DISCUSSION PAPER

February 2013 ■ RFF DP 13-05

Regulating Greenhouse Gases from Coal Power Plants under the Clean Air Act

Joshua Linn, Erin Mastrangelo, and Dallas Burtraw

A version of this paper will be published in the *Journal of the Association of Environmental and Resource Economists* 2014 (forthcoming).

1616 P St. NW
Washington, DC 20036
202-328-5000 www.rff.org

RESOURCES FOR THE FUTURE

Regulating Greenhouse Gases from Coal Power Plants under the Clean Air Act

Joshua Linn, Erin Mastrangelo, and Dallas Burtraw

Abstract

The Clean Air Act has assumed the central role in US climate policy, directing the development of regulations governing greenhouse gas emissions from existing coal-fired power plants. This paper examines the operation of coal-fired generating units over 25 years to estimate the marginal costs and potential magnitude of emissions reductions from improving their efficiency. We find that a 10 percent increase in coal prices causes a 0.2 to 0.5 percent heat rate reduction, broadly consistent with engineering assessments. We also find that coal prices have a significant effect on utilization. The results are used to compare cost-effectiveness of alternative policies.

Key Words: efficiency, regulation, greenhouse gas, carbon dioxide, coal, performance standards

JEL Classification Numbers: L94, Q54

Contents

Regulating Greenhouse Gases from Coal Power Plants under the Clean Air Act

Joshua Linn, Erin Mastrangelo, and Dallas Burtraw*

1. Introduction

There has been considerable debate over the costs and effectiveness of energy efficiency investments, such as improving fuel economy of passenger vehicles or retrofitting buildings with better-insulated windows. On the one hand, many estimates suggest that low-cost and even negative cost opportunities exist across the economy, where the market value of the energy savings outweighs the investment cost. On the other hand, many analysts are skeptical of these assertions, arguing that if such opportunities were available, firms and consumers would take advantage of them.

Many of the optimistic estimates are based on case studies or engineering assessments of particular technologies. Previous analysis has identified several reasons why such assessments may be incomplete. First, there may be costs that the analyst does not observe and that hinder adoption. Second, technologies, particularly new ones, may be less effective than expected or not used as expected. Third, missing data on the extent to which the technologies have already entered the market may cause an overestimate of available efficiency opportunities. Fourth, there may be a rebound effect, in which adopting energy-efficient technology reduces its cost of operation and increases its use. Underestimating the rebound effect could lead to an overestimate of emissions reductions caused by technology adoption. For the most part, however, there is little direct evidence on these possibilities, and the controversy remains.

Recent policy developments heighten this debate. Since a legislative approach to climate policy stalled in the US Congress, the Clean Air Act (CAA) has assumed the central role in the development of regulations that will reduce greenhouse gas (GHG) emissions. The US

* Linn (linn@rff.org) is a fellow and Burtraw (burtraw@rff.org) is a the Darius Gaskins Senior Fellow at Resources for the Future. Mastrangelo (ekmastrangelo@gmail.com) is an analyst at the Federal Energy Regulatory Commission, and previously at Resources for the Future, where she conducted most of her research. We thank participants at the USAEE, APPAM, and InfraDay conferences along with Aaron Bergman, David A. Evans, Alex Marten, Bryan Mignone, Stephen Puller, and Margaret Taylor for helpful comments, and Louis Preonas and Matt Woerman for outstanding research assistance. The authors received funding from the Bechtel Foundation, the US Environmental Protection Agency's National Center for Environmental Economics, and the Department of Energy. This paper does not reflect the official views of any of the funding entities.

Environmental Protection Agency (EPA) has been developing emissions rate standards for sectors of the economy including passenger vehicles and new industrial sources. EPA also is expected to introduce performance standards for existing stationary sources, such as electricity generators and industrial facilities—an approach that is nearly unprecedented. Such standards raise the possibility of achieving carbon dioxide (CO_2) emissions reductions rapidly due to the low capital costs of the measures being contemplated, which is particularly important given the US commitment to reduce emissions under the 2009 Copenhagen Accord.

Coal-fired electricity generators account for about one-third of annual US CO_2 emissions. EPA estimates that 2 to 5 percent efficiency improvements may be achieved on average at these facilities,[1] yielding annual emissions reductions comparable to those expected from efficiency standards for new passenger vehicles sold from 2012 through 2016.[2] The novelty and potential of these electricity sector standards raise two questions: (1) what are the available abatement opportunities; and (2) what are the costs of reducing emissions? Answering these questions requires addressing each of the issues above: estimating technological potential, technology costs, and the rebound effect.

This paper focuses on existing coal-fired electricity generation units. We analyze the actual efficiency of the entire fleet of coal units in the United States, where efficiency is measured as electricity generated per unit of heat input. We assess abatement opportunities and costs by observing how market and regulatory incentives affect the energy efficiency of coal plants. We use the results to compare the cost-effectiveness of alternative energy efficiency policies.

We make two contributions to the literature. First, other studies, for example Metcalf and Hassett (1999) and Linn (2008), have analyzed the effect of fuel prices on energy efficiency in other sectors, but this is the first study on the electricity sector. Second, previous research has

[1] EPA (2008), p. 16. With no change in utilization, a 5 percent efficiency improvement would produce emissions reductions of 90 million tons per year, about 1.6 percent of total US emissions in 2009. Drawing on estimates from Sargent & Lundy, LLC (2009) that we discuss below, these reductions could be achieved at a cost of \$10.74 to \$63.91 per ton CO_2 before accounting for the value of saved energy.

[2] This estimate is calculated by taking the 992 million tons in cummulative emissions reductions expected from pasenger vehicles introduced between 2012 and 2016 (http://environment.about.com/od/environmentallawpolicy/a/obama-sets-new-fuel-efficiency-standards.htm) and dividing by the expected 13 year lifetime of passenger vehicles in the United States resulting in average emissions reductions of 76 million tons per year. The calculation does not consider the effect of CAFE on the ownership and operation of existing passenger vehicles.

shown that heterogeneous abatement costs and opportunities and the rebound effect cause flexible standards to be more cost effective than traditional standards, and cause emissions prices to be more cost effective than flexible standards (a traditional standard requires all units to meet the target whereas a flexible standard allows units to exceed the standard and sell credits to units that do not meet the standard) (Holland et al. 2009). This paper is the first to study the electricity sector using econometric techniques and the first to estimate abatement costs, the rebound effect, and heterogeneity parameters in an internally consistent manner and compare the alternative policies.

We first assess abatement opportunities from efficiency improvements by examining heterogeneity in the efficiency of existing coal units. The analysis is performed using a unique panel data set of coal-fired generation units for the years 1985–2009. The data include monthly fuel input, generation, and coal prices by generation unit for nearly all US coal plants, and the units in the sample account for 95 to 98 percent of total coal generation in each year. We use a generation unit's heat rate (the ratio of heat input to electricity generated) to measure efficiency; heat rate is approximately proportional to the rate of CO_2 emissions per unit of electricity generation.

We show that there is considerable heterogeneity and a substantial right-hand (positive) tail in the heat rate distribution. Specific technical factors help explain heterogeneity across units, including boiler design, size, and vintage, and features such as pollution control equipment and cogeneration. After controlling for these factors, fleetwide emissions rate reductions of up to 6 percent may be technically feasible by improving performance up to a 90th percentile emissions-rate benchmark. This estimate does not account for costs, and we consider it an upper bound, given current technology.

To compare the costs of alternative policies, we specify and estimate the key parameters in an electricity sector model. In principle, several types of policies could be used to incentivize heat rate improvements, and the costs may vary across the alternative policies. We compare the costs of four policy alternatives: a traditional (inflexible) performance standard, a flexible performance standard, an emissions tax and a fuel tax. To make this comparison we observe that (1) cost-effectiveness depends largely on potentially heterogeneous abatement costs and the rebound effect; and (2) coal prices mimic the incentives created by a CO_2 emissions price (i.e., an emissions cap or tax) or some types of performance standards. Demonstrating the first point requires a brief description of the policies.

A traditional standard requires improvements at all units with emissions rates in excess of a target without regard to relative cost effectiveness. However, units that decrease heat rates to meet a traditional standard also experience a rebound effect because the lower heat rate reduces the marginal cost of generating electricity.[3] In contrast, a flexible emissions rate standard sets a benchmark emissions rate and allows firms to over-comply and sell credits to firms that under-comply, thus it promotes the lowest-cost efficiency improvements rather than requiring improvements at all units initially exceeding the target. The flexible standard has two effects on generation: it imposes an opportunity cost on heat rates by effectively adding to the cost of fuel, and it provides an output subsidy through the allocation of credits based on generation. The opportunity cost provides a disincentive for generation while the output subsidy provides an incentive to increase generation. Hence, unlike the traditional standard the flexible standard creates incentives to adopt energy efficiency technology at all units. Moreover, unlike the traditional standard the effect of the flexible standard on generation is ambiguous.

A CO_2 emissions or fuel tax raises the cost of using fuel, thereby creating an incentive to adopt energy efficient technology. By raising fuel costs, they also create an incentive for firms to reduce generation so they create the smallest rebound effect; hence, the emissions or fuel tax would require the smallest overall reduction in heat rates to achieve a given emissions target. In short, the relative cost-effectiveness of the policies depends on the cost of improving heat rates and the magnitude of the rebound effect.

Because data on energy efficiency technology adoption are not available, we focus on the response of heat rates to changes in coal prices. A simple model demonstrates that we can estimate the cost of adopting technology by examining the empirical relationship between coal prices and heat rates. We show that, conditional on the utilization of the unit, there is a one-to-one correspondence between the level of energy efficiency technology and the unit's heat rate. Similar to the CO_2 policies, an increase in the price of coal increases the opportunity cost for heat rates, conditional on utilization. Using the same panel data set as for the analysis of abatement opportunities, we find that a 10 percent coal price increase, corresponding to a tax on CO_2 emissions of about \$1.64 per ton, reduces heat rates by 0.2 to 0.5 percent, depending on the estimation procedure. A change in coal prices commensurate with a \$10 per ton tax on CO_2 emissions (representing a coal price increase of about 60 percent) would stimulate a 1 to 2

[3] The analysis in this paper focuses on the short run and does not consider retirements. Consistent with this focus, in the policy analyis we consider relatively small efficiency improvements of 1-2 percent.

percent heat rate reduction (holding fixed utilization). This range of estimates encompasses the estimates suggested in the engineering literature but includes the possibility of somewhat lower costs than have been estimated. We note that the overall efficiency improvements of 2 to 5 percent discussed in the engineering literature correspond to the change in heat rate resulting from an increase in coal prices of more than two standard deviations—that is, out of sample. We also obtain a significant relationship between coal prices and utilization. A 10 percent increase in coal prices reduces utilization by 2 to 6 percent.

We use a stylized model of the electricity sector to simulate the effects of four energy efficiency policies: a traditional emissions rate standard, a flexible standard, a coal Btu tax (roughly equivalent to a coal emissions tax), and a fossil fuel emissions tax. We find that because of the narrower focus of the performance standards and the greater rebound effect, more investment in heat rate technology is required under the performance standards than the taxes to achieve a given emissions reduction. This raises the relative costs of the standards, but overall, the costs approximate the engineering estimates.

This paper is organized as follows. Section 2 provides a brief background on the regulation of existing coal units under the CAA. Section 3 discusses the operation of coal-fired units in the US electricity system. Section 4 describes the data and summarizes heterogeneity in the heat rates across individual units. Section 5 describes the electricity sector model and motivates the empirical focus on the effects of coal prices on heat rates and utilization. Section 6 presents the empirical strategy and Section 7 presents the estimation results. Section 8 uses the estimation results to compare cost-effectiveness across policies, and Section 9 concludes.

2. The Clean Air Act

The modern CAA was passed in 1970 and conveys broad authority to EPA to develop regulations to mitigate harm from air pollution. In 2007 the Supreme Court affirmed this authority with respect to the regulation of GHGs *(Massachusetts v. EPA)*.[4] Subsequently, the agency made a formal, science-based determination that GHGs are dangerous to human health and the environment. This “endangerment finding” compelled the agency to mitigate that harm.

[4] 549 US 497 (2007).

In 2011 EPA implemented regulations affecting CO_2 emissions from passenger vehicles, medium-duty trucks, and heavy-duty trucks.[5] The agency also implemented regulations for construction permitting (New Source Review, NSR) for major new and modified sources, such as power plants and industrial facilities.[6] The third anticipated EPA regulatory action is the development of performance standards for GHGs affecting the *operation* of stationary facilities.[7] EPA has a long history of setting performance standards for new sources, but performance standards for existing sources are nearly unprecedented.[8] The first standards, expected in 2013, will target new steam boilers at power plants fueled with coal, oil, and natural gas. Subsequently the EPA is required to begin development of similar standards at existing facilities in the same source category. These sources represent more than one-third of GHG emissions in the United States (EIA 2011).

In principle, power plants could reduce emissions in many ways, including fuel switching, making incremental changes to efficiency, or adopting new, energy-efficient technology. Indications from EPA (2008) are that the regulations will encourage improvements in the efficiency of power plants without requiring large-scale substitution among fuels or technologies. The efficiency of facilities varies substantially, possibly indicating potential opportunities for improvements at the least efficient facilities. EPA suggested that average efficiency might be improved by 2 to 5 percent at moderate cost.[9] However, one might conjecture that incentives already exist to reduce costs at these facilities, and the presence of heterogeneity across plants might reflect variation in technological, geographic, and economic factors that make operational improvements more expensive in some situations. To date, there has been no comprehensive examination of the actual opportunities for efficiency improvements, or the magnitude and cost of potential emissions reductions.

---

[5] Beginning with model year 2012, average fuel economy of cars and light trucks improves by 5 percent per year to a fleet average of 35.5 miles per gallon (mpg) in 2016. New regulations are expected to extend this target to 54.5 mpg by 2025.

[6] Implementing regulations including the definition of best available control technology are in development.

[7] Standards under §111(b) of the CAA apply to new sources (these are termed New Source Performance Standards), and those under §111(d) apply to existing sources.

[8] Existing sources regulated under other provisions of the CAA are not eligible for regulation under §111(d) (Richardson et al. 2011).

[9] Equivalent additional reductions could be achieved if facilities were allowed to co-fire with biomass under the assumption that waste biomass was used and its combustion led was carbon neutral.

3. Coal Plant Operations

At a simplified level, a coal-fired power plant consists of one or many boilers that burn pulverized coal to produce steam from water. The boilers are connected to one or many generation turbines, which spin from the pressure of the steam to produce electricity. A condensing and cooling system collects and recycles the steam.

3.1 Determinants of Efficiency

Efficiency is often measured by the heat rate—the amount of heat input, in million British thermal units (mmBtu), required to generate one megawatt-hour (MWh) of electricity. A lower heat rate represents a more efficient unit. A generating unit can improve efficiency by reducing the amount of fuel required for a specific level of generation. A percentage improvement in heat rate is nearly equivalent to an equal percentage improvement in the emissions rate in terms of the change in CO_2 emissions.[10] The difference stems from the small variation in carbon per Btu across coal varieties.

The heterogeneity in heat rates across coal-fired generation units can partly be explained by technical characteristics determined at the time of plant construction that cannot be changed without a major overhaul. This category includes size, age, firing type, and the technology employed. Higher efficiency is generally associated with plants that are used more heavily because efficient units are less costly to operate.

A second factor is how the boiler is used. The relationship between the heat rate and utilization is nonlinear, as efficiency tends to be lower at very low and very high levels of utilization. Routine decisions regarding plant operations affect efficiency, and many of these decisions result from constraints or factors that weigh against efficiency concerns (e.g., variation in demand, voltage regulation, or system reliability). Units with lower utilization may be ramped up and down more frequently, which requires additional fuel input as temperature in the boiler fluctuates. The result could involve efficiency losses at least partly outside the control of plant decision makers.

[10] The product of the rate of emissions per mmBtu and the heat rate gives the rate of emissions per MWh. For example, combustion of Powder River Basin low-sulfur subbituminous coal emits about 212.7 lb CO_2 per mmBtu of heat content. A plant with a heat rate of 10 mmBtu/MWh has a CO_2 emissions rate of about 2,127 lb CO_2/MWh (212.7 * 10).

