

The Roles of Energy Markets and Environmental Regulation in Reducing Coal-Fired Plant Profits and Electricity Sector Emissions

Joshua Linn (University of Maryland and Resources for the Future)

Kristen McCormack (Harvard University)*

March 2018

Abstract

Between 2005 and 2015, US electricity sector emissions of nitrogen oxides, 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 that accurately reproduces unit operation, emissions, and retirement, we find that electricity consumption and gas prices account for nearly all the coal plant profitability decline and resulting retirements. Environmental regulations explain most of the emissions reductions but had little effect on coal plant profitability and retirement.

Key words: nitrogen oxides, shale gas, renewables, energy efficiency, cap-and-trade

JEL codes: L94, Q5, Q4

* Linn is an associate professor at the University of Maryland and a senior fellow at Resources for the Future (linn@rff.org) and McCormack (kmccormack@g.harvard.edu) is a PhD candidate at Harvard University. We thank David Evans, Catherine Hausman, and participants at the Association of Environmental and Resource Economists Summer Conference for comments.

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. Underlying these changes are the tightening stringency and broadening scope of NO_x emissions caps that the US Environmental Protection Agency (EPA) administers. Coinciding with the emissions trends, many coal-fired plants became unprofitable and about one-third of coal-fired plants retired, raising concerns about electricity service reliability (i.e., the ability to provide electricity to consumers on demand; DOE 2017).

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 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. Asserting that the emissions caps have improved the environment and human health, many adherents of the first view favor tightening the emissions caps in light of lower than expected compliance costs due to declining costs of natural gas and renewables. The second view argues 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 announced its intention to reduce regulation and support 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 estimate the effects of environmental regulation on the coal-fired fleet. Several recent articles examine the statistical relationship among natural gas prices, wind generation, fossil fuel-fired generation, and emissions (e.g., Cullen and Mansur 2017; Fell and Kaffine, forthcoming; Linn and Muehlenbachs 2016; Holladay and LaRiviere 2017; Johnsen et al. 2016). 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. The long run effects on emissions, coal plant profits, and retirements may differ from the short run effects. On the one hand, in the long run low gas prices may raise natural gas plant investment, potentially compounding the short-run effects on coal plant profits and retirements. On the other hand, in the long run gas prices interact with the emissions caps. Emissions caps for NO_x were binding in the mid-2010s with emissions credit prices trading at several hundred dollars per ton, implying that gas prices and wind generation do not affect emissions in the long run (i.e., for the regions covered by the caps) but do affect emissions credit prices. For example, low gas prices would reduce emissions credit prices, reducing costs for coal plants and opposing the short run effects of gas prices on coal plant profits. That is, the short run and marginal effects of natural gas prices and renewables –which have been the focus of the literature–may differ from the long run effects, either positively or negatively.¹

Moreover, while several studies have compared expected and realized costs of sulfur dioxide emissions reductions under the Acid Rain Program (e.g., Carlson et al. 2000; Ellerman et al.

¹ Houser et al. (2017) compare the effects of electricity consumption, natural gas prices, and renewables on coal consumption between 2006 and 2016. Rather than comparing counterfactuals they make back-of-the-envelope calculations from Energy Information Administration projections, 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. DOE (2017) argues that natural gas prices are the most important factor explaining coal plant retirements, but provides little evidence supporting this conclusion other than the timing of events.

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 it established the caps, EPA expected that NO_x emissions caps would cost the sector at least \$3 billion per year (2005 dollars), representing a large share of overall estimated costs of federal environmental regulations of the electricity sector.² These costs fall largely on the coal-fired fleet (Linn 2010), implying roughly a 10 percent cost increase for those units, with most of the costs incurred by the older and higher-emitting coal-fired units. Moreover, the Mercury and Air Toxics Standards (MATS) impose emissions standards for pollutants other than NO_x, and were expected to cost several times more than the NO_x caps. Thus, it seems possible that environmental regulation could have substantially reduced profitability of coal-fired units, leading to their retirement. Only Fowlie and Muller (2013) analyze the costs of achieving NO_x emissions caps, but they consider only the early portion of the program and they do not evaluate whether either view of the electricity sector trends is correct.

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

We consider three market shocks: natural gas prices, renewables generation, and electricity consumption. Largely because of the rise of production from shale formations, natural gas prices in 2015 were 30 percent lower than projections made in 2005 predicted them to be. Improved wind generator performance and subsidies caused wind generation in 2015 to be 10 times higher

² Between 2003 and 2015, EPA implemented the emissions caps in three phases, as described in Section 2. The agency reports costs of complying with each phase (EPA 1998, 2005, and 2009). For the latter two phases the costs are combined with the costs of achieving the sulfur dioxide caps. As a conservative estimate, the cost number in the main text uses only the cost estimate from EPA (1998).

³ The United States contains three major interconnections, across which there is little available transmission. Throughout the paper, East refers to the eastern interconnection, which spans the Great Plains to the East Coast.

than had been expected. Because of the 2008–9 economic recession and other factors, 2015 electricity consumption was 20 percent below 2005 expectations. For convenience, we refer to differences between 2015 realized outcomes and 2005 projections of those outcomes as energy market shocks, noting that policies have contributed to them.

We use a new operational and investment model of the eastern US power system to explore the causes of the observed decline in emissions and coal-fired plant profitability. The model includes 3,500 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 natural gas prices.

We model the NO_x emissions caps that were adopted between 2005 and 2015, which require that emissions in 2015 equal about half of 2005 levels. Using projections made in 2005 of electricity consumption, wind generation, and fuel prices, we estimate that without shocks, abatement costs would have been about \$2.9 billion per year, which roughly agrees with ex-ante EPA assessments.⁴

The three shocks collectively reduced regulatory costs from \$2.9 billion to \$0.4 billion per year (86 percent) and reduced coal-fired plant profits by 89 percent. The shocks reduced coal use by 41 percent, with the electricity consumption shock explaining the majority of that decline. These shocks explain nearly all of the coal-fired plant retirements observed between 2005 and

⁴ We cannot compare our estimated costs directly with EPA estimates because the EPA cost estimates cannot be combined, as noted above. The agency reports costs of complying with each of three regulatory phases, but the costs are estimated relative to different baselines, making it inappropriate to add the three cost estimates.

2015. After accounting for these shocks, the emissions caps had a negligible effect on coal-fired plant profits and retirements. Because NO_x emissions from part of the East are not subject to the emissions caps, we report separate results for uncapped regions. In those regions, the market shocks reduced NO_x emissions by half and coal-fired plant profits by 64 percent. These results therefore confirm the first of the two views, that factors other than NO_x emissions caps explain most of the decline in the profits of coal-fired plants and the resulting retirements. Although we do not model environmental regulations other than the NO_x emissions caps, particularly MATS, in Section 6 we argue that including them would not affect this finding.

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 directly affect emissions as long as the emissions caps continue to bind.

Like the findings in the recent literature (e.g., Fell and Kaffine, forthcoming), our results confirm the importance of natural gas prices on coal- and gas-fired generation. In contrast to the empirical literature, we show that the natural gas price shock had little effect on NO_x emissions, demonstrating the importance of accounting for long-run interactions between emissions caps and market shocks. That is, 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 paper builds on the extensive literature that has used structural models of the electricity sector to address economic and environmental questions (e.g., Borenstein et al. 2002; 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.

2. Background

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

2.1 Overview of regulating electricity sector NO_x emissions

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), EPA established air quality standards for NO_x and ground-level ozone that reflect the maximum ambient level of the pollutant to protect human health and welfare. States submit plans to demonstrate their strategies for meeting the standards. The CAA also authorizes EPA to create emissions standards for certain sources.

Between the passage of the CAA in 1970 and the late 1980s, these regulations and state plans proved to be ineffective at reducing NO_x emissions in absolute terms. 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—the fact that, because of prevailing winds, emissions in one location can affect air quality hundreds of miles away. Indeed, the CAA created incentives for firms to construct tall smokestacks at their power

plants, which improved local air quality but exacerbated downwind air quality problems (Burtraw and Palmer 2003).

In the 1980s, policymakers became increasingly aware of the contributions of NO_x and sulfur dioxide to acid rain. The shortfalls of the initial regulations and 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, 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, EPA created the Clean Air Interstate Rule to replace the NO_x Budget Trading Program in 2009 and 2010. The new 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. However, in 2008, the US Court of Appeals ruled that the Clean Air Interstate Rule was “fundamentally flawed.” The court was concerned that the regional emissions caps could not prevent a situation in which generation units in a state purchase and use so many emissions credits that the state does not achieve its

required emissions reduction. The Clean Air Interstate Rule remained in place while the Cross-State Air Pollution Rule (CSAPR) was developed and underwent judicial review. The CSAPR program restricts cross-state credit trading to prevent the situation the court highlighted, and it began capping NO_x emissions from 27 states and the District of Columbia in 2015. Thus, over time the NO_x emissions caps have expanded geographically and increased in stringency.

2.2 Data

The main source of data is the EPA Continuous Emissions Monitoring System (CEMS). The data set comprises nearly all fossil fuel-fired units operating in the eastern interconnection. Using 2005–15 CEMS data, we compile hourly fuel consumption and generation; hourly emissions of NO_x, sulfur dioxide, and carbon dioxide; and unit characteristics for each fossil fuel-fired generation unit. Unit characteristics include the state in which the unit is located, whether the unit has specific NO_x 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.

2.3 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 national NO_x emissions between 1990 and 2015. Emissions declined by 75 percent during this period, with most of the decline occurring after 2000. For comparison, the figure shows that sulfur dioxide emissions declined by 85 percent between 1990 and 2015.

In this paper, we focus on the eastern interconnection, which accounts for about 90 percent of national electricity sector NO_x emissions. Between 2000 and 2015, NO_x emissions in the East declined by nearly 75 percent, mirroring the national trend during that period.

