



Contents lists available at ScienceDirect

Journal of Environmental Economics and Management

journal homepage: www.elsevier.com/locate/jeeem



Plant vintages, grandfathering, and environmental policy

Garth Heutel*

Department of Economics, University of North Carolina Greensboro, PO Box 26165, Greensboro, NC 27402, USA

ARTICLE INFO

Article history:

Received 22 October 2009

Available online 8 August 2010

Keywords:

Clean Air Act

Sulfur dioxide

Method of simulated moments

Vintage-differentiated regulation

Electric power plants

ABSTRACT

Environmental regulations that grandfather existing plants, by not holding them to the same strict standards as new plants, may have the unintended consequence of retarding new investment. If new plants are cleaner, then this effect may increase pollution in the short run. I develop a dynamic model of a facility's decisions over scrapping and abatement, which depend on capital depreciation, profitability shocks, and environmental policy. Using data from fossil fuel fired boilers at electric power plants, I estimate the structural parameters of the model and assess the impact of grandfathering in the Clean Air Act on sulfur dioxide emissions. Counterfactual policy simulations show that an increase in the stringency of performance standards would have led to a decrease in investment in new boilers. However, this does not lead to increased emissions, since there is less investment in dirtier coal boilers as compared to relatively cleaner oil or natural gas boilers.

© 2010 Elsevier Inc. All rights reserved.

0. Introduction

Regulations often contain grandfathering provisions, where facilities already built or workers already employed at the time of passage are not subject to the new standard. While the reasoning for such provisions may relate to fairness, such as a wish not to “change the rules in the middle of the game,” they often come with unintended consequences. By giving different incentives to grandfathered agents and non-grandfathered agents, the regulations can lead to unexpected outcomes. The federal Clean Air Act (CAA) and its New Source Performance Standards (NSPS) for major sources of air pollution are examples. By mandating that any new pollution sources (including power plants) meet strict pollution control standards, the rule may keep older facilities from scrapping and being replaced by newer ones. If older facilities are dirtier than newer ones, then this may create a “perverse” effect that increases pollution.

The purpose of this paper is to determine how grandfathering provisions in environmental policy affect both investment in and pollution from fossil fuel fired electric power plants. I develop a dynamic model of each facility's decisions about investment in new capital and choice of abatement techniques. These decisions are affected by the relative profitability of new capital, the costs of upgrading, and environmental regulations. Newer capital pollutes less, and hence stricter environmental policy without grandfathering provides an extra incentive to upgrade. Yet stricter environmental policy with grandfathering may provide a disincentive to upgrade. Using data from 1985 to 1995 on U.S. electric power plants, I estimate the parameters of the model. Finally, I simulate the estimated model to determine how grandfathering in the CAA impacts emissions and investment.

An early examination of a grandfathered environmental policy is [1], who looks not at stationary air pollution sources but at automobiles. He finds that stricter vehicle emissions standards, which apply only to new cars and hence effectively

* Fax: +1 336 334 5580.

E-mail address: gaheutel@uncg.edu

grandfather old cars, lead to a perverse short-term increase in emissions. [2] find that environmental regulations increase the age of capital but not the level of emissions for electric utilities, while [3] finds no significant difference in capital vintage between states with and without grandfather provisions for manufacturing plants in commercial printing and paint manufacturing. [4,5] also find perverse effects of grandfathering in the electric power industry and in manufacturing plants in New York state, respectively. Finally, [6] find that grandfathering in the CAA decreases the capital expenditures of coal-fired power plants but has no effect on their operating costs, fuel efficiency, or emissions.

In this paper's model, firms face a discrete choice of whether or not to scrap and upgrade their capital. This model contains two important extensions to discrete choice scrapping models to accommodate this industry and this policy. First, since the regulations are grandfathered, the model contains an additional state variable representing a facility's grandfathered status. Investment affects this status (if you scrap your plant, you lose its grandfathered status), so the policy impacts dynamic investment decisions. Second, instead of just considering a single binary decision (scrap or do not scrap), the model includes decisions over pollution abatement and plant type. These are important avenues for plants to respond to environmental policy, and without including them as options, plants are constrained in how they can react to policy changes.

This paper also extends the literature that studies the effects of grandfathering provisions in the CAA, including NSPS and New Source Review (NSR), by developing the first model that is structural and dynamic. A structural model allows for simulation of countless counterfactual policy scenarios. A dynamic model makes explicit the dynamic programming problem that utilities are solving, and it captures the effect that the cross-sectional distribution of plant vintages and grandfathering statuses has on the impact of the policy. Although this paper uses the model to retrospectively analyze the CAA's effect on the electricity industry through the mid-1990s, the model can be adapted to study and make forecasts about other industries and other policy environments in which grandfathering plays a role.

Grandfathering, or vintage-differentiated regulation, appears frequently outside of the CAA, both in environmental and non-environmental laws. Corporate average fuel economy (CAFE) standards and emissions rate standards apply only to new cars, so old cars are effectively grandfathered. The Clean Water Act and the Safe Drinking Water Act both set differential standards for water treatment plants based on when they went into operation. Furthermore, grandfathering is by no means absent from current policy debate.¹ Fire sprinklers are required in new buildings, but existing buildings are often not required to have them unless they are renovated. Zoning ordinances generally do not apply to businesses or homes built before the ordinance went into effect. Given the prevalence of grandfathering, it is important to study its effects, including any potentially perverse or counterproductive effects.

I find a significant effect of grandfathering in environmental regulations on both emissions and investment. Using the model to simulate the CAA and counterfactuals, I find that if grandfathering provisions were eliminated in 1985, emissions from power plants would be 60% less than the baseline level in 1995. This is because a large majority of plants in the cross section are grandfathered, and eliminating grandfathering amounts to a large policy shift. If the stringency of the grandfathered CAA standards was weakened in 1985, investment in new boilers would rise and yet emissions would also rise. The rise in new investment occurs for the same reason that strengthening a grandfathered standard can reduce new investment. When the grandfathered regulations are weakened, grandfathered plants lose some value in their grandfathered status, and hence their disincentive to invest falls. However, emissions do not fall, because the increased new investment occurs in relatively dirty coal boilers as opposed to relatively clean oil or natural gas boilers. Thus, the structural model that includes choice of boiler type helps explain results from previous studies, including [2,6], finding that grandfathered regulations inhibit new plant investment but have no perverse effect on emissions.