Plant managers control several other factors that affect heat rates. Techniques, management, or technology may improve the efficiency of the plant by targeting the major components of the coal combustion process: oxygen, temperature, and pressure. Excessive deviations in any of these areas may decrease efficiency through waste or shortfalls. For example, Beer (2007) and Rosen and Dincer (2003) explain that enhanced coal-feeding systems or grinding the coal more finely can reduce excess air in the boiler and increase efficiency by reducing heat loss. Maintenance and performance testing are also critical for identifying and preventing losses.

Although technical adjustments can be made to improve heat rate, a firm's institutions, such as training or commitment to goals, are also important. Industry experts identify as potential barriers weak support from management and a lack of expertise or onsite engineers dedicated to heat rate improvement (DOE/NETL 2009).

Regulations can also affect heat rates. Some measures that improve efficiency may trigger NSR, under which the firm must demonstrate that the efficiency improvement would not create or exacerbate air quality violations. NSR can raise the cost of improving efficiency if the firm has to install pollution abatement equipment. In addition, the market environment may influence investment and operational behavior of firms. Firms subject to greater competition in wholesale power markets may have greater awareness of and incentive to minimize costs (Fabrizio et al. 2007). Firms may vary in their ability to access capital when investment opportunities exist.

Table 1 describes factors that affect heat rates and hypotheses from the literature on how the factors affect heat rates. Appendix 1 includes a more complete literature review.

3.2 Engineering Estimates of Costs and Heat Rate Improvements

According to an engineering analysis prepared for EPA, the cost of reducing heat rates varies widely across options (Sargent & Lundy 2009). In one example, for a 200-MW unit, cleaning the condenser improves heat rates up to 0.07 mmBtu/MWh and costs a modest $30,000 per year.[11] Accounting for utilization and coal prices, this would be profitable for a typical 200-MW unit in 2008. Other improvements may be less cost-effective, however. A new fan system to

[11] The range of identified reductions could be achieved at a cost of $10.74-$63.91 per ton of CO_2 before accounting for the value of reduced fuel expenditures.

control air flow in the flue gas system can cost $6 million in capital investment plus $50,000 per year in operating and maintenance costs for heat rate reductions up to 0.05 Btu/kWh for a 200-MW unit.

Each potential measure at a typical plant improves the heat rate between 0.01 and 0.1 mmBtu/MWh (compared with an average heat rate of about 11 mmBtu/MWh) and improvements may be additive. Importantly, significant analysis and expertise are required to find the optimal combination of upgrades and techniques, if any, for each specific plant. There are no comprehensive data, however, on the extent to which such technologies are already in use.

4. Data, Summary Statistics, and Opportunities for Heat Rate Improvements

To study the efficiency of coal-fired power plants, we assemble a comprehensive annual panel data set of coal-fired generating units in the United States from 1985 through 2009. The dataset combines several public sources and contains a uniquely detailed set of characteristics and efficiency of coal-fired units over time. Because of data limitations, 2006 and 2007 are excluded from the panel. Appendix 2 provides further information about data sources.

The annual heat rate is calculated from reported monthly heat input and generation at a boiler and corresponding generator. Because some units fire multiple fuels, heat input from fuels other than coal is included in the calculation of heat rate; however, close to 99 percent of total annual heat input is from coal in any given year in the panel. Each year contains about 1,000 units, although the sample size is somewhat larger after 2001 because of changes in the reporting requirements of the primary data source. The sample contains about 340 gigawatts (GW) of capacity and accounts for 97 percent of US coal-fired electricity generation and emissions in 2009.

Although we do not observe heat rate investments in the data, the data set includes other variables that influence heat rate. Characteristics of the plant, such as vintage, size, boiler firing type, and cogeneration capacity, help determine the unit’s physical ability to achieve a particular level of efficiency. We also include variables that describe how the plant is used and managed. These include utilization rates, the presence of pollution controls, choice of fuel including coal type, regulatory environment, and ownership structure. Finally, we merge plant-by-year prices of the delivered price of coal. Note that the coal price is the annual average price of coal because it is calculated by dividing expenditure by heat content of the delivered coal. Most of the analysis uses the overall price of coal, including contract and spot market purchases, and some of the analysis includes contract and spot market prices separately.

Table 2 and Figures 1–3 summarize the data. Table 2 aggregates the data to five-year time periods and reports means of several variables with standard deviations in parentheses. Figure 1 shows the distribution of heat rates in 2008. The horizontal axis is the average annual heat rate and the vertical axis is the heat input, which maps into the unit's net electricity generation. The figure displays a right-hand tail, indicating that many of the least efficient units have relatively little electricity generation. Figure 2 disaggregates the data and illustrates differences in heat rates across three prominent firing types. Figure 3 reports the distribution of heat rates after controlling for several characteristics, including firing type, cogeneration, capacity and fuel type. More specifically, the figure plots the residuals of a regression of the unit's heat rate on the indicated control variables; the variables are added sequentially to generate the different plots (in each case, the sample mean heat rate is added to the residuals before plotting). Even after accounting for these characteristics, one observes substantial heterogeneity, which is the focus of our analysis.

We use the observed heat rate heterogeneity to place an upper bound on the available abatement opportunities. For this calculation, we put aside abatement costs and consider what levels of abatement—that is, percentage emissions reductions—would be possible under alternative traditional performance standards. We first consider a uniform standard equal to the 90th percentile of efficiency (corresponding to the 10th percentile of heat rates) and calculate the percentage emissions reduction across all units in the sample. This calculation implicitly assumes that it is technically feasible for all units to achieve a heat rate equivalent to the most efficient 10th percentile, but this may not be possible for some units whose inherent features, such as firing type, size, or other characteristics, are inflexible. Therefore, we also analyze traditional standards that are set according to the firing type or other attributes of the generation unit that cannot be changed.

Table 3 reports the results of these calculations. Each column reports a different performance standard that is determined according to the generator attributes in the column heading. In Panel A we assume that the standard does not affect utilization, and in Panel B we assume a 20 percent rebound effect (as estimated below). We find that the percentage emissions reduction is about 6 to 7 percent assuming no rebound effect, and 5 to 6 percent assuming a 20 percent rebound effect. Somewhat surprisingly, the reductions are fairly insensitive across the different columns.

5. Comparing Policies for Improving Emissions Rates at Existing Generators

This section introduces the model we use to compare the investment, emissions, and costs of four policies that aim to reduce emissions rates at existing coal generators: a traditional standard, a flexible standard, a tax on coal consumption, and a CO^2 emissions tax. Firms choose a heat rate technology to maximize profits. Using the model's first-order conditions, we show that the policies create different outcomes because of their effects on heat rate technology and utilization.

5.1 Profit Maximization in the Absence of CO_2 Regulation

We begin with a static, one-period, model of the electricity sector that contains a large number of coal electricity generation units, each of which maximizes profits and takes prices as given. The sector consists of a set of markets that are indexed by the subscript r (referring to region). We first consider the case in which there is no CO_2 regulation.

Each firm in the market owns a generation unit, which we index by i. The firm has already chosen the capacity of the unit, K_i, and the analysis focuses on the optimal choice of heat rate technology conditional on the capacity and exogenous variables. At the beginning of the period the firm chooses a heat rate technology and then sells electricity into a competitive wholesale market in which there is a backstop natural gas generator that supplies electricity at a constant marginal cost (below, we turn to the case in which the firm operates in a regulated environment and in simulation analysis we consider price-responsive gas supply). Let p_r equal the price of electricity in region r in dollars per MWh and p_r^f equal the delivered price of coal in dollars per mmBtu.

To simplify the exposition, we assume that the heat rate is fixed after the firm has chosen the heat rate technology, or in other words, that there is a direct correspondence between technology and heat rate. This assumption abstracts from the fact that, for a fixed technology, there is typically a negative relationship between heat rate and utilization. Therefore, we relax this assumption in the empirical analysis.

The cost of choosing heat rate h_i is $c(h_i, \gamma_i)K_i^{\rho}$, where $c' < 0$, $c'' > 0$, $\gamma_i > 0$, and $0 < \rho < 1$. Because this is a static one-period model, the cost represents the annualized capital cost of choosing heat rate h_i. There are two important features of this cost function. The first is that the cost scales less than linearly with the unit's capacity, so that reducing the heat rate is relatively less costly at larger units. The second feature is that the cost function includes the unit-specific constant γ_i, which allows for the possibility that costs vary across units that have the same capacity. For example, it may be less costly to reduce the heat rates for units with one firing type rather than another.

The marginal cost of generating electricity includes fuel costs, which equal the price of coal multiplied by the heat rate, $(p_r^f h_i)$, plus nonfuel operations and maintenance costs, m_r, which is the capacity-weighted average for coal units in the region.

The unit's generation output is equal to its capacity multiplied by the capacity utilization rate, u. The utilization rate is a function of fuel costs and unit type: $u_i = u(p_r^f h_i, \theta_i)$, where θ_i is a unit-specific constant. For simplicity we ignore operations and maintenance costs in the utilization decision. Utilization is decreasing in fuel costs because a unit with higher marginal costs is less competitive with other generation units in the market. Thus, a unit with a higher heat rate produces less electricity. The constant θ_i is proportional to the quality of the unit, and it may reflect unit-specific characteristics such as managerial quality or its idiosyncratic economic value, given its location in the transmission system. Units with a higher value of θ_i produce more electricity.

Thus, the model allows for two types of unit-specific heterogeneity. First, because of the constant γ_i, the cost of reducing the heat rate may differ for units that have the same capacity. Second, because of the constant θ_i, two units with the same heat rate and fuel price may nonetheless have different utilization rates. The wide heterogeneity in utilization and in heat rates, which the previous section showed persisting even after controlling for unit characteristics, demonstrates the importance of introducing both types of heterogeneity in the model.

The firm maximizes profits by choosing the heat rate. The profit maximization problem is

$$\max_{h_i}(p_r - m_r - p_r^f h_i)u(p_r^f h_i, \theta_i)K_i - c(h_i, \gamma_i)K_i^\rho \tag{A1}$$

The first-order condition for the heat rate is

$$\frac{\partial u}{\partial h}(p_r - m_r - p_r^f h_i) - p_r^f u(p_r^f h_i, \theta_i) - \frac{\partial c}{\partial h}K_i^{\rho-1} = 0 \tag{A2}$$

If utilization is relatively insensitive to heat rate (consistent with the empirical evidence), the first term on the left-hand side of equation (2) is small in magnitude and an increase in the price of coal reduces heat rates. The net effect of a coal price increase on utilization is ambiguous, depending on whether the decrease in heat rate more than offsets the increase in the coal price.

5.2 Traditional Heat Rate Standard

Suppose that, before firms have chosen heat rate technology, the government mandates a heat rate ceiling such that all units have to achieve this standard through efficiency investments. To simplify the expressions, we normalize the CO_2 emissions rate from coal generation to equal 0.1 ton of CO_2 per mmBtu of heat input. Because the emissions rate is normalized to 0.1, the heat rate standard is synonymous with an emissions rate standard.

Units respond in one of two ways. Units that, in the absence of regulation, would have had a heat rate at or below the standard ($h_i \le \hat{h}_i$) make no change (note that the standard could be unit-specific). Units that would have otherwise had a heat rate above the standard ($h_i > \hat{h}_i$) make investments to lower their heat rates to equal the standard incurring a cost: $c\left(\hat{h}_i\right) - c\left(h_i\right)$.[12] These units are expected to increase utilization to capitalize on their lower heat rates: $u\left(p_r^f \hat{h}_i, \theta_i\right)$. The units' endogenous decisions under the traditional standard determine the aggregate change in emissions.

5.3 Flexible Heat Rate Standard

Instead of imposing a traditional heat rate standard, suppose the government sets a heat rate standard of $\bar{h}$. The standard is chosen so that the aggregate emissions reduction is the same for the traditional and flexible heat rate standards.

If the unit exceeds the standard, it generates credits equal to the difference between its heat rate and the standard, multiplied by the electricity it generates. The profit maximization problem is

$$\max_{h_i}(p_r - m_r - p_r^f h_i)u[p_r^f h_i + P(h_i - \bar{h}), \theta_i]K_i - P(h_i - \bar{h})K_i u[p_r^f h_i + P(h_i - \bar{h}), \theta_i] - c(h_i, \gamma_i)K_i^\rho \quad \text{(A3)}$$

where the price of the credits is P, which is exogenous to the firm. The first-order condition is

$$\frac{\partial u}{\partial h}[p_r - m_r - p_r^f h_i - P(h_i - \bar{h})] - (p_r^f + P)u - \frac{\partial c}{\partial h}K_i^{\rho-1} = 0 \quad \text{(A4)}$$

[12] We assume that the standard is not so stringent that profits are negative and firms exit the market. Section 9 discusses the possibility of firm exit.

Comparing equation (4) and the first-order condition in the absence of regulation, equation (2), and continuing to assume that magnitude of the first term on the left-hand side is small, we observe that the flexible standard creates an incentive for a firm to reduce its heat rate.

5.4 Coal Btu Tax

Suppose instead the government imposes a tax, τ , on coal heat input, or equivalently a tax on coal Btu. This results in the following profit maximization problem

$$\max_{h_i}[p_r - m_r - (p_r^f + \tau)h_i]u[(p_r^f + \tau)h_i, \theta_i]K_i - c(\gamma_i, h_i)K_i^\rho \qquad \text{(A5)}$$

The first-order condition for the heat rate is

$$\frac{\partial u}{\partial h}[p_r - m_r - (p_r^f + \tau)h_i] - (p_r^f + \tau)u - \frac{\partial c}{\partial h}K_i^{\rho-1} = 0 \qquad \text{(A6)}$$

A comparison of equations (2) in the absence of carbon policy and (6) shows that the Btu tax affects the firm's heat rate and utilization the same as an increase in the price of coal. Consequently, imposing the Btu tax reduces heat rates and has an ambiguous effect on utilization. Comparing equation (6) with equation (4) for the flexible heat rate standard, we observe that the Btu tax reduces utilization (i.e., holding the heat rate constant), whereas the emissions rate standard increases or decreases utilization, depending on whether the heat rate is lower than or greater than the standard. This comparison demonstrates the potential importance of the rebound effect, which is the change in utilization arising from the efficiency improvement.

5.5 CO_2 Emissions Tax

An emissions tax on all fossil fuel is modeled similarly to the Btu tax on coal because the CO_2 emissions rate is normalized to equal 0.1 ton of CO_2 per mmBtu of heat input. However, the emissions tax is applied to the backstop natural gas technology (at an assumed emissions rate of 0.05 tons of CO_2 per mmBtu) as well as to coal generation, and consequently it increases the wholesale price of electricity as well as the coal price. Otherwise, the profit maximization problem and first-order condition are identical to those for the Btu tax.

5.6 Summary

Comparing the first-order conditions across the regulatory scenarios shows how the policies affect a firm's optimal heat rate. All four policies create incentives for a firm to reduce its heat rate. The policies have differing effects on utilization, however. The traditional standard increases utilization at units that initially do not meet the standard and have to make investments to reduce heat rates, but it does not raise the total cost of using fuel. Compared to the traditional standard, the flexible standard creates less of an incentive to increase utilization for units that do not meet the standard because it raises the cost of using fuel, but the flexible standard has a positive effect on units that initially satisfy or overcomply with the standard. The taxes create less incentive to increase utilization than the standards because the cost of using coal increases, regardless of the firm's heat rate.

6. Strategy for Estimating the Effects of Coal Prices on Heat Rates and Utilization

The model in the preceding section provides a simple framework for comparing the four policies. The comparison depends on the cost of heat rate technology, on how fuel costs affect utilization, and on the unit-specific heterogeneity parameters. This section discusses how we estimate technology costs and the effect of fuel costs on utilization.

