unit has a shorter remaining lifetime over which costs are annualized (Fowlie, 2010).

Average abatement costs are defined as the ratio of total abatement costs to abatement:

$$\frac{K_i^a}{(e_i - e')\, g_i} + \frac{M^a e'}{(e_i - e')}, \tag{1}$$

where SCR operating costs are $M^a e' g_i$. Average abatement costs increase with the unit's capital costs and decrease with its generation level. Units with higher uncontrolled emissions rates have lower average costs.

If a firm installs abatement equipment and reduces its emissions below its credit allocation, the firm can sell excess credits to a firm whose emissions exceed its credit allocation. Each state's emissions credit market is perfectly competitive, and there is a market-clearing credit price, $\tau_s \geq 0$, for each state, s. In each state, aggregate emissions cannot exceed the state's emissions cap. Because the credit market is competitive, firms install abatement equipment if their average abatement costs do not exceed the emissions credit price. The price therefore adjusts so that in equilibrium, credit demand equals credit supply.

The fact that g_i affects average abatement costs implies that expected generation affects the decision to install abatement equipment. Because of the assumption of perfect foresight, the firm makes its abatement decision knowing the value of g_i.

A few states have a summer emissions cap but not an annual emissions cap. In these states, firms make abatement decisions as described above, except that they compute average abatement costs using generation during summer months and do not operate SCR in nonsummer months to avoid operating costs. Many states have both annual and summer caps, and there are separate credit prices for annual and summer emissions. Firms in these states install SCR if the average abatement costs are less than either the annual or the summer emissions price. For reasons explained in Section 4, we use observed compliance decisions to estimate the costs of MATS.

□ **Phase 3: Hourly operation.** The operational stage of the model represents a steady state. We characterize hourly operation over a single year, and that year is repeated into the infinite future. Revenues and costs are discounted back to the retirement and construction phase of the model.

We build a unit commitment model that introduces constraints affecting a unit's minimum generation level and its ability to vary generation across hours. A standard unit commitment model (e.g., Castillo and Linn, 2011; Wang and Hobbs, 2016) includes stochastic electricity demand and unit outages, fixed costs of starting up and shutting down, and constraints on changes in generation level across hours. Our unit commitment model is simplified for tractability, approximating a unit commitment model's main features. We first describe the assumptions and the market equilibrium, and then explain how the model approximates uncertainty, fixed costs, and constraints on changing generation across hours.

A unit's generation costs include both fuel costs and nonfuel costs. Fuel costs equal the price of fuel (p_{ih}), in dollars per million British thermal units (mmBtus), multiplied by the unit's heat rate (h_i), in mmBtus per megawatt hour (MWh) of generation. The price of fuel varies across units because of fuel type and regional fuel price variation, and across hours because of monthly changes in fuel prices. The nonfuel costs (n_i), in dollars per MWh, include costs of labor and materials and vary across units but not across hours. For simplicity, the heat rates and nonfuel costs do not depend on the level of generation, and marginal costs are given by $m_{ih} = h_i p_{ih} + n_i + e_i \tilde{\tau}_s$. Note that marginal costs depend on the emissions costs, $e_i \tilde{\tau}_s$, where $\tilde{\tau}_s$ is the sum of the annual and summer emissions credit prices (summer credit prices equal zero in nonsummer months; e' replaces e_i for units with SCR).

Each coal and large natural gas-or oil-fired unit has a minimum generation level, $\underline{g}_i$, such that if the unit is operating, it cannot operate below that level. All units have a maximum generation level, $\overline{g}_{ih}$. The maximum generation level varies across units and hours because of time-varying factors such as transmission-constraints, and the minimum level varies across units.

Total generation is exogenous to the model, which includes generation from nuclear, hydroelectric, renewables, and fossil. Following Bushnell, Chen, and Zaragoza-Watkins (2014), among others, we assume that generation from nuclear, hydroelectric, and renewables does not respond to electricity prices. The lack of available data necessitates this assumption, although we note that it is particularly reasonable for nuclear and renewables.[13]

Next, we turn to the market equilibrium. We assume that the market is perfectly competitive and that firms treat the equilibrium price as being independent of the generation.[14] We distinguish among three types of hours: the peak hour, when electricity demand reaches its daily maximum; near-peak hours, which are within six hours of the peak hour; and off-peak hours, which include all other hours in the same day.[15]

At the beginning of each day, a system operator determines the peak hourly aggregate fossil generation for the day. The operator solicits bids, where a unit's bid includes a generation level and minimum price above which the unit generates the specified amount. The operator ranks the bids in order of increasing price and accepts bids to meet the forecast peak aggregate generation. Below, we explain how the firm chooses its minimum price bid.

Firms whose bids are accepted for peak hour generation must operate their units above their minimum levels, $g_{ih} \geq \underline{g}_i$, during near-peak hours of the same day. Except for firms owning small gas- and oil-fired units, firms whose bids are not accepted for peak hour generation cannot generate in near-peak hours. This structure prevents units from turning off or shutting down repeatedly during near-peak and peak hours. Because small gas- and oil-fired units are exempted from these constraints, those units may turn on and off multiple times during a day.

Firms that generate in the peak hour submit two-part bids for generating in the near-peak hours. If the near-peak price exceeds the firm's marginal operating costs, the firm generates at the maximum level, $\overline{g}_{ih}$. If the price is below the unit's marginal costs, the firm generates at the minimum level, $\underline{g}_i$. The operator accepts bids such that total supply equals aggregate fossil generation. The equilibrium price in near-peak hours equals the marginal costs of the highest-cost unit that operates above its minimum level.

During near-peak hours, certain units may earn negative profits. For example, consider a unit that operates during a peak hour. During near-peak hours, when aggregate fossil generation is below the peak, the electricity price may lie below the marginal costs of the unit. Because the unit cannot operate below its minimum level, the unit must operate during those hours even if the electricity price is less than its marginal costs. However, the firm anticipates the negative profits when it submits its bid for peak hour generation and submits a price sufficient to recover its losses during nonpeak hours. Therefore, the firm submits a peak hour price that is greater than

[13] These technologies have low marginal operating costs and therefore generate as much electricity as technologically possible. Hydroelectric plants, on the other hand, can be dispatched to some extent subject to environmental and other constraints. However, in the East, hydroelectric plants accounted for just 3% of power generation in 2005. The exogeneity assumption, therefore, has little effect on the main results. For eastern nuclear units, between 2005 and 2015, we observe little variation in annual capacity factors and few trends in monthly capacity factors, supporting the exogeneity assumption.

[14] In models that allow for the possibility that US electricity markets are imperfectly competitive (e.g., Borenstein, Bushnell, and Wolak, 2002; Mansur, 2007), firms account for the effect of their generation on equilibrium prices and restrict their generation to increase prices. However, models now commonly assume perfect competition (e.g., Borenstein and Holland, 2005; Blanford, Merrick, and Young, 2014; Zhou, 2016), reflecting expansion of the geographic scope of wholesale power markets and other factors that have increased competition.
The modelling does not distinguish units at plants that are subject to cost of service regulation. In principle, operation and retirement decisions could differ for regulated plants. For example, an owner of a regulated plant may be able to persuade regulators to cover costs even if the plant is otherwise unprofitable. However, the model appears to predict generation equally well for both types of units. Moreover, about 63% of the units that the model predicts retire (and that actually retire) are subject to cost of service regulation. These results suggest that the perfect competition assumption does not systematically affect the performance of the model in predicting behavior of regulated units.

[15] The choice of six hours for nonpeak hours is somewhat arbitrary. We have recalibrated the model using a range of assumptions on the nonpeak period for coal- and gas-fired units. The assumptions described in the text and in the Appendix yield the best fit between observed and simulated hourly generation. However, the main conclusions about the importance of market shocks appear to be robust to these assumptions.

its marginal operating costs. Dynamic models with startup or shutdown costs, such as Bushnell, Mansur, and Saravia (2008), similarly yield equilibria in which firms bid prices below marginal costs in certain hours. The gap between the peak bid and marginal costs is greater for high marginal cost firms than for low-cost firms (all else equal), because the high-cost firms must recover greater losses incurred during near-peak hours. In equilibrium, the peak price exceeds the marginal costs of the highest-cost unit operating. This is an equilibrium, because units whose costs exceed the highest-cost unit that actually operates would earn negative profits across peak and near-peak hours if they were to operate.[16]

During off-peak hours, the equilibrium is determined according to economic dispatch. Units operate at $\overline{g}_{ih}$ if the equilibrium price exceeds their marginal costs, and they operate at zero otherwise. The operator stacks the units in order of increasing marginal costs and selects the price such that combined generation equals aggregate fossil generation.

This model differs from a standard economic dispatch model in several important ways. In a dispatch model, a unit's decision to operate in a particular hour does not affect its decision to operate in other hours; there is no distinction between peak, near-peak, and off-peak hours. Consequently, units are assumed to operate at their maximum generation level if price exceeds marginal costs, and they do not operate at all otherwise. As demand varies across hours, units start up and shut down such that supply equals demand, and a unit's generation varies along the extensive but not the intensive margin. In contrast, in the stylized unit commitment model, if a firm operates during the peak hour, it must operate during all near-peak hours of the same day. This constraint captures the effects on unit operation of startup and shutdown costs because firms typically avoid incurring these costs multiple times each day.[17] Importantly, in the unit commitment model, exogenous factors (e.g., fuel prices) may affect both intensive and extensive generation margins. In contrast, in a dispatch model, exogenous factors affect only the extensive margin. Furthermore, in the commitment model, electricity prices during peak hours exceed marginal costs of operating units (during other hours, price equals marginal costs of the highest-cost unit operating above its minimum level, just as in a dispatch model). As the Appendix explains, the maximum generation level, $\overline{g}_{ih}$, captures transmission constraints or other factors that may cause balkanized dispatch within the eastern interconnection. For example, transmission constraints between regions that arise when aggregate demand is high may prevent a unit from operating at full capacity, which would be captured by $\overline{g}_{ih}$. Finally, as the Appendix explains, we allow for uncertainty in unit availability and aggregate fossil fuel-fired generation.[18]

□ **Parameter assumptions and solution algorithm.** Whenever possible, we use observational data to populate and parameterize the model. The set of units at the beginning of phase 1 (retirement and new construction) includes all CEMS units that operated in 2005. Potential entrants include all units that actually entered between 2005 and 2015 and units that entered construction planning prior to 2005 but did not actually enter the market. This allows us to use

[16] A firm's peak-hour bid is determined by the losses incurred during the near-peak hours. In turn, the near-peak losses depend on the firm's exogenous variable costs and the equilibrium price, which it takes as exogenous. For a near-peak hour, the equilibrium price is determined by the marginal costs of the highest-cost unit that operates above its minimum capacity. That is, the bid of a firm that operates at its minimum generation level does not affect the equilibrium price because that firm is inframarginal. Because the price rather than the firm's bid determines the firm's losses, those losses do not depend on its bid during the near-peak hour. Consequently, the near-peak bid does not affect profits, making the peak bid unique.

[17] In the simulations, many natural gas- or oil-fired units can start up or shut down multiple times within the peak and near-peak periods. We have also considered versions of the model that prevent more than one startup and shutdown for each unit and day, which yield similar results to those reported here.