1. Model

I develop a model to estimate the impact of grandfathering in the Clean Air Act of 1970 on the electric power generating industry.² The CAA is an appropriate policy to consider because of its explicit grandfathering of existing sources. Section 111 of the law gave EPA the power to set binding emissions standards on all new sources of emissions—the New Source Performance Standards. Regulation of existing plants was left up to the states and is likely to be less strict. The plants were grandfathered for reasons of efficiency (it was costlier to retrofit existing plants than new ones), equity (it was unfair to “change the rules of the game mid-stream” by regulating existing plants), and politics (potential facilities has less clout than existing ones).³

I consider the behavior of a profit-maximizing single-plant firm. This model could be generalized to multi-plant firms, but the assumption that each individual plant's maximization decision is independent leads to identical results. The basis

¹ The US energy bill passed in 2007 contains an ethanol mandate with a requirement that the ethanol produced over its lifetime achieves at least a 20 percent reduction in greenhouse gases compared to gasoline. However, any biorefinery built or under construction before the bill's signing is exempt from that requirement; this could include up to 13.7 billion gallons of production capacity out of the 15 billion gallons of corn ethanol envisioned by the bill [21].

² Technically, the CAA was passed in 1963 and amended in 1970, but the 1970 amendments are often referred to as the “Clean Air Act of 1970” because they contained the bulk of the regulations.

³ Shavell [26] discusses when it might be socially optimal to grandfather, and Stavins [27] discusses the effects of grandfathering of environmental policy in many contexts.

Table 1

Summary of variables in dynamic model.

State variables	Choice variables	Transition process
v —Boiler age, integer	z —Replacement decision, discrete ($z=0$ if no replacement, $z=1$ if replaced with a coal boiler, $z=2$ if replaced with non-coal boiler)	$v' = v + 1$ if $z=0$ $v' = 1$ if $z > 0$
b —Boiler type, binary ($b=1$ for coal, $b=0$ for non-coal)	x —Build scrubber decision, binary	$b' = b$ if $z=0$ $b' = 1$ if $z=1$ $b' = 0$ if $z=2$
s —Scrubber present, binary	c —Use clean coal decision, binary	$s' = 1$ if ($s=1$ and $z=0$) or ($x=1$) $s' = 0$ otherwise
u —Grandfathered status, binary ($u=0$ if grandfathered, $u=1$ if not grandfathered)	m —Intensity of use decision, continuous	$u' = u$ if $z=0$ $u' = 1$ if $z > 0$
p —Price differential between high- and low-sulfur coal, continuous		p' determined exogenously
A —productivity shock, continuous		A' determined exogenously
B —emissions shock, continuous		B' determined exogenously

of this model is the discrete choice scrapping model [7], extended in two important ways. First, in addition to the choice of scrapping, plants can choose abatement investment to reduce emissions and can choose among different types of capital. Second, plants can be grandfathered from environmental policy.

Five nonrandom and two random state variables describe the plant at each period. First, the plant's age is v . Second, a plant can be of two types: coal ($b=1$) or non-coal ($b=0$).⁴ Third, a coal plant can have a flue gas desulfurization unit, or "scrubber," to reduce its emissions of sulfur dioxide (SO_2). If so, then $s=1$, otherwise $s=0$. Fourth, a plant can be grandfathered from environmental policy. Let $u=0$ if the plant is grandfathered and exempt from the policy, and $u=1$ if not. Fifth, a plant faces a price differential, p , between the prices it faces for using low-sulfur coal and high-sulfur coal.⁵ Finally, two random state variables, A and B , represent a productivity shock and a shock to emissions, respectively. The state variables are summarized in the first column of Table 1.

Each plant in each period faces decisions over four variables. First, it makes a discrete decision whether or not to update its production technology, that is, whether or not to scrap and replace with a new plant of age 1.⁶ If updating, it can be replaced with a coal or a non-coal plant. Let the decision variable $z=0$ if no replacement, $=1$ if replaced with a coal plant, and $=2$ if replaced with a non-coal plant. The second decision faced by plants is whether to add a scrubber. This decision is possible only for coal plants. Let $x=1$ if the plant chooses to add a scrubber and $x=0$ otherwise. Third, also available for coal plants only, is the opportunity to use low-sulfur "clean" coal as a means of reducing SO_2 emissions. Let $c=1$ if the plant chooses to use clean coal and 0 otherwise. Fourth, the plant faces a continuous decision over the intensity of use. This is measured by the amount of fuel input, m , used by the plant in a period. The second column of Table 1 summarizes the four choice variables.

A plant's single-period profit function is given by

$$Af(v, m, b) - F_c 1(z=1) - F_{nc} 1(z=2) - xG - pcm - \tau ue/m.$$

The first term, $Af(v, m, b)$, represents the reduced-form revenue of a plant of age v , operating at intensity m , and of type b . Here A is a multiplicative productivity shock. By reduced-form revenue I mean that the utilities are implicitly optimizing given supply and demand curves and whatever regulations they face, including rate of return regulations on electricity prices. Given the conditions under which these utilities operate, some function exists that they are trying to optimize, and it varies by capital age, type, and intensity of use. The next two terms capture fixed adjustment costs, F_c and F_{nc} , that are faced only if a plant is replaced with a coal plant ($z=1$) or a non-coal plant ($z=2$), respectively. G is the fixed cost of adding a scrubber, paid only when a scrubber is added ($x=1$). The per-unit cost differential between clean and dirty coal is p , so that

⁴ The data set includes fossil fuel fired steam-electric power plants powered by coal, oil, and natural gas. I lump together oil and natural gas because their emissions of sulfur dioxide are negligible compared to coal plants; coal accounts for 97% of SO_2 emissions from this group.