6.1 The Relationship between Heat Rate Technology and Measured Heat Rate

The cost of fuel input creates an incentive for firms to adopt heat rate-improving technology. Ideally, we could estimate the cost of adopting such technology by investigating empirically the relationship between coal prices and technology adoption. However, because data on technology adoption are not available, we instead focus on the effect of coal prices on heat rates, which are expected to change with technology adoption. This presents the challenge that heat rates are not fixed for a given heat rate technology, and in fact heat rates depend on utilization, as discussed above. However, if we condition the analysis on utilization, there is a direct relationship between heat rate technology and heat rate. That is, if two units have the same utilization but different heat rates (and are otherwise identical), they have a different heat rate technology; we refer to the unit with the lower heat rate as having a "better" heat rate technology. Figure 4 demonstrates this point by plotting heat rates against utilization for two levels of heat rate technology, T. The figure shows that if the units have the same utilization, the unit with the lower heat rate has the better technology. Therefore, conditioning on utilization, there is a direct relationship between heat rate and heat rate technology.

6.2 Estimating the Effect of Coal Prices on Heat Rates

We begin by focusing on the average effect of coal prices on efficiency. Based on the model in the previous section, we show that we can recover technology costs by estimating the relationship between coal prices and heat rates. We use the full panel data set and a linear regression model in which heat rate is the dependent variable that measures efficiency, and the major independent variables are coal prices and unit characteristics.

Equation (2), the first-order condition from the firm's heat rate choice in the absence of regulation, is the basis for the estimating equation. We begin by assuming a functional form for the heat rate cost function

$$c(\gamma_i, h_i) = -\gamma_i \frac{\alpha}{1+\alpha} h_i^{(1+\alpha)/\alpha} \tag{A7}$$

where $\alpha < 0$ and $\gamma_i > 0$ (below we discuss the implications of this particular functional form). We assume further that the first term in equation (2) is small compared with the other terms. Taking logs of equation (2) and allowing the firm to adjust heat rate in response to temporal fuel price variation yields the following relationship between the heat rate of unit i in time t and the price of coal, the unit's utilization, and other unit characteristics

$$\ln h_{it} = \alpha \ln p_{it}^f + \ln u_{it} + f(X_{it}) + \varepsilon_{it} \tag{A8}$$

where X_{it} is a vector of characteristics and parameters that affect the cost of adopting a heat rate technology and the equilibrium heat rate. According to the model, the vector includes the unit's capacity, K_i; the scale parameter ρ; and the heterogeneity parameter in the cost function γ_i. The coefficient α is to be estimated and ϵ_{it} is a random error term.

Because equation (8) uses a log-linear relationship between coal prices and efficiency, we can interpret α as the elasticity of the heat rate to the coal price. Based on equation (2), we expect that high coal prices increase the incentive to adopt technology that improves efficiency and reduces heat rates, so α should be negative. Importantly for the policy comparison below, equation (2) suggests that α is inversely related to the cost of reducing the heat rate.

We observe that the functional form in equation (7) plays an important role in deriving the linear estimating equation. We view the constant elasticity approximation as reasonable given that the policy simulations consider small heat rate improvements that are of similar magnitude to those observed in the data. Besides the constant elasticity assumption, an important feature of equation (7) is that the unit-specific constant, γ_i, and the capacity, K_i, enter multiplicatively as do the unit's capacity in equation (2). Consequently, the effect of fuel prices on heat rates can be separated from the effect of γ_i or K_i on heat rates. We can partially relax this assumption by including interactions of the fuel price with characteristics of the generation unit such as capacity.

The extent to which this predicted relationship will be observed at individual units depends on unit-specific characteristics and other factors, including the assumption that the firm maximizes profits. If the firm does not minimize costs perfectly, which may be the case in a regulated environment for example, the data should reveal a relatively weaker relationship between heat rates and fuel price for generating units owned by that firm. Hence, the economic relationship observed in the data may overestimate the cost of responding to regulations that require an improvement in emissions rates or heat rates.

For two reasons, we estimate equation (8) by weighted least squares using the unit's generation as weights, as opposed to estimating the equation by ordinary least squares. First, there is likely to be measurement error in both heat rates and coal prices, where the former would increase standard errors and the latter would bias the coefficients. Appendix 2 describes the steps we take to reduce measurement error in heat rates, but measurement error for both variables is likely to persist. We weight by generation based on the observation that heat rates and coal prices are much less variable at larger and more heavily utilized units.[13] The second reason we weight by generation is motivated by our policy focus; the policies aim to reduce generation-weighted average heat rates (i.e., emissions), and hence we seek an estimate of the generation-weighted average cost of reducing heat rates.

The theoretical model and discussion of the data and power plant characteristics illustrate three major challenges to estimating equation (8): reducing measurement error in heat rates, controlling for utilization, and controlling for the generator characteristics that affect heat rate adoption in $f(X_{it})$. We just discussed weighting observations to reduce measurement error, and to further reduce measurement error we aggregate observations to five-year time periods. Each unit is observed a maximum of five time periods over the sample 1985–2009.

An increase in coal prices decreases utilization because it makes the relative cost of operating coal plants more expensive. The decreased utilization raises heat rates, for reasons discussed earlier. Failing to control for utilization would bias the estimate of α toward zero; that

[13] We obtain similar results using alternative weighting schemes, such as the unit's capacity, and other approaches to reducing measurement error such as dropping observations for generators below a particular size or utilization cutoff.

is, a simple regression of heat rates on coal prices would combine the positive effect of a change in utilization with the negative effect of a change in efficiency. Since we are interested in the latter, it is important that we control for the effect of coal prices on heat rates due to utilization. To the extent we fail to do so would overestimate the cost of reducing heat rates.

To reduce the possibility of such bias, equation (8) includes an extensive set of controls for utilization. Equation (8) allows for a constant elasticity of heat rate to utilization, which was assumed in the model from the previous section. However, the relationship between coal prices and utilization may not be log-linear; that is, including the log of utilization may imperfectly control for utilization and bias other estimated coefficients. Therefore, we also include a large set of variables that account for the potentially nonlinear relationship between utilization and efficiency, particularly at low levels of utilization. Utilization controls in X_{it} include the unit's utilization rate (generation divided by the maximum possible generation over the time period). We also include a set of fixed effects for the number of months in the five-year period that the unit operated below 10 percent or below 30 percent of rated capacity.

Despite the extensive set of utilization controls, a potential concern is that because they are based on monthly data they may not account for the nonlinear effects on the heat rate of stopping or starting a generator. However, the control variables in X_{it} were selected based on an analysis of hourly data in which we control directly for stops and starts. From that analysis, we concluded that these monthly utilization variables predicted heat rates nearly as well as hourly utilization variables such as variables measuring stops and starts. Furthermore, the results are unaffected if we use high-order polynomials in utilization rather than the log of utilization.

As the model suggests, many variables affect heat rates besides fuel prices and utilization. However, controlling for utilization in equation (3) avoids the need to include other variables that affect heat rates via utilization, such as natural gas prices; below we show that the results are similar if we control for natural gas prices and other determinants of utilization.[14]

However, there remain many variables that affect heat rates independently of utilization. Based on equation (2) and equation (7), two such variables are the unit's capacity, K_i, as well as the unit-specific constant, γ_i. There may be other, potentially time-varying factors that also

[14] There may be an interaction between natural gas prices and coal prices. A coal plant that faces high coal prices and low natural gas prices would have an even greater incentive to reduce its heat rate than if natural gas prices were higher. We do not find strong evidence for this mechanism during the sample period.

affect heat rates. For example, the plant or firm may have better managers, or the unit may operate with less variation in its utilization rate. For shorthand, we refer to the unit's type as representing the cumulative effect of such factors on the unit's heat rate; a unit with a high type has a low heat rate, all else equal. This introduces a significant challenge because the unit's type is not directly observed in the data.

We take two approaches to controlling for time-invariant factors that enter equation (2) or equation (7), including components of the unit's type that do not vary over time. Below, we show results that control for fixed unit characteristics by including the variables for the technical characteristics that were analyzed earlier: state, rated capacity, firing type, primary fuel type, and whether the unit is a cogenerator.[15] As an alternative specification, we include fixed effects for each generating unit.

We take several approaches to controlling for time-varying components of the unit's type. The utilization variables control for type to the extent that high-type units have higher utilization rates because they are more efficient; that is, there is a monotonic relationship between type and utilization (Olley and Pakes 1996). Several other included variables are likely to be correlated with type: age, vintage, percentage of coal in total heat input (to account for units that use a small amount of natural gas or biomass), ownership type, and whether the unit has selective catalytic reduction (for nitrogen oxides) or flue gas desulfurization (for sulfur dioxide).[16] The ownership type fixed effects control for the fact that different ownership types may have different incentives to improve heat rates. We also report results from a specification that includes firm by time period fixed effects, which control for changes over time in management quality at the parent company level. Throughout, we assume that $f(X_{it})$ is a linear function of the time-invariant and time-varying factors, although in a few cases we allow for interactions among a subset of these factors.

Given the control variables included in equation (8), the main remaining concern is that coal prices may be correlated with omitted, particularly time-varying, unit characteristics, even

---

[15] More specifically, the control variables include a set of fixed effects for the unit's firing type and fuel type. Deciles are computed for rated capacity across the units in the sample, and equation (3) includes fixed effects for each decile.

[16] Equation (8) includes a set of age fixed effects, for 6-10 years, 11-20 years, 21-30 years, 31-40 years, 41-50 years, 51-60 years and above 60 years, where 0-5 years is the omitted category. We do not control for vintage, although the specifications that include unit fixed effects control for vintage.

after conditioning the analysis on utilization. In particular, coal markets are not perfectly competitive (Busse and Keohane 2004), and certain owners of coal plants may be able to negotiate more favorable prices. This could result in a negative correlation between the unobserved type and the observed coal price. In the estimation we take several approaches to address this possibility, such as including firm by time period fixed effects and instrumenting for the unit's coal price.

6.3 Estimating the Effect of Coal Prices on Utilization

To estimate the rebound effect stemming from an improvement in heat rate, we assume the following utilization function

$$u(p_r^f h_i, \theta_i) = \theta_i (p_r^f h_i)^\beta \tag{A9}$$

Note that fuel costs and the heat rate have the same effect on utilization by assumption, which is reasonable given the importance of marginal generation costs in determining utilization. Therefore, we are interested in the effect of coal prices on utilization, holding heat rates constant. We estimate the reduced-form utilization function

$$\ln u_{it} = \beta \ln p_{it}^f + g(X_{it}) + \varepsilon_{it} \tag{A10}$$

The dependent variable is the log of utilization, and the independent variables are the same as in equation (8), except that we omit the utilization controls. Because we are interested in the short-run effect of fuel prices on utilization—that is, holding fixed heat rates, we use annual observations rather than five-year time periods. The main challenge to estimating equation (10) is that utilization and heat rate are jointly determined; that is, changes in coal prices affect both heat rates and utilization rates. In equation (8) we account for this problem by assuming the effect of the heat rate on utilization is small. We could include heat rate on the right-hand side of equation (10) but doing so would bias the estimated coefficients because of the endogeneity of the heat rate. Instead, we omit heat rate from equation (10) and interpret the coefficient β as the effect of coal prices on utilization after accounting for the change in heat rate technology. In practice, this quantity should be very similar to the quantity we are trying to estimate, which is the effect of coal prices on utilization holding fixed heat rate technology; the latter should be larger in magnitude. The elasticity of heat rates to coal prices turns out to be about –0.02. This suggests we may underestimate the effect of coal prices on utilization, but probably not by a large amount. Using annual observations rather than five-year time periods to estimate equation (10) further reduces the likelihood that omitting the unit's heat rate biases the estimate of β.

7. Estimation Results

This section first presents the estimated effects of coal prices on heat rates and utilization. We discuss the magnitude and robustness of the estimates, as well as interactions between New Source Review and coal prices. We then discuss the abatement costs of efficiency improvements that the estimates imply. The section concludes by reporting the estimated rebound effect.

7.1 Effect of Coal Prices on Heat Rates

7.1.1 Main Results

Table 4 reports estimates of equation (8), where the dependent variable is the unit's log heat rate over the five-year time period. Panel A does not include unit fixed effects, and Panel B includes unit fixed effects. The main coefficient of interest is α, which is the coefficient on the log price of coal and is interpreted as the elasticity of the heat rate to the price of coal.[17] The regression includes a large number of other control variables for technical characteristics and utilization, as well as state and time period fixed effects.

Column 1 shows unweighted results and column 2 shows weighted results. Without fixed effects, the estimate of α is -0.043 (unweighted) and -0.049 (weighted); in both cases, the estimate is significant at the 1 percent level. The point estimate can be interpreted as an elasticity; a one standard deviation increase in coal prices reduces heat rate by about 1.5 percent.

Adding unit fixed effects decreases the estimate to -0.011 (unweighted) and -0.023 (weighted). The weighted estimate is significant at the 1 percent level but the unweighted estimate is not statistically significant. The weighted estimate implies that a one standard deviation coal price increase causes a 0.8 percent heat rate reduction.[18] Recall that we weight observations by generation to address measurement error in heat rates and coal prices and to interpret estimates as the generation-weighted average. The results in the first two columns

[17] The coal price variable reflects the cost of using coal. Therefore we add to the delivered coal price the associated sulfur dioxide emissions multiplied by the sulfur dioxide permit price, adjusting for whether the unit is connected to a scrubber and whether the unit participated in the Acid Rain Program.

[18] The calculation of the change in heat rate for a one standard deviation price increase uses the standard deviation of observed coal prices across the entire estimation sample. In Panel B, α is identified by within-unit coal price variation. Therefore, a more appropriate calculation may use the standard deviation of coal prices after removing unit-specific means. The standard deviation is smaller than for the full sample ($0.50/mmBtu instead of $0.62/mmBtu), and a one standard deviation increase corresponds to a 0.6 percent heat rate decrease.

support the first rationale for weighting. Because the unit fixed effects absorb much of the coal price variation, adding unit fixed effects aggravates attenuation bias caused by measurement error.[19] Therefore, weighting observations by generation to reduce bias from measurement error should cause the magnitude of the coefficient to increase more for the regression with unit fixed effects than for the regression without unit fixed effects; in fact, we observe exactly this pattern.[20]

There are several possible explanations for the fact that adding unit fixed effects decreases the estimate by a substantial amount (i.e., comparing column 2 in Panels A and B). First, the estimate with fixed effects could correspond to a short run estimate because it is estimated using within-unit variation; the estimate without fixed effects may reflect efficiency improvements that require major changes to capital equipment that are difficult to implement quickly but can be observed on a cross-sectional basis. However, another explanation for the discrepancy suggests caution with making this interpretation: the specification that does not include unit fixed effects may not fully control for time-invariant unit characteristics. Given these considerations, we consider column 2 in Panel B as the baseline specification.

We show that, overall, the results are robust to other specifications and functional form assumptions. Equation (8) imposes a log-linear relationship between coal prices and heat rates. Although this functional form is a first-order approximation to a more complicated relationship, it is somewhat arbitrary. The results are similar using other functional forms, for example, estimating the equation using the heat rate and coal price in levels rather than in logs. Figure 5 provides further confirmation that the log-linear approximation is reasonable. Residuals are constructed from a regression of log heat rate and log coal price on the same independent variables as column 2 in Panel B, except that log coal price is not included as an independent variable. The figure plots the residuals as well as the fitted values from a regression of the heat rate residuals on the coal price residuals (the slope of the fitted values equals α). The figure does

[19] More specifically, if we regress coal prices on the same controls as are included in column 1, the root mean square error is about 25 percent smaller including unit fixed effects instead of state fixed effects.

[20] We have tried several other approaches for reducing measurement error in the heat rate and coal price. First, if we weight by the unit's capacity rather than generation the estimates are similar to the baseline, although slightly smaller. Second, because some of the data are imputed for units with capacity less than 100 MW, we estimate equation (8) without weighting but restricting the sample to units with capacity of at least 100 MW. The estimate of the coal price coefficient is -0.014 and is significant at the 5 percent level. Finally, Table 7 reports results if we instrument for the coal price.

not indicate the presence of a nonlinear relationship between the heat rate and coal price residuals, which supports the linearity assumption in equation (8).

Given the fixed costs and the potential importance of managerial quality or technical information in implementing heat rate improvements, heat rates at units owned by larger firms may be more responsive to coal prices than units owned by smaller firms. Column 3 restricts the sample to the 20 largest firms, based on firms' total coal capacity in 2009.[21] The magnitudes are similar to those in column 2 although the standard errors are larger (likely reflecting the smaller sample) and the estimate is significant at only the 10 percent level in Panel B. Thus, we do not find strong evidence that units owned by larger firms respond more to coal prices.