[18] We incorporate two types of uncertainty. First, units may be unavailable because of unplanned outages or maintenance. We include an exogenous probability that the unit is unavailable for a particular day. Second, the system operator introduces a reserve requirement to account for the fact that peak aggregate fossil generation is forecasted with error. We introduce a reserve margin, $r > 0$. The system operator accepts bids for peak-hour generation such that total generation of accepted bids is equal to $1 + r$ multiplied by forecast peak aggregate fossil generation.

the heat rate and capacity of actual and planned units to characterize the potential entrants in the model. For the abatement phase, we estimate capital costs of installing SCR from EPA (2010). In the operational phase, fuel prices are constructed from EIA data, and aggregate fossil generation is computed as the sum of observed generation across all fossil fuel-fired units. Profits in the operational phase are discounted to 2005 using a 10% discount rate (this discount rate is typical for the industry, and results are similar using alternative discount rates). The Appendix describes the methodology for constructing unit attributes and other parameter values, including nonfuel costs and transmission constraints.

Turning to the solution algorithm, we solve the model iteratively, beginning with the operation phase and an initial guess of the equilibrium emissions credit prices, and assuming that there are no retirements. The operation stage includes a single year to represent the steady state, and it can be solved each day of the year by first determining which units operate in the peak period; only those units operate during the near-peak hours of the same day. In each hour, the equilibrium price is determined such that supply equals aggregate generation and is subject to all operating constraints of the units.

In the abatement phase, given an assumed emissions credit price, units install abatement equipment if the unit's average abatement cost is no greater than the credit price. The credit price is increased from zero until the emissions cap is satisfied. Capital investments are annualized and subtracted from each unit's operating profits computed in the operation stage.

If profits of all units are positive but electricity prices are not high enough to induce entry, the equilibrium is determined. If at least one unit has negative profits, the unit with the lowest profits is assumed to retire (or not to enter), and the model is resimulated using the smaller set of generation units and the credit prices estimated from the previous iteration.[19]

□ **Validation of the abatement and hourly operation stages and comparison with dispatch model.** Section 5 validates the full model by comparing predicted and observed coal-fired plant retirement decisions. In this section, we focus on validating the abatement and operational phases of the model by comparing model outputs with observed behavior. Although the scenarios described in the next section use 2015 as the representative year, in this subsection, we simulate the operation stage for each year between 2005 and 2015. This allows us to compare predicted and observed outcomes across the range of fuel price, demand, and renewables conditions that occurred during these years.

To validate the abatement cost estimates, we compare estimated abatement costs of units that do and do not install SCR, expecting the installers to have lower costs. Figure 6 uses data and assumptions from the EPA to plot the estimated density functions of average abatement costs for units that do not have SCR in 2005, separating units that install SCR between 2005 and 2015 from units that do not. Abatement costs, in dollars per ton of NO_x, are defined as in equation (1). The figure shows that units with higher average abatement costs were less likely to install SCR. If other factors predict SCR installation, or if there is a nonlinear relationship between annualized capital costs and installation, we could improve the model by incorporating these other factors in the installation decision. We find that, conditional on average abatement costs, other unit attributes, such as age, do not predict SCR installation.[20] Unobserved factors, such as compliance with local regulations and EPA enforcement of New Source Review, may explain SCR installation. These

[19] The exit rule does not account for the fact that one unit's exit can affect another unit's profits, raising the possibility of multiple equilibria if we accounted for strategic behavior in exit. We have considered other exit rules, such as randomly choosing a unit with negative profits or choosing the unit whose profits increase most from another unit's exit (this is feasible only in simulations that include a small number of units with negative profits). Because these alternatives yield similar results, we use the simpler exit rule.

[20] Fowlie (2010) finds that traditionally regulated units are more likely to install SCR. However, after using more recent data and conditioning on estimated costs, we do not find a strong correlation between SCR adoption and regulatory status. Therefore, we do not distinguish between regulated and unregulated plants in the abatement phase.

FIGURE 6

ESTIMATED DENSITY FUNCTIONS OF ANNUALIZED AVERAGE ABATEMENT COSTS FOR SELECTIVE CATALYTIC REDUCTION

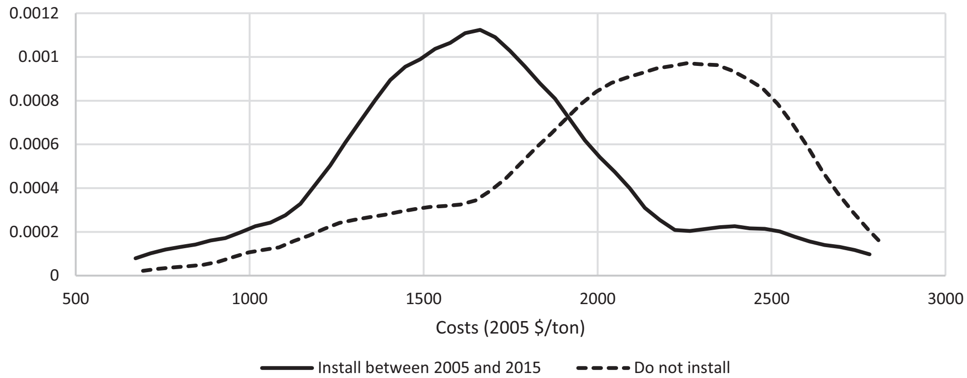

Notes: For each coal-fired unit operating in 2005 that does not have SCR, capital costs of installing SCR are computed using the unit's rated capacity, heat rate, and observed emissions rate and cost assumptions from the EPA Integrated Planning Model. Capital costs are annualized using the unit's estimated remaining lifetime. Emissions abatement is equal to the product of the unit's simulated 2005 generation and the change in emissions rate from installing SCR. Average abatement costs, in dollars per ton of NO_x emissions, equal the annualized capital costs divided by emissions abatement. The figure plots the estimated density function of average abatement costs separately for units that do and do not install SCR between 2005 and 2015.

factors are exogenous to the model. Below, we discuss estimated abatement costs in context of the simulated and observed credit prices.

Turning to the performance of the hourly operation phase, we compare observed generation and emissions with levels predicted by the unit commitment model as well as the levels predicted by an economic dispatch model. The dispatch model uses the same data as the unit commitment model and does not include the constraints on minimum generation levels or distinguish among peak, near-peak, and off-peak hours.

We compare the simulated outcomes using the commitment and dispatch models for each year between 2005 and 2015. For selected years, Panel A of Table 2 shows that the percentage of coal-fired generation predicted by the unit commitment model matches the observed percentages more closely than do the predictions of the dispatch model (results for other years are available on request; percentage differences between the simulated and the observed emissions are reported in curly brackets). Across all years between 2005 and 2015, the mean absolute deviation is about 3 percentage points for the unit commitment model and 20 percentage points for the dispatch model. The dispatch model overpredicts cross-year changes in the coal-fired percentage. For example, the dispatch model predicts a 36.8 percentage point reduction between 2005 and 2015, whereas the observed change was 21.9 percentage points. This suggests that the dispatch model would overstate the effects of the natural gas shock on coal plant profits and retirements. By comparison, the commitment model predicts a 21.9 percentage point decrease.

Panels B through D in Table 2 compare observed and simulated aggregate emissions for selected years. The table shows that the unit commitment model outperforms the dispatch model in every year, and typically by a wide margin.

The unit commitment model also outperforms the dispatch model when comparing unit-level outcomes. Above, we noted the observed variation over time along the extensive and intensive margins. By construction, the dispatch model predicts capacity factors, conditional on operation, equal to one. In contrast, the unit commitment model predicts capacity factors between zero and one because of the minimum and maximum generation constraints. Table 3 shows that the unit

TABLE 2 Observed and Simulated Generation and Emissions

Year	Observed	Dispatch Model Simulation	Unit Commitment Model Simulation
Panel A: Percentage of coal in total generation			
2005	79.9	96.9	81.3
		{21.28}	{1.75}
2008	80.2	97.7	82.2
		{21.87}	{2.57}
2012	62.9	63.1	62.2
		{0.37}	−{1.11}
2015	58.0	60.1	59.4
		{3.69}	{2.38}
Panel B: Nitrogen oxides emissions (million tons)			
2005	2.94	3.42	2.97
		{16.28}	{1.11}
2008	2.39	2.83	2.46
		{18.39}	{2.66}
2012	1.28	1.39	1.35
		{9.06}	{5.97}
2015	1.00	1.09	1.08
		{8.97}	{7.53}
Panel C: Sulfur dioxide emissions (million tons)			
2005	9.05	10.70	9.31
		{18.16}	{2.82}
2008	6.65	8.26	7.01
		{24.08}	{5.39}
2012	2.75	2.99	3.00
		{8.70}	{9.06}
2015	1.79	1.98	1.97
		{10.43}	{9.67}
Panel D: Carbon dioxide emissions (million tons)			
2005	1923	2092	1928
		{8.75}	{0.22}
2008	1872	2055	1875
		{9.74}	{0.12}
2012	1630	1637	1626
		{0.43}	−{0.26}
2015	1521	1517	1514
		−{0.30}	−{0.48}

Notes: The first column reports the observed percentage of total fossil fuel-fired generation that is from coal-fired generation in Panel A and the emissions in millions of tons in Panels B–D. The right two columns report the corresponding simulated outcomes using the dispatch and commitment versions of the model. The percentage difference between simulated and observed values is reported in curly brackets.

commitment model approximates the observed changes along the intensive margin for both coal- and natural gas-fired units and across time periods.

Regarding the extensive margin, each annual simulation includes units that are observed to generate electricity in that year. Therefore, the unit commitment model would ideally predict positive generation for each unit in each year. In practice, Table 4 shows that the percentages of units predicted to have zero generation are close to zero, whereas the dispatch model predicts zero generation for about 5% of units on average.

TABLE 3 Observed and Simulated Capacity Factors

	Mean Capacity Factor Conditional on Positive Generation					
	2005			2015		
	Observed	Simulation Using Dispatch Model	Simulation Using Unit Commitment Model	Observed	Simulation Using Dispatch Model	Simulation Using Unit Commitment Model
Coal	0.78	1.00	0.79	0.70	1.00	0.78
Gas	0.70	1.00	0.66	0.89	1.00	0.75

Note: For the years indicated in the column headings and fuel types indicated in the row headings, the table reports the observed and simulated mean capacity factors conditional on positive generation, where means are weighted by the unit's rated capacity.

TABLE 4 Percentage of Units with Zero Annual Generation

Year	Simulation Using Dispatch Model	Simulation Using Unit Commitment Model
2005	1.80	0.00
2008	4.92	0.06
2012	7.52	0.01
2015	5.13	0.05

Note: For the versions of the model indicated in the column headings and years indicated in the row headings, the table reports the share of units that have zero simulated annual generation.

Figure 7 further confirms the superiority of the commitment model by plotting simulated against observed annual generation for the two versions of the model. The predicted values for the commitment model are more similar to observed values than they are for the dispatch model. If we regress simulated on observed generation, the R^2 is typically about 0.9 for the commitment model and 0.7 for the dispatch model. Consistent with the results in Table 4, Figure 7 shows that the dispatch model is more likely to predict zero generation than is the unit commitment model.