⁵ [29] finds evidence of a discrete jump in coal prices right at the CAA-imposed standard of 1.2 lb/MMBtu.

⁶ This choice is assumed to be binary rather than continuous. Doms and Dunne [20] and Nilsen and Schiantarelli [24] find that investment is mostly "lumpy." Entry and exit are not modeled. The vast majority of the electric utilities are present in every year, indicating that the utility-level industry makeup is roughly constant and entry and exit are not prevalent in this industry. Capacity growth, on the other hand, is prevalent, and it can be accommodated in the model in the same way as can technological growth. See the discussion later in this section.

when clean coal is used ($c=1$), p is paid per unit of coal m .⁷ This price differential p is allowed to vary by plant and over time.

The last term in the expression for profit, $-\tau ue/m$, represents the implicit costs of environmental policy. A plant's emissions is given by e . Environmental policy is modeled by a virtual tax rate τ levied on emissions intensity e/m . The tax payment $\tau e/m$ is paid only if the plant is not grandfathered, that is, $u=1$. If the plant is grandfathered, then $u=0$ and no environmental tax is paid.

The NSPS are performance standards, not an emissions tax, though here they are modeled as a tax. These standards create an implicit price on emissions, and utilities respond to the NSPS the same as they would respond to a tax ([8–10]). Importantly, the tax is levied not on the quantity of emissions e but on the emissions intensity, e/m . The NSPS have always been targeted towards emissions intensity rather than absolute emissions. For example, the NSPS for new coal plants set in the 1970 CAA was 1.2 lbs of SO₂ per MMBtu of heat input. By modeling the NSPS as a tax on emissions intensity, plants cannot limit the impact of this tax by reducing operating intensity only while maintaining the same ratio of emissions to fuel input. Thus, it more closely matches the regulatory environment these firms faced; no evidence exists suggesting that plants met their NSPS through reduced operating intensity.

Different specifications of the policy have different implications for firm behavior. Many observers have claimed that regulations rewritten in 1978 effectively required the installation of scrubbers on all new coal plants. In a sensitivity analysis, I consider policy modeled as a tax on the level of emissions rather than emissions intensity.

Emissions, like revenue, are a function of plant age, type, and intensity of use, but also of abatement activity: $e=Bg(v,c,s,m,b)$, where B is a random shock to emissions. Like a total factor productivity shock, this shock could arise from an innovation in abatement technology. It is expected that older plants emit more, plants using scrubbers or clean coal emit less, and a higher operating intensity leads to higher emissions; all of these properties are found in the data.

The final column of Table 1 describes the transition process between states. If a plant of age v updates, its age next period (v') becomes 1. If it does not update its next-period age is $v+1$. The plant type next period b' is determined by the replacement decision. The indicator for the presence of a scrubber next period, s' , equals one if either a scrubber was present this period and the plant was not replaced or a scrubber was built this period. If a plant updates, then it loses its grandfathering status so $u'=1$. The last three state variables, p , A , and B , are determined exogenously by an as-yet unspecified Markov process, so that the expectation of next periods' values of these variables is taken as a function of this process.

Plants maximize discounted expected lifetime profits over an infinite horizon $E_0 \sum_{t=0}^{\infty} \beta^t y_t$ where E_0 is the expectation operator at $t=0$, β is the discount factor, and y_t is profit in period t . The plant's choice can thus be written as a dynamic programming problem:

$$V(v,s,u,b,p,A,B) = \max_{\substack{m \in [0,\bar{m}] \\ z \in \{0,1,2\} \\ x \in \{0,1\} \\ c \in \{0,1\}}} \left\{ \begin{array}{l} Af(v,m,b) - F_c 1(z=1) - F_{nc} 1(z=2) - u\tau Bg(v,c,s,m,b)/m \\ -xG - pcm + E_{A',B',p'|A,B,p} \beta V(v',s',u',b',p',A',B') \end{array} \right\},$$

where the variables are determined according to the transition processes described in Table 1. Even when parameterizing the revenue function, the emissions function, and the process governing the evolution of the shocks and p , this model cannot be solved analytically. Thus, the model is solved numerically through value function iteration (VFI).⁸

While this model does not explicitly incorporate technological growth, it can be derived from a model with an exogenous growth rate that is factored out to make the problem stationary [11]. Suppose that the productivity of a new plant increases at the rate κ each year. Older capital is less productive both absolutely, because of real depreciation at rate, say, δ_r , and relatively, because of the improved technology of newer capital. Thus "capital depreciation" in this model encompasses both physical depreciation (δ_r) and obsolescence (κ). This model cannot separately identify δ_r from κ , but it includes both. Since I am concerned with analyzing policy by simulating counterfactuals, this does not present a problem so long as these two forms of depreciation remain constant under the various counterfactuals. Thus, the model does not accommodate *endogenous* technological growth, which may be induced by changes in environmental policy. A counterfactual strengthened environmental policy may induce additional growth in abatement technology, and thus the counterfactual simulation here will overstate emissions from that counterfactual. So, long-run predictions of the model should be viewed with more caution than short- or medium-run predictions, where technological change cannot yet respond to policy.

While federal emissions regulations have been grandfathered since the 1970 CAA, Title IV of the 1990 Clean Air Act Amendments (CAAA) established an emissions trading program for SO₂ emissions among electric utilities beginning in 1995 [12]. Because of this regulatory change, I do not estimate the model on data after 1995. This policy was not

⁷ The cost differential of switching to low-sulfur coal is modeled as a variable cost. It is possible to also include a fixed cost that is incurred when first switching coal types. Different types of coal can have different ash contents, different difficulties of pulverization, and other different characteristics that may require capital or fixed costs to accommodate the boiler. However, those capital costs are dwarfed by the differential in fuel cost [25, p. 158].