Recall that we aggregate the data from annual to five-year time periods to reduce measurement error. For this reason we expect to obtain smaller coefficients (in magnitude) using annual observations. Column 4 shows this to be the case.[22]

7.1.2 Persistence

Ideally, we would estimate the long-run effect of coal prices on heat rates because this relationship would provide insight into the effects of potentially permanent CO_2 policies.

Because the regressions include state and period fixed effects, we estimate α using deviations in coal prices from state averages and from the period averages. If such deviations were less than fully persistent, we would underestimate the long-run effect of coal prices on heat rates. We use several approaches to address this issue and find some evidence that the baseline estimates understate the long-run effect of coal prices on heat rates.

We begin by estimating the persistence of coal prices. Panel A of Table 5 reports a series of regressions in which the log price of coal is the dependent variable and independent variables include the one-period lag coal price, state fixed effects, and period fixed effects. Column 1 uses the overall price of coal based on total contract and spot market purchases, column 2 uses the contract price, and column 3 uses the spot price. The estimates are all much less than one, which suggests that all prices exhibit some mean reversion.

---

[21] A firm refers to a parent company; for example, Southern Company is a single firm, which owns Alabama Power, Georgia Power, and several other subsidiaries.

[22] Another possible explanation for the smaller magnitude in column 4 is that firms may take more than one year to adopt heat rate technology in response to coal prices.

Panel B shows that the coefficients on lagged prices are much higher if the state fixed effects are removed. The same is not true if period fixed effects are removed (not reported), which suggests that if the estimates in Table 4 are less than long-run estimates, removing state fixed effects in the baseline regression would increase the estimated magnitude of α.

Column 1 in Table 6 repeats the specifications in column 2 of Table 4, except that the average coal price is separated into spot and contract prices. Given the relatively low persistence of these prices in the data, we would expect these coefficients to be smaller in magnitude than the estimate for the overall price of coal. This pattern is observed, in fact, and the estimates on the spot and contract prices are essentially zero.

The specifications in Panel B of Table 5 indicate that omitting state fixed effects from equation (8) should result in larger coefficients for the spot and contract prices, and to a lesser extent for the overall price. Columns 2 and 3 of Table 6 report the results of this exercise, which conform to expectations.[23]

Another approach to assessing whether we estimate long-run responses is to use the coal price in the first year of each period rather than the average price in each period. Suppose firms take several years to implement heat rate improvements. If coal prices happen to change late in the five-year periods in our sample, we would underestimate the effect of coal prices on heat rates over a five-year period. To investigate this possibility, column 4 replaces the average coal price over the time period with the price in the first year. The results are similar to the baseline, as is the case if we estimate equation (8) using the price in the third year of each time period (not reported). We would expect to estimate a smaller coefficient using the price in the final year of each period rather than the average price, which is the case (not reported).[24]

Finally, we attempt to incorporate persistence directly in equation (8). We estimate the same regression as column 1 of Table 5 separately for each region of the North American Electric Reliability Corporation (NERC), and we define a persistence variable as the coefficient on the lagged coal price. Column 5 of Table 6 adds to the baseline the interaction of the

---

[23] Because of differences in the persistence of contract and spot prices, heat rates at units that use coal purchased under long term contracts may respond differently to coal prices than units that rely more on spot market purchases. In fact, the estimates are similar to the baseline if we restrict the sample to units that use coal mostly purchased under long term contracts, suggesting that this is not a significant concern.

[24] A related concern is that the five-year time periods are somewhat arbitrary. We obtain very similar results using 3-year, 4-year or 10-year time periods.

persistence variable with the coal price. The interaction coefficient would be negative if more persistent coal price shocks had a larger effect on heat rates. The estimate is positive, but it is small and not statistically significant; in short, there is not sufficient cross-region variation in persistence to identify an effect. Taken together, Table 6 provides some evidence that we underestimate the effect of a permanent change in coal prices on heat rates.

7.1.3 Potential Coal Price Endogeneity

As noted above, coal prices may be correlated with unobserved and time-varying generation unit attributes, such as managerial quality or negotiations over coal prices. Table 7 reports specifications that address this possibility. Overall, we find that the results are robust.

In column 1 of Table 7 we instrument for the coal price to address the possibility that some firms have market power and are able to negotiate lower coal prices. We interact two pre-sample variables—the plant's coal price in 1984 and the plant's share of coal from the Western region in 1984—with a set of period fixed effects. For the regression that includes unit fixed effects, the instruments are valid if changes in a firm's market power (or measurement error or any other time-varying unit attribute that affects coal prices) are uncorrelated with the initial level of the omitted variable. The instruments are jointly strong predictors of the unit's coal price, and Panel B shows that although the estimate is larger than the baseline it is not statistically different. The result suggests that time-varying firm or unit-level variables do not bias the estimate towards zero.

Unobserved unit characteristics are likely to be correlated with utilization. Column 2 restricts the sample to units with high utilization, or more specifically, units with a median utilization rate above 0.5 across the five time periods. The estimates of α are similar to the baseline.

Plant or generation unit entry and exit could be correlated with unobserved unit characteristics. Estimating equation (8) and omitting unit fixed effects would result in biased estimates in that case. Column 3 controls for entry and exit by restricting the sample to a balanced panel. The estimate of α is close to the baseline.

Unobserved unit characteristics are likely to be correlated at the firm level. Some firms may be more efficient than others because of ownership structure, regulatory environment, or other factors. Column 4 allows for time-varying unobserved firm characteristics by including firm-period fixed effects. The results are again similar to the baseline and to the instrumental variables results in column 1, which also attempt to control for time-varying firm characteristics.[25]

The discussion has focused on unit, plant, or firm variables that vary over time and affect heat rates, but there may also be aggregate shocks that affect utilization in ways that we do not control for in equation (8). Column 5 includes several variables measured by state and year to proxy for such shocks: generation capacity, generation, gross state product, employment, and population. All variables are included in logs, and there are three separate variables for capacity and generation: total, coal, and natural gas. Because the capacity and generation variables are not available before 1990, the regressions do not include the first time period. Adding these variables to the baseline specification does not significantly affect the estimate, although they are somewhat smaller. The results are not affected if we add natural gas prices to the baseline regression (not reported).

7.1.4 Regulatory and Market Incentives

Regulatory and market incentives may affect the relationship between coal prices and heat rates. We have noted that some measures that might improve efficiency may be considered a major modification to the plant and thereby trigger a permitting process for NSR. This process may impose additional costs on plant owners and may provide a formidable barrier to making investments to improve heat rates. To examine the influence of NSR, we divide our data into two time periods. Before 1998 there was little concern about NSR enforcement proceedings, but that year EPA initiated information requests that signaled a potentially more aggressive stance. Keohane et al. (2007) note that by October of that year, the electricity trade press began reporting the possibility of EPA enforcement, and a year later, in November 1999, the Department of Justice initiated the first of a series of enforcement actions. We consider the first regime through 1998 as one in which energy efficiency investments would be considered and implemented as part of routine maintenance and would not be expected to trigger an enforcement action. During

[25] Previously we noted that if firms are not profit maximizing, the estimate of α is likely to be biased to zero because firms do not choose heat rates based on expected costs and benefits. It is also possible that some firms "satisfice" or face high capital costs, in which case they may choose only the most profitable heat rate improvements rather than those with positive net present value. In such a case, equation (8) would be modified to include a vector of the coal prices and attributes of each of the other units owned by the same firm. Note that if all coal prices of other units have the same coefficient, including firm-year interactions would account for their affect (and likewise for other unit characteristics). But this may not hold in practice, and we also estimate a specification that includes the coal prices of other units owned by the firms, which does not affect the main results.

the second regime (1999–2009) we assume that energy efficiency investments would be subject to scrutiny because they might increase utilization of a plant and its emissions, and hence would be likely to trigger NSR review.

We use the two regimes to assess the interactions between NSR and heat rates. We expect heat rates to respond more to coal prices in the first regime when heat rate improvements would not have been expected to trigger NSR. We could estimate a separate α for the two regimes, expecting the estimate for the first regime to be larger in magnitude. Many other regulatory and market changes occurred between the first and second regimes, however, which could confound the analysis. Instead, we use annual data and estimate a separate α for each year of the sample; otherwise the specification is identical to the baseline in Table 4. Figure 6 plots the coefficient estimates and 95 percent confidence intervals. The figure provides only suggestive evidence for an NSR effect because it reveals a positive trend in α (heat rates becoming less sensitive to coal prices) over most of the sample period; in other words, factors other than NSR may have caused the upward trend. Without better data we cannot reach a stronger conclusion.[26]

Section 3 suggested that privately owned firms might respond more to economic incentives. We examine the influence of ownership structure with a dummy variable interacted with log coal price. The sign of the interaction coefficient is consistent with the hypothesis that privately owned firms are more responsive, separately considering investor-owned utilities and nonutility generators, but it is not statistically significant (not reported). Similarly, we do not observe large differences across federal, state, municipal, and co-op ownership types.

Beginning with the 1992 Energy Policy Act, competition began to affect the industry in a stronger way and different types of private ownership emerged. These changes suggest that privately owned units, and in particular nonregulated privately owned units, would respond more to coal prices after around 1995. For example, Fabrizio et al. (2007) find that the transition to market-oriented environments in this period had led to the greatest efficiency gains at investor-owned plants in states that restructured their electricity markets. Chan et al. (2013) find a two percent improvement in heat rates at investor-owned coal plants following the introduction of competition in states that restructured since the mid-1990s. Our results are not affected if we

[26] Whether adding pollution abatement controls triggers NSR also varies over time. This suggests examining changes in heat rates after pollution equipment is added, but unfortunately, there are not enough units that add equipment in the different regimes to test this hypothesis.

control for the effect of restructuring on average heat rates by adding to the baseline specification the interaction of a set of state dummies with a post-1995 dummy. Although there are modest differences in heat rates across ownership types, we have estimated several specifications that allow for different responses to coal prices across ownership types and generally do not find significant differences, either for units in the same period that have a different ownership type or for units with a particular ownership type across periods. The pre- and post-1995 distinction is somewhat crude, however, and the interaction between regulatory environment and coal prices is a topic for future research.

7.2 Effect of a Carbon Tax on Heat Rates

The estimation results can be used to estimate abatement costs at existing coal units from improving heat rates. We use the baseline estimates in Table 4 to calculate the percentage reduction in heat rate from a $10 per ton of CO_2 tax on coal (i.e., assuming the tax is fully passed through to delivered coal prices and holding fixed utilization). We estimate heat rate reductions of 1 to 2 percent, where the lower estimate derives from the specification with fixed effects and the upper estimate from the specification without fixed effects. As the previous analysis showed, this range is quite robust to a variety of alternative specifications. The range overlaps engineering estimates, although the upper end of the range is higher than the engineering estimates. For some of the robustness checks, the effect of the coal price on heat rates was even larger than in the baseline specification, which suggests that the baseline estimate may, if anything, be biased toward zero. On the other hand, Figure 6 suggests that the responsiveness of heat rates to coal prices may have diminished over the sample period at the same time as other policy changes were observed. Hence, interpretation of our results should consider the interaction with other relevant policies.

7.3 Effect of Coal Prices on Utilization

We use equation (10) to estimate the effect of coal prices on utilization. The dependent variable is the log utilization rate by unit and year, and the independent variables are the same as in equation (8) except that the utilization controls are omitted (and we do not weight observations by generation).

Tables 8 and 9 report the estimation results. Columns 1 and 2 repeat the specifications from Table 4. The coefficient on the log price of coal is interpreted as the elasticity of the utilization rate to the coal price. Across specifications, the estimated elasticity is typically around

–0.4, which suggests that a 10 percent price increase would cause a 4 percent reduction in the utilization rate (from a mean utilization rate of 0.61 across all years in the sample).

The remaining columns in Table 8 and all of Table 9 document the robustness of these results. One specification in Table 8 is particularly noteworthy. The effect of coal prices on utilization depends on other factor prices, particularly natural gas prices. If natural gas prices are high relative to coal, a coal price increase may have a small effect on coal unit utilization because the increase would not affect the competitiveness of coal units compared with gas units. In fact, we do observe a smaller effect of coal prices on coal unit utilization when natural gas prices are high. For example, column 4 of Table 8 uses observations from the years 2001–2009, when natural gas prices were higher than in the previous decade, and the reported elasticities are much smaller than in column 1, which uses the entire sample. The results suggest that the price of natural gas relative to the price of coal affects utilization, which is confirmed in other unreported specifications (e.g., replacing the coal price in equation (10) with the ratio of the coal to natural gas price); other factors could explain the results, of course, and we treat the results with some caution.

Table 9 reports a statistically significant effect of coal prices on utilization across specifications that parallel the heat rate specifications in Table 7. Overall, we estimate an elasticity of utilization to coal prices of about –0.2 to –0.6. Because of our interest in estimating the cost-effectiveness of future policies, we prefer the estimates from the most recent time period which are at the lower end of this range (i.e., the specification in column 4 of Table 8, which uses observations from 2001–2009).[27]

8. The Cost-Effectiveness of Policy Alternatives

We use the simple model from Section 5 to illustrate the policy implications of the empirical estimates. One paper (Burtraw et al. 2011b) has previously represented endogenous efficiency improvements at existing coal plants within a simulation model using engineering estimates of investment opportunities and costs. This section addresses similar questions but does so by incorporating the econometric estimates into a simple model of the electricity sector. We compare the cost-effectiveness of four policies for reducing CO_2 emissions from existing

[27] The rebound effect could vary across regions because of differences in coal supply curve price elasticities or other factors. If we estimate a separate rebound effect by NERC region the estimates are not statistically signficiant and the standard errors are large, which likely reflects the limited within-region coal price variation.

coal units. We consider two forms of a technology standard, a traditional heat rate standard and a flexible standard, plus a Btu tax on coal and an emissions tax on fossil fuel. We briefly summarize these policy alternatives, and then we summarize the model and report the results.

8.1 Policies

We characterize a traditional heat rate standard as imposing a maximum heat rate ceiling (mmBtu/MWh) requiring all facilities to achieve that standard or to retire. (The heat rate standard has a close analogue in an emissions rate standard in tons of CO_2/MWh.) Facing a traditional standard, a facility operator evaluates the net profit of the necessary improvements to enable continued operation, and α serves as a measure of the opportunity cost of those improvements. Improving heat rates would reduce fuel costs and increase utilization, as captured by β.

A flexible heat rate performance standard sets a benchmark (e.g., a uniform standard or one that varies by generator type) and allows firms that over-comply with the benchmark to transfer efficiency credits to those who fail to comply. Facilities with economic advantages make relatively greater investments and transfer efficiency credits to other facilities. Because the cost of efficiency improvements varies across facilities, and the flexible standard equates marginal costs across facilities, flexibility allows for greater total emissions reductions to be achieved for the same aggregate cost as under a traditional standard.

Introducing a market-determined opportunity cost for generation units is one feature that a flexible performance standard has in common with a coal Btu tax or an emissions price that might be introduced through either an emissions cap or a tax. However, a flexible performance standard is different from cap-and-trade in two important ways. First, a standard does not cap emissions. Overall emissions are able to grow with increased production, even as that production becomes more efficient. Second and following from the first, with the flexible performance standard the regulator does not explicitly allocate credits. Instead, the regulator sets a benchmark emissions rate for each unit. Credits are earned for electricity production at the benchmark rate; that is, the regulator implicitly allocates credits through the assignment of benchmark rates. Facilities surrender credits at their actual emissions rates. Facilities with emissions rates below the benchmark earn surplus credits with each unit of generation, while facilities with rates above the benchmark earn a deficit that would be filled with transfers from other facilities. Hence, the flexible performance standard can be understood to encompass two instruments in one policy: it imposes an opportunity cost on heat rate, providing an incentive for heat rate improvements, and it provides an output subsidy equivalent to the value of credits earned for each unit of electricity generation.[28] Thus, the flexible standard introduces an inherent rebound effect for units that over-comply by giving them added incentive to increase output. The consequence is that a flexible performance standard is likely to result in greater utilization of the regulated facilities than under a Btu tax or an emissions price, and conceivably greater utilization than in the absence of regulation.