We assume that the transmission constraints are the same across counterfactuals. These assumptions could bias the simulation results if the operation constraints do not accurately approximate transmission congestion in the baseline scenario, or if transmission constraints would affect units differentially across the counterfactuals considered. Figures A5 through A7 compare predicted and observed outcomes at more disaggregated levels than in Figure 7, illustrating results by North American Electric Reliability Corporation (NERC) region and fuel type. Overall, the model predicts outcomes accurately for these subsets of generation units, and across the range of operating conditions considered in the counterfactuals. The agreement between observation and data supports the validity of our approach to approximating transmission constraints.

Thus, over the range of conditions observed between 2005 and 2015, when fuel prices, renewables, and consumption varied considerably, the unit commitment reproduces outcomes more accurately than does the dispatch model. Note that to avoid overfitting the model, we estimate the model parameters using observations across the entire 2005–2015 period rather than estimating the parameters during subperiods. The fact that we use the entire period to estimate the parameters, and that the model performs well in all subperiods, indicates that the model can accurately predict outcomes across the range of scenarios described next.

4. Scenarios

■ We use the model to quantify the costs of reducing NO_x emissions and to estimate the effects of market shocks on those costs, as well as on coal plant profits and retirements. Each scenario includes a set of initial generation units that were operating in the year 2005. The scenarios use the year 2015 to represent the steady state and differ in their projected fuel prices, electricity

FIGURE 7

SIMULATED VERSUS OBSERVED ANNUAL UNIT GENERATION (MILLION MWH) [Color figure can be viewed at wileyonlinelibrary.com]

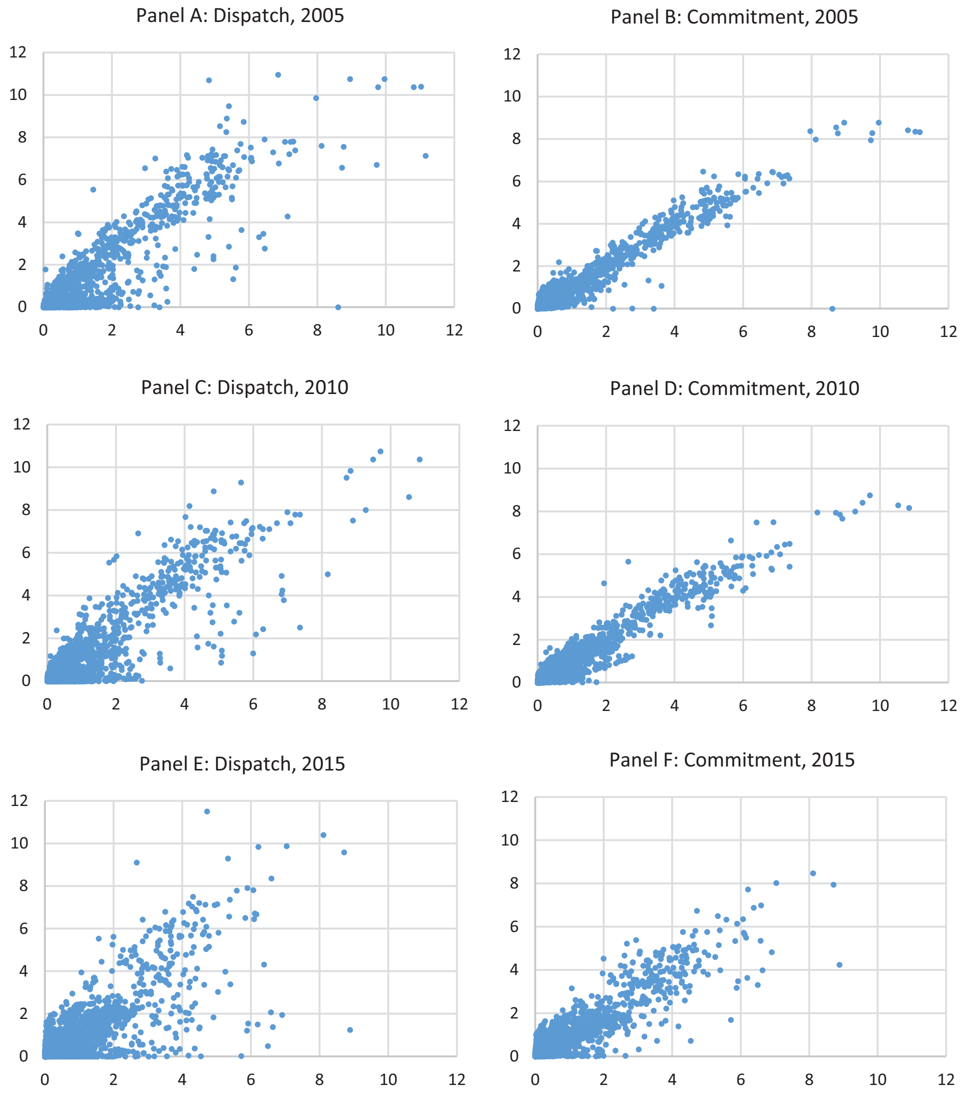

Notes: Two versions of the hourly operation model, dispatch and stylized unit commitment, are run using fuel prices and aggregate fossil fuel-fired generation for the years 2005, 2010, and 2015. In the dispatch version, units are dispatched each hour according to marginal costs. The unit commitment model includes stochastic unit availability, a reserve margin, minimum and maximum generation levels, and daily unit commitment. Panels A, C, and E plot simulated against observed generation for each unit, using the dispatch version. Panels B, D, and F plot simulated against observed generation for each unit, using the commitment version.

consumption, renewables generation, and environmental regulation. This section defines the scenarios that we analyze in the next section.

□ **Baseline.** The year 2005 represents the initial unit construction-retirement and abatement stages of the model. The year 2005 is chosen for reasons of data availability and regulatory history. As the geographic extent of emissions regulation expanded in the 2000s, so too did the coverage of the CEMS data. By 2005, CEMS included nearly all units that were eventually covered by the end of our data period, 2015. The year 2005 also represents the second year of the NO_x Budget Trading Program, which was the first time the regional cap-and-trade system expanded beyond the Northeast. Therefore, the 2005–2015 period contains most of the NO_x emissions reductions that the emissions programs have required.

Recall that a planning-style model, such as this one, is useful for comparing steady states. The operating phase consists of a single year of operating conditions, which is repeated to the infinite future and discounted to 2005. For several reasons, we use the most recent year for which we have data, 2015, to characterize the steady-state operating conditions. First, 2015 is the first year for which the CSAPR emissions caps applied, and these require deeper and geographically broader emissions reductions than the previous NO_x emissions caps. Second, the 2015 data allow us to compare steady states that use the observed outcomes with steady states using 2005 projections of 2015 outcomes. By comparing the steady states, we can evaluate the effects of the differences between projected and realized outcomes, which we refer to as shocks. Alternatively, we could define shocks using EIA's Annual Energy Outlook (AEO) projections from 2005 and 2015 and simulate the hourly phase for each year through 2030, discounting back to 2005; doing so does not affect the main conclusions.[21]

The baseline fuel prices and consumption growth are based on EIA projections from the 2005 AEO. For the eastern interconnection, the 2005 AEO projected a 23% increase in consumption, a 15% increase in wind generation, and a 19% decrease in the real price of natural gas between 2005 and 2015.

□ **Consumption, wind, and natural gas price shocks.** We define four scenarios around the consumption, wind, and natural gas price shocks. Year 2015 consumption in the eastern interconnection turned out to be 9% lower than 2005 consumption, versus the 23% increase between 2005 and 2015 that the EIA projected. The first scenario uses the observed electricity consumption growth rather than the projected level from the baseline scenario, which amounts to assuming that firms correctly anticipate the shock in 2005. Note that this scenario uses the projected wind and nonrenewables generation to compute aggregate fossil generation. We expect that lower electricity consumption reduces equilibrium electricity prices and generation of all units, reducing emissions and potentially causing the retirement of high-cost units.

The second shock is for wind generation. In 2005, the EIA projected a 15% increase in wind generation from the eastern interconnection between 2005 and 2015, but wind generation in 2015 was 16 times higher than it was in 2005. The second scenario uses the observed 2015 wind generation rather than the level of wind generation the EIA had projected in 2005. We do not include solar power generation, which has accounted for a negligible share of generation in the East, even in 2015. This shock reduces electricity prices and has the largest negative effect on generation from high-cost units that operate during hours when wind generation is high.

The third shock is that natural gas prices turned out to be lower than the EIA projected. This scenario uses observed 2015 fuel prices rather than the projected prices from the baseline scenario, replacing the 19% projected price decrease with the observed 50% price decrease. The fuel price shock includes the effects of shale gas as well as other demand and supply shocks in natural gas markets. The fuel price shock reduces generation costs for gas-fired units.

The first three scenarios treat each of the three shocks individually, and the fourth scenario has all three shocks simultaneously. Because three shocks occur simultaneously in the fourth scenario, this scenario illustrates their combined long-run effects. These scenarios allow us to

[21] The alternative approach also includes perfect foresight of the shocks. The baseline uses the EIA forecasts from the AEO 2005 for each year from 2005 through 2030. The hourly operational model is simulated each year and profits are discounted back to 2005. For the other scenarios, we use the observed values of the variables between 2005 and 2015, and the AEO 2015 projections for years 2005 through 2030. For example, in the wind scenario, we use observed wind generation in the eastern interconnection for the years 2005 through 2015, and the AEO 2015 forecast of wind generation for 2015 through 2030. In other words, we assume that the firm has perfect foresight for 2005 through 2015 and that the AEO 2015 correctly forecasts outcomes after 2015.

quantify the effects of each shock on emissions and generator profits in a hypothetical situation that does not include environmental regulations.

□ **Emissons caps and MATS.** We define four regulation scenarios. The first uses parameter assumptions from the baseline scenario and the 2015 CSAPR emissions caps. After simulating this scenario, we check that the summer emissions caps are not exceeded; if they are, we model the annual and summer caps jointly.[22] Comparing this scenario with the baseline allows us to estimate the expected costs of the NO_x regulations, given the EIA projections made in the 2005 AEO. This scenario corresponds to an analysis the EPA might have made had it created the 2015 CSAPR caps in 2005, without the intermediate caps under the preceding programs.

The second regulation scenario includes the 2015 CSAPR emissions caps and the consumption shock. The third regulation scenario includes CSAPR as well as the consumption, wind generation, and fuel price shocks. Comparing the three CSAPR scenarios allows us to quantify the effects of the shocks on the costs of CSAPR as well as on generator profits.

The final scenario includes MATS in addition to the three shocks and CSAPR. The fact that MATS was implemented after the market shocks simplifies the modelling, because we can model MATS using observed compliance decisions. We estimate abatement costs from the same EPA sources that we use for the SCR costs. This approach allows us to characterize the effects of MATS on emissions, profits, and retirements conditional on the market shocks and CSAPR.[23]

5. Results

■ In this section, we first compare the baseline scenario with the scenarios that include shocks to electricity consumption, wind generation, and fuel prices. Finally, we consider the scenarios that include environmental regulations.