⁸ An appendix that described this method can be found at JEEM's online archive of supplementary materials, which can be accessed at <http://www.aere.org/journals/>.

grandfathered; plants were required to participate regardless of their vintage. Therefore, the analysis here should be viewed as a retrospective analysis of NSPS for SO₂, since those standards exist alongside the permit trading market. The extent to which plants respond in anticipation of a forthcoming policy (for example, by installing a scrubber in 1992 so that their emissions can be reduced when the policy takes effect in 1995) can be captured in this model by incorporating dynamic policy.⁹ Another potential concern is the restructuring of electricity markets, though I suspect that this will be unimportant during my sample period (1985–1995), since restructuring did not begin until the late 1990s.¹⁰

A different grandfathered policy created in the 1977 CAAA, New Source Review (NSR), applies to both new plants and existing plants making major modifications. Thus, a unit could be built before 1970 and hence be grandfathered from NSPS, but have made major modifications that triggered NSR and hence be subject to federal environmental regulations. As identified by Bushnell and Wolfram [6], this additional aspect of NSR creates a second distortion from grandfathering: in addition to keeping older plants from scrapping, it keeps them from making modifications that might trigger NSR, although these modifications may increase productivity and reduce emissions. The second distortion is unlikely to actually apply to any units in my sample period (1985–1995), since it has been widely documented that there had been no vigorous enforcement of this aspect of NSR prior to the late 1990s.¹¹

2. Data

The econometric analysis is done at the boiler level as opposed to the plant level. Each fossil fuel fired power plant can have multiple boilers that could be of different vintages, have different grandfathered statuses, and use different abatement strategies. The data come from three sources. The Energy Information Administration's Form EIA-767 contains information on each boiler, including age, intensity of use, and abatement choices, and is used to calculate the emissions function for coal boilers and the moments for the structural estimation. Form EIA-423 contains plant-level fuel prices for both clean and dirty coal as well as oil and natural gas. The EPA's Emissions and Generation Resource Integrated Database (eGRID) provides emissions data used in estimating the emissions function for non-coal boilers.

Emissions data are not available for the years studied (1985–1995). To estimate the emissions function, I use two different strategies, one for coal boilers and one for non-coal boilers. For coal boilers, information from the Form EIA-767 is used to calculate each boiler's annual emissions from its amount of coal used, average sulfur content of coal, and abatement technology. An engineering equation is used to calculate emissions from these data on a mass-balance basis.¹² Given these emissions estimates (in tons of SO₂), the emissions function $g(v,m,s,c,b)$ for coal boilers is estimated using OLS. These emissions calculations are only available for coal boilers. For non-coal boilers, I use measured emissions data from eGRID.¹³ While measured emissions data ought to be preferable to those calculated from an engineering equation, the measured eGRID data are only available from 1996 to 2000, after the sample period.¹⁴

Data on boiler age, total fuel input, use of clean coal and scrubbers, and grandfathering status come from the EIA-767 data files.¹⁵ These data cover steam-electric plants (including fossil fuel fired plants) with a generator nameplate rating of 10 or more megawatts. I use data from 1985 to 1995, which contain approximately 1800 boilers per year.

Boiler age is calculated from the inservice date reported, or the year in which the boiler went into operation. The grandfathered status of a boiler, that is, whether it is subject to NSPS, is reported in the data set.¹⁶ In 1985, 87% of boilers are grandfathered from NSPS. Fig. 1 presents a histogram of the vintage distribution of all operating boilers in the 1985 EIA-767 data file, categorized according to grandfathered status. While virtually all boilers built before 1970 are grandfathered, and all built after 1978 are not, during the intermediate period variation exists in grandfathering status even for boilers of the same age.

The EIA-767 data also provide information on abatement activity: scrubbers and clean coal. In 1985, 15% of coal boilers used a scrubber. The data include the average sulfur content of the coal used, which can be converted to pounds of sulfur per MMBtu of heat input. Though this value is continuous, for simplicity I model the choice of coal type as a binary variable; it is either clean (low-sulfur) or dirty (high-sulfur). I choose the cut-off point between the two coal types to be

⁹ A dynamic policy incorporating uncertainty can be modeled by through a state variable that captures the policy and evolves according to a Markov process. See [19]. If firms are responding in advance to the policy change, then the baseline estimates of the implicit emissions tax may be biased downwards.

¹⁰ Only 12 states had begun implementing electricity market restructuring before the end of 2000, and as of September 2001 only 2.6% of total customers (residential and non-residential) were served by competitive electricity suppliers (http://www.eia.doe.gov/cneaf/electricity/chg_str/restructure.pdf).

¹¹ Two recent papers also study the impact of NSR on existing plants [6,23].

¹² See [23]. For boilers without scrubbers, total emissions is the product of the total tons of coal, the average sulfur content of coal, and a boiler-specific emissions factor. For scrubbed boilers, the equation also considers scrubber-specific emissions reduction factors.

¹³ Available online at www.epa.gov/cleanenergy/egrid.

¹⁴ Data on emissions of nitrogen oxides (NO_x) and other pollutants are also included; I focus here solely on SO₂.

¹⁵ Available at <http://www.eia.doe.gov/cneaf/electricity/page/eia767.html> for 1996–2005, but the data used here for earlier years were obtained by request from the EIA.

¹⁶ At least three different types of NSPS can apply to the boilers, and each boiler is asked which type of NSPS it faces. While the standards differ somewhat, for simplicity I consider only a binary grandfathered policy.