8.2 Model

Section 5 describes our static model of the electricity market. The model contains a representation of all coal-fired power plants in operation in 2009 organized into 22 regions. Each firm in a region owns a single unit and sells electricity into a competitive wholesale market. There is perfect substitutability among facilities within a region but no transmission between regions. Every coal plant is paired with a backstop natural gas combined-cycle power plant that provides residual generation necessary to meet a fixed level of demand. The marginal costs of the natural gas plants vary by location and are calibrated to equal those of the coal plant in the baseline. Gas supply is calibrated to adjust to changes in coal prices and generation consistent with the coal generator utilization elasticity β. Profits for the power plants come from the difference between the wholesale power price and the sum of operating costs and fuel costs.

The analysis focuses on the optimal choice of heat rate conditional on the capacity and exogenous variables. We use the model in section 5 to solve for the equilibrium with and without each of the four policies in place. We use equations (7) and (9) for the heat rate cost and utilization functions, and choose baseline parameter values of $\alpha = -0.023$ and $\beta = -0.2$, based on the regression of heat rate and utilization on coal prices. As noted above, we find some evidence that the estimate of α may be biased toward zero but also that the estimate may have diminished over the sample period.

The model contains two parameters that introduce unit-specific heterogeneity. The first is θ_i, which allows for the possibility that units with the same coal price and heat rate may have different utilization rates. We obtain a value of θ_i based on the unit's observed utilization, coal

[28] One can view this as equivalent to an emissions trading program where the cap is determined endogenously as the product of the benchmark emissions rate and output, and allocation is performed on the basis of output.

price, and heat rate. The second heterogeneity parameter, γ_i , allows the cost of improving heat rates to vary across generators because of factors such as age or firing type. We calculate γ_i from the first-order condition, equation (2).

The cost of each policy to regulated units depends on (1) the cost of changing heat rates and (2) the change in operating profits (i.e., the difference between the electricity price and marginal costs multiplied by generation). In other words, we estimate the short-run costs of the policies, in the absence of entry, exit, or changes in capacity.

The effectiveness of each policy is equal to the change in electricity sector CO_2 emissions. Therefore, we need to characterize the effect of each policy on the rest of the sector. We make the simplifying assumptions that the policies do not affect total electricity generation in the system, and that any changes in total coal generation are offset by natural gas generation. The natural gas emissions rate is assumed to be 0.05 tons of CO_2 per mmBtu. The wholesale electricity price is determined by the marginal cost of a natural gas combined-cycle unit at each location. Their marginal costs depend on utilization, the price of natural gas, and operations and maintenance costs for combined-cycle units in the region. The national average marginal cost is $38/MWh.

We use the first-order conditions above to solve for the heat rate of each unit under each policy. We then use the utilization function to estimate generation for each unit. Finally, we calculate the change in total power plant operating profits and emissions by summing across units.

We measure the rebound effect as the change in system level emissions due to the change in utilization that arises from heat rate improvements at the generating units. To estimate the change in emissions in the absence of the rebound effect we evaluate the new equilibrium with utilization determined by the baseline heat rates. See Appendix 3 for details.

8.3 Policy Effects on Efficiency Investments and Fuel Switching

The traditional standard is calibrated to achieve a 1 percent reduction in the average heat rate, and the other policies are calibrated to achieve the same emissions reductions that result. Under the traditional standard, each unit must achieve a heat rate of 10.80 mmBtu/MWh or reduce its heat rate by 10 percent, whichever results in a higher heat rate. The changes in heat rates that occur in the simulations are of similar magnitude to the changes observed in the data. In the baseline 22 percent of units have heat rates above this standard. Achieving the standard results in a 0.88 percent reduction in emissions across coal and gas units, illustrating that the rebound effect erodes 13 percent of the emissions reduction associated with efficiency improvements. Coal generation increases of 0.19 percent are concentrated entirely at facilities that have initial heat rates above the standard. These results are presented in Table 10.

To achieve a comparable 0.88 percent reduction in emissions under the flexible standard requires an emissions rate benchmark of 10.28 mmBtu/MWh, and a price for a tradable performance standard credit of $1.01/mmBtu. The rebound effect due to the increase in utilization erodes 18 percent of emissions reductions that would otherwise result from the heat rate improvement, illustrating the important role of the output subsidy under a tradable performance standard. Heat rate improvements occur across the distribution of coal plants.

Under the taxes, the emissions reductions are achieved with an emissions tax of $1.17 per metric ton CO_2 and a coal Btu tax of $0.12 per mmBtu. As expected, the coal tax requires a higher tax rate (after converting to a common metric) than the more broadly based emissions tax. Under the taxes, the rebound effect is much smaller, less than 2 percent in each case. However, it is noteworthy that the rebound is slightly greater under the emissions tax because it applies to natural gas as well as coal. Hence, the change in relative prices is less under the emissions tax than under the coal Btu tax.

Although calibrated to achieve comparable emissions reductions as the inflexible standard, the flexible standard also yields a nearly identical decline in heat rates (1.01 percent). In contrast, the tax policies yield a reduction in heat rates of 0.12 percent. The flexible standard leads to an order of magnitude greater annual investment in heat rate improvements than the taxes, and the inflexible standard results in nearly another order of magnitude greater investment than the flexible standard.[29] Because of these investments, the cost of generation falls by 0.7 percent for the flexible standard and by 0.8 percent for the inflexible standard; in contrast generation cost is nearly unchanged under the taxes (not including the tax burden). The standards result in an increase in coal generation whereas the taxes lead to an increase in natural gas generation. In sum, we find these policies are relatively most effective with respect to their specific targets. The relatively narrowly focused performance standards lead to greater

[29] This result differs from Burtraw et al. (2011b), where retirement of existing facilities can occur.

improvements in energy efficiency and less fuel switching. The price-based policies are effective across all eligible channels to achieve comparable emissions reductions.[30]

We have conducted sensitivity analysis using a larger utilization elasticity ($\beta = -0.4$ instead of -0.2), using a larger responsiveness of heat rates to coal prices ($\alpha = -0.049$, the estimated value without fixed effects and weighted by generation, instead of -0.023) and using more stringent policies (heat rate improvements of 2 or 3 percent as the benchmark instead of 1 percent). For the larger utilization elasticity, to achieve a given emissions reduction with the standards there is more investment (compared to the low elasticity simulation) because of the larger rebound effect. The greater responsiveness to coal prices reduces the cost of compliance for the flexible standard by more than half and by substantially more for the inflexible standard. The investment costs reported in Table 10 change roughly linearly with the stringency of the policies. The rebound effect for the flexible standards grows with stringency for the flexible policy only. Across the alternative simulations, the conclusions are qualitatively similar to those from the simulations reported in Table 10 (results are available upon request).

9. Conclusions

Currently, the Clean Air Act is the most important federal policy for addressing climate change. EPA has issued a draft standard for regulation of new fossil-fired electricity generating units and if that is finalized EPA will be required to regulate greenhouse gas emissions from existing coal-fired electricity steam boilers. Engineering case studies have identified possible reductions in emissions through efficiency improvements at zero or moderate cost that could amount to 1.6 percent of total US emissions in 2009. However, a substantial literature raises doubts that widespread and unrealized cost-effective efficiency opportunities exist in general. Heretofore, empirical information was lacking about the actual magnitude and cost of these potential efficiency improvements across the fleet of existing generating units. This information should be central to rulemakings under the CAA, which in this case will rely on technical estimates of potential opportunities as well as costs for the determination of a standard for these

[30] Welfare comparisons across the policies are not strictly valid because we assume the marginal cost of natural gas generation is equal to that for coal generation at every location and that consumer demand is perfectly inelastic. Consequently under the tax policies that increase average gas utilization, we underestimate the costs of that substitution. Under the standards, which increase coal utilization due to the rebound effect, we overestimate the the costs of the policy on average. In our baseline we find the welfare costs of emissions reductions to be $5.57 per ton under the flexible standard and $216 per ton under the inflexible standard.

facilities. Furthermore, this information should be important to the characterization of other policy alternatives, such as legislative proposals for cap-and-trade or emissions taxes. For example, federal agency analysis of proposed cap-and-trade legislation (HR 2454) that passed the House of Representatives in 2009 had no information about potential efficiency improvements at existing coal boilers, and consequently it may have overestimated the abatement cost for the economy as a whole (EIA 2009). Moreover, engineering estimates of potential improvements in other existing stationary sources suggest potential emissions reductions that add up to more than 6 percent of US emissions, without accounting for fuel switching (Burtraw et al. 2011a). Analysis similar to the one in this paper could possibly verify or refute those potential opportunities.

We compile a unique data set covering about 97 percent of US coal-fired units over 25 years of operation and use the data to estimate the opportunities to reduce emissions rates by improving the efficiency of these plants. We note that the introduction of a performance standard will impose opportunity costs similar to costs associated with fuel use; hence, the effective price of fuel will change in response to the regulation.

We find strong evidence that heat rates (a measure of efficiency and a proxy for emissions) respond to changes in fuel prices. For example, a change in coal prices commensurate with a $10 tax on CO_2 emissions would stimulate a 1 to 2 percent reduction in heat rates and emissions rates. This range of costs encompasses the estimates suggested in the engineering literature; indeed it includes the possibility of somewhat lower costs than have been suggested.

An important consideration in whether these efficiency improvements lead to emissions reductions is the degree to which generating units respond to improvements in heat rates by increasing their utilization. If they do, one would observe a rebound in emissions that would erode some of the emissions reductions that would otherwise occur. In fact, we find a significant effect of coal prices on utilization. A 10 percent coal price increase would decrease utilization 2-6 percent (holding fixed heat rates).

We evaluate these econometric estimates in a simulation model with inelastic demand. Natural gas generation is used to meet residual demand, and we solve for changes in heat rate and utilization at coal facilities under a variety of policy scenarios. We find that a performance standard leads to substantially greater investment in efficiency improvements than taxes (or other policies that set an emissions price, such as cap-and-trade) that allow for broader compliance options, including fuel switching. This outcome is consistent with the theory that broad-based, incentive-based policies should be more efficient than performance standards, even after accounting for the opportunity for efficiency improvements. We also find the rebound effect to be 13- 18 percent under the standards, which reduces the emissions savings by roughly the same percentage. The rebound effect under the taxes is an order of magnitude less than under the performance standards.

We have noted a few possible directions for future work. Because the econometric model is static, we do not allow for plant retirements in the simulations. Introducing dynamics would enable an evaluation of more aggressive performance standards than those considered here.

We have provided some suggestive evidence that the economic and regulatory environments affect the relationship between fuel prices and heat rates. Future work could investigate further the effect of competition, NSR, and other policies on the adoption of energy efficiency technology.

In conclusion, substantial long-term reductions in GHG emissions from the power sector will require greater use of nonemitting sources (renewables, nuclear), lower-emitting sources (natural gas), or post-combustion control of carbon. However, we find evidence that there exist opportunities to reduce emissions at existing facilities at low cost in the short run. These reductions could contribute importantly to meeting international commitments as articulated in the Copenhagen Accord.

References

Beer, J. M. 2007. High Efficiency Electric Power Generation: The Environmental Role. *Progress in Energy and Combustion Science* 33: 107–134.

Burtraw, D., A. Fraas, and N. Richardson. 2011a. Greenhouse Gas Regulations under the Clean Air Act: A Guide for Economists. *Review of Environmental Economics and Policy* 5(2): 293–313.

Burtraw, D., M. Woerman, and A. Paul. 2011b. Retail Electricity Price Savings from Compliance Flexibility in Greenhouse Standards for Stationary Sources. *Energy Policy* 10.1016/j.enpol.2011.11.036.

Busse, M. R., and N. O. Keohane. 2004. Market Effects of Environmental Regulation: Coal, Railroads and the 1990 Clean Air Act. CSEM WP 137, Berkeley, CA, Center for the Study of Energy Markets, University of California Energy Institute, UC Berkeley.

Chan, H. R., H. Fell, I. Lange, and S. Li. 2013. Efficiency and Environmental Impacts of Electricity Restructuring on Coal-Fired Power Plants. Cornell University Department of Economics.

Fabrizio, K. R., N. L. Rose, and C. D. Wolfram. 2007. Do Markets Reduce Costs? Assessing the Impact of Regulatory Restructuring on US Electric Generation Efficiency. *American Economic Review* 97(4): 1250–1277.

Fleishman, R., R. Alexander, S. Bretshneider, and D. Popp. 2009. Does Regulation Stimulate Productivity: The Effect of Air Quality Policies on the Efficiency of US Power Plants. *Energy Policy* 37(11): 4574–4582.

Holland, Stephen P., Jonathan E. Hughes, and Christopher R. Knittel. 2009. Greenhouse Gas Reductions under Low Carbon Fuel Standards? *American Economic Journal: Economic Policy*, 1(1): 106–146.

Keohane, N., E. T. Mansur, and A. Voynov. 2007. Averting Regulatory Enforcement: Evidence from New Source Review. NBER Working Paper Series. Working Paper 13512, Cambridge, MA, National Bureau of Economic Research.

Linn, J. 2008. Energy Prices and the Adoption of Energy-Saving Technology. *The Economic Journal* 118: 1986–2012.

Metcalf, G. E., and K. A. Hassett. 1999. Measuring the Energy Savings From Home Improvement Investments: Evidence from Monthly Billing Data. *Review of Economics and Statistics* 81: 516–528.

Olley, S., and A. Pakes. 1996. The Dynamics of Productivity in the Telecommunications Equipment Industry. *Econometrica* 64(6): 1263–1298.

Richardson, N., A. Fraas, and D. Burtraw. 2011. Greenhouse Gas Regulation under the Clean Air Act: Structure, Effects, and Implications of a Knowable Pathway. *Environmental Law Reporter* 41: 10098–10120.

Rosen , M. A., and I. Dincer. 2003. Survey of Thermodynamic Methods to Improve the Efficiency of Coal-Fired Electricity Generation. *Proceedings of the Institution of Mechanical Engineers, Part A: Journal of Power and Energy* 217: 63–73.

Sargent & Lundy, LLC. 2009. Coal-Fired Power Plant Heat Rate Reductions.

Stavins, R. N. 2006. Vintage-Differentiated Environmental Regulation. *Stanford Environmental Law Journal* 25(1): 29–63.

US Department of Energy (DOE) and National Energy Technology Laboratory (NETL). 2009. *Opportunities to Improve the Efficiency of Existing Coal-Fired Power Plants*. Technical Workshop Report, Rosemont, IL.

———. 2010. *Improving the Thermal Efficiency of Coal-Fired Power Plants in the United States*. Baltimore, MD.

US Energy Information Administration (EIA). 2009. Energy Market and Economic Impacts of HR 2454, the American Clean Energy and Security Act of 2009. SR/OIAF/2009-05, Washington, DC, Energy Information Administration.

———. 2011. *Annual Energy Outlook 2011*. DIE.EUA-0383(2011), Washington DC.

US Environmental Protection Agency (EPA). 2008. Technical Support Document for the Advance Notice of Proposed Rulemaking for Greenhouse Gases: Stationary Sources. Washington DC.