□ **Consumption, wind, and natural gas price shocks.** The first column of Table 5 reports summary statistics from the baseline scenario. Recall that the baseline scenario uses EIA projections made in 2005 of fuel prices, renewables generation, and electricity consumption. Panel A shows the generation percentages by fuel type, with coal accounting for 74% of total generation and natural gas for 22% (oil accounts for the remaining 4%).

Column 2 in Table 5 reports the simulation results if we use the observed consumption rather than the projected consumption. The consumption shock has little effect on percentages of coal and natural gas in total generation and reduces capacity factors (Panel B) and profits (Panel C) for both natural gas- and coal-fired generation. Panel D shows that the consumption shock causes 1.1 GW of coal-fired plant retirements. Panel E indicates that it reduces NO_x emissions by 34%.

The wind scenario in column 3 uses the observed wind generation level, which was higher than the EIA projection. The increase in wind generation reduces capacity factors of both coal- and natural gas-fired units (Panel B) and by a larger percentage for natural gas than coal; consequently, the share of natural gas in total generation decreases slightly (Panel A). Wind generation has a larger effect on coal-fired plant profits than on natural gas-fired plant profits (Panel C) and reduces emissions by about 6%.

Column 4 shows that the lower natural gas prices, relative to projected prices, cause a substantial shift from coal- to natural gas-fired generation (Panel A). Because natural gas-fired

[22] We find no cases in which a state's summer cap is binding and the annual cap is not binding. Because SCR constitutes a lumpy investment, a single unit adding SCR or exiting could cause emissions to fall below the cap. In those cases, we assume that the SCR units are operated sufficiently often (i.e., less than 100% of hours) such that emissions exactly equal the cap. This is consistent with the observation that many units turned off SCR when natural gas prices and consumption growth fell in the early 2010s.

[23] In MATS, the EPA assesses compliance based on the emissions controls technologies installed, rather than by directly monitoring emissions. Our assumption that any controls installed after 2011 (when MATS was proposed) were installed because of MATS may cause us to overestimate the costs of MATS if they were installed because of state-level regulations.

TABLE 5 Effects of Shocks to Demand, Wind Generation, and Fuel Prices

	(1) Baseline	(2) Observed (lower) Consumption Growth	(3) Observed (higher) Wind Generation	(4) Observed Fuel Prices	(5) Observed Consumption, Wind, and Fuel Prices
Panel A: Generation percentage by fuel type					
Coal	74	75	75	63	58
Natural gas	22	23	21	36	42
Panel B: Mean capacity factor by fuel type					
Coal	0.77	0.66	0.73	0.66	0.50
Natural gas	0.29	0.23	0.25	0.46	0.42
Panel C: Mean annual operating profits by fuel type (million 2005 $ per megawatt)					
Coal	0.32	0.22	0.29	0.09	0.03
Natural gas	0.08	0.05	0.07	0.09	0.05
Panel D: Retirements (gigawatts of capacity)					
Coal		1.11	0.00	0.00	31.61
Natural gas		0.00	0.94	0.00	0.00
Oil		0.00	0.00	0.00	0.84
Panel E: Annual emissions (million tons)					
Nitrogen oxides	2.87	1.90	2.70	2.51	1.54
Sulfur dioxide	10.68	7.14	10.10	8.97	5.46
Carbon dioxide	2516	1864	2352	2375	1661

Notes: Each column reports the results of the scenario indicated in the column heading. See text for scenario definitions. Capacity factor and profits are capacity-weighted.

units often determine the electricity price, in many hours the electricity price falls in proportion to the heat rate of the marginal natural gas-fired unit. Consequently, the decrease in natural gas prices reduces equilibrium electricity prices, consistent with empirical evidence (Linn and Muehlenbachs, 2018). Profits of coal-fired units decrease because of the lower capacity factor and electricity price. Profits of natural gas-fired units increase only slightly because the decrease in equilibrium electricity prices offsets the increase in capacity factors (Panel B), and because low gas prices induce entry of gas-fired units. Because coal-fired units have higher emissions rates than natural gas-fired units, NO_x emissions decline by about 13%. The gas price shock causes fewer retirements than the consumption shock (Panel D).

Column 5 combines the consumption, wind, and fuel price scenarios. Comparing the results in columns 2 through 5 shows that all three shocks reduce coal capacity factors and profits. The consumption shock has a larger effect on emissions than the fuel price shock, but the fuel price shock has a larger effect on coal profits than the consumption shock. The three shocks combined cause 31 GW of coal-fired capacity retirements. The fact that the individual shocks cause few retirements but the combined shocks cause the retirements of more than 10% of the fleet implies that it is the combination of the shocks that is so damaging to coal-fired profits.

Figure A8 provides additional context for these results by plotting the estimated density functions for profits in the baseline, consumption, fuel price, and combined scenarios. If a unit retires, it does not operate and its profits are zero. Individually, the consumption and fuel price scenarios cause the distributions to shift leftward by substantial amounts. However, because coal-fired units were quite profitable in the absence of the shocks, neither of these shocks is sufficient by itself to cause profits to fall to zero. It is the combined effect of the shocks that causes the

TABLE 6 Emissions Cap Scenarios

	(1) Baseline	(2) State Emissions Caps	(3) Column (2) Plus Observed Consumption	(4) Column (3) Plus Observed Wind Generation and Fuel Prices	(5) Column (4) Plus MATS
Panel A: Generation percentage by fuel type					
Coal	74	73	75	56	55
Natural gas	22	23	23	44	45
Panel B: Mean capacity factor by fuel type					
Coal	0.77	0.77	0.66	0.48	0.47
Natural gas	0.29	0.29	0.23	0.44	0.45
Panel C: Mean annual operating profits by fuel type (million 2005 $ per megawatt)					
Coal	0.32	0.33	0.21	0.03	0.04
Natural gas	0.08	0.10	0.05	0.05	0.06
Panel D: Retirements (gigawatts of capacity)					
Coal		0.00	1.21	35.02	40.66
Natural gas		0.00	0.00	0.00	0.00
Oil		0.00	0.00	0.60	0.11
Panel E: Annual emissions (million tons)					
Nitrogen oxides	2.87	1.77	1.51	1.22	1.23
Sulfur dioxide	10.68	10.63	7.09	5.21	5.14
Carbon dioxide	2516	2511	1862	1638	1629
Panel F: Abatement costs					
Emissions price (2005 $/ton)		3109	1483	980	1101
Annualized costs (2005 billion $)		2.94	0.90	0.41	0.36

Notes: Each column reports the results of the scenario indicated in the column heading. See text for scenario definitions.

retirements, which is consistent with Fell and Kaffine (2018), who show that natural gas prices and wind power interact positively with one another to reduce coal-fired generation.

□ **Emissions caps and MATS.** Table 6 reports the simulation results for the three scenarios that include CSAPR and the fourth scenario that includes CSAPR and MATS, with the baseline repeated in column 1 for convenience. Column 2 includes the baseline assumptions and introduces both summer and annual emissions caps. The caps raised the marginal cost of coal-fired generation more than they raised the costs of natural gas-fired generation, which causes the generation share of coal to decrease slightly. The caps do not cause any retirements. The emissions caps decrease aggregate emissions by 38%, at an annual cost of $2.9 billion (all reported dollar numbers are in 2005 dollars).

Comparing columns 2 and 3 in Panel F shows that the consumption shock reduces the cost of the emissions caps by almost two thirds. The consumption shock reduces eastern NO_x emissions for two reasons. First, some fossil fuel-fired plants in the East are not subject to the cap, and the lower consumption reduces generation and emissions from the unregulated plants. Second, for

FIGURE 8

ESTIMATED DENSITY FUNCTIONS OF PERCENTAGE CHANGE IN PROFITS OF COAL-FIRED UNITS THAT CONTINUE OPERATING AFTER 2015, BETWEEN ALL AND BASELINE SCENARIOS

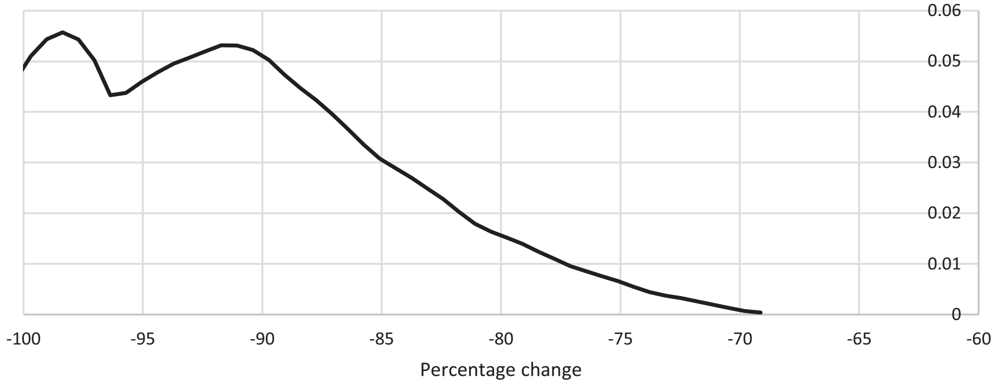

Notes: For each coal-fired unit, the percentage change in profits between the scenario that includes all market shocks and CSAPR caps and the baseline scenario is calculated. The figure plots the estimated density function of percentage changes for coal-fired units that operated in 2005 and 2015, according to the CEMS data.

states that have a binding summer cap but a nonbinding annual cap, lower consumption reduces fossil generation and emissions in nonsummer months.

Column 4 adds the wind and fuel price shocks to the scenario in column 3. In combination, the three shocks reduce the cost of the emissions cap by 86%, to \$0.4 billion per year. Comparing columns 1, 2, and 4 suggests that the three shocks explain nearly all the reduction in coal operating profits. Comparing column 4 in Tables 5 and 6 shows that, given the presence of the three market shocks, CSAPR caused 3.4 GW of coal-fired plant retirements.

Recall that we model the entire eastern interconnection. Because CSAPR covers most but not all of the eastern interconnection, the model includes some units that do not participate in CSAPR. Table A2 shows that the three shocks affect emissions differently for CSAPR states than for non-CSAPR states. For CSAPR states, CSAPR reduces NO_x emissions by 41%. Adding the market shocks to CSAPR further reduces emissions, but this additional reduction is smaller than the effect of CSAPR. For non-CSAPR states, the three market shocks reduce NO_x emissions by half.

Column 5 of Table 6 adds MATS to column 4. MATS has a small incremental effect on coal-fired plant profits and retirements. Just 20% of the total retirements in column 5 are caused by CSAPR and MATS (i.e., comparing column 5 in Tables 5 and 6). Note that the natural gas and coal generation shares in column 4 match observed 2015 levels for the eastern interconnection, confirming the accuracy of the simulation model.