Total NO_x emissions from the East, in tons, equal the total generation multiplied by the average rate of emissions, in tons per megawatt hour (MWh) of generation. Therefore, reductions in total generation or average emissions rates could explain the declining emissions. Figure 2 shows that total generation in the East increased steadily between 2001 and 2007, at about 2 percent 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 the expanded use of energy efficiency, electricity generation in 2015 was slightly lower than generation in 2009. The figure illustrates that fossil fuel-fired generation, which accounts for nearly all NO_x emissions from the electricity sector, experienced a similar leveling off of generation growth after 2007. The fact that fossil fuel-fired generation in 2015 was the same level as in 2001 implies that changes in average emissions rates, and not total fossil generation, explain the emissions decline.

The average emissions rate could decline because of reductions in emissions rates at individual units or because of a generation shift to lower-emitting fuels. Coal-fired units have steadily adopted technology that reduces emissions, such as selective catalytic reduction (SCR), which reduces emissions rates by roughly 90 percent. Generation shifts have also contributed to the decline in average emissions rates. Coal-fired units typically have higher emissions rates than gas-fired units. Figure 3 depicts the shift from coal- to gas-fired generation that occurred between 2000 and 2015, which reduced the average emissions rate across fossil generation units. The increase in the wind generation share, from close to zero in 2000 to 4 percent in 2015, further reduced emissions. In short, changes in unit emissions rates and shift from coal to cleaner fuels contribute 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 in generation shares. Figure 4 shows that about 90 gigawatts (GW) of coal-fired capacity retired between 2005 and 2015, which accounts for almost one-third of the initial capacity. Table 1 compares the attributes of coal-fired units that retire between 2005 and 2015, with those that continue operating. The retiring units tend to be smaller, older, less efficient, and less heavily utilized than the continuing units.

Figure 5 illustrates that after 2008, natural gas prices declined relative to coal and oil prices, reducing the relative cost of using natural gas to generate electricity. Cullen and Mansur (2017), Fell and Kaffine (forthcoming), and others have shown that lower natural gas prices increased generation from natural gas-fired units and decreased generation from coal-fired units.

We summarize the recent developments in the eastern electricity system by comparing projections of the electricity system made in 2005 and 2015. In Figure 6, we compare projections that the EIA made in the 2005 Annual Energy Outlook, with outcomes between 2005 and 2015. Compared with the 2005 projections of 2015 outcomes, in 2015 natural gas prices were 25 percent lower, generation from renewables was 2.5 times higher, and total electricity sector generation was 15 percent lower. These differences represent the unanticipated changes in fuel prices, renewables generation, and aggregate electricity consumption that occurred during the 10-year period. We use the term *shocks* to describe the difference between observed 2015 outcomes and 2005 projections of 2015 outcomes. The projections through 2025 show that these shocks were permanent rather than temporary.

2.4 A few stylized facts about generator emissions and hourly operation

This subsection describes several patterns in unit-level emissions and generation that we observe in the CEMS data; an objective of our model is to reproduce these patterns. In principle, 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, Appendix Figure 1 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 2005. The observed variation in NO_x emissions rates within fuel types suggests that accurately 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 2 demonstrates that generation units experienced changes along both the extensive and intensive margins between 2005 and 2015. 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 shows that between 2005 and 2008, coal units operated 85 percent of all hours on average. Between 2009 and 2015, coal units operated 73 percent of all hours, representing a 12 percentage point decline along the extensive margin. Across the same two years, natural gas-fired unit operation increased by 7 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 6 percentage points for coal. For natural gas, conditional capacity factors increased by 5 percentage points.

In short, the data indicate substantial variation over time and across aggregate fossil generation levels 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 modeling 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 models of the electricity system, such as Bushnell et al. (2014), with attributes of the reduced-form model in Davis and Hausman (2016). The model reproduces observed unit-level operation and emissions, as well as abatement and plant retirement decisions.

3.1 Overview of model structure

The model consists of three phases: retirements and new construction, pollution abatement investment, and hourly operation. We find an equilibrium by solving the model backwards. The hourly operation phase is simulated over a single representative year (i.e., 8,760 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.⁵ Intuitively, an increase in natural gas prices reduces profits of coal-fired units in the operation phase, particularly the inefficient coal-fired units, and potentially causing the retirement of those units. Because pollution abatement has high fixed costs and low marginal operating costs, typically higher-emitting, larger, and higher capacity factor units install abatement equipment, making it more likely for smaller and lower capacity 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.

⁵ 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, such an approach is computationally infeasible.

Second, while 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 and expected future state variables (e.g., Mansur 2007), prices are not available in much of the eastern United States, particularly in the Southeast.⁶ Consequently, a dynamic model would omit much of the eastern emissions.

3.2 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 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 are negative, where profits equal the difference between discounted revenues and costs. Costs include pollution abatement costs, variable costs, and fixed operations and maintenance costs. The model includes 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 synonymous with 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 is a standard metric in the electric power sector and 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 include permitting and construction costs, as well as costs and revenue from the abatement and operation phases.

⁶ 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 forthcoming).

3.3 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. Similar to 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.⁷

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 abatement equipment. For simplicity, we focus on SCR, noting that the main conclusions would be similar if we included other abatement technologies such as low-NO_x burners. 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

⁷ The modeled emissions program includes a few simplifications relative to CSAPR. First, CSAPR allows a limited amount of emissions credit trading across states. To reduce computational burden, we assume that there is free emissions credit trading within states but no trading across states. Second, we assume that each state's credit market is perfectly competitive. This assumption is consistent with EPA analysis of CSAPR, but in practice, firms may have market power in credit markets. Note that the planning model structure avoids the need to model credit banking.

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.⁸

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 MWh of generation) in the absence of the abatement controls, 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).

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

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

where SCR capital costs are K_i^a and 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 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 are greater than 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 as long as their average abatement costs do not exceed the emissions credit price. The price therefore adjusts so that in equilibrium, credit demand equals credit supply.

⁸ 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).

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 the SCR in non-summer months to avoid operating costs. Many states have both annual and summer caps; in these cases, 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 summer emissions price.

3.4 Phase 3: Hourly operation

The operational stage of the model represents a steady state. We characterize hourly operation over a single year (i.e., 8,760 hours), and that year is repeated into the infinite future. Revenues and costs are discounted back to the retirement and construction phase of the model. This setup is typical of the planning models cited above.

We build a unit commitment style 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 a unit's generation level across hours. Unfortunately, it is computationally infeasible to combine the first two model phases with an hourly unit commitment model for the entire eastern interconnection. Therefore, we build a simplified unit commitment model for tractability, approximating a unit commitment model's key 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 temporal 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, \bar{g}_{ih} . The maximum generation level varies across hours and units, and the minimum level varies across units.

We introduce an hourly cap on fossil generation, which is exogenous to the model. Recall that aggregate fossil generation excludes generation from nuclear, hydroelectric, and renewables. Following Bushnell et al. (2014), among others, we assume that generation from these technologies 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. These technologies have very low marginal operating costs and therefore generate as much electricity as technically possible.⁹ Hydroelectric plants, on the other hand, can be dispatched to some

⁹ For eastern nuclear units, between 2005 and 2015 we observe little variation in annual capacity factors and few trends in monthly capacity factors, supporting this exogeneity assumption.

extent subject to environmental and other constraints. However, in the East, hydroelectric plants accounted for just 3 percent of power generation in 2005, and this fact, combined with the limited dispatchability of hydroelectric generation, suggests that the exogeneity assumption has little effect on the main results.

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.¹⁰ We distinguish among three types of hours: the peak hour, when electricity demand reaches its daily maximum; near-peak hours, which are within 6 hours of the peak hour; and off-peak hours, which include all other hours in the same day.¹¹

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

¹⁰ Some models in the literature allow for the possibility that US electricity markets are imperfectly competitive (e.g., Borenstein et al. 2002; Mansur 2007). In such an environment, firms account for the effect of their generation on equilibrium prices and restrict their generation to increase prices. However, it has become increasingly common in the literature to assume perfect competition (e.g., Borenstein and Holland 2005; Blanford et al. 2014; Zhou 2016), reflecting expansion of the geographic scope of wholesale power markets and other factors that have increased competition among firms. Although cases of imperfect competition may persist in US electricity markets, the perfect competition assumption is a commonly used approximation.

¹¹ The main conclusions are unaffected if we use longer spans for the near-peak hours.

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, \bar{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 non-peak hours. Therefore, the firm submits a peak hour price that is greater than its marginal operating costs. Dynamic models with startup or shutdown costs, such as Bushnell et al. (2008), similarly yield equilibriums 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.

During off-peak hours, the equilibrium is determined according to economic dispatch. Units operate at \bar{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.

Finally, 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. Many electricity systems include spinning reserves, which are units that are available to supply electricity in the event that other units are unexpectedly unavailable or if realized aggregate fossil generation exceeds forecast aggregate fossil generation. We include spinning reserves by introducing 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.

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 do not operate at all otherwise. As demand varies across hours, units start up and shut down so 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 startup and shutdown costs multiple times each day.¹²

Importantly, in the unit commitment model, exogenous factors (e.g., fuel prices) may affect both the 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 like in a dispatch model). Finally, we allow for uncertainty in unit availability and aggregate fossil fuel-fired generation.

3.5 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 the attributes of actual and planned units to characterize the attributes of 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 percent 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.

¹² Recall that small 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 yields similar results to those reported in the paper.

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; have similar age, generating capacity, and heat rate; and yet one unit 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.¹³

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—i.e., ignoring the intensive margin.

¹³ The calculation of average capacity factors excludes days in which 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.

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 due to 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.¹⁴ Appendix Figure 2 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, which ensures consistency with the fixed cost assumptions that are also from EIA.

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 constraint. 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 \bar{g}_{ih} , which is the maximum generation level by unit and hour.

¹⁴ 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.