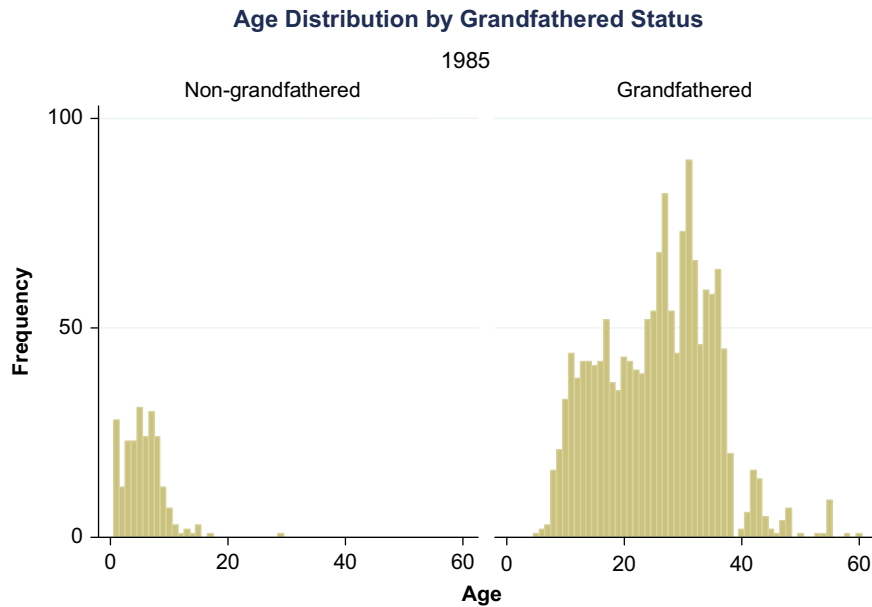


Fig. 1. Data source is EIA-767, 1985. The x-axis is the age of the boiler; the y-axis is the number of boilers of that age. Grandfathered boilers are those not subject to any New Source Performance Standard.

1.2 lb/MMbtu [12], the value of the performance standard for new boilers set in the 1970 CAA. Based on this definition, 34% of coal boilers use clean coal in 1985.

This differential is clearly due to heterogeneous price differentials between clean and dirty coal. The EIA-423 data contain plant-level information on total quantity and average price of annual coal consumption.¹⁷ The data are given for the price actually paid for the coal actually received by the plants. A single plant can use coal from multiple sources, some of which may be high-sulfur and some of which may be low-sulfur. For plants that use both types of coal, I am able to construct an average price paid for each of the two types. Many plants, though, only use one type of coal. For these plants, I use the state average of coal prices for the other type of coal. For some states, no plants use a particular type of coal, and I use the average price for the census region of a particular type of coal for plants in these states. For those plants with multiple boilers, I assign each boiler the same plant-level value of coal prices.

Coal price data are available for coal plants only. Since the model allows for non-coal boilers to scrap and be replaced by coal boilers, I also want coal prices available to these boilers. Thus for non-coal boilers I assign coal prices based on the state or regional average coal prices.

The last item obtained from the EIA-767 data set is the annual amount of fuel input used by a boiler. For each fuel, I convert the fuel amounts into the amount of heat energy (in MMBtus). While 58% of all boilers are coal in 1985, they represent over 90% of the total amount of fuel input. Most of the remaining 10% is from gas boilers.

In addition to the differential fuel use by type of boiler, the age of the boiler also makes a large difference in the operating intensity. Boilers built after 1970 are used much more intensively. The grandfathered boilers represent a large fraction of the industry by number (85%) but only a small fraction by fuel use (7%). Boiler efficiency is declining in age, averaging 86.4% for the youngest decile of boilers and 84.0% for the oldest decile.

I choose two sets of moments to identify the parameters in the structural portion of the estimation. First, I use a set of moments from each year in the sample period. Table 2 includes moments from the first year, 1985, although they will not be used in the estimation, since the 1985 data are used to set the initial distribution of boilers. Also, I omit 1995, the last year for which I have data, since one of the moments requires data from the following year to construct it. For each year in 1985–1994, I calculate the fraction of boilers grandfathered, the fraction that uses coal, and of the coal boilers, the fraction that has a scrubber and the fraction that uses clean coal. These annual data moments are meant to identify changes in the industry over the different years and match how boiler characteristics change with time. Second, I aggregate the boilers from all ten years of data and divide them into five categories based on their age and grandfathered status: less than or equal to 15 years old, between 16 and 25 years old and grandfathered, between 16 and 25 years old and not grandfathered, between 26 and 35 years old, and greater than 35 years old. The age group between 16 and 25 years old is divided into grandfathered and non-grandfathered boilers since it is those ages of boilers for which there is variation in the grandfathered status. For each category, I evaluate the fraction of boilers scrapping, the fraction of boilers that uses coal,

¹⁷ Available at: http://www.eia.doe.gov/cneaf/electricity/cq/cq_sum_backissues.html.

Table 2

Data moments used in MSM estimation.

Annual moments					
	1985	1986	1987	1988	1989
Fraction grandfathered	0.8736018	0.8659091	0.8609502	0.859366	0.8571429
Fraction coal boilers	0.5755079	0.5869809	0.6020702	0.604611	0.6030006
Fraction of coal with scrubber	0.1354962	0.1453155	0.1451767	0.1458532	0.1492823
Fraction of coal using clean coal	0.3435115	0.3518164	0.3906399	0.4566254	0.3617225
	1990	1991	1992	1993	1994
Fraction grandfathered	0.8564946	0.8510029	0.8485023	0.8463768	0.8435092
Fraction coal boilers	0.5996563	0.6087963	0.6101104	0.6153396	0.6230586
Fraction of coal with scrubber	0.1509074	0.1625475	0.1685714	0.1712655	0.1783317
Fraction of coal using clean coal	0.3352436	0.3612167	0.3314286	0.3644148	0.4074784
Categorical moments					
	Age ≤ 15	15 < Age ≤ 25 and grandfathered	15 < Age ≤ 25 and not grandfathered	25 < Age ≤ 35	Age > 35
Fraction scrapping	0.0052647	0.0091001	0.0233645	0.0233645	0.0705286
Fraction of coal with scrubber	0.4264165	0.1077819	0.5617978	0.0556046	0.0382166
Fraction of coal using clean coal	0.5283297	0.3499792	0.511236	0.3415711	0.2845011
Fraction coal boilers	0.7680136	0.5695662	0.5615142	0.5695375	0.5675839
Fraction in category	0.1630	0.1927	0.0144	0.2726	0.3573

Notes: Data source is EIA-767, 1985–1995.

of the coal boilers, the fraction with scrubbers and the fraction using clean coal, and the fraction of boilers in each category. These categorical data moments are used to identify the impact of both age and policy on boiler operating decisions.