Figures and Tables

Figure 1. Heat Input vs. Heat Rate (2008)

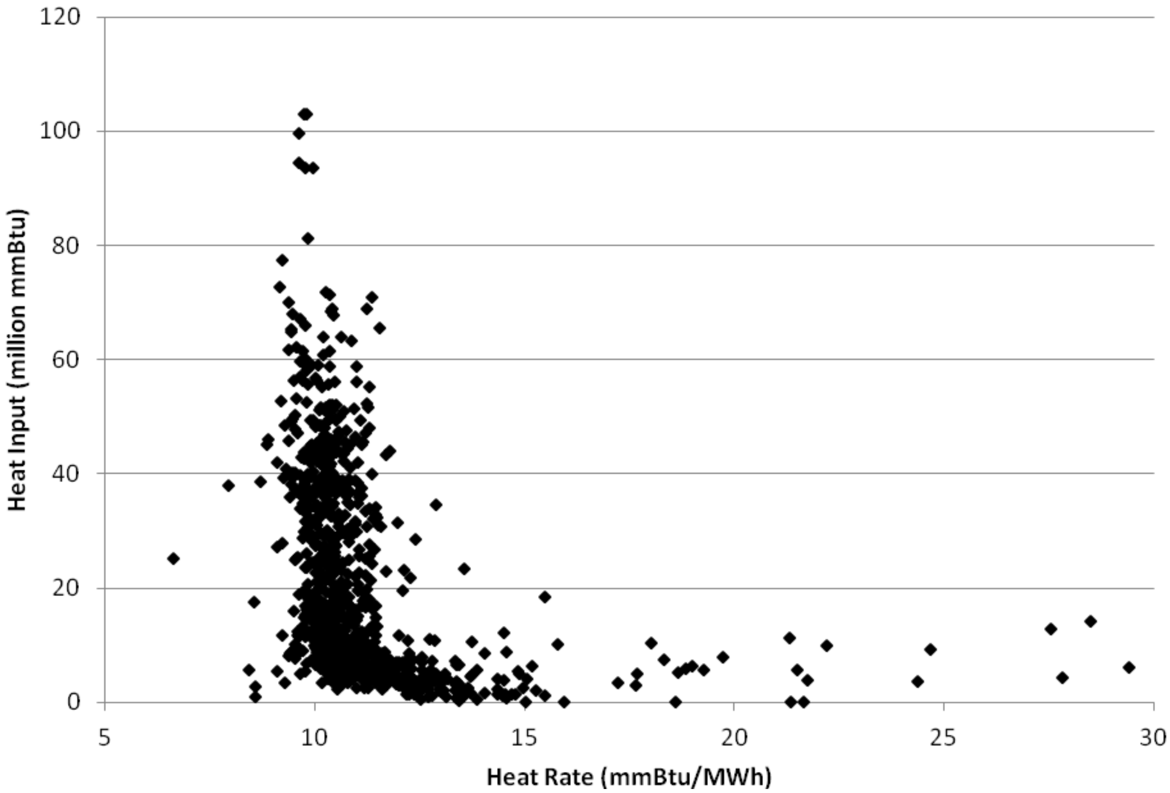

Notes: The figure plots the heat input (million megawatt hours, MWh) against the heat rate (million Btu, mmBtu, per MWh) for all units in the sample in 2008. The vertical line indicates the generation-weighted mean heat rate of 10.38 mmBtu/MWh.

Figure 2. Estimated Heat Rate Distribution by Firing Type

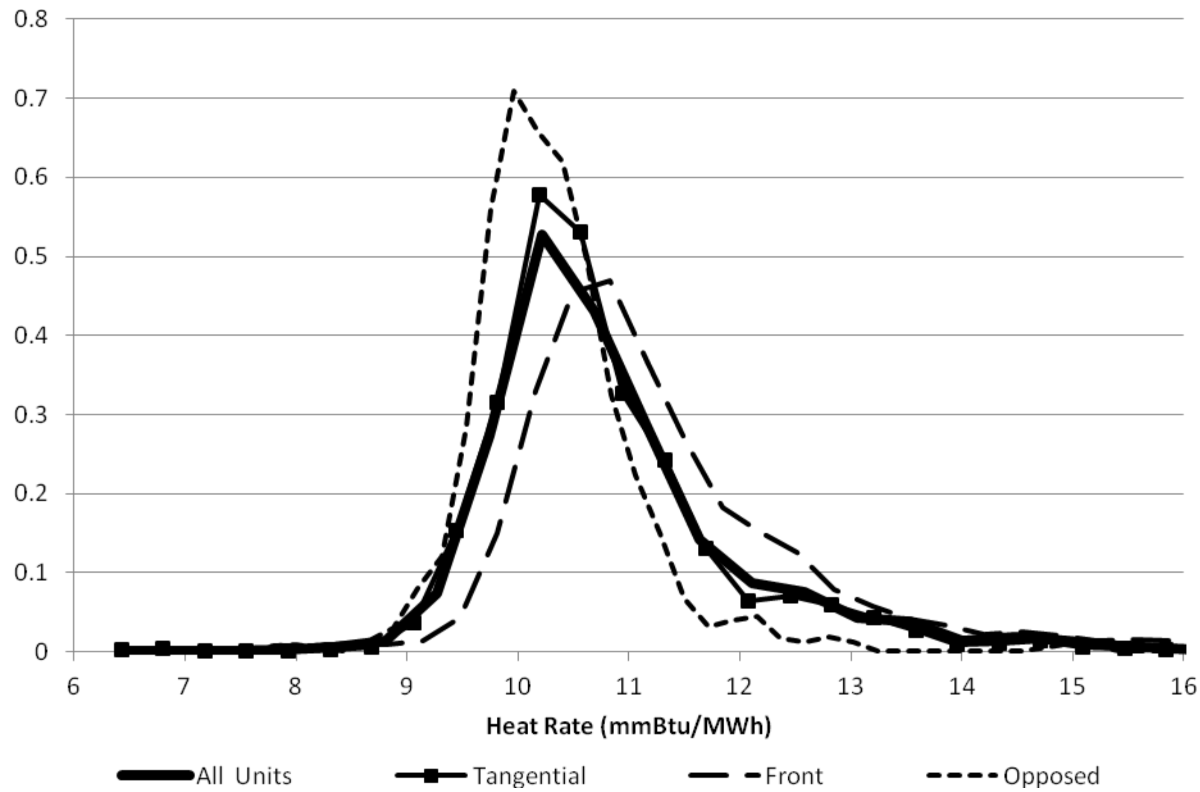

Notes: The figure plots the estimated heat rate kernel density function for all units and for the three most common firing types in the sample using observations from 2008.

Figure 3. Estimated Heat Rate Distribution Controlling for Technology Variables

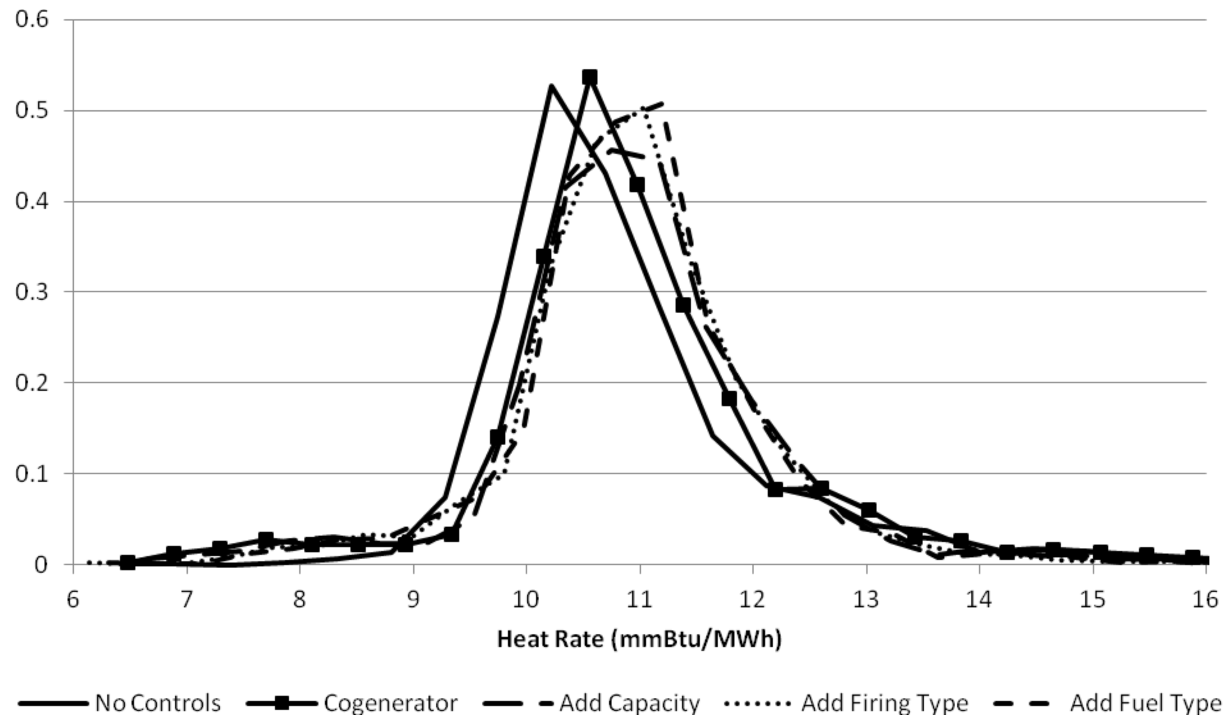

Notes: The figure plots the estimated kernel density function of heat rates for 2008. The first series plots the function using the heat rates observed in the data set. To construct the other plots, each unit's heat rate is regressed on the indicated control variables. The figure plots the residuals after adding the mean heat rate across units in the sample.

Figure 4. Utilization, Heat Rates, and Technology

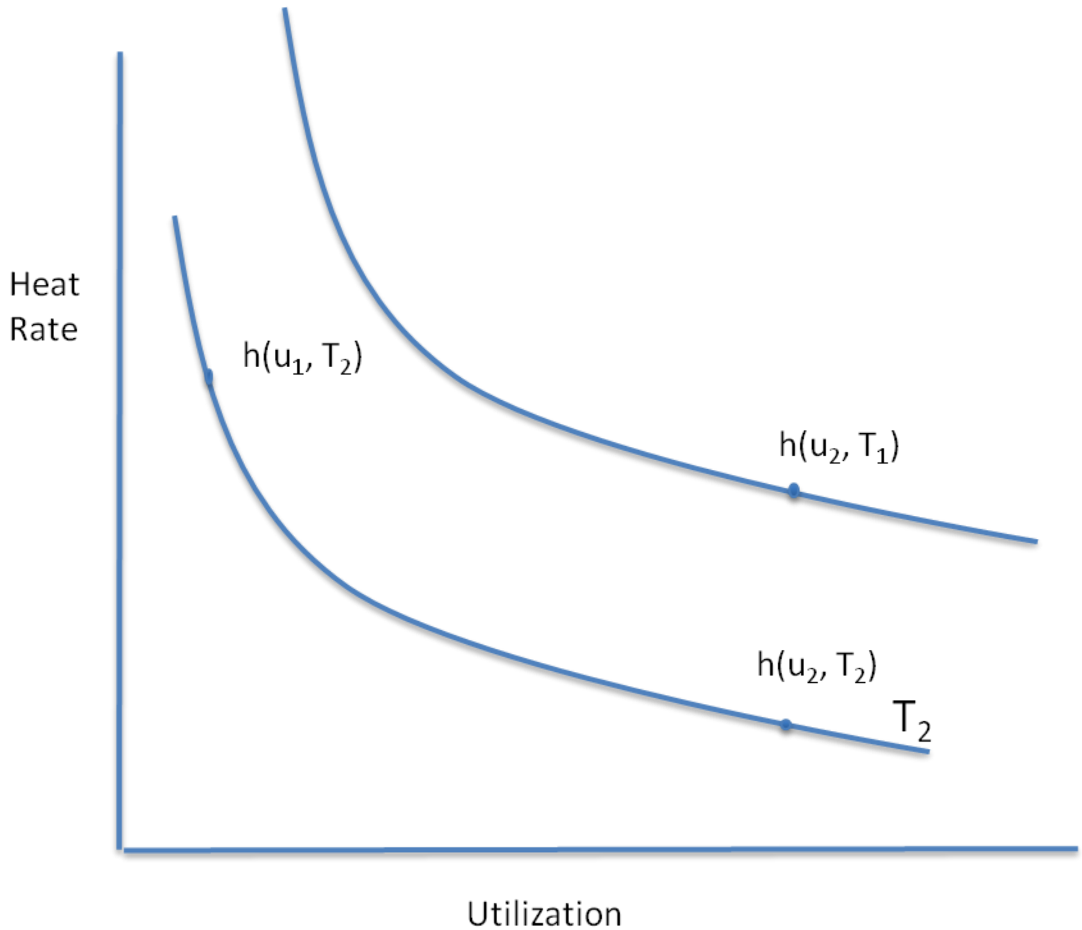

Notes: Heat rate, h, is a function of utilization, u, and heat rate technology, T. The figure plots two hypothetical technologies, T_1 and T_2. The figure shows that at the same level of utilization, u_2, the heat rate using technology T_2 is lower than the heat rate using technology T_1.

Figure 5. Heat Rate vs. Coal Price Residuals

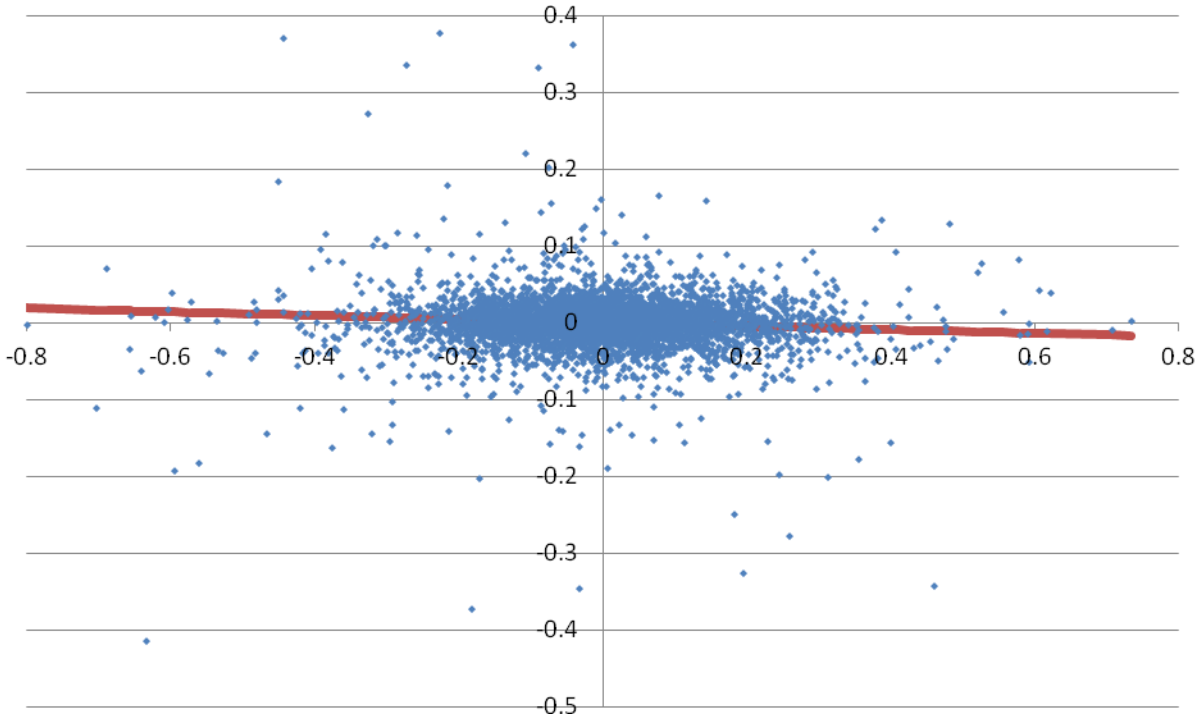

Notes: Heat rate and coal prices are obtained from a regression of log heat rate and log coal price on the control variables in column 2 of Panel B in Table 4, omitting the log coal price as an independent variable. The figure plots the heat rate residuals against the coal price residuals. The red line shows the fitted values from a regression of the heat rate residuals on the coal price residuals. The slope of the fitted values is the same as the estimated coefficient in Table 4.