As in Davis and Hausman (2016), we assume that the counterfactuals we consider in Section 5 do not affect \bar{g}_{ih} .¹⁵

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 (8,760 hours) 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 re-simulated using the smaller set of generation units and the credit prices estimated from the previous iteration.¹⁶

¹⁵ Firms may choose not to operate at full capacity during certain hours if they desire to maintain some capacity available 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 modeled in Section 5 do not affect the amount and source of capacity that is withheld for reserves.

¹⁶ 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 were to account for strategic behavior in exit. We have considered other exit rules, such as randomly choosing the exiting unit from units with negative profits or choosing the unit whose profits increase the most from another unit's exit (this is feasible only in simulations that include a small number of units with negative profits). These alternatives yield similar results to those reported in the paper; because of this, we use the simpler exit rule.

3.6 Validation of the abatement and hourly operation stages

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 the year 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 7 uses data and assumptions from 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 and 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. Note that 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. Unobserved factors, such as compliance with local regulations and EPA enforcement of New Source Review, may explain SCR installation. These factors are exogenous to the model.

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. In the dispatch model a unit's ability to generate in one hour does not depend on its generation in any other hour; i.e., there is no distinction between 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 3 shows that the percentage of coal-fired generation predicted by the unit commitment style model matches the observed percentages more closely than do the predictions of the dispatch model (results for other years are available upon request; percentage differences between the simulated and 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 3 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 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 4 shows that the unit 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 5 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 percent of units on average.

Figure 8 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*-squared is typically about 0.9 for the commitment model and 0.7 for the dispatch model. Consistent with the results in Table 5, Figure 8 shows that the dispatch model is more likely to predict zero generation than is the unit commitment model.

We approximate operation constraints imposed by transmission congestion using the maximum generation levels for each unit. We assume that these 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. Figure 9 and Appendix Figures 4 and 5 compare predicted and observed outcomes at more disaggregated levels than in Figure 8, illustrating results by 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 2005–15 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, supports the ability of the model to 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 from one another according to the projected fuel prices, electricity consumption, renewables generation, and regulation. This section defines the scenarios that we analyze in the next section.

4.1 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 the year 2005, CEMS includes 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 2005–15 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 projections of 2015 outcomes that were made in 2005. 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 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 (results available upon request).

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 percent increase in consumption, a 15 percent increase in wind generation, and a 19 percent decrease in the real price of natural gas between 2005 and 2015.

4.2 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 percent lower than 2005 consumption, as compared with the 23 percent increase between 2005 and 2015 that EIA projected. The first scenario uses the observed electricity consumption growth rather than the projected level from the baseline scenario. 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 costs units.

The second shock is for wind generation. In 2005, EIA projected a 15 percent 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 EIA had projected in 2005. We do not include solar power generation, which accounts 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 EIA projected. The third scenario uses observed 2015 fuel prices rather than the projected prices from the baseline scenario, replacing the 19 percent projected price decrease with the observed 50 percent price decrease. The fuel price shock includes the effects of shale gas as well as other demand and supply shocks in natural gas markets. This shock reduces generation costs for gas-fired units, reducing electricity prices and generation from high cost coal-fired units (Linn and Muehlenbachs, forthcoming).

The first three scenarios include each of the three shocks individually, and the fourth scenario includes all three shocks simultaneously. Because three shocks occur simultaneously in the fourth scenario, this scenario illustrates the long-run effects of the shocks; as noted above, this class of simulation model does not make it possible to represent the transitional dynamics that arise from the fact that the consumption shock began prior to the fuel price shock. These scenarios allow us to quantify the effects of each of the shocks on emissions and generator profits in a hypothetical situation that does not include the emissions caps.

4.3 Emissions caps

We define three emissions 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.¹⁷ 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 EPA might have made had it created the 2015 CSAPR caps in 2005, without the intermediate caps under the NO_x Budget Trading Program or Clean Air Interstate Rule. Thus, the emissions scenario includes nearly all of the emissions reductions required under the three phases of EPA emissions caps.

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

Note that the scenarios do not include other emissions regulations besides the CSAPR NO_x caps, such as sulfur dioxide emissions caps or the MATS. Thus, the three CSAPR scenarios correspond to a hypothetical in which only NO_x emissions regulations changed after 2005. This allows us to isolate the effects of the NO_x caps and to estimate the effects of the market shocks on the costs of these emissions caps. In Section 6, we discuss the potential interactions between the CSAPR NO_x caps and MATS.

¹⁷ We do not find any 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 percent of hours) so that emissions exactly equal the cap. This is consistent with the observation that many units turned off their SCR when natural gas prices and consumption growth fell in the early 2010s.

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. Subsequently, we estimate the effects of the NO_x caps and show how the shocks to consumption, wind generation, and fuel prices affected the costs of the NO_x caps, as well as coal plant profits and retirements.

5.1 Consumption, wind, and natural gas price shocks

The first column of Table 6 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 percent of total generation and natural gas for 22 percent (oil accounts for the remaining 4 percent).

Column 2 in Table 6 reports the simulation results if we use the observed consumption rather than 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 indicates that the consumption shock reduces NO_x emissions by 34 percent. The consumption shock reduces coal consumption by 24 percent (not shown).

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 percent.

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 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, forthcoming). 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, emissions decline by about 13 percent.

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 does, but the fuel price shock has a larger effect on coal profits than the consumption shock does. Combined, the three shocks reduced coal consumption by 41 percent (not shown).

5.2 Emissions caps

Table 7 reports the simulation results for the three scenarios that include CSAPR, repeating the baseline in column 1 for convenience. Column 2 includes the baseline assumptions and introduces both summer and annual emissions caps. The caps cause a small amount of coal-fired retirements, which causes the generation share of coal to decrease slightly. The emissions caps decrease aggregate emissions by 38 percent, at an annual cost of \$2.9 billion (all reported dollar numbers are in 2005 dollars).

Comparing columns 2 and 3 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 states that have a binding summer cap but a nonbinding annual cap, lower consumption reduces fossil generation and emissions in non-summer 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 percent, to \$0.4 billion per year. Comparing columns 1, 2, and 4 suggests that the three shocks explain nearly all of the reduction in coal operating profits. Note that the natural gas and coal generation shares in column 4 match observed 2015 levels for the eastern interconnection, further confirming the accuracy of the simulation model.

The three shocks affect emissions differently for CSAPR states than for non-CSAPR states. In Appendix Table 1, column 1 reports the estimated emissions in the baseline scenario. The remaining columns show the change in emissions relative to the baseline. For CSAPR states, CSAPR reduces NO_x emissions by 41 percent. Adding the market shocks to CSAPR further reduces emissions, but this additional reduction is smaller than the effect of CSAPR.

Panels B and C show that CSAPR slightly reduces sulfur dioxide and carbon dioxide emissions because of the small reduction in coal-fired generation (as we discuss later, the sulfur dioxide emissions caps are not binding in 2015). In CSAPR states, the consumption, wind, and fuel price shocks have comparatively larger effects than CSAPR on sulfur dioxide and carbon dioxide emissions.

In non-CSAPR states, CSAPR has zero direct effect on NO_x emissions. The consumption shock reduces NO_x emissions by 25 percent in non-CSAPR states, compared with 8 percent in CSAPR states. Thus, relative to a scenario that includes only CSAPR, the three market shocks primarily reduce coal plant profits and compliance costs in CSAPR states and reduce coal plant profits and NO_x emissions in non-CSAPR states. The shocks reduce sulfur dioxide and carbon dioxide emissions in both CSAPR and non-CSAPR states.

To provide further information about the simulation results as well as validation of the entire model, Figure 10 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 and the CSAPR caps. The percentage change is -100 percent 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 actually continue operating, the model correctly predicts this decision 97 percent of the time. For the units that actually retire between 2005 and 2015, the model predicts percentage profit changes of -100 percent for all units—i.e., the model correctly predicts retirement for all units that actually retire.

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

6. Conclusions

Between 2005 and 2015, NO_x emissions from the US electricity sector decreased by about 8 percent per year, and emissions in 2015 were just two-thirds what they were 10 years prior. Over the same period, firms retired 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 have used a new computational model to assess whether either view is correct. The model covers 3,500 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 outcomes more accurately than a standard economic dispatch model and matches 98 percent of observed retirements of coal-fired generation units.

We find that market shocks have larger effects than regulation on emissions, coal consumption, and 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 percent of the decline in NO_x emissions and 99 percent 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. This importance of the natural gas price shock is consistent with the empirical literature (e.g., Holladay and LaRiviere 2017), and we believe that the literature has not previously quantified the importance of the consumption shock.

The analysis focuses on NO_x emissions regulation and considers neither the sulfur dioxide caps nor MATS. Although these regulations did not affect NO_x emissions directly, they could have affected emissions indirectly by causing retirements of coal-fired plants. In principle, this effect could have reduced the compliance costs of the NO_x regulation, similarly to the market shocks considered here.

However, the sulfur dioxide caps are unlikely to have affected NO_x emissions. Unlike NO_x emissions credit prices, which have been several hundred dollars per ton in the mid-2010s, sulfur dioxide emissions credit prices have been close to zero in the 2010s. These prices suggest that the sulfur dioxide emissions caps have not been binding, likely because of the market shocks and the fact that complying with MATS reduced sulfur dioxide emissions (Burtraw et al. 2012). Consequently, the sulfur dioxide caps have not affected long-run profits or sulfur dioxide emissions.

On the other hand, MATS could have affected coal plant retirements. However, Burtraw et al. (2012) conclude that MATS had a much smaller effect on retirements than did natural gas prices.¹⁸ Moreover, according to EPA, the expected costs of MATS were two to four times greater than the expected costs of NO_x emissions caps that we model. Using the factor of four, the effect of MATS on coal-fired plant profits would have been one-sixth the combined effect of the consumption, wind, and fuel price shocks. Thus, accounting for MATS would not affect the conclusion that market forces rather than regulation explain the declining coal-fired plant profitability and rising retirements.