To provide further information about the simulation results as well as validation of the entire model, Figure 8 illustrates the effects of the market shocks on coal-fired plant profits. For each coal-fired plant in the baseline scenario, we compute the percentage change in profits between the baseline scenario and the scenario that includes the three market shocks, CSAPR, and MATS. The percentage change is –100% for units that retire in the latter scenario. We plot the estimated density function of the percentage profit change for units that actually continue operating after 2015. For units that the model predicts will continue operating, the model correctly predicts this decision 96% of the time. All units that the model predicts to retire by 2015 are observed to retire (i.e., stop operating) by 2015. Note that this comparison includes only units that stopped operating by 2015, and not units that were planned to retire after 2015.

The figure also shows that even among the units that continue operating, profits decline by at least 70% for all units. Thus, the effects of the market shocks on profits were widespread across the coal fleet.

As we have noted, several of the modelling assumptions of CSAPR and MATS may cause us to overstate compliance costs. It may appear surprising that we accurately predict retirements despite this overestimate. The explanation is that the estimated compliance costs are about 0.6% of the total change in profits caused by the three market shocks; that is, CSAPR and MATS account for a negligible portion of the overall change in profitability of coal-fired plants. Because CSAPR and MATS explain such a small share of the profits changes, even if we overestimate CSAPR and MATS costs, we still correctly predict retirement decisions for nearly all plants. The two regulations cause 9 GW of retirements because the market shocks push many plants close to retirement.

6. Conclusions

■ Between 2005 and 2015, NO_x emissions from the US electricity sector decreased by about 8% per year, and emissions in 2015 were just two thirds what they were 10 years prior. Over the same period, firms prepared to retire about one third of coal-fired plant capacity. The causes of those emissions reductions have been the source of intense controversy in the public debate over environmental regulation. One view is that market shocks have reduced emissions and coal-fired plant profits, and that environmental regulation has reduced emissions substantially while having a relatively small effect on coal plant profits. The other view is that environmental regulation is the primary driver of declines in emissions and coal plant profits. Although the economics literature suggests that market forces have reduced coal-fired generation and profits on the margin, there is no direct evidence of the aggregate effects of market forces, or a comparison of market forces and environmental regulation.

We use a new computational model to assess whether either view is correct. The model covers 3500 fossil fuel-fired generation units in the eastern US electricity system and consists of three phases: unit construction and retirement, pollution abatement, and hourly operation. The operational phase approximates dynamic operating constraints and unit availability, as well as transmission congestion. The model reproduces observed changes in the extensive and intensive generation margins, fuel consumption, and emissions more accurately than a standard economic dispatch model and matches 97% of observed retirements that occurred by 2015.

We find that market shocks have larger effects than regulation on coal-fired plant profits. The consumption shock is about as important as the fuel price shock, both of which are more important than the wind generation shock. Combined, the market shocks explain 82% of the decline in NO_x emissions and 99% of the decline in coal-fired plant profits. The consumption shock explains a large share of the overall reduction in coal-fired plant profits, albeit a smaller share than the fuel price shock. The consumption shock reduces emissions 2.5 times more than does the fuel price shock, suggesting that both shocks played important roles in reducing NO_x compliance costs and in causing coal plant retirements. The importance of the natural gas price shock is consistent with the empirical literature (e.g., Holladay and LaRiviere, 2017). We believe that the literature has not previously quantified the importance of the consumption shock.

The three market shocks reduced the costs of CSAPR by 86%. Conditional on the market shocks, CSAPR and MATS explain 20% of retirements. As we have noted in the article, we make several simplifying assumptions about CSAPR and MATS that likely cause us to overestimate the costs of those regulations; relaxing the assumptions would strengthen our conclusions about the primacy of the market shocks in explaining retirements.

Appendix

The Appendix provides additional details on the model formulation and parameter estimation. We begin by defining the set of existing units at the beginning of phase 1 (retirement and new construction), as well as a set of potential entrants. Existing units are fossil fuel-fired units with positive generation in CEMS in the year 2005. Potential entrants include units that actually entered the system between 2005 and 2015, and those that were being planned in 2005 but did not actually enter. According to EIA 860, about 83 GW of new coal- and natural gas-fired capacity began operating between

2005 and 2015, and an additional 9 GW of coal- and natural gas-fired units were either in planning or construction in 2005 but did not actually begin operating before 2015. Capital costs and fixed operations and maintenance costs for each fuel type are from the EIA 2005 AEO. We have also attempted to estimate capital costs for new natural gas-fired units based on entry decisions of these units, yielding estimated capital costs similar to the EIA estimates.

Most of the unit characteristics are from the EPA. For each unit, the EPA data include state, NERC region, fuel type, rated capacity, initial year of operation, and whether the unit is connected to SCR. Fowlie (2010) analyzes several NO_x abatement technologies in addition to SCR, but most of these were widely installed at coal-fired units at the beginning of our sample. For that reason, we exclude these technologies from our analysis, as well as selective noncatalytic reduction, which few plants have installed.

We use EPA (2010) to estimate K_i^a, which is the capital cost of SCR. EPA (2010) documents the assumptions that are made in the agency's electricity sector model. For emissions abatement equipment such as SCR, the EPA has constructed a model that estimates the installation costs as a function of the unit's generation capacity, heat rate, and other factors. Because these factors vary across units, installation costs also vary. We annualize the capital costs assuming a 25-year lifetime of the equipment and a maximum 60-year life of the plant.

For phase 3 (hourly unit operation), we compute each unit's emissions rates of NO_x, sulfur dioxide, and carbon dioxide by computing the generation-weighted average across hours in 2005. Using the 2005 data yields the emissions rates at the beginning of the simulation period—that is, before subsequent abatement decisions are made. We compute an average heat rate for each unit using fuel consumption and generation from 2005 through 2015. Using the 11 years of data yields an average heat rate across a wide range of operating levels, accounting for the fact that heat rates tend to be higher at very low or high levels of operation. The nonpeak hours are within six hours of the peak hour in the corresponding day.

We obtain delivered fuel prices from EIA Forms 423 and 923. To reduce measurement error and concerns about the potential correlation among plant-level fuel prices and plant attributes that are not incorporated in the model, rather than using plant-level prices, we use average prices by NERC region and month. The plants used to construct the prices include publicly available data for traditionally regulated plants and proprietary data for unregulated plants (Cicala, 2015; Linn and Muehlenbachs, 2018). Future profits are discounted at a rate of 10%.

Aggregate fossil fuel-fired generation is computed as the sum of observed generation across all fossil fuel-fired units. Aggregate fossil generation is equal to consumer demand plus transmission line losses, net of generation from nonfossil technologies such as nuclear (i.e., we assume that demand is perfectly price-inelastic, which is a common assumption when modelling wholesale markets, such as by Bushnell, Chen, and Zaragoza-Watkins, 2014).

We use observed operation of each unit to estimate a minimum operating constraint. For each unit, we set the minimum generation level equal to the fifth percentile of generation observed across all hours between 2005 and 2015 in which the unit operates with positive generation. Figure A4 plots the estimated density functions of the distributions of minimum generation levels, separately for coal and gas units. For some units, particularly oil-fired and small gas-fired units, this level is close to zero. For most coal units, this minimum level corresponds to 30%–50% of rated capacity, which is consistent with assumptions made in many power system operational models in the academic literature (e.g., Castillo and Linn, 2011) and with assumptions used by industry.

Because we do not observe total hourly wind generation in the eastern interconnection, we estimate hourly wind generation using Pennsylvania, New Jersey, Maryland Interconnection (PJM) and EIA data. We use the average hourly wind generation capacity factor from the PJM market for the year 2015. The EIA 923 data include monthly wind generation for each wind generator in the eastern interconnection. We compute the monthly capacity factor for each wind plant using the total capacity from EIA 860.

We assume that each wind-powered plant has the same hourly variation within a month as in PJM. Using this assumption, we scale the hourly PJM capacity factor for each wind plant to match the monthly average capacity factor computed from the EIA data. Then, we use the wind plant's capacity to estimate hourly generation for the entire year. Estimated total hourly wind generation equals the cross-plant sum of estimated hourly wind generation.

□ **Estimation of nonfuel costs and transmission constraints.** Here, we discuss two particular parameterization challenges. First, whereas fuel costs can be estimated from observed data, nonfuel costs are not included in available data. Most computational models of the electric power system include *ad hoc* assumptions about nonfuel costs. For example, many researchers assume that nonfuel costs do not vary across units within a fuel type.

We extend the logic of Davis and Hausman (2016) to circumvent the data limitations. They argue that observed deviations from economic dispatch are due to transmission constraints. We extend this argument by observing that nonfuel costs and transmission congestion affect unit-level hourly generation in different ways from one another. Nonfuel costs affect the extensive margin—whether the unit is operating—at all levels of aggregate fossil generation. In contrast, transmission congestion affects generation at high levels of aggregate fossil generation, when the unit owner would like to operate the unit at full capacity but cannot do so because of transmission congestion. To illustrate this distinction, consider two particular coal-fired units in our data that are located in the same state and are similar in age, generating capacity, and heat rate; one unit nevertheless has a capacity factor twice that of the other unit over periods of moderate aggregate fossil generation. Because these differences in utilization rates occur at moderate aggregate fossil generation levels, they are not likely to be explained by transmission congestion (which should be most important at high levels of aggregate generation). Rather, differences in nonfuel costs are a likely explanation for the observed differences in utilization.[24]

FIGURE A1

EASTERN INTERCONNECTION AND CSAPR [Color figure can be viewed at wileyonlinelibrary.com]

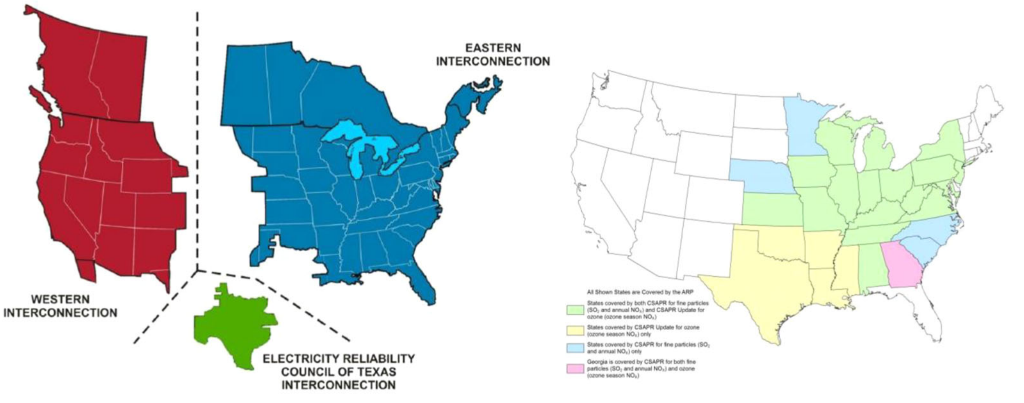

Notes: The map on the left shows the three major interconnections. The map on the right shows the states included in the Cross State Air Pollution Rule (CSAPR). The map on the left is from the North American Reliability Council, and the map on the right is from the Environmental Protection Agency (EPA).