The top half of Table 2 displays the annual moments. The fraction of boilers grandfathered declines slightly over the period from 87% to 84%. Of coal boilers, the fraction using scrubbers increases from 14% to 18%. The percentage using clean coal goes from 34% in 1985 to 41% in 1994, but not monotonically; this fraction peaks at 46% in 1988. The fraction of all boilers that are coal boilers increases slightly from 58% to 62%.

The second panel of Table 2 displays the categorical moments. While the annual moments do not differ significantly from year to year, the categorical moments do strikingly differ by category. The first row shows that the fraction of generators scrapping is very low. The percentages in the first four categories are all less than three percent, and they show that a very small number of boilers are scrapping. For example, the value of 0.5265% for boilers younger than 15 years old comes from 18 boilers scrapping out of 3419 in that age category over the ten years. The patterns in the data are conformable to intuition. One would expect that grandfathered boilers are less likely to scrap than non-grandfathered boilers, and that is just what is seen comparing the second and third categories. As boilers get older, they are more likely to scrap; 7% of boilers older than 35 years old scrap.

The moments related to abatement activity are consistent with boilers responding to grandfathered policy. Grandfathered units are less likely to have a scrubber or to use clean coal. The rate at which boilers adopt either of these abatement measures decreases with boiler age. The youngest boilers are more likely to be coal fired, but there is no difference in the fraction of coal boilers between the other age categories. About 35% of the boilers are more than 35 years old, while about 16% are younger than 16 years old.

3. Estimation strategy and results

The model is estimated in three steps. First, the emissions function $g(v,m,s,c,b)$ is estimated with OLS, after imposing a functional form. This equation is an emissions production function; given the plant's age, abatement equipment, type, and choice of operating intensity, a level of emissions will be produced. Second, I estimate the revenue function $f(v,m,b)$ using panel data estimation allowing for boiler-specific effects and autocorrelated error terms. Finally, since it is impossible to estimate all of the parameters in this way, I estimate the rest of the parameters using the method of simulated moments (MSM) ([13–16]). While all of the parameters could be estimated using MSM, estimating a subset of parameters using other methods when possible is preferable because it saves computational time and the identification of those parameters is more straightforward.

Because the emissions function, g , is a function of plant type, b , I estimate a separate emissions function for coal and non-coal plants. The estimating equation for coal boilers is

$$\ln(\text{emis}) = \beta_0 + \beta_1 \ln(\text{age}) + \beta_2 \ln(m) + \beta_3 \text{lsc} + \beta_4 \text{scrubber} + \beta_5 \text{lsc} \times \text{scrubber} + \varepsilon.$$

Table 3
Emissions function estimation results.

	Coal				Non-coal				
	(1) lnemis	(2) lnemisrate	(3) lnemisrate	(4) lnemisrate	(5) lnemisrate	(6) lnemis	(7) lnemisrate	(8) lnemisrate	(9) lnemisrate
lnage	0.128*** (0.0102)	0.138*** (0.00850)	0.165*** (0.0102)	0.0504*** (0.0166)	0.297*** (0.0239)	0.166 (0.148)	−0.189 (0.136)	0.438** (0.190)	0.0258 (0.141)
scrubber	−1.308*** (0.0328)	−1.306*** (0.0330)	−1.239*** (0.0373)	−1.213*** (0.0622)	−1.458*** (0.0384)				
lsc	−1.059*** (0.00961)	−1.058*** (0.00958)	−0.997*** (0.0120)	−0.907*** (0.0129)	−0.924*** (0.0169)				
lsc_scrubber	0.285*** (0.0435)	0.283*** (0.0435)	0.130*** (0.0494)	0.385*** (0.0826)	0.680*** (0.0537)				
lnmmbtus	0.994*** (0.00404)					1.207*** (0.0308)			
nonattainment			−0.0964*** (0.0190)				−0.697 (0.495)		
Coal–gas ratio			−0.129 (0.0972)				−0.236 (1.176)		
Coal–oil ratio			−0.0700 (0.0597)				−0.125 (1.824)		
Capacity				−1.7e−7*** (3.03e−08)					0.000226*** (4.87e−05)
Heat rate				5.19e−07 (5.25e−07)					−3.39e−7*** (9.74e−08)
Hours				−1.0e−5*** (2.44e−06)					6.5e−5*** (1.64e−05)
Constant	−6.697*** (0.0892)	−6.832*** (0.0333)	−6.791*** (0.0578)	−6.592*** (0.0665)	−7.720*** (0.0930)	−17.03*** (0.837)	−13.57*** (0.602)	−14.98*** (1.023)	−14.45*** (0.624)
Observations	10,837	10,837	6530	6209	5154	2940	2940	1446	2940
R-squared	0.937	0.806	0.819	0.722	0.734	0.719	0.646	0.579	0.650

Notes: Coal emissions data in columns 1–4 are calculated based on EIA-767, 1985–1995. Coal emissions data in column 5 and non-coal emissions data in columns 6–9 are from eGrid, 1996–2000. Columns 3 and 8 include only boilers in states that had not begun electricity market restructuring by 2000. The dependent variable is the log of tons of SO₂ emissions or the log of tons of SO₂ emissions per MMBtus heat input. All columns include unreported state and year fixed effects. Robust standard errors are in parentheses.

*** $p < 0.01$.

** $p < 0.05$.

The age of the boiler, the amount of fuel input (m), whether the boiler uses low-sulfur coal (lsc) and whether the boiler has a scrubber are all included, as is the interaction term between the last two abatement methods. The estimating equation for non-coal boilers is identical, save that the lsc and $scrubber$ terms (and their interaction) are not included. I do not estimate the emissions shock B because I do not have actual emissions data and thus do not have to accommodate variation in the data.