7. References


- Blanford, G., J. Merrick, and D. Young. 2014. A Clean Energy Standard Analysis with the US-REGEN Model. Special issue, *Energy Journal* 35.
- Borenstein, S., Bushnell, J.B., and F.A. Wolak. 2002. Measuring Market Inefficiencies in California's Restructured Wholesale Electricity Market. *American Economic Review* 92: 1376–1405.
- Borenstein, S., and S Holland. 2005. On the Efficiency of Competitive Electricity Markets with Time-Invariant Retail Prices. *RAND Journal of Economics* 36: 469–93.
- Burtraw, D., D. A. Evans, A. Krupnick, K. Palmer, and R. Toth. 2005. Economics of Pollution Trading for SO₂ and NO_x. Discussion paper 05-05. Washington, DC: Resources for the Future.

¹⁸ In principle, MATS could have affected fuel prices or electricity consumption by reducing relative demand for coal and increasing generation costs. These effects would be included in the fuel price and consumption scenarios, although they are likely to be small given the conclusions of Burtraw et al. (2012).

- Burtraw, D., and K. Palmer. 2003. The Paparazzi Take a Look at a Living Legend: The SO₂ Cap-and-Trade Program for Power Plants in the United States. Discussion paper 03-15. Washington, DC: Resources for the Future.
- Burtraw, D., K. Palmer, A. Paul, and M. Woerman. 2012. Secular Trends, Environmental Regulations, and Electricity Markets. *Electricity Journal* 25: 35–47.
- Bushnell, J., Y. Chen, and M Zaragoza-Watkins. 2014. Downstream Regulation of CO₂ Emissions in California’s Electricity Sector. *Energy Policy* 64: 313–23.
- Bushnell, J., E. Mansur, and C. Saravia. 2008. Vertical Arrangements, Market Structure, and Competition: An analysis of Restructured US Electricity Markets. *The American Economic Review* 98: 237–66.
- Carlson, C., M. Cropper, K. Palmer, and D. Burtraw. 2000. SO₂ Control by Electric Utilities: What are the Gains from Trade? *Journal of Political Economy* 108: 1292–1326.
- Castillo, A., and J. Linn. 2011. Incentives of Carbon Dioxide Regulation for Investment in Low-Carbon Electricity Technologies in Texas. *Energy Policy* 39: 1831–44.
- Cicala, S. 2015. When Does Regulation Distort Costs? Lessons from Fuel Procurement in U.S. Electricity Generation. *American Economic Review* 105: 411–44.
- Cullen, J., and E. Mansur. 2017. Inferring Carbon Abatement Costs in Electricity Markets: A Revealed Preference Approach Using the Shale Revolution. *AEJ: Economic Policy* 9 (3): 106–33.
- Cullen, J. A., and S. S. Reynolds. 2016. The Long Run Impact of Environmental Policies on Wholesale Electricity Markets: A Dynamic Competitive Analysis. University of Arizona Working Paper.
- Davis, L., and C. Hausman. 2016. Market Impacts of a Nuclear Plant Closure. *AEJ: Applied Economics* 8: 92–122.
- DOE (US Department of Energy). 2017. Staff Report to the Secretary on Electricity Markets and Reliability.
- Ellerman, A. D., P. L. Joskow, R. Schmalensee, J.-P. Montero, and E. M. Bailey. 2000. *Markets for Clean Air: The U.S. Acid Rain Program*. Cambridge: Cambridge University Press.
- EPA (US Environmental Protection Agency). 1998. Regulatory Impact Analysis for the NO_x SIP Call, FIP, and Section 126 Petitions.
- EPA. 2005. Regulatory Impact Analysis for the Final Clean Air Interstate Rule.
- EPA. 2010. Documentation for EPA Base Case v.4.10 Using the Integrated Planning Model. EPA #430R10010. Washington, DC: EPA.
- EPA. 2011. 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.
- Fell, H., and D. Kaffine. Forthcoming. The Fall of Coal: Joint Impacts of Fuel Prices and Renewables on Generation and Implications for Policy. *AEJ: Economic Policy*.
- Fell, H. and J. Linn. 2013. Renewable Electricity Policy, Intermittency, and Cost-Effectiveness. *Journal of Environmental Economics and Management* 66: 688–707.
- Fowlie, M. 2010. Emissions Trading, Electricity Restructuring, and Investment in Pollution Abatement. *American Economic Review* 100: 837–69.
- Fowlie, M., and N. Muller. 2013. Market-Based Emissions Regulation When Damages Vary across Sources: What Are the Gains from Differentiation? NBER Working Paper No. 18801. Cambridge, MA: National Bureau of Economic Research.


- Holladay J. S., and J. LaRiviere. 2017. The Impact of Cheap Natural Gas on Marginal Emissions from Electricity Generation and Implications for Energy Policy. *Journal of Environmental Economics and Management* 85: 205–27.
- Houser, T., J. Bordoff, and P. Marsters. 2017. Can Coal Make a Comeback? New York: Columbia Center on Global Energy Policy.
- Johnsen, R., J. LaRiviere, and H. Wolff. 2016. Estimating Indirect Benefits: Fracking, Coal, and Air Quality. Discussion paper 10170. Bonn, Germany: Institute for the Study of Labor (IZA).
- Linn, J. 2010. The Effect of Cap-and-Trade Programs on Firms' Profits: Evidence from the Nitrogen Oxides Budget Trading Program. *Journal of Environmental Economics and Management* 59: 1–14.
- Linn, J., and L. Muehlenbachs. Forthcoming. The Heterogeneous Impacts of Low Natural Gas Prices on Consumers and the Environment. *Journal of Environmental Economics and Management*.
- Mansur, E. 2007. Do Oligopolists Pollute Less? Evidence from a Restructured Electricity Market. *Journal of Industrial Economics* 55: 661–89.
- Stavins, R. 2005. The Effects of Vintage-Differentiated Regulation. Discussion paper 05-12. Washington, DC: Resources for the Future.
- Wang, B., and B. F. Hobbs. 2016. Real-Time Markets for Flexiramp: A Stochastic Unit Commitment-Based Analysis. *IEEE Transactions on Power Systems* 31: 846–60.
- Zhou, Y. 2016. Emission Responses to Carbon Pricing in Dynamic Electricity Markets. Job Market Paper. Storrs: University of Connecticut.

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.


Figure 2. Total and Fossil Fuel-Fired Generation in the Eastern Interconnection, 2001–15




Note: The figure plots total net generation and fossil fuel-fired generation, in billion MWh, from EIA Form 920.

Figure 3. Generation Share by Energy Source for Eastern Interconnection, 2000–15

Panel A: Coal, gas, and nuclear




Panel B: Hydro, wind, and other




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.

Figure 4. Cumulative Retirements and Capacity Additions in the East, 2005–15



Note: Cumulative retirements and capacity additions by fuel type, in gigawatts (GW), are computed from EIA Form 860 for the years 2005 through 2015.

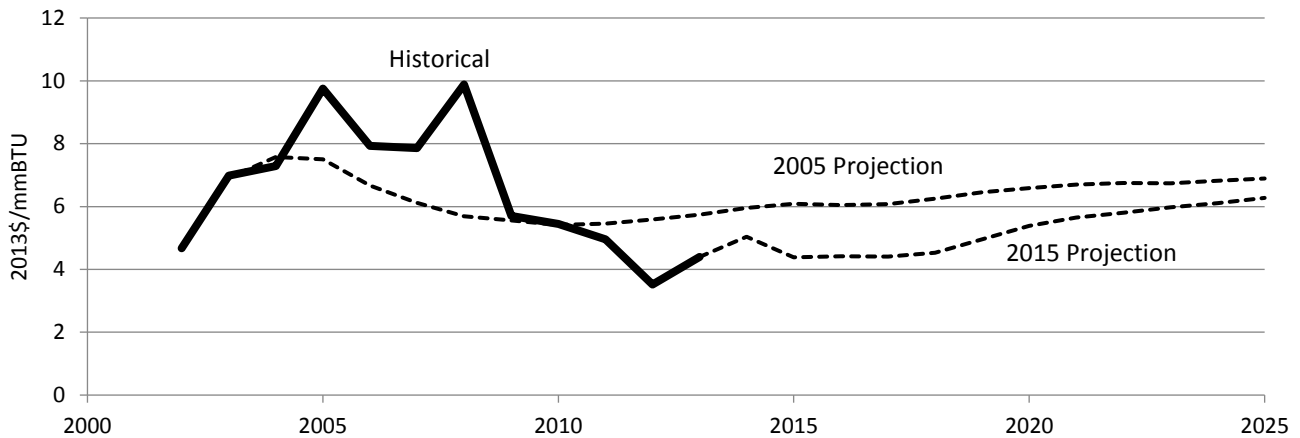
Figure 5. Fuel Prices for Eastern Interconnection, 2000–2015




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.