Figure 6. Effect of Coal Prices on Heat Rates by Year

Panel A: Without Fixed Effects

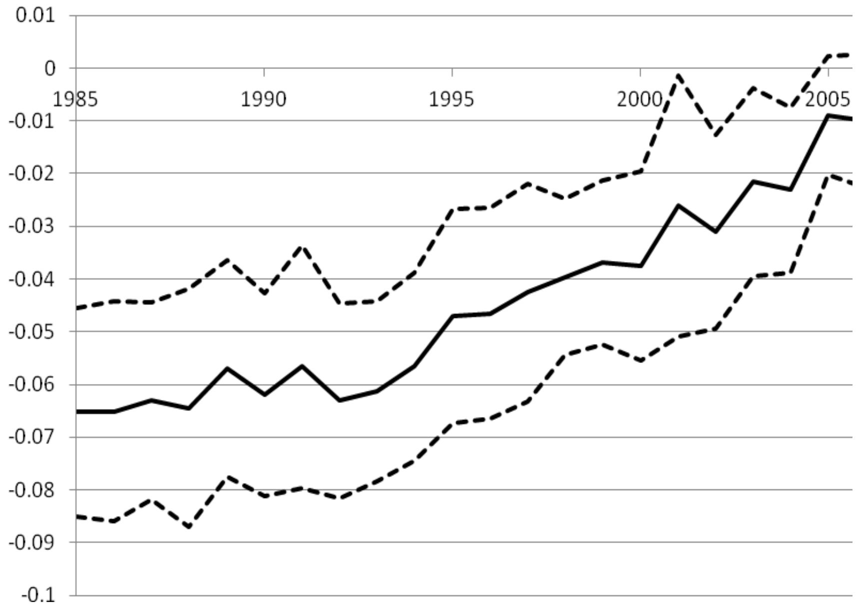

Panel B: With Fixed Effects

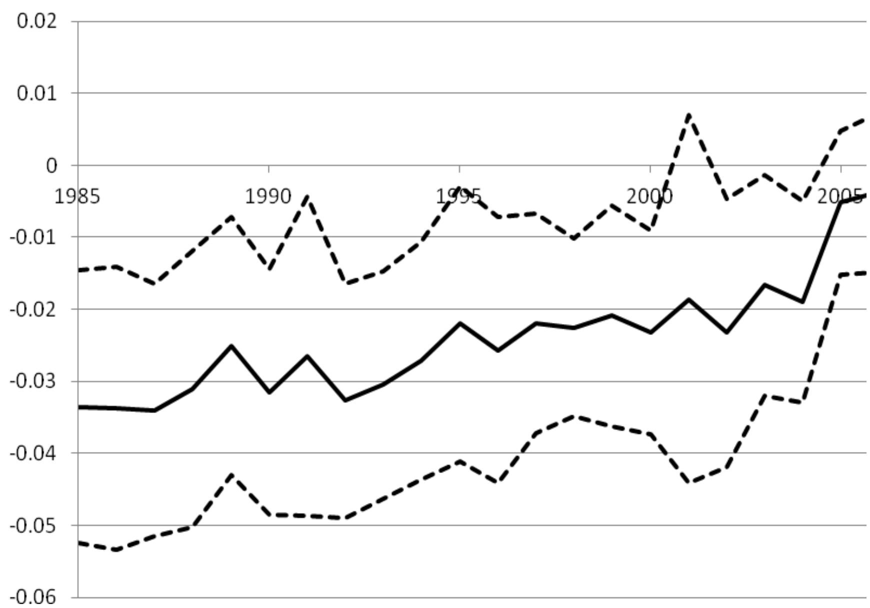

Notes: The log coal price is interacted in a set of year fixed effects, and the interactions are added to the specification in column 1 of Table 4. The figure plots the effect of coal prices on heat rates for each year as the sum of the coefficient on the coal price and the coefficient on the corresponding year-coal price interaction. The dashed curves show the 95 percent confidence intervals. In Panel A unit fixed effects are not included in the regression, and in Panel B, fixed effects are included. The vertical lines indicate the end of the first and second NSR regimes (see text).

Table 1. The Determinants of Heat Rates Associated with How Boilers Are Used

Determinant	*Hypotheses*
Cogeneration	Traditional measures of heat rate do not account for the thermal output produced by cogenerators and are misleading.
Utilization	Higher utilization rates are associated with lower heat rates because it is less costly to run a more efficient plant and there are efficiency losses associated with varying utilization and with starting up and shutting down the unit.
Pollution Controls	Pollution controls impose a heat rate penalty because of the heat input or electricity required to run them. Also, investment in pollution controls may crowd out other investments.
Fuel Choice	Low sulfur coal is associated with increased efficiency because it reduces the need for pollution controls. However, low sulfur (sub-bituminous) coal also has a higher rate of CO_2 emission per BTU than high sulfur (bituminous) coal. Coal type may also signal an economic decision based on location and availability.
Regulation and Incentives	Fear of triggering NSR for other pollutants may delay efficiency improvements. Fuel cost adjustment clauses allow plants to pass through fuel costs to customers, which may reduce incentives for making efficiency improvements. Competition in wholesale markets incentivizes efficiency improvements.
Ownership	Competition in wholesale power markets provides an incentive for efficiency improvements.

Table 2. Summary Statistics by Time Period

	Number of Observations	*Age (years)*	*Capacity (MW)*	*Utilization Rate*	*Scrubber*	*Heat Rate (mmBtu/MWh)*	*Coal Price ($/mmBtu)*
1985-1989	4858	22.2 (11.1)	317.5 (258.0)	0.49 (0.20)	0.13	10.9 (1.5)	1.57 (0.42)
1990-1994	4926	26.7 (11.4)	322.2 (261.8)	0.51 (0.21)	0.14	10.8 (1.5)	1.46 (0.39)
1995-1999	4911	31.5 (11.6)	324.8 (266.3)	0.57 (0.19)	0.17	10.8 (1.5)	1.29 (0.34)
2000-2004	4846	35.6 (12.1)	320.9 (265.1)	0.62 (0.17)	0.20	11.0 (1.8)	1.32 (0.36)
2005-2009	2893	39.4 (12.9)	332.8 (273.3)	0.60 (0.20)	0.29	11.1 (2.1)	2.12 (0.84)

Notes: The table preports summary statistics for all units in the sample using annual observations. The table reports the number of observations in each time period and the means of the variables indicated in the column headings; standard deviations are in parentheses. Age is reported in years and capacity in MW. Utilization rate is total generation divided by generation if the unit operates at full capacity throughout the time period. The scrubber variable is an indicator for whether the unit is connected to a scrubber. Heat rat is reported in mmBtu/MWh. Coal price is in dollars per mmBTU.

Table 3. Estimated Abatement Opportunities for Traditional Emissions Rate Standards

Uniform Standard	*Standard by Firing Type*	*Standard by Firing Type and Size*	*Standard by Firing Type, Size, and Scrubber*	*Standard by Firing Type, Size, Scrubber, and Cogenerator*
Panel A: Percent Abatement without Rebound Effect				
6.38	6.19	7.19	6.71	6.45
Panel B: Percent Abatement with Rebound Effect				
5.19	5.02	5.83	5.44	5.23

Notes: The table reports abatement opportunities under different hypothetical emissions rate standards. The calculations do not account for abatement costs or technical constraints, but capture only the heterogeneous operating performance of existing coal generators. The standards are defined as the 10th percentile heat rate for the indicated categories (i.e., 90 percent of units initially have a higher heat rate than the standard). Size categories are assigned based on the unit's quartile of generation capacity. Scrubber and cogenerator categories are assigned based on whether the unit has a scrubber or whether the unit is a cogenerator. Panel A assumes that the standards do not affect generation and Panel B assumes that a 10 percent heat rate decrease causes a 2 percent generation increase.

Table 4. Effect of Coal Prices on Heat Rates

Dependent Variable: Log Heat Rate				
	(1)	(2)	(3)	(4)
		Panel A: No Unit Fixed Effects		
Log Coal Price (α)	-0.043 (0.009)	-0.049 (0.008)	-0.056 (0.010)	-0.039 (0.006)
Number of Observations	5,132	5,132	2,523	22,434
R-Squared	0.75	0.68	0.75	0.56
		Panel B: Unit Fixed Effects		
Log Coal Price (α)	-0.011 (0.008)	-0.023 (0.008)	-0.020 (0.011)	-0.018 (0.006)
Number of Observations	5,132	5,132	2,523	22,434
R-Squared	0.95	0.90	0.91	0.75
Specification	Unweighted	Baseline, weighted by generation	Include 20 largest firms	Single-year time periods

Notes: The table reports estimates of equation (3), in which α is the coefficient on the log coal price. Standard errors, in parentheses, are clustered by state. The unit of observation is a generation unit by 5-year time period from 1985-2009 in columns 1-3; column 4 uses annual observations. The dependent variable is the log of heat input divided by generation over the time period. Log coal price is the log of the average price of coal for the corresponding plant and time period. Utilization is the total generation of the unit over the time period divided by the unit's capacity multiplied by the number of hours in the time period. All specifications include utilization; separate indicator variables for whether the unit is a cogenerator, whether the unit has SCR, and whether the unit has a scrubber; age group fixed effects; vintage decile fixed effects; period fixed effects; firing type fixed effects; fuel type fixed effects; state fixed effects; and ownersip type fixed effects. Deciles are constructed for the unit's rated capacity, and all specifications include a set of fixed effets for each capacity decile. All specifications include a set of fixed effects for the number of months in the time period that the unit operates at below 10 percent or below 30 percent of rated capacity. Panel B also includes unit fixed effects. Each column reports the indicated specification. Columns 2-4 weight observations by the unit's annual generation. Column 3 includes the 20 largest firms, where firms are ranked by total coal capacity in 2009. Column 4 uses annual observations rather than five-year time periods.

Table 5. Coal Price Persistence

	Dep var: log coal price	*Dep var: log contract price*	*Dep var: log spot price*
	(1)	(2)	(3)
		Panel A: Include State Fixed Effects	
Log Price	0.712 (0.024)	0.494 (0.151)	0.596 (0.084)
Number of Observations	3,975	3,805	3,595
R-Squared	0.83	0.37	0.43
		Panel B: Omit State Fixed Effects	
Log Price	0.898 (0.033)	0.738 (0.104)	0.730 (0.081)
Number of Observations	3,975	3,805	3,595
R-Squared	0.79	0.28	0.33

Notes: Standard errors, in parentheses, are clustered by state. The unit of observation is a generation unit by 5-year time period from 1985-2009. The dependent variable is the average price of coal in column 1, the contract price in column 2, and the spot price in column 3. Each regression includes the one-period lag of the corresponding price. All specifications include period fixed effects, and Panel A also includes state fixed effects.

Table 6. Persistence Results

	Dependent Variable: Log Heat Rate				
	(1)	(2)	(3)	(4)	(5)
	Panel A: No Unit Fixed Effects				
Log Coal Price		-0.064 (0.008)		-0.039 (0.006)	-0.049 (0.040)
Log Contract Price	0.001 (0.002)		-0.002 (0.002)		
Log Spot Price	0.000 (0.001)		-0.003 (0.002)		
Log Coal Price X Peristence					0.017 (0.052)
Number of Observations	4,732	5,132	4,732	4,890	4,527
R-Squared	0.69	0.59	0.61	0.69	0.70
	Panel B: Unit Fixed Effects				
Log Coal Price				-0.014 (0.005)	-0.107 (0.079)
Log Contract Price	0.000 (0.002)				
Log Spot Price	0.001 (0.001)				
Log Coal Price X Peristence					0.111 (0.102)
Number of Observations	4,732			4,890	4,527
R-Squared	0.91			0.90	0.90
Specification	Separate spot and contract prices	Drop state fixed effects	Separate spot and contract prices without state fixed effects	Use coal price in first year of time period	Interact coal price with NERC region price persistence

Notes: The table reports the same specifications as column 2 of Table 4, except as noted in the bottom row of the table. The first column replaces the log average coal price with the log spot price and log contract price. Column 2 repeats column 1 of Table 2 except that state fixed effects are omitted. Column 3 repeats column 1 omitting state fixed effects. Column 4 uses the coal price in the first year of the corresponding time period in place of the average price over the period. For the five NERC regions with the most coal units, the coal price persistence is estimated using the same specification as column 1 in Panel A of Table 5. Column 5 includes the interaction of the log coal price with the estimated persistence.

Table 7. Time-Varying Omitted Variables

	Dependent Variable: Log Heat Rate				
	(1)	(2)	(3)	(4)	(5)
			Panel A: No Unit Fixed Effects		
Log Coal Price	-0.113	-0.051	-0.052	-0.062	-0.038
	(0.017)	(0.008)	(0.008)	(0.009)	(0.008)
Number of Observations	5,132	3,556	4,580	4,989	4,113
R-Squared	0.69	0.67	0.69	0.84	0.69
			Panel B: Unit Fixed Effects		
Log Coal Price	-0.038	-0.025	-0.022	-0.032	-0.015
	(0.019)	(0.009)	(0.008)	(0.010)	(0.007)
Number of Observations	5,132	3,556	4,580	4,989	4,113
R-Squared	0.90	0.88	0.88	0.94	0.92
Specification	Instrument for coal price	Include high utilization units	Balanced panel	Include firm X period fixed effects	Add state X period controls

Notes: Standard errors, in parentheses, are clustered by state. See Table 2 for sample and variable construction. All specifications are the same as in column 2 of Table 4 except as noted. Column 1 instruments for the coal price using interactions of a set of year fixed effects with the plant's log coal price in 1984 and the plant's share of coal purchased from the west in 1984. Column 2 includes units whose median utilization across all periods exceeds 0.5. Column 4 includes units that operate in all time periods. Column 4 includes firm by period fixed effects. Column 5 includes total capacity, coal capacity, natural gas capacity, total generation, coal generation, natural gas generation, gross state product, employment, and population by state and period, where all variables are in logs.

Table 8. Effect of Coal Prices on Utilization Rates

Dependent Variable: Log Utilization Rate				
	(1)	(2)	(3)	(4)
		Panel A: No Unit Fixed Effects		
Log Coal Price (β)	-0.458 (0.066)	-0.423 (0.075)	-0.401 (0.065)	-0.271 (0.055)
Number of Observations	22,403	11,340	5,132	6,760
R-Squared	0.85	0.86	0.84	0.86
		Panel B: Unit Fixed Effects		
Log Coal Price (β)	-0.397 (0.042)	-0.427 (0.062)	-0.396 (0.074)	-0.184 (0.049)
Number of Observations	22,403	11,340	5,132	6,760
R-Squared	0.93	0.92	0.94	0.95
Specification	Baseline	Include 20 Largest Firms	5-year Intervals	Include 2001-2009

Notes: The table reports estimates of equation (4), in which β is the coefficient on the log coal price. Standard errors, in parentheses, are clustered by state. Specifications in columns 1 and 2 are identical to columns 1 and 3 in Table 4, except that the dependent variable is the log of the utilization rate, utilization controls are omitted, and observations are annual rather than aggregated to five-year time periods. Column 3 uses 5-year time intervals. Column 4 is the same as column 1 except that it includes only observations from 2001-2009.

Table 9. Utilization: Time-Varying Omitted Variables

	Dependent Variable: Log Utilization Rate					
	Panel A: No Unit Fixed Effects					
	(1)	(2)	(3)	(4)	(5)	(6)
Log Coal Price	-0.561 (0.117)	-0.260 (0.033)	-0.390 (0.043)	-0.445 (0.060)	-0.362 (0.072)	-0.459 (0.066)
Number of Observations	22,403	15,543	15,387	21,884	17,511	22,247
R-Squared	0.85	0.90	0.88	0.89	0.85	0.85
	Panel B: Unit Fixed Effects					
	(1)	(2)	(3)	(4)	(5)	
Log Coal Price	-0.499 (0.110)	-0.291 (0.040)	-0.388 (0.040)	-0.365 (0.047)	-0.255 (0.042)	-0.396 (0.042)
Number of Observations	22,403	15,543	15,387	21,884	17,511	22,247
R-Squared	0.93	0.93	0.93	0.95	0.94	0.93
Specification	Instrument for coal price	Include high utilization units	Balanced panel	Include firm X period fixed effects	Add state X period controls	Control for natural gas prices

Notes: Standard errors, in parentheses, are clustered by state. Specifications in columns 1-5 are identical to Table 7, except that the dependent variable is the log of the utilizaiton rate, utilization controls are omitted, and observations are annual rather than aggregated to five-year time periods. Column 6 adds to the baseline from Table 8 the log of the state's natural gas price.