Results are presented in Table 3. The first four columns are the results for coal boilers from the calculated emissions data from 1985 to 1995. Unreported state and year dummies are included in all columns of Table 3 to account for unobserved heterogeneity in the regulatory environment at the state level, as well as trends in emissions picked up by the year effects. Older plants emit more, once controlling for total heat input and abatement activity. The elasticity of emissions with respect to age is about 0.13. Note that this is measuring both physical depreciation and technological obsolescence; older boilers are both absolutely and relatively dirtier. The use of scrubbers and the use of low-sulfur coal significantly decrease emissions, as expected, and the interaction term is positive (since a scrubber has less to “scrub” when low-sulfur coal is already being used). The elasticity of emissions with respect to heat input is almost exactly one. This suggests that emissions are linear in heat input, which is used in identifying the revenue function off the first-order condition.¹⁸

The next column of Table 3 uses the log of the emissions rate, in tons of SO₂ per MMBtu heat input, as the dependent variable. This estimating equation is

$$\ln(emis/m) = \beta_0 + \beta_1 \ln(age) + \beta_2 scrubber + \beta_3 lsc + \beta_4 lsc \times scrubber + \varepsilon.$$

The emissions rate is also higher for older boilers, with an elasticity of about 0.14. Scrubbers and low-sulfur coal have the expected effects.

Columns 3–5 consider alternate specifications for coal boilers. In each of these columns the dependent variable is the log of the emissions rate. Column 3 includes an indicator variable equal to one if the boiler's county is not in attainment of

¹⁸ One can also regress emissions on a polynomial of heat input. The non-linear terms are significant but orders of magnitude smaller than the linear term.

the Clean Air Act's air quality standards for SO₂. It also includes a measure of relative prices of three fossil fuel inputs, averaged at the state-year level from EIA-423 data. Column 3 also removes all observations from states that had begun electricity restructuring by 2000, to eliminate any potential bias from utilities' anticipatory behavior of restructuring. Boilers in nonattainment counties have a significantly lower emissions rate, and fuel prices do not enter significantly. The coefficient on the log of boiler age is 0.17, only slightly higher than in column 2.

Column 4 includes three additional variables related to boiler technology or intensity. These are the capacity of the boiler's plant (in megawatts), the heat rate of the boiler's plant (in BTU/kilowatt-hours), and the number of hours the boiler is in operation over the year. Plant capacity and hours of operation enter negatively, while heat rate is insignificant. The coefficient on the log of boiler age is still significantly positive, but falls to 0.05. Lastly for coal boilers, column 5 performs the same regression as column 2 but uses the actual emissions data from eGrid (available only from 1996 to 2000) rather than the estimated emissions data from EIA-767 (from 1985 to 1995). The coefficient on the log of boiler age is still significant but is larger than in column 2.

The regression results for non-coal boilers are presented in columns 6–9 of Table 3. Column 6 regresses the log of total emissions on the log of age and the log of total heat input, as well as unreported state and year fixed effects, and column 7's dependent variable is the log of the emissions rate. Column 8 includes county-level attainment status, relative fuel prices, and omits states that began restructuring by 2000; column 9 includes capacity, heat rate, and hours. The non-coal regressions do not tell a consistent story. Only in column 8 is there a positive and significant coefficient on age (as expected).

This unusual result is not worrisome for two reasons. First, the contribution of non-coal boilers to SO₂ emissions is small; coal plants are responsible for 97% of total SO₂ emissions from this sector. If the non-coal regression results in this stage are off, it will not substantially affect the simulations. Second, the absence of a positive elasticity between age and emissions for non-coal plants is possible. Technological innovation in abatement of SO₂ emissions from power plants has focused mainly on coal plants because of their dominance in this pollutant. Since the age elasticities in Table 3 include the effect of obsolescence from technological growth, the lack of a positive age effect may signal this lack of technological growth in this area. I use the point estimates from columns 2 and 7 in the MSM estimation and simulations.

In the second stage of the estimation, I exploit the fact that two of the choice variables are chosen in a static, not dynamic, problem. Both m , the operating intensity, and c , the sulfur content of coal input, affect the single-period profit function but no future state variables and thus not the next period value function.¹⁹ The optimal decision over these two variables can be made by just looking at the single-period profit function. In the case of m , a continuous choice variable, a first-order condition arises

$$Af_m(v, m, b) - \frac{d}{dm} [u\tau Bg(v, c, s, m, b)/m] - pc = 0.$$

The first term in the equation represents the marginal benefit to a plant of an additional unit of fuel input from increased revenue. The middle term represents the costs from the environmental policy as measured by a marginal increase in the virtual emissions tax paid, paid only by non-grandfathered plants ($u=1$). The last term is the marginal cost of using clean rather than dirty coal.

The middle term of the first-order condition disappears when the emissions function g is linear in fuel input m , as in the emissions estimates above. This simplifies the moment condition. It depends on the parameterization of f , defined as follows:

$$f(v, m, b = \text{coal}) = d_{0c}m + d_{1c}vm - \frac{1}{2}m^2$$

$$f(v, m, b = \text{non-coal}) = d_{0nc}m + d_{1nc}vm - \frac{1}{2}m^2.$$

This parameterization captures an important dimension of the behavior of these boilers: older plants are used less intensively. Under this parameterization, this behavior of plants can be captured when d_{1c} and d_{1nc} are negative.²⁰

Given this parameterization, the first-order condition for fuel use for non-coal boilers is $d_{0nc} + d_{1nc}v - m = 0$, while the first-order condition for coal boilers is $d_{0c} + d_{1c}v - m - pc = 0$. Notice that the emissions function for coal plants has the additional variable c for low-sulfur coal use. A coal boiler optimizes over whether to use clean or dirty coal ($c=1$ or 0), but after doing so the boiler chooses a level of fuel input which satisfies the first-order condition above, for either value of c .

The reduced-form revenue function abstracts from supply and demand curves and from the strategic behavior of firms operating under regulators. Because markets as well as regulations differ by time and location, I include in the revenue function estimation state- and year-fixed effects.

These results are presented in Table 4. Columns 1–4 present the results for non-coal boilers, and columns 5–8 present the results for coal boilers. Columns 1 and 5 are simply OLS regressions with standard errors clustered at the boiler level;

¹⁹ This neglects effects from long-term contracts for coal, where utilities cannot choose an optimal coal type each period since they are locked in to a contract [22]. This should not present a significant distortion, since utilities are optimizing over contract terms and, on average, the outcome of that behavior should follow the outcome of choosing coal types each year.