Figure 6. Comparison of 2005 and 2015 EIA Projections


Panel A: Natural gas price



Panel B: Renewable energy generation




Panel C: Electricity sector generation



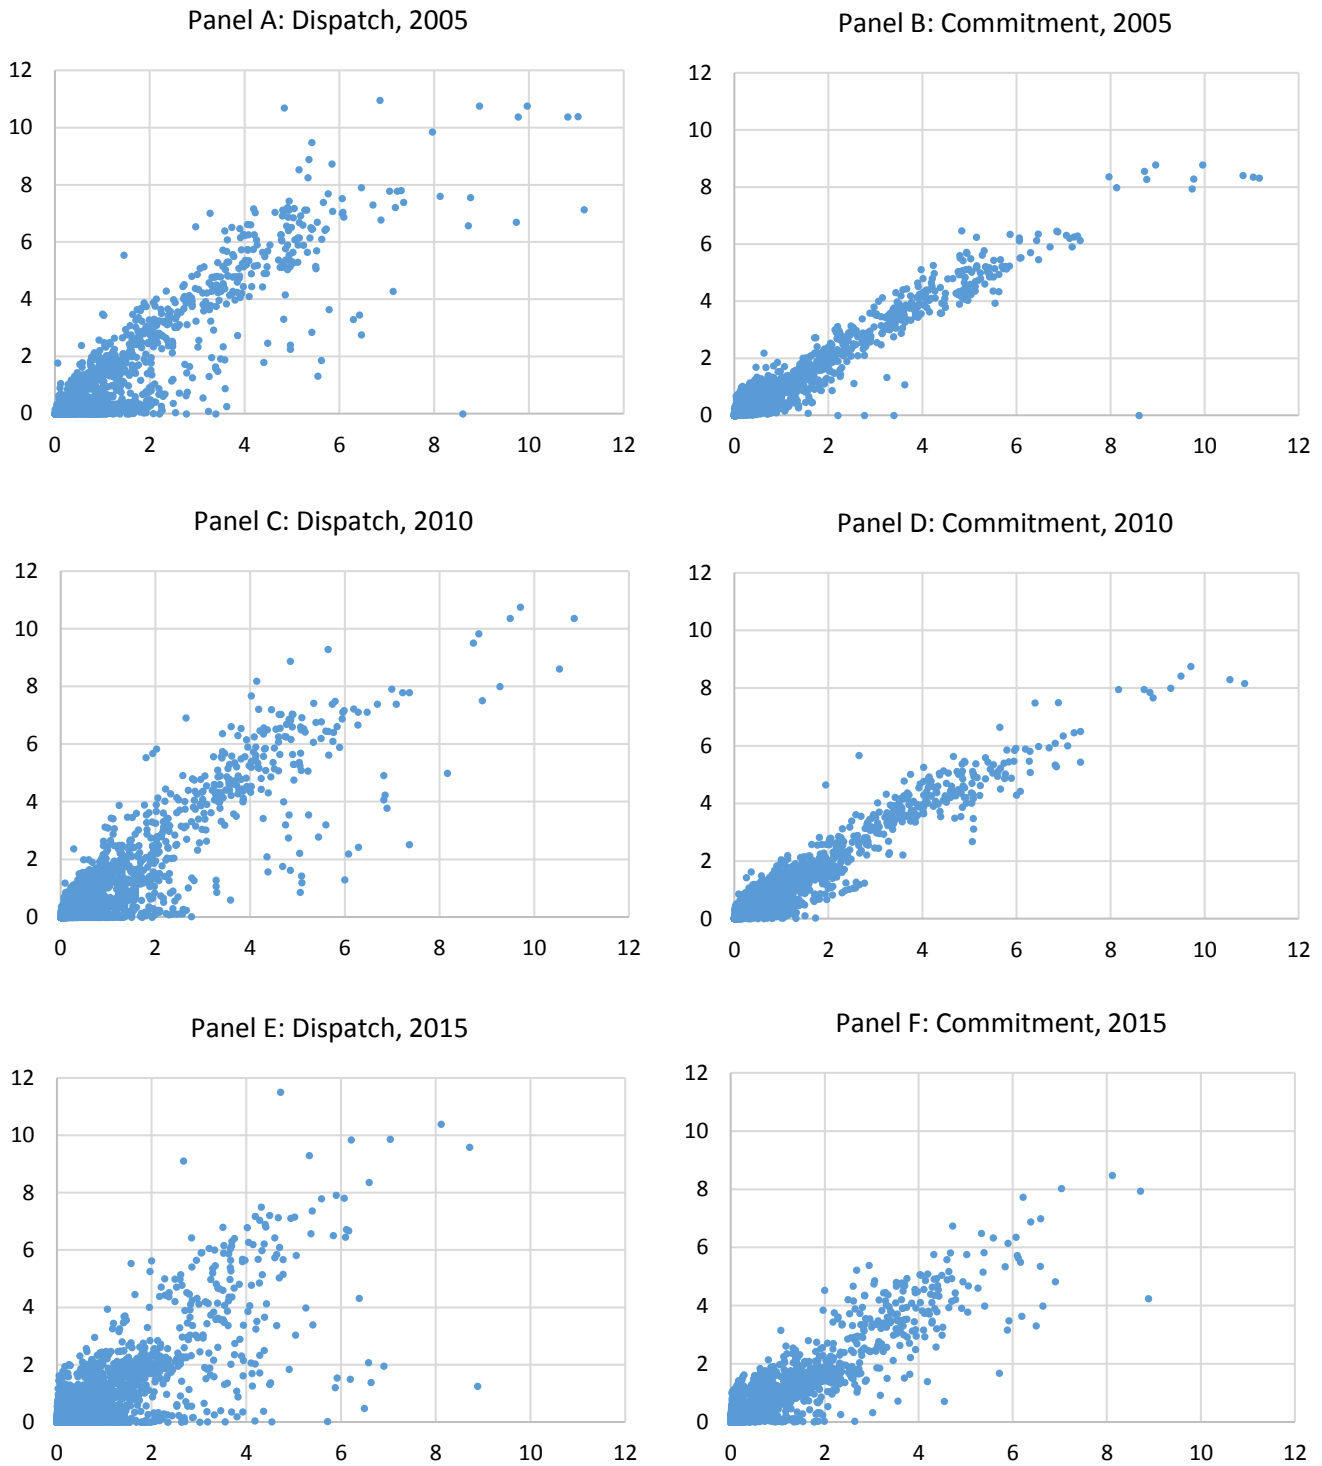
Notes: The dashed lines show projections from the Annual Energy Outlook (AEO) for the indicated years. Solid lines show the estimated historical price or generation level using data reported in the AEOs between 2005 and 2015. For example, historical electricity sector generation for the year 2002 is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

Figure 7. Estimated Density Functions of Annualized Average Abatement Costs for Selective Catalytic Reduction



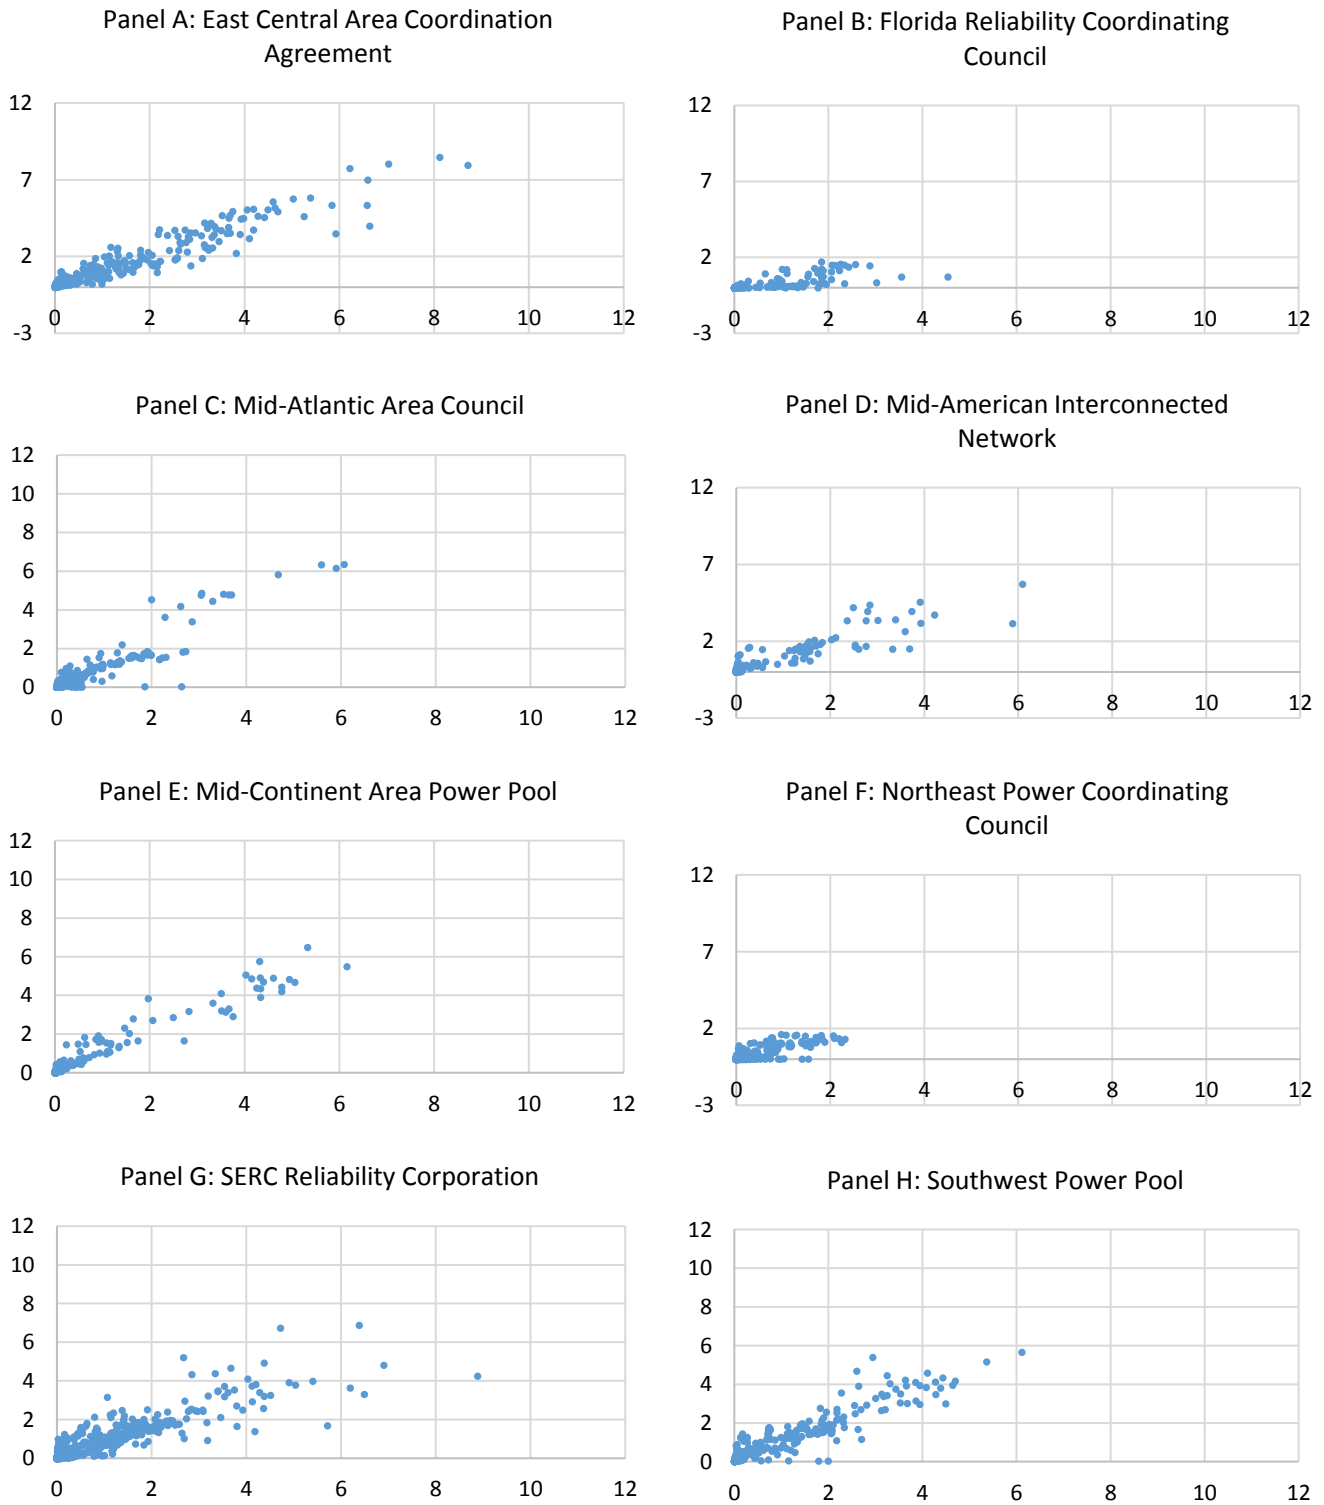
Notes: For each coal-fired unit operating in 2005 that does not have SCR, capital costs of installing SCR are computed using the unit's rated capacity, heat rate, and observed emissions rate and cost assumptions from the EPA Integrated Planning Model. Capital costs are annualized using the unit's estimated 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.