TABLE A1 Changes in Unit Operation By Fuel Type, 2005–2008 versus 2009–2015

	Probability of Positive Generation			Capacity Factor Conditional on Positive Generation		
	2005–2008	2009–2015	Change	2005–2008	2009–2015	Change
Coal	0.85	0.73	−0.12	0.76	0.71	−0.06
Gas	0.26	0.33	0.07	0.61	0.65	0.05

Notes: The table uses data on hourly operation across all coal- and natural gas-fired units in the CEMS data (the natural gas units include steam, combined cycle, and large turbines). The table reports the probability a unit has positive generation and the capacity factor conditional on positive generation, across units indicated in the row headings and years indicated in the column headings. Probabilities and capacity factors are weighted by the unit's rated capacity.

Based on this reasoning, we estimate nonfuel costs using a simple regression. Focusing on hours in which aggregate generation lies between the 30th and 70th percentiles, for each unit, we compute the average share of hours the unit generates—that is, we ignore the intensive margin. During these hours, we expect transmission constraints not to bind, in which case, if we observe a unit operating less than would be predicted by its fuel costs, we would infer that the unit has high nonfuel costs. We omit days in which the unit does not operate to account for situations in which the unit is unavailable because of maintenance or other reasons. Using a separate sample for each fuel type, we regress the operation probability on the unit's heat rate, using the estimated residual and heat rate coefficient to estimate nonfuel costs.[25] Figure A3 shows the distribution of estimated nonfuel costs. The mean costs by fuel type are adjusted to match the mean costs by fuel type in the EIA National Energy Modeling System, to ensure consistency with the fixed-cost assumptions that are also from the EIA.[26] For example, if we predict a unit's nonfuel costs to be 50% higher than the mean of the

[24] The calculation of average capacity factors excludes days when the units did not operate, because of which differences in maintenance needs cannot explain the differences in capacity factors. Differences in supply of ancillary services across the two units could explain the differences in capacity factors, which we account for in the cost estimation as described in the next footnote.

[25] As noted above, capacity factors could vary across units for reasons other than nonfuel costs such as the provision of ancillary services. In that case, these services would be equivalent to having a negative nonfuel cost. We assume that the provision of these services does not vary across policy scenarios described in Section 4.

[26] The adjustments are about $4 per MWh for each fuel type. These adjustments are small compared with the fuel costs, which tend to be several times larger—and an order of magnitude larger for natural gas. Moreover, the estimated variation in nonfuel costs is substantially smaller than the variation in fuel costs; the coefficient of variation for nonfuel costs is about one quarter that for fuel costs. In practice, the level and variation of nonfuel costs have little effect on the main results: the results are similar if we simply use EIA estimates of nonfuel costs or if we do not adjust our nonfuel costs to match the means of the EIA nonfuel costs.

FIGURE A2

ESTIMATED EMISSIONS RATE DISTRIBUTIONS BY FUEL TYPE, 2005

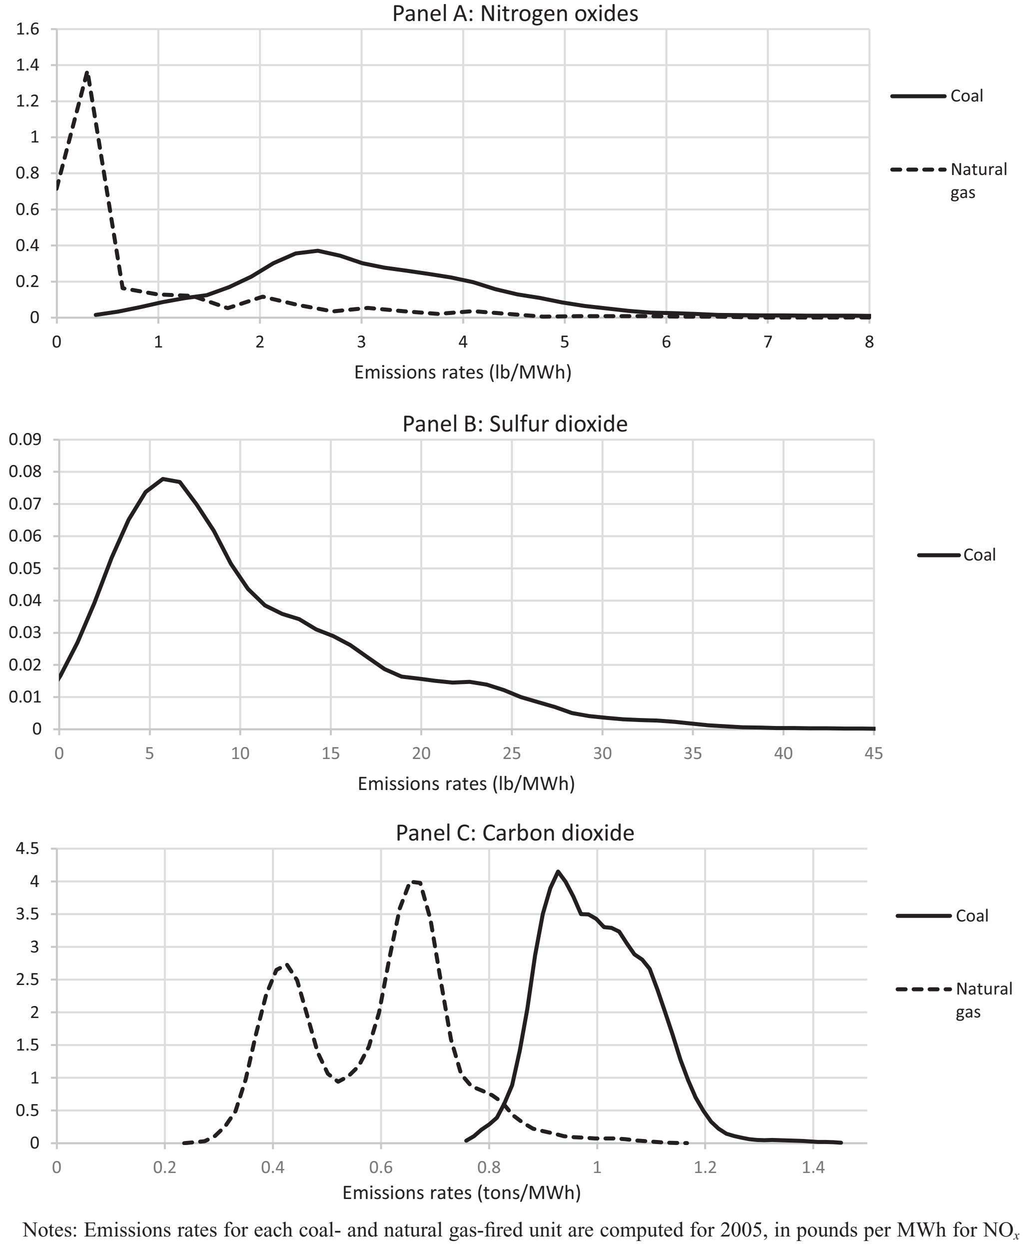

Notes: Emissions rates for each coal- and natural gas-fired unit are computed for 2005, in pounds per MWh for NO_x (Panel A) and sulfur dioxide (Panel B) and in tons per MWh for carbon dioxide (Panel C). A unit's emissions rate equals total annual emissions divided by total annual generation, including hours with positive emissions and generation. The figure plots estimated density functions of the emissions rates by fuel type.

predicted costs across all units, we adjust that the unit's nonfuel costs such that the costs are 50% higher than the EIA costs.

The second challenge is to account for transmission congestion. We use the intensive margin to estimate the constraints that congestion or other unit-specific factors place on the operation of a unit. Once a unit is operating, we would expect it to operate either at its lowest available level (if marginal costs exceed the electricity price) or at its highest possible level (if the electricity price exceeds marginal costs). Therefore, observing the unit operating at less than full capacity, but above its minimum level, implies that the unit is facing transmission congestion or some other operating

TABLE A2 Emissions for CSAPR and Other States

	(1) Baseline	(2) State Emissions Caps	(3) Column (2) Plus Observed Consumption	(4) Column (3) Plus Observed Wind Generation and Fuel Prices	(5) Column (4) Plus MATS
Panel A: Changes in nitrogen oxides emissions (million tons)					
CSAPR states	2.67	−1.10	−1.31	−1.56	−1.53
Other states	0.20	0.00	−0.05	−0.10	−0.11
Total emissions	2.87	−1.10	−1.36	−1.65	−1.64
Panel B: Changes in sulfur dioxide emissions (million tons)					
CSAPR states	10.26	−0.05	−3.46	−5.22	−5.26
Other states	0.42	0.00	−0.13	−0.25	−0.29
Total emissions	10.68	−0.05	−3.59	−5.47	−5.54
Panel C: Changes in carbon dioxide emissions (million tons)					
CSAPR states	2368.27	−5.59	−613.36	−835.27	−840.18
Other states	147.81	0.79	−41.05	−43.07	−46.60
Total emissions	2516.08	−4.80	−654.41	−878.34	−886.77

Notes: Column 1 reports the emissions in the baseline scenario, and columns 2–5 report comparisons between the scenarios indicated in the column headings and the baseline scenario.

FIGURE A3

ESTIMATED DENSITY FUNCTIONS OF NONFUEL OPERATING COSTS

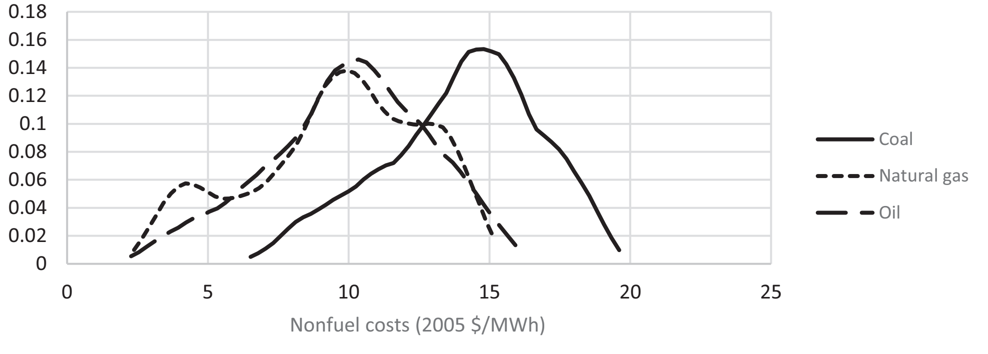

Notes: Nonfuel operating costs are estimated as described in the text. The figure plots the estimated density function of nonfuel costs (in 2005 $/MWh) by fuel type.

constraint. For example, inefficiencies may arise in dispatch because of market rules or regulation, which would cause actual dispatch to differ from merit-order dispatch. Our reduced-form approach accounts for these factors because we use the unit's observed maximum generation levels, which reflect those factors, rather than assuming that units can run at full capacity in all hours.

These constraints may vary with the overall level of aggregate fossil generation, and we compute deciles of the aggregate fossil generation distribution. For each decile and unit, we determine the 95th percentile of generation during hours that fall within the decile. This calculation determines $\overline{g}_{ih}$, which is the maximum generation level by unit and hour. As in Davis and Hausman (2016), we assume that the counterfactuals we consider in Section 5 do not affect $\overline{g}_{ih}$.[27]

[27] Firms may choose not to operate at full capacity during certain hours if they desire to maintain some capacity for reserve markets. Therefore, the estimated hourly maximum capacity factors may reflect this consideration as well as transmission constraints. Implicitly, we assume that the scenarios modelled in Section 5 do not affect the amount and source of capacity that is withheld for reserves.