Table 10. Effects of Policies That Achievean 0.88 Percent Emissions Reduction

	Traditional Standard	*Flexible Standard*	*Coal Btu Tax*	*Emissions Tax*
Percent Change in Heat Rates	-1.00	-1.01	-0.12	-0.12
Percent Change in Coal Generation	0.19	0.20	-1.18	-1.16
Percent Change in Coal Emissions	-0.82	-0.81	-1.29	-1.29
Coal Capacity Factor (percent)	60.16	60.17	59.34	59.35
Rebound in Emissions per Change in Emissions	0.13	0.18	0.02	0.02
Change in Investment Costs (million $)	3,777.5	417.6	38.5	41.4
Change Investment Costs/Change Emissions ($/ton)	-240.81	-26.62	-2.45	-2.64

Notes: The table reports results of simulations of the electricity market model. See Section 5 for a detailed description of the model. Each column reports the results of a separate policy scenario. The scenarios are calibrated to achieve the same total emissions reduction. Under the inflexible standard units must achieve a heat rate of 10.80 mmBtu/MWh or reduce heat rate by 10 percent, whichever results in a higher heat rate. The flexible standard sets a benchmark heat rate of 10.28 mmBtu/MWh. The coal Btu tax imposes a tax of $0.12 per mmBtu of coal. The emissions tax is $1.17 per ton of CO_2 emissions. Each row reports the change in the indicated variable as compared to the baseline (no policy) case. Coal capacity factor is the ratio of generation of the coal units to the generation if all units operated at full capacity. The rebound in emissions is the change in emissions resulting from the change in heat rates, holding fixed the effective coal price under each policy (see Appendix 3 for details on the rebound calculation). Investment costs include the annualized capital costs of heat rate improvements using a fixed charge rate of 0.1.

Appendix 1. Determinants of Heat Rates

This section provides a more detailed review of the literature on heat rates than is found in the main text.

Cogeneration

Cogenerators, or combined heat and power generators, are facilities that recycle heat to produce both electricity and useful thermal energy. The additional energy captured is typically used for manufacturing processes or central heating. Cogeneration is regarded as a highly efficient technology for many applications, yet traditional heat rate calculations are misleading because heat input per unit of generation does not account for the useful thermal output. Furthermore, the manager of a cogeneration facility can adjust the proportion of heat versus electricity based on the customers' needs, heating fuel and electricity prices, and regulation.

Utilization

As noted, more efficient plants typically are used more heavily. Moreover, higher utilization tends to reduce heat rates. Conversely, less efficient plants are available to be ramped up and down more frequently, which requires additional fuel input because the temperature in the boiler fluctuates, and which further amplifies their relatively high heat rates.

Pollution Controls

Because fossil fuel–fired electric power plants emit a variety of pollutants, they are subject to numerous environmental regulations at the local, state, and federal levels. Regulations exist for sulfur dioxide, nitrogen oxides, particulate matter, mercury, and acid gases. Reducing emissions to comply with these regulations can affect efficiency. Fleishman et al. (2009) found that implementing pollution controls may "crowd out" investment in productivity improvements if firms face capital constraints. On the other hand, several papers have studied the trade-offs between pollution controls and productivity of a plant in the power sector as well as other manufacturing industries that face similar environmental regulations. These studies conclude that inputs and pollution can be reduced without sacrificing productivity (!!! INVALID CITATION !!!).

Fuel Choice

Coal-fired power plants generally use one or a combination of three types of coal: bituminous, subbituminous, and lignite. The choice depends on relative prices and characteristics. Coal types vary in heat, ash, and sulfur per pound and are priced accordingly; higher heat content justifies a higher price. Because of regulations on emissions of sulfur dioxide, low-sulfur coal is also priced higher. Efficiency may also increase with low-sulfur coal as it negates the need for energy-intensive postcombustion pollution control (scrubbers). Transportation costs also have a large effect on delivered coal prices. Because certain coal types are more common to specific regions of the country, location may influence a plant's choice of fuel.

Regulation and Incentives

Regulation may also affect efficiency. Coal-fired power plants have demonstrated useful lives that are much longer than many anticipated when they were originally constructed. Ellerman (1998) notes that these life extensions are due to new electronics in the boilers and other features that improve plant efficiency and longevity. However, some measures that might improve efficiency and reduce CO_2 emissions may be considered a major modification to the plant and thereby trigger New Source Review (NSR) for other pollutants under the Clean Air Act. Permitting under NSR requires a site-specific and technology-based review of the control technology proposed by the source and a demonstration that the plant will not create or exacerbate violations of air quality standards in the surrounding area (Richardson et al. 2011). Consequently, NSR can raise the cost of efficiency improvements by requiring the installation of pollution abatement equipment.

At least three papers have studied the effect of NSR on plant operations. Bushnell and Wolfram (2010) and Stavins (2006) found that NSR causes older plants—and therefore potentially less efficient plants—to operate more. Also, anecdotal evidence suggests that plant owners have deferred cost-effective improvements that might initiate a permitting process (Keohane et al. 2007).

Other regulatory practices may insulate plant operators from the cost of continuing the inefficient operation of existing plants. For example, state-level fuel cost adjustment clauses allow firms to pass fuel costs through to retail electricity ratepayers (DOE/NETL 2010). These provisions are common in regions of the country that operate under cost-of-service regulation, including regulated investor-owned utilities and publicly owned utilities. Such provisions are intended to remove from shareholders the risk from fluctuations in fuel prices. Depending on how they are structured, fuel cost pass-through provisions could reduce the incentive to make efficiency improvements if cost savings would be offset by lower retail prices.

Knittel (2002) examined various efficiency incentive programs run by electric utilities. Programs for generator efficiency modifications improve efficiency more than other types of programs. Automatic fuel cost pass-through, in which an increase in fuel cost passes directly to the consumer without a rate hearing, reduces plant efficiency.

Since 1978, the electricity sector has been restructured gradually so that electricity generators increasingly operate in competitive environments. Many studies have focused on the effects of electricity restructuring in the United States. Much of the literature has examined the effect of restructuring on electricity prices, market performance, or market power (Joskow 1997; Borenstein et al. 2002; Bushnell et al. 2008), but several have considered the relationship between regulation and technical efficiency. Joskow (1974) and Hendricks (1975) showed that regulatory structures can provide significant incentives for public utilities to reduce costs. As noted previously, Fabrizio et al. (2007) found that restructuring reduces employment and fuel consumption at firms throughout the affected wholesale power markets, including those that were not restructured.

Ownership

The firm's ownership structure may affect efficiency improvements. Ownership types include private utilities (investor owned or privately owned), public utilities (municipal, state, or federal), and cooperatives. Because of different objective functions, degrees of principal-agent conflicts, and exposure to market forces or other factors, different ownership types may not place the same emphasis on improving efficiency (all else equal).

Ownership structure is also closely related to the regulatory environment. Following deregulation in many states, privately owned utilities were required or encouraged to divest generation assets and act solely as transmitters and distributers of electricity. Publicly owned utilities were largely unaffected by this trend. States that remained under regulation tended to remain dominated by vertically integrated utilities, although nonutilities may still participate in the wholesale market. Bushnell and Wolfram (2005) report greater efficiency at units that were divested from utility to nonutility; however, similar improvements were observed at nondivested units that were subject to other efficiency incentives. Thus, restructuring or changing regulatory incentives can encourage efficiency investments. On the other hand, Fischlein et al. (2009) find large differences in the efficiency of investor-owned utilities as opposed to municipalities, rural cooperatives, or district power providers. A possible explanation is that publicly owned companies might be insulated from market-driven incentives to improve efficiency, but this is widely disputed, and publicly owned companies have broader performance criteria than do privately owned firms that may affect efficiency measures.

Appendix 2. Data

Data Sources

The main data source, which defines our universe of units, is the Energy Information Administration's form 767 (EIA-767) and successor forms. This government-mandated survey collects boiler- and generator-level information from fossil fuel–fired electric power plants with nameplate capacity greater than 10 megawatts (MW). All units that fire any coal during each year were included in the panel data set.

The EIA-767 provides most of the variables of interest, including monthly heat input, monthly generation, boiler vintage, boiler technology, nameplate capacity, location, and fuel type. Other data sources provide additional variables or fill in missing observations. EIA forms 860 and 861 provide information on ownership and generator characteristics. EIA form 423 and 923 provide data on fuel expenditure and quantity by fuel type, plant and month, from which we construct average delivered coal prices by plant and year. The publicly available fuel data include only regulated plants. We supplement the public data with data obtained from EIA covering unregulated plants for the years 2002-2009. EPA's National Electric Energy Data System (NEEDS) includes variables for the existence and vintage of environmental controls at individual units.

Issues with the Data

Coverage

The panel data set covers the years 1985 through 2009. However, the sample does not include the years 2006 and 2007 because EIA discontinued form 767 after 2005 and did not collect unit-specific data on the successor forms until 2008.

Another issue that affects coverage is the change in reporting requirements for EIA-767 over time. Prior to 2001, only regulated plants with nameplate capacity greater than 10 MW were required to report. Between 2001 and 2003, the survey expanded to include unregulated plants, but they reported generation only if nameplate capacity exceeded 100 MW. After 2003, all plants greater than 10 MW—both regulated and unregulated—reported all variables.

Heat Rate Adjustments

Because the annual heat rates are calculated using sums of monthly data, measured annual heat rates can be highly influenced by missing data or unreasonably high or low heat rates in a single month. These "outlier" heat rates are not possible according to engineering estimates and usually arise because of a problem in the data, such as negative generation, missing heat input or generation, or reporting error.[31] Often, these outlier monthly heat rates contributed to a misleading figure for the annual heat rate.

To address this issue and avoid dropping full years of data, we adjust the annual heat rates by dropping months that contain negative generation, missing data, or outlier heat rates. We define an outlier monthly heat rate as being more than two standard deviations above or below the monthly mean across all years for the unit. This ensures that only heat rates outside normal engineering limits were dropped. The annual heat rates are calculated using the remaining months.

Boiler-to-Generator Correspondence

The last constraint on the data stems from the boiler-to-generator correspondence at individual plants. Heat input is measured at the boiler level, whereas generation is measured at the generator level. In most cases, a single boiler is connected to a single generator, and calculating heat rate is straightforward. When many boilers connect to a single generator, we aggregate heat input across boilers. Other characteristics of that group of boilers are aggregated or averaged. Similarly, when many generators are linked with a single boiler, we aggregate generation to a single generator. This technique allows us to calculate heat rates for the units in our sample that do not have a one-to-one boiler-to-generator correspondence. However, the trade-off is that for some units, the vintage, size, technology, and fuel represent an average or dominant characteristic of the group rather than an actual unit. In general, the configurations with

[31] Negative generation occurs when the plant generates electricity for use at the plant but does not supply any electricity to end users.

multiple boilers or generators account for less than 4 percent of the observations.[32] Omitting these units in the estimation does not affect the results.

Appendix 3. Measuring the Rebound Effect

We characterize the rebound effect as the change in emissions due to the change in utilization in response solely to changes in heat rate improvements at the generating units. To estimate the change in emissions in the absence of the rebound effect, we evaluate the new equilibrium with utilization determined by the baseline heat rates. Against this estimate, we measure the rebound effect as the further change that occurs when utilization is based on the improved heat rates.

To simplify notation, let $j,k=\{0,1\}$ denote the baseline indicated by zero and the new equilibrium indicated by 1. We suppress notation for the individual unit type and describe utilization as a function of only the price of fuel and the heat rate: $u_{jk}=u\left(p_j^f h_k\right)$. We describe emissions from coal-fired units as a function of utilization and the heat rate: $e_c\left(u_{jk},h_k\right)$. For example, we represent emissions at the new effective fuel prices and heat rates but based on utilization that would occur at the new prices and the old heat rate as $e_c\left(u_{10},h_1\right)$. Utilization of natural gas units adjusts to the change in utilization of coal units to satisfy residual demand, and emissions from natural gas units is represented by $e_g\left(u_{jk}\right)$. With this notation, we measure the rebound effect as the difference between emissions that would occur if utilization depended on the baseline heat rate and the actual emissions, normalized by the former measure:

$$
\begin{aligned}
R &= \frac{\left\{e_c\left(u_{10},h_1\right)+e_g\left(u_{10}\right)-e_c\left(u_{00},h_0\right)-e_g\left(u_{00}\right)\right\}-\left\{e_c\left(u_{11},h_1\right)+e_g\left(u_{11}\right)-e_c\left(u_{00},h_0\right)-e_g\left(u_{00}\right)\right\}}{\left\{e_c\left(u_{10},h_1\right)+e_g\left(u_{10}\right)-e_c\left(u_{00},h_0\right)-e_g\left(u_{00}\right)\right\}} \\
&= \frac{e_c\left(u_{10},h_1\right)+e_g\left(u_{10}\right)-e_c\left(u_{11},h_1\right)-e_g\left(u_{11}\right)}{e_c\left(u_{10},h_1\right)+e_g\left(u_{10}\right)-e_c\left(u_{00},h_0\right)-e_g\left(u_{00}\right)}
\end{aligned}
\tag{11}
$$

[32] The percentage is higher after 2004 because of the different reporting requirements across years. Prior to 2004, smaller or unregulated units did not report generation and these units are less likely to have a one-to-one correspondence.

Appendix References

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

Boyd, G. A., and J. D. McClelland. 1999. The Impact of Environmental Constraints on Productivity Improvement in Integrated Paper Plants. *Journal of Environmental Economics and Management* 38: 121–142.

Bushnell, J. B., and C. Wolfram. 2005. Ownership Change, Incentives and Plant Efficiency: The Divestiture of US Electric Generation Plants. Center for the Study of Energy Markets, University of California Energy Institute, UC Berkeley, CSEM WP 140.

Bushnell, J., and C. Wolfram, 2010. Enforcement of Vintage Differentiated Regulations: The Case of New Source Review. Iowa State University Department of Economics Working Paper No. 10027.

Bushnell, J. B., E. Mansur, and C. Saravia. 2008. Vertical Arrangements, Market Structure, and Competition: An Analysis of Restructured US Electricity Markets. *American Economic Review* 98(1): 237–266.

Ellerman, A. D. 1998. Note on the Seemingly Indefinite Extension of Power Plant Lives, a Panel Contribution. *Energy Journal* 19(2): 129–132.

Fabrizio, K. R., N. L. Rose, and C. D. Wolfram. 2007. Do Markets Reduce Costs? Assessing the Impact of Regulatory Restructuring on US Electric Generation Efficiency. *American Economic Review* 97(4): 1250–1277.

Fischlein, M., T. M. Smith, and E. J. Wilson. 2009. Carbon Emissions and Management Scenarios for Consumer-Owned Utilities. *Environmental Science and Policy* 12: 778–790.

Fleishman, R., R. Alexander, S. Bretshneider, and D. Popp. 2009. Does Regulation Stimulate Productivity: The Effect of Air Quality Policies on the Efficiency of US Power Plants. *Energy Policy* 37(11): 4574–4582.

Joskow, P. L. 1974. Inflation and Environmental Concern: Structural Change in the Process of Public Utility Price Regulation. *Journal of Law and Economics* 17(2): 291–327.

Joskow, P. L. 1997. Restructuring, Competition and Regulatory Reform in the US Electricity Sector. *Journal of Economic Perspectives* 11(3): 119–138.

Keohane, N., E. T. Mansur, and A. Voynov. 2007. Averting Regulatory Enforcement: Evidence from New Source Review. *NBER Working Paper Series.* Working Paper 13512, Cambridge, MA, National Bureau of Economic Research.

Knittel, C. R. 2002. Alternative Regulatory Methods and Firm Efficiency: Stochastic Frontier Evidence from the US Electricity Industry. *Review of Economics and Statistics* 84: 530–540.

Richardson, N., A. Fraas, and D. Burtraw. 2011. Greenhouse Gas Regulation under the Clean Air Act: Structure, Effects, and Implications of a Knowable Pathway. *Environmental Law Reporter* 41: 10098–10120.

Rosen , M. A., and I. Dincer. 2003. Survey of Thermodynamic Methods to Improve the Efficiency of Coal-Fired Electricity Generation. *Proceedings of the Institution of Mechanical Engineers, Part A: Journal of Power and Energy* 217: 63–73.

Shadbegian, R. J., and W. B. Gray. 2006. Assessing Multidimensional Performance: Environmental and Economic Outcomes. *Journal of Productivity Analysis* 26: 213–234.

Stavins, R. N. 2006. Vintage-Differentiated Environmental Regulation. *Stanford Environmental Law Journal* 25(1): 29–63.