Figure 8. Simulated versus Observed Annual Unit Generation (million MWh)




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.

Figure 9. Simulated versus Observed Annual Unit Generation by NERC Region, 2015
(million MWh)



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 10. 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.

Table 1. Summary Statistics of Eastern Coal-Fired Units

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

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.

Table 2. Changes in Unit Operation by Fuel Type, 2005-2008 vs. 2009-2015

	<u>Probability of positive generation</u>			<u>Capacity factor conditional on positive generation</u>		
	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 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.

Table 3. Observed and Simulated Generation and Emissions

<u>Year</u>	<u>Observed</u>	<u>Dispatch model simulation</u>	<u>Unit commitment model</u>
<u>Panel A: Percentage of coal in total generation</u>			
2005	79.9	96.9 {21.28}	81.3 {1.75}
2008	80.2	97.7 {21.87}	82.2 {2.57}
2012	62.9	63.1 {0.37}	62.2 -{1.11}
2015	58.0	60.1 {3.69}	59.4 {2.38}
<u>Panel B: Nitrogen oxides emissions (million tons)</u>			
2005	2.94	3.42 {16.28}	2.97 {1.11}
2008	2.39	2.83 {18.39}	2.46 {2.66}
2012	1.28	1.39 {9.06}	1.35 {5.97}
2015	1.00	1.09 {8.97}	1.08 {7.53}
<u>Panel C: Sulfur dioxide emissions (million tons)</u>			
2005	9.05	10.70 {18.16}	9.31 {2.82}
2008	6.65	8.26 {24.08}	7.01 {5.39}
2012	2.75	2.99 {8.70}	3.00 {9.06}
2015	1.79	1.98 {10.43}	1.97 {9.67}
<u>Panel D: Carbon dioxide emissions (million tons)</u>			
2005	1,923	2,092 {8.75}	1,928 {0.22}
2008	1,872	2,055 {9.74}	1,875 {0.12}
2012	1,630	1,637 {0.43}	1,626 -{0.26}
2015	1,521	1,517 -{0.30}	1,514 -{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.

Table 4. Observed and Simulated Capacity Factors

Mean capacity factor conditional on positive generation						
<u>2005</u>			<u>2015</u>			
	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 5. Percentage of Units with Zero Annual Generation

<u>Year</u>	<u>Simulation using dispatch model</u>	<u>Simulation using unit commitment model</u>
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.

Table 6. Effects of Shocks to Demand, Wind Generation, and Fuel Prices

	(1)	(2)	(3)	(4)	(5)
	Baseline	Observed (lower) consumption growth	Observed (higher) wind generation	Observed fuel prices	Observed consumption, wind, and fuel prices
<u>Panel A: Generation percentage by fuel type</u>					
Coal	74	75	75	63	59
Natural gas	22	23	21	36	41
<u>Panel B: Mean capacity factor by fuel type</u>					
Coal	0.77	0.66	0.73	0.66	0.51
Natural gas	0.28	0.23	0.25	0.45	0.40
<u>Panel C: Mean annual operating profits by fuel type (million 2005 dollars per unit of capacity)</u>					
Coal	0.34	0.24	0.30	0.10	0.04
Natural gas	0.09	0.05	0.08	0.09	0.04
<u>Panel D: Annual emissions (million tons)</u>					
Nitrogen oxides	2.87	1.90	2.69	2.49	1.53
Sulfur dioxide	10.68	7.13	10.07	8.91	5.46
Carbon dioxide	2,516	1,872	2,351	2,372	1,683

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.

Table 7. Emissions Cap Scenarios

	(1)	(2)	(3)	(4)
	Baseline	State emissions caps	Emissions caps with observed consumption	Emissions caps with observed consumption, wind generation, and fuel prices
<u>Panel A: Generation percentage by fuel type</u>				
Coal	74	73	75	56
Natural gas	22	23	23	44
<u>Panel B: Mean capacity factor by fuel type</u>				
Coal	0.77	0.77	0.66	0.49
Natural gas	0.28	0.28	0.25	0.43
<u>Panel C: Mean annual operating profits by fuel type (million 2005 dollars per unit of capacity)</u>				
Coal	0.34	0.34	0.22	0.04
Natural gas	0.09	0.09	0.23	0.05
<u>Panel D: Annual emissions (million tons)</u>				
Nitrogen oxides	2.87	1.77	1.51	1.23
Sulfur dioxide	10.68	10.64	7.09	5.28
Carbon dioxide	2,516	2,511	1,869	1,651
<u>Panel E: Abatement costs</u>				
Emissions price (2005 \$/ton)		3,304	2,409	920
Annualized costs (2005 billion \$)		2.92	0.90	0.42

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

Data Appendix

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 that 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 the year 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 Annual Energy Outlook. 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 EPA. For each unit, the EPA data include state, North American Electric Reliability Corporation (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. 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 non-peak hours are within 6 hours of the peak hour in the corresponding day.