FIGURE A4

ESTIMATED DENSITY FUNCTIONS OF MINIMUM CAPACITY FACTOR FOR COAL- AND GAS-FIRED UNITS

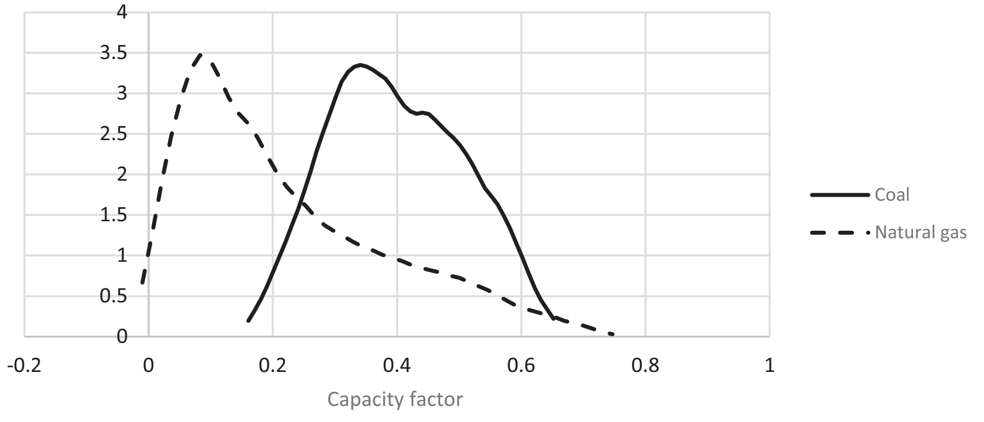

Notes: Minumum capacity factor is the ratio of the unit's minimum generation level to its rated capacity. The figure plots the estimated density functions of minimum capacity factor for coal- and gas-fired units.

FIGURE A5

SIMULATED VERSUS OBSERVED ANNUAL UNIT GENERATION BY NERC REGION, 2005 (MILLION MWH) [Color figure can be viewed at wileyonlinelibrary.com]

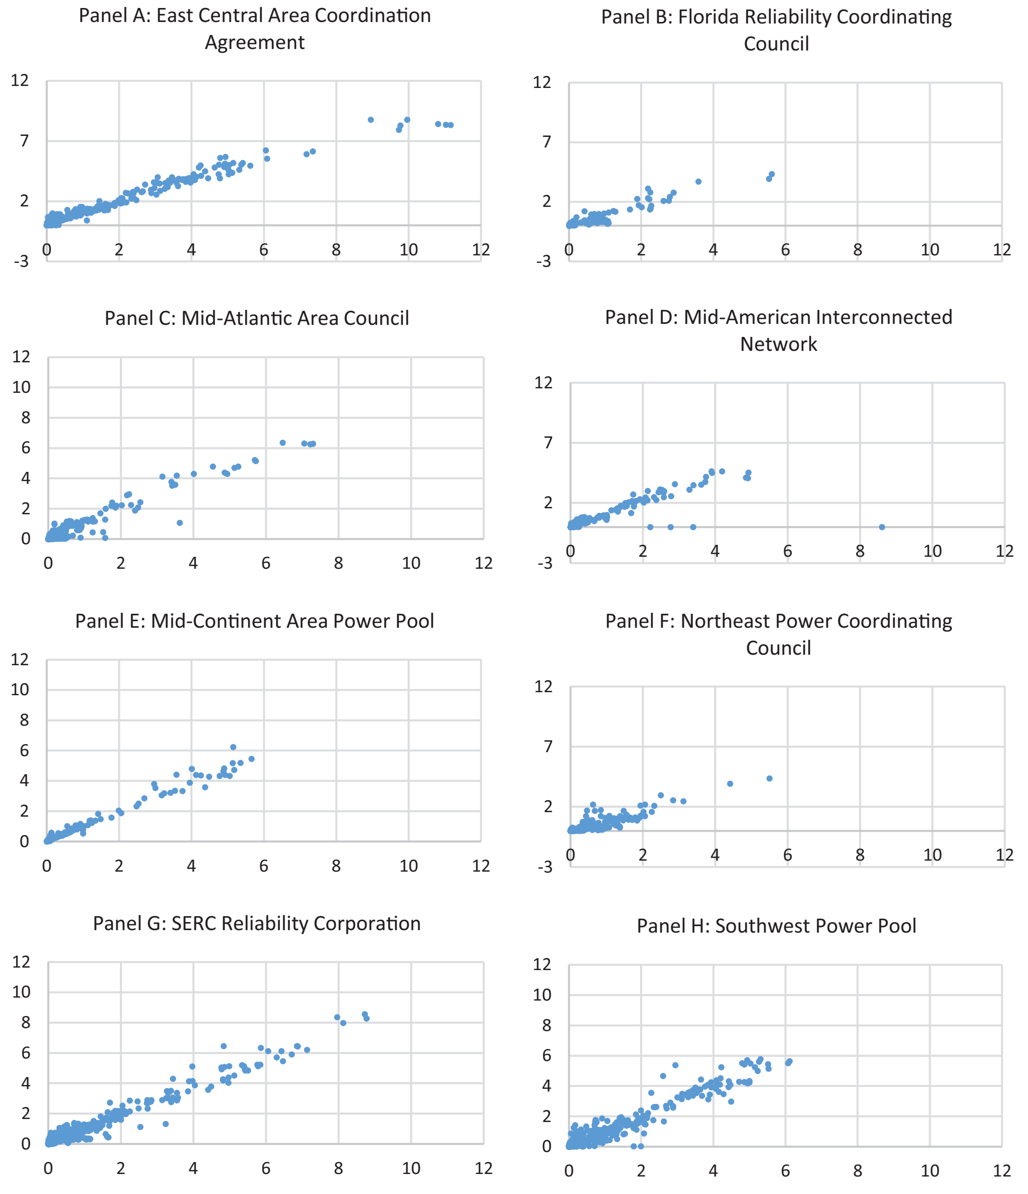

Notes: The figure plots the simulated against observed generation using the same results as in Panel F of Figure 8 (2015, unit commitment model). Each panel includes units in the indicated NERC region.

FIGURE A6

SIMULATED VERSUS OBSERVED ANNUAL UNIT GENERATION BY NERC REGION, 2015 (MILLION MWH) [Color figure can be viewed at wileyonlinelibrary.com]

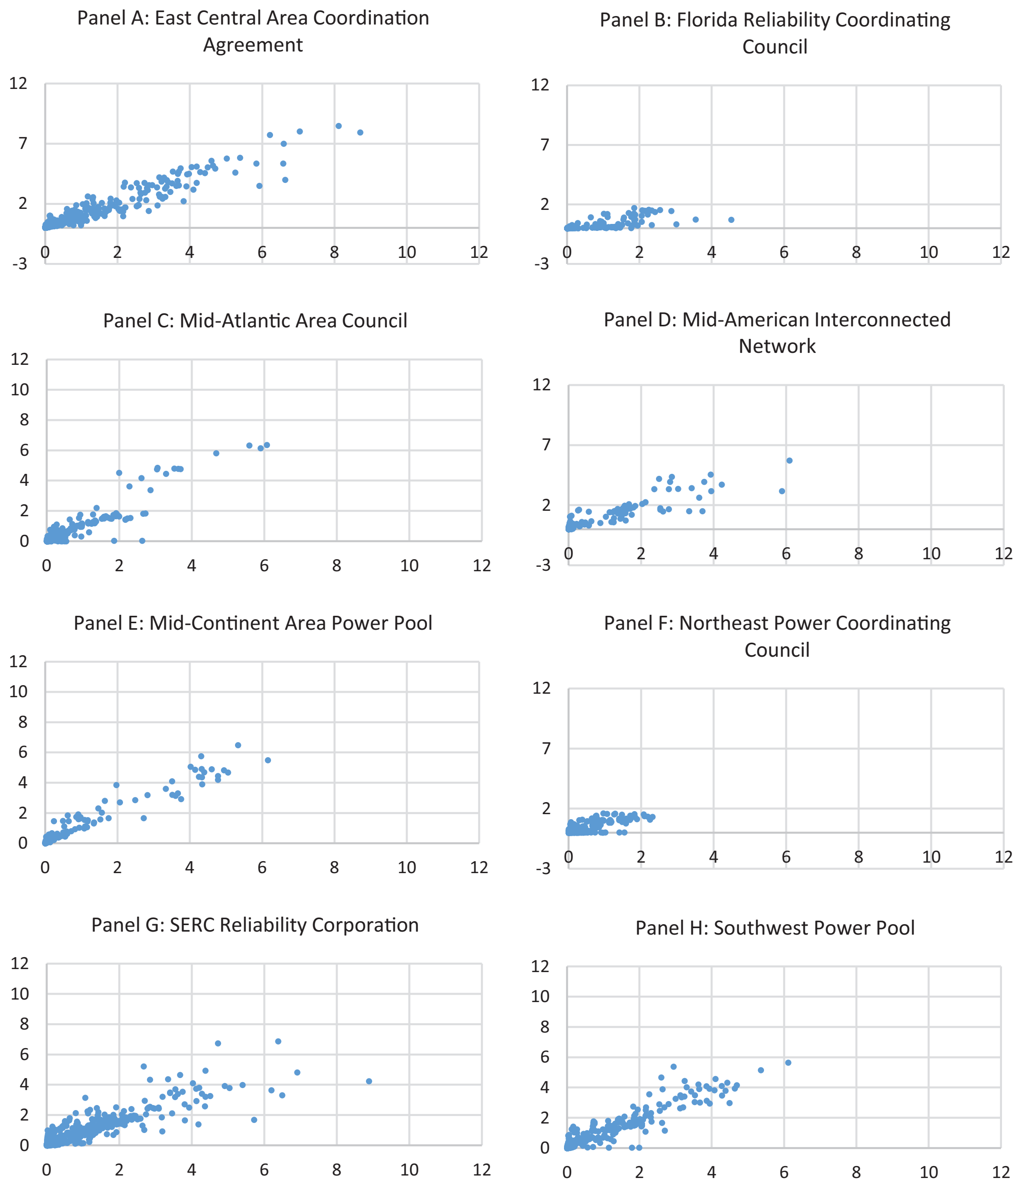

Notes: The figure plots the simulated against observed generation using the same results as in Panel F of Figure 8 (2015, unit commitment model). Each panel includes units in the indicated NERC region.

FIGURE A7

SIMULATED VERSUS OBSERVED ANNUAL UNIT GENERATION BY FUEL TYPE, 2015 (MILLION MWH)
[Color figure can be viewed at wileyonlinelibrary.com]

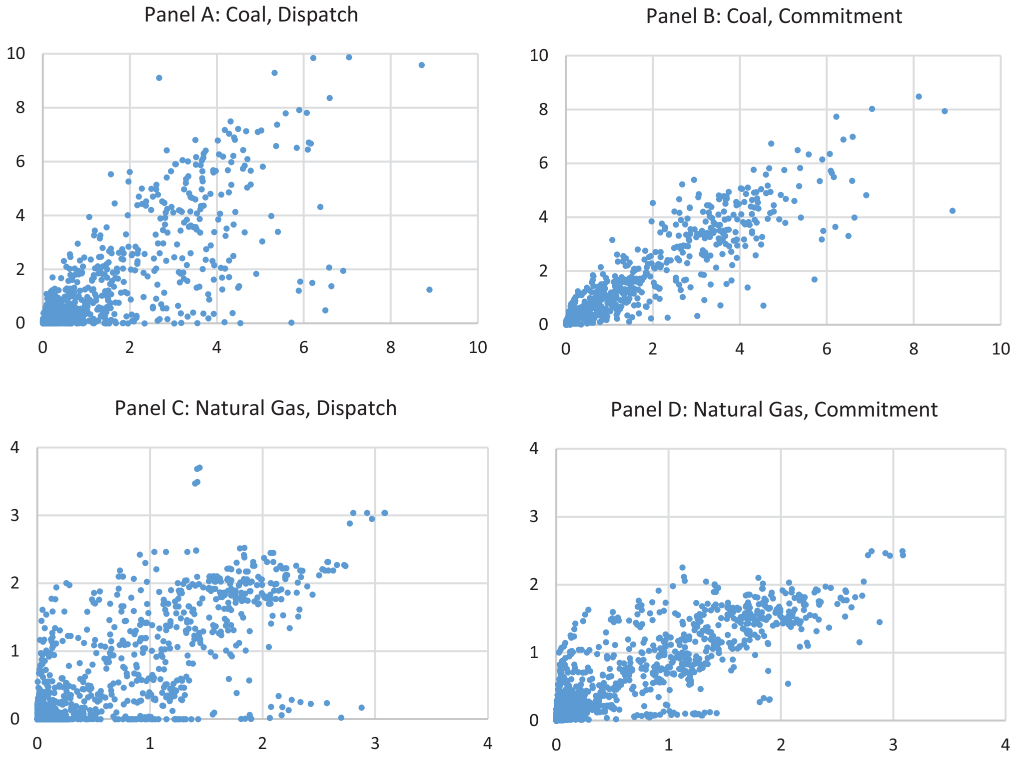

Notes: The dispatch and commitment simulations are the same as those shown in panels E and F of Figure 8. The panels in this figure show simulation results for the indicated fuel type

FIGURE A8

ESTIMATED DENSITY FUNCTIONS OF PROFITS OF COAL-FIRED UNITS

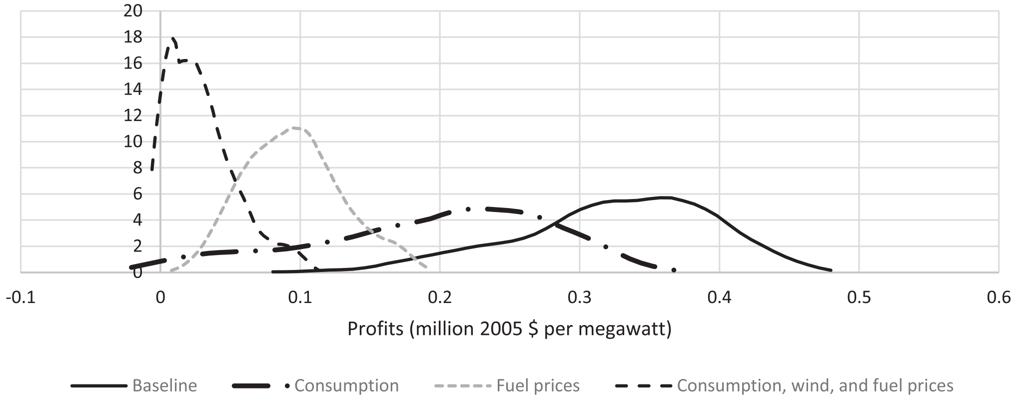

Notes: The figure plots the estimated density function of profits per megawatt of capacity from the indicated scenarios. Zero profits indicates that the unit is retired.

References

Blanford, G., Merrick, J., and Young, D. "A Clean Energy Standard Analysis with the US-REGEN Model." Special issue, *Energy Journal*, Vol. 35 (2014).

Borenstein, S., Bushnell, J.B., and Wolak, F.A. "Measuring Market Inefficiencies in California's Restructured Wholesale Electricity Market." *American Economic Review*, Vol. 92 (2002), pp. 1376–1405.

Borenstein, S. and Holland, S. "On the Efficiency of Competitive Electricity Markets with Time-Invariant Retail Prices." *RAND Journal of Economics*, Vol. 36 (2005), pp. 469–493.

Burtraw, D., Evans, D.A., Krupnick, A., Palmer, K., and Toth, R. "Economics of Pollution Trading for SO_2 and NO_x." Discussion Paper no. 05-05, Washington DC: Resources for the Future, 2005.

Burtraw, D. and Palmer, K. "The Paparazzi Take a Look at a Living Legend: The SO_2 Cap-and-Trade Program for Power Plants in the United States." Discussion Paper no. 03–15, Washington DC: Resources for the Future, 2003.

Bushnell, J., Chen, Y., and Zaragoza-Watkins, M. "Downstream Regulation of CO_2 Emissions in California's Electricity Sector." *Energy Policy*, Vol. 64 (2014), pp. 313–323.

Bushnell, J., Mansur, E., and Saravia, C. "Vertical Arrangements, Market Structure, and Competition: An Analysis of Restructured US Electricity Markets." *American Economic Review*, Vol. 98 (2008), pp. 237–266.

Carlson, C., Cropper, M., Palmer, K., and Burtraw, D. "SO2 Control By Electric Utilities: What Are the Gains from Trade?" *Journal of Political Economy*, Vol. 108 (2000), pp. 1292–1326.

Castillo, A. and Linn, J. "Incentives of Carbon Dioxide Regulation for Investment in Low-Carbon Electricity Technologies in Texas." *Energy Policy*, Vol. 39 (2011), pp. 1831–1844.

Cicala, S. "When Does Regulation Distort Costs? Lessons from Fuel Procurement in U.S. Electricity Generation." *American Economic Review*, Vol. 105 (2015), pp. 411–444.

Cullen, J. and Mansur, E. "Inferring Carbon Abatement Costs in Electricity Markets: A Revealed Preference Approach Using the Shale Revolution." *AEJ: Economic Policy*, Vol. 9 (2017), pp. 106–133.

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

Curtis, E.M. "Who Loses under Cap-and-Trade Programs? The Labor Market Effects of the NOx Budget Trading Program." *Review of Economics and Statistics*, Vol. 100 (2018), pp. 151–166.

Davis, L. and Hausman, C. "Market Impacts of a Nuclear Plant Closure." *AEJ: Applied Economics*, Vol. 8 (2016), pp. 92–122.

Department of Energy (DOE). Staff Report to the Secretary on Electricity Markets and Reliability. Washington DC, 2017.

Ellerman, A.D., Joskow, P.L., Schmalensee, R., Montero, J.-P., and Bailey E.M. *Markets for Clean Air: The U.S. Acid Rain Program*. Cambridge: Cambridge University Press, 2000.

Environmental Protection Agency (EPA). Regulatory Impact Analysis for the NOx SIP Call, FIP, and Section 126 Petitions. Washington DC, 1998.

———. Regulatory Impact Analysis for the Final Clean Air Interstate Rule. Washington DC, 2005.

———. Documentation for EPA Base Case v.4.10 Using the Integrated Planning Model. EPA 430R10010. Washington DC, 2010.

———. Regulatory Impact Analysis for the Federal Implementation Plans to Reduce Interstate Transport of Fine Particulate Matter and Ozone in 27 States; Correction of SIP Approvals for 22 States. Washington DC, 2011.

Fell, H. and Kaffine, D. "The Fall of Coal: Joint Impacts of Fuel Prices and Renewables on Generation and Implications for Policy." *AEJ: Economic Policy*, Vol. 10 (2018), pp. 90–116.

Fell, H. and Linn, J. "Renewable Electricity Policy, Intermittency, and Cost-Effectiveness." *Journal of Environmental Economics and Management*, Vol. 66 (2013), pp. 688–707.

Fowlie, M. "Emissions Trading, Electricity Restructuring, and Investment in Pollution Abatement." *American Economic Review*, Vol. 100 (2010), pp. 837–869.

Fowlie, M. and Muller, N. "Market-Based Emissions Regulation When Damages Vary across Sources: What Are the Gains from Differentiation?" Working Paper no. 18801, Cambridge, MA: National Bureau of Economic Research, 2013.

Holladay J.S. and LaRiviere, J. "The Impact of Cheap Natural Gas on Marginal Emissions from Electricity Generation and Implications for Energy Policy." *Journal of Environmental Economics and Management*, Vol. 85 (2017), pp. 205–227.

Houser, T., Bordoff, J., and Marsters, P. "Can Coal Make a Comeback?" New York: Columbia Center on Global Energy Policy, 2017.

Johnsen, R., LaRiviere, J., and Wolff, H. "Estimating Indirect Benefits: Fracking, Coal, and Air Quality." *Journal of the Association of Environmental and Resource Economists* (forthcoming).

Linn, J. "The Effect of Cap-and-Trade Programs on Firms' Profits: Evidence from the Nitrogen Oxides Budget Trading Program." *Journal of Environmental Economics and Management*, Vol. 59 (2010), pp. 1–14.

Linn, J., Burtraw, D., and McCormack, K. "The Supreme Court's Stay of the Clean Power Plan: Economic Assessment and Implications for the Future." *Environmental Law Reporter*, Vol. 46 (2016), pp. 10759–10772.

Linn, J. and Muehlenbachs, L. “The Heterogeneous Impacts of Low Natural Gas Prices on Consumers and the Environment.” *Journal of Environmental Economics and Management*, Vol. 89 (2018), pp. 1–28.

Mansur, E. “Do Oligopolists Pollute Less? Evidence from a Restructured Electricity Market.” *Journal of Industrial Economics*, Vol. 55 (2007), pp. 661–689.

Reguant, M. “Complementary Bidding Mechanisms and Startup Costs in Electricity Markets.” *Review of Economic Studies*, Vol. 81 (2014), pp. 1708–1742.

Ryan, S. “The Costs of Environmental Regulation in a Concentrated Industry.” *Econometrica*, Vol. 80 (2012), pp. 1019–1061.

Shapiro, J. and Walker, R. “Why Is Pollution from U.S. Manufacturing Declining? The Roles of Environmental Regulation, Productivity, and Trade.” *American Economic Review* (forthcoming).

Stavins, R. “The Effects of Vintage-Differentiated Regulation.” Discussion Paper no. 05–12, Washington DC: Resources for the Future, 2005.

Wang, B. and Hobbs, B.F. “Real-Time Markets for Flexiramp: A Stochastic Unit Commitment-Based Analysis.” *IEEE Transactions on Power Systems*, Vol. 31 (2016), pp. 846–860.

Zhou, Y. “Emission Responses to Carbon Pricing in Dynamic Electricity Markets.” Job Market Paper. Storrs: University of Connecticut, 2016.