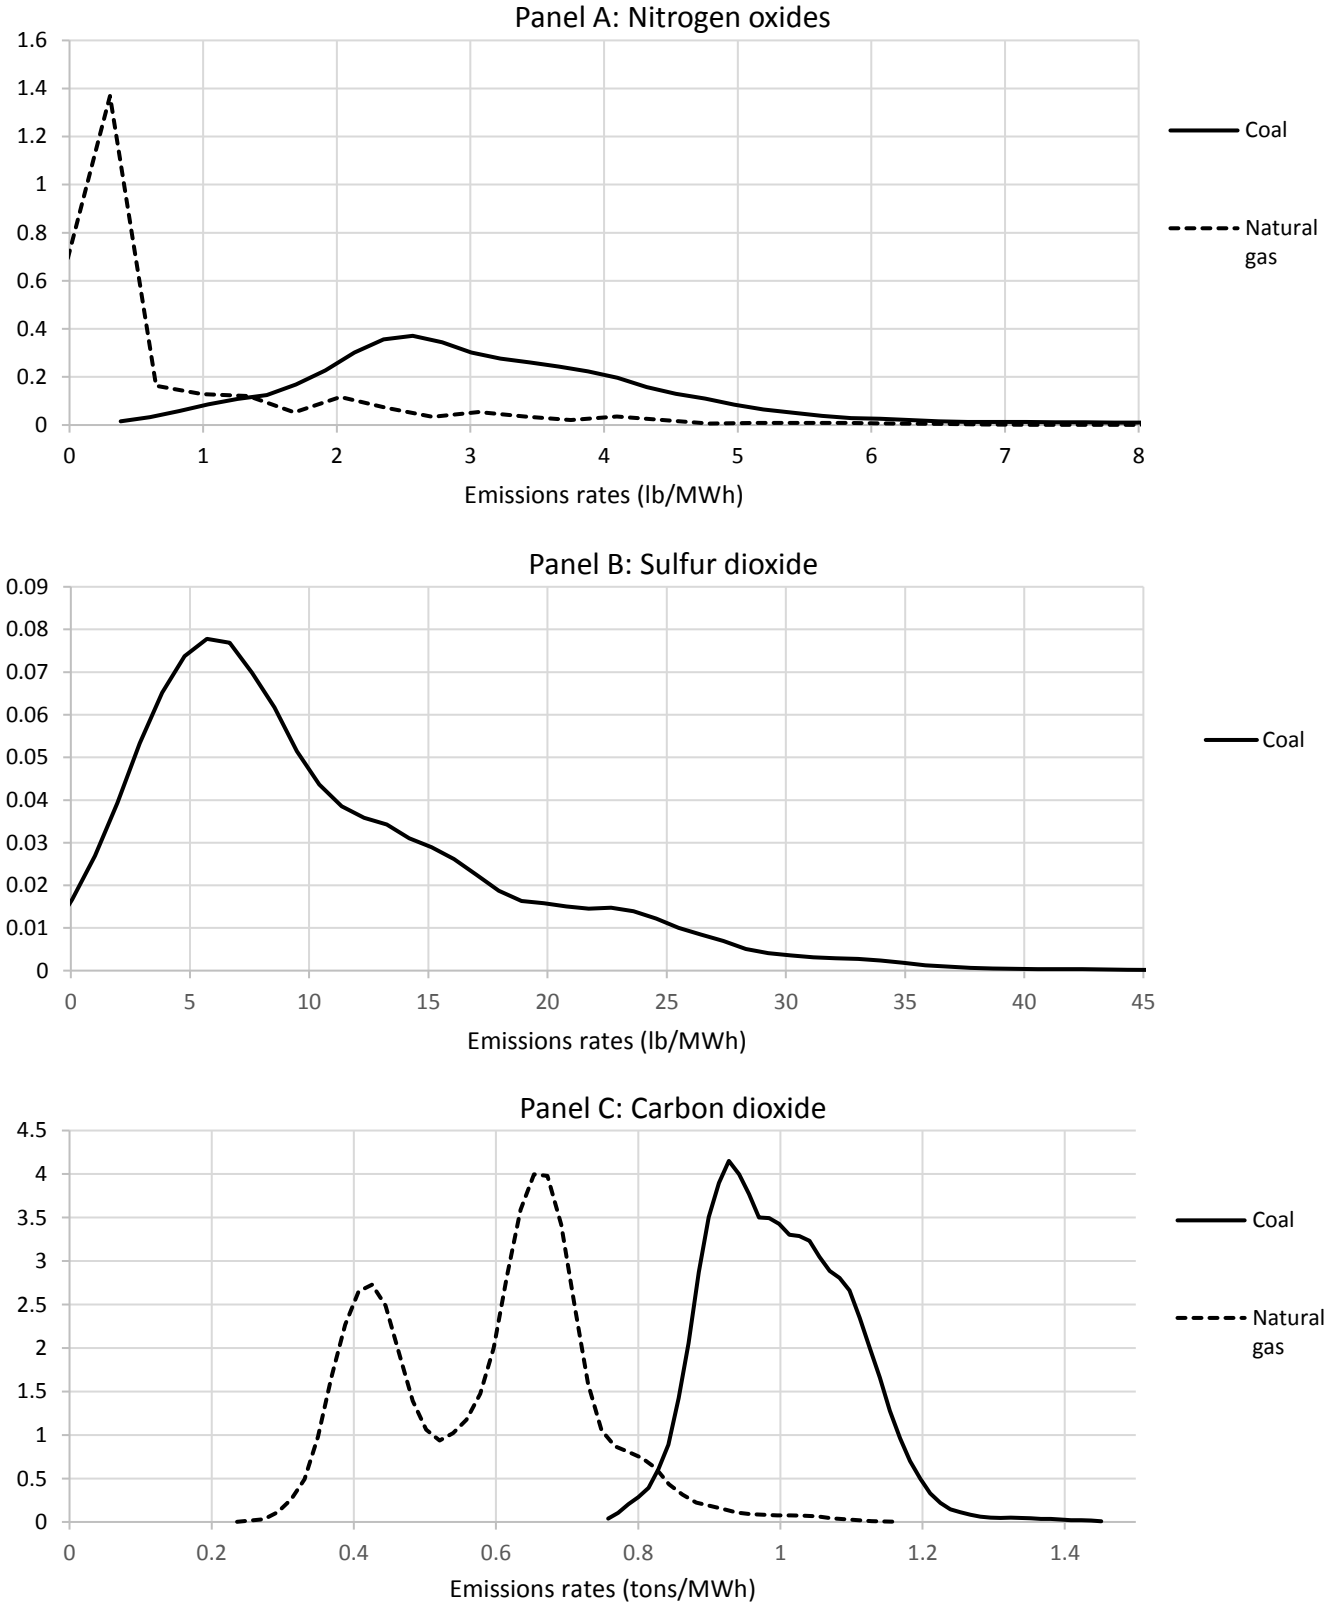
We obtain delivered fuel prices from EIA Forms 423 and 923. To reduce measurement error and concerns about the potential correlation among plant-level fuel prices and plant attributes that are not incorporated in the model, rather than using plant-level prices, we use average prices by NERC region and 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 2016). Future profits are discounted at a rate of 10 percent.

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 non-fossil technologies such as nuclear (i.e., we assume that demand is perfectly price-inelastic, which is a common assumption when modeling wholesale markets, such as by Bushnell et al. 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 5th percentile of generation observed across all hours between 2005 and 2015 in which the unit operates with positive generation. Appendix Figure 3 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 percent 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.

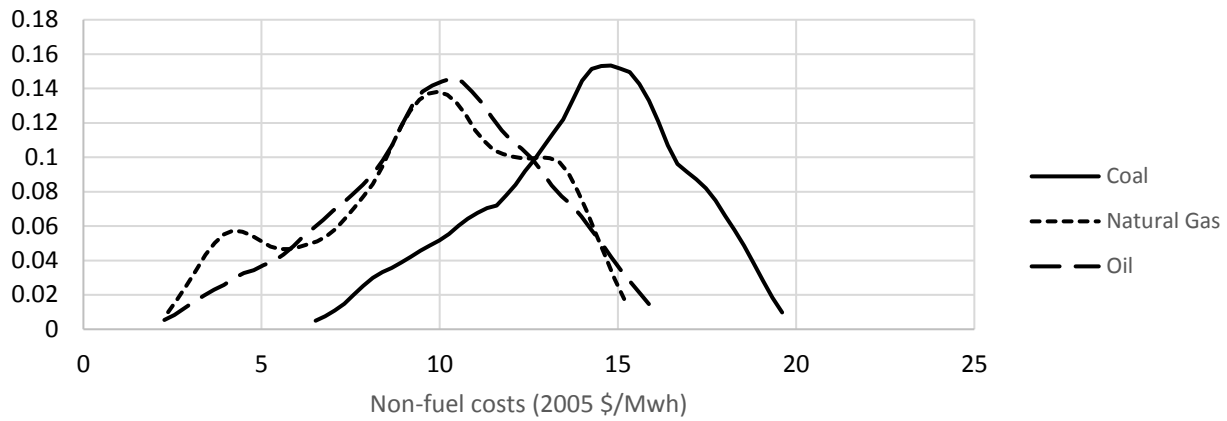
Appendix Figures and Tables

Appendix Figure 1. 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). The figure plots estimated density functions of the emissions rates by fuel type.

Appendix Figure 2. Estimated Density Functions of Nonfuel Operating Costs



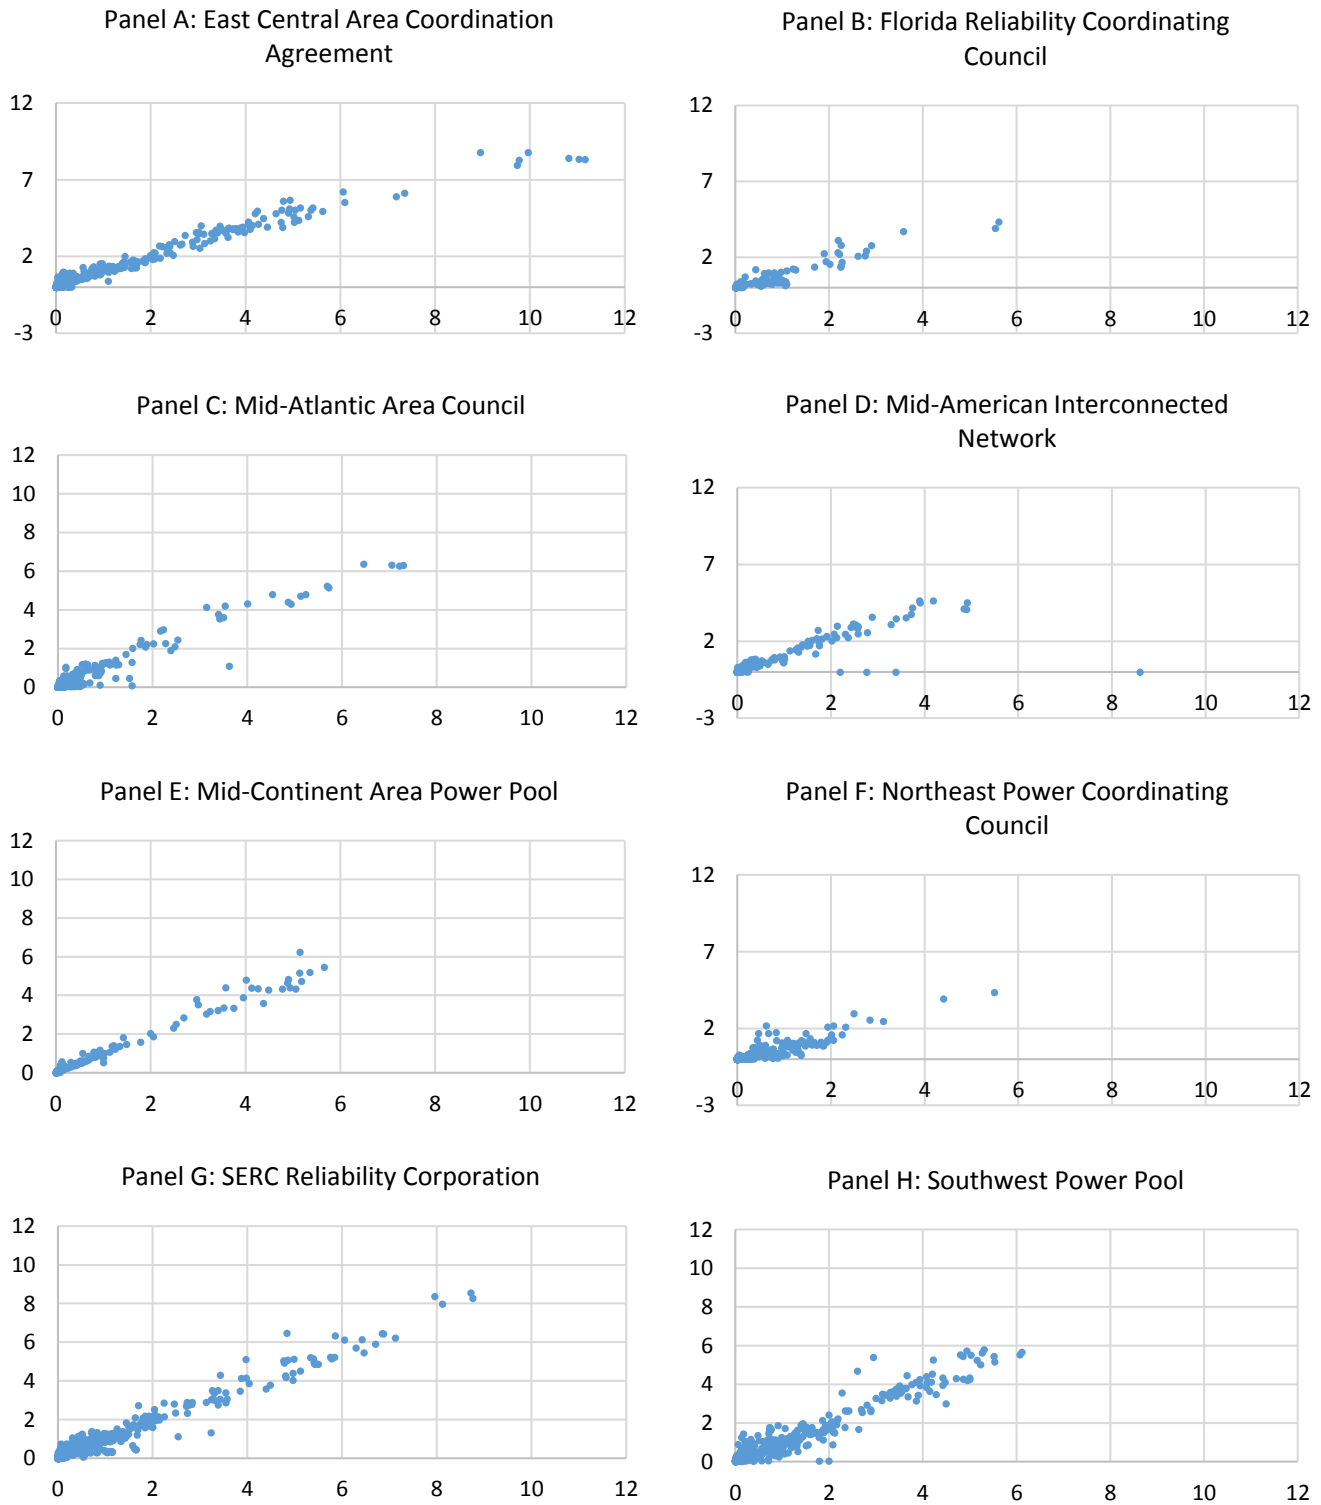
Notes: Nonfuel operating costs are estimated as described in the text. The figure plots the estimated density function of nonfuel costs (in 2005 \$/MWh) by fuel type.

Appendix Figure 3. Estimated Density Functions of Minimum Capacity Factor for Coal and Gas-Fired Units



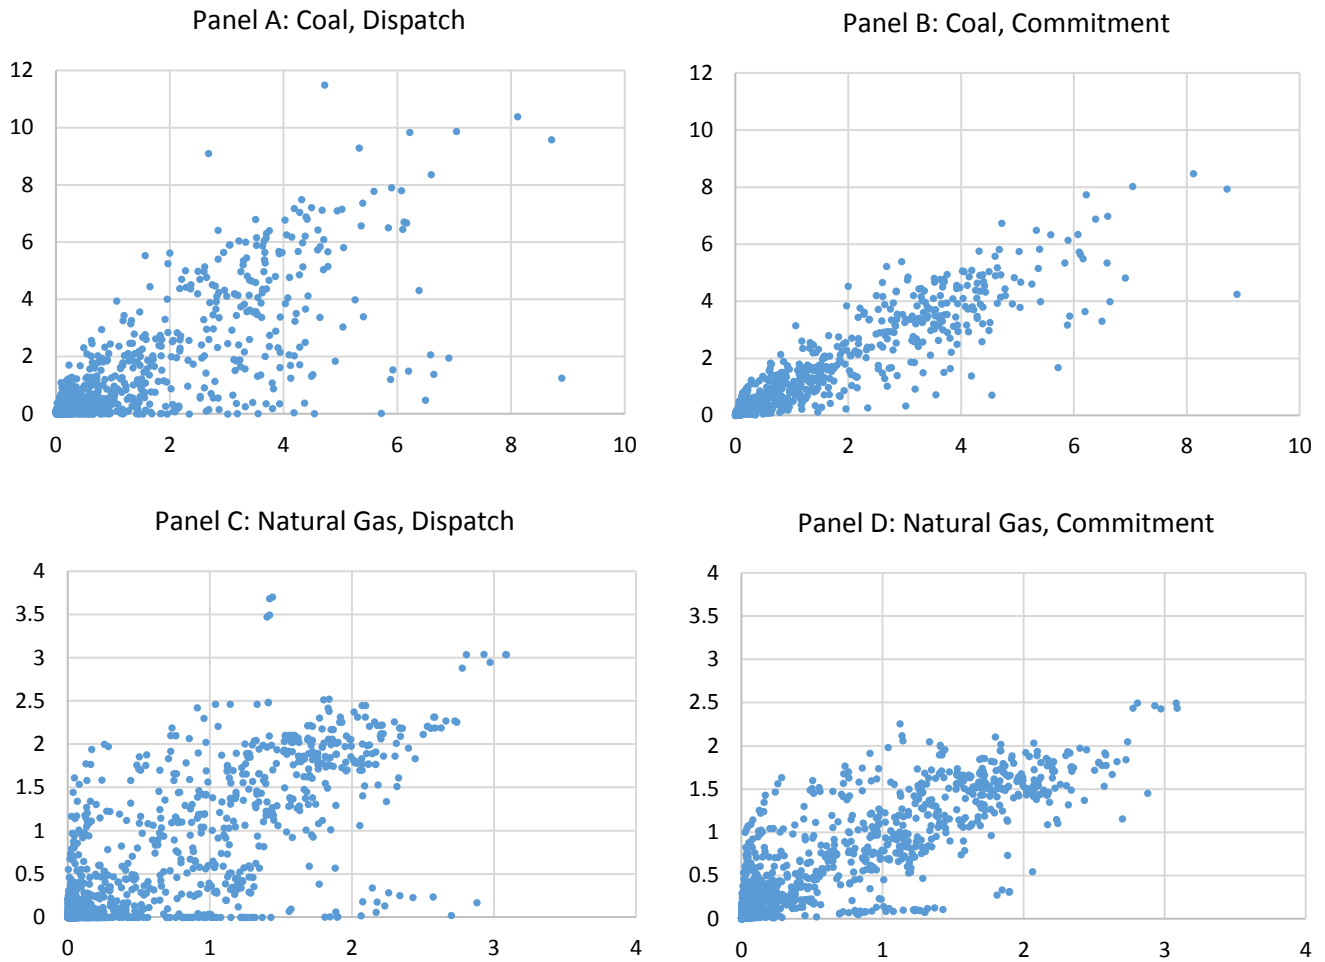
Notes: Minimum capacity factor is the ratio of the unit's minimum generation level to its rated capacity. The figure plots the estimated density functions of minimum capacity factor for coal and gas-fired units.

Appendix Figure 4. Simulated versus Observed Annual Unit Generation by NERC Region, 2005 (million MWh)



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.

Appendix Figure 5. Simulated versus Observed Annual Unit Generation by Fuel Type, 2015 (million MWh)



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

Appendix Table 1. Emissions for CSAPR and Other States

	(1)	(2)	(3)	(4)
	Baseline	State emissions caps	Emissions caps with observed consumption	EMISSIONS caps with observed consumption, wind generation, and fuel prices
<u>Panel A: Changes in nitrogen oxides emissions (million tons)</u>				
CSAPR states	2.67	-1.10	-1.31	-1.54
Other states	0.20	0.00	-0.05	-0.10
Total emissions	2.87	-1.10	-1.36	-1.64
<u>Panel B: Changes in sulfur dioxide emissions (million tons)</u>				
CSAPR states	10.27	-0.04	-3.46	-5.15
Other states	0.41	0.00	-0.13	-0.25
Total emissions	10.68	-0.04	-3.59	-5.40
<u>Panel C: Changes in carbon dioxide emissions (million tons)</u>				
CSAPR states	2,368.98	-5.38	-606.01	-822.25
Other states	147.11	0.71	-40.72	-42.55
Total emissions	2,516.09	-4.67	-646.73	-864.80

Notes: Column 1 reports the emissions in the baseline scenario, and columns 2-4 report comparisons between the scenarios indicated in the column headings and the baseline scenario.