²⁰ Identification requires some normalization of the profit function; I therefore fix the coefficient on m^2 .

Table 4
Revenue function estimation results.

	Non-coal				Coal			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Age	-1.8e+06*** (256,324)	-891,911*** (223,868)	-920,688*** (233,285)	-777,489*** (130,785)	-9.31e+6*** (496,721)	-6.74e+6*** (1.599e+06)	-6.72e+6*** (1.603e+06)	-3.59e+6*** (226,540)
Coal-gas ratio			7.018e+06 (7.495e+06)				1.502e+07 (1.204e+07)	
Coal-oil ratio			-9.773e+06 (1.044e+07)				-2.10e+7*** (6.154e+06)	
$p \times c$					-124,463 (115,035)	-95,069*** (36,650)	-92,138** (36,725)	-43,314 (35,468)
Constant	9.31e+7*** (8.937e+06)	7.07e+7*** (7.240e+06)	7.61e+7*** (7.926e+06)	2.657e+07 (5.880e+07)	4.91e+8*** (1.442e+07)	3.80e+8*** (2.724e+07)	3.87e+8*** (2.756e+07)	3.30e+8*** (1.040e+08)
Observations	3096	3096	2907	3096	3760	3760	3760	3760
R-squared	0.347	0.334	0.322	0.331	0.527	0.510	0.510	0.419
Number of boilers	-	720	704	720	-	975	975	975
Estimation Error structure	OLS Clustered at boiler	RE Clustered at boiler	RE Clustered at boiler	RE AR(1)	OLS Clustered at boiler	RE Clustered at boiler	RE Clustered at boiler	RE AR(1)

Notes: Data source is EIA-767, 1985–1995. Relative fuel price ratio data are from EIA-423, 1985–1995. Robust standard errors in parentheses.

*** $p < 0.01$.

** $p < 0.05$.

the significant negative coefficient on age indicates that older boilers are used less intensively. For coal boilers, the coefficient on pc is negative, as expected, though only significantly so in columns 6 and 7.

Columns 2 and 6 use a random-effects specification, to pick up an unobserved boiler-specific effect. These specifications result in a coefficient on age that is still significantly negative but has a smaller magnitude, by about a factor of one-half for non-coal boilers and a factor of two-thirds for coal boilers. Columns 3 and 7 keep the random-effects specification and also add state-year relative fuel prices as regressors. Given that many boilers and plants belong to a utility's fleet and are dispatched on a least-cost basis, relative fuel prices may matter. The coefficient on age does not change after inclusion of the price variables. Finally, in columns 4 and 8 the error is allowed to be autocorrelated according to an AR(1) process; again the coefficient of interest is largely unchanged. The preferred specification for the parameter values comes from columns 4 and 8, and those estimated parameters are used in the simulations.

Finally, in the third stage I use MSM to estimate the remaining parameters of the model. This method is described in more detail in the Appendix. These parameters include the fixed costs for building a coal and non-coal boiler, F_c and F_{nc} , and the fixed cost of building a scrubber G . The implicit tax on emissions intensity τ is also estimated in this stage. Ideally, this could be calibrated based on the policy. Because the complex CAA is modeled simply as an emissions tax, where the value of the tax is the shadow price on pollution intensity that the policy creates, this value is unknown. Hence, it falls into the parameter set to be estimated. The productivity shock A is idiosyncratic.²¹ It is assumed to be multiplicative and distributed log-normally with median one. That fixes its first parameter, μ , to zero, while σ_A^2 is estimated. The shock is persistent and evolves according to a Markov process. A boiler has probability P of having the same productivity shock in the next period. With probability $1-P$, the next-period productivity shock is randomly chosen from the log-normal distribution.

The annual discount rate β is set at 0.95.²² This leaves six parameters to be estimated with MSM: $[F_c, F_{nc}, G, \tau, \sigma_A^2, P]$. The data moments that are used in MSM are described in the following section and listed in Table 2.

The MSM estimation is based on an explicitly dynamic structural model, which differs from estimation techniques used in previous literature.²³ As a result, the mechanisms by which policy affects firm behavior can be better understood.

As with any estimation but in particular with structural estimation, one must pay careful attention to identification of the model's parameters. In the MSM estimation stage, I identify the implicit tax on emissions using variation in the behavior of grandfathered versus non-grandfathered facilities. Grandfathering is based on age, so it is difficult to separate an age effect from a grandfathering effect. I do so by exploiting the fact that of the facilities aged 15–25 years, some are grandfathered and some are not. It is unfortunately impossible to determine the source of this variation from the data. The age is based on the date the plant went online, but the grandfathering status for some plants may have been partly based

²¹ I have experimented with adding a common component to the productivity shock to capture business cycle effects, using an additional set of data moments based on business cycles. However, the common component of the shock was not found to be significant.

²² Though β could be estimated, it may be difficult to identify due to collinearity with the fixed costs of replacement F_c and F_{nc} [7, p. 1023].

²³ [28] estimates a structural model of the cement industry's response to the 1990 CAA amendments, but not grandfathering.

Table 5

Structural parameter estimation results.

F_c	1803.34 (1.61)
F_{nc}	84.95 (0.60)
G	61.41 (0.45)
τ	2989.77 (258.62)
σ_A^2	0.1104 (0.4851)
P	0.8727 (0.1207)

Notes: Standard errors are in parentheses. Estimates come from simulated method of moments, matching moments from Table 2.

on the date of construction. For some facilities, the grandfathering status could be determined by the utilities bargaining with regulators to adjust the official online date.

Clearly, grandfathering status was not randomly assigned; the question is how or if this determination creates a bias. Suppose that of the plants within that age group, those with the highest expected abatement costs were more likely to be grandfathered, say as a result of lobbying. The higher use of abatement among non-grandfathered plants is thus partly due to the unobservable difference in abatement costs, though the estimation strategy attributes it to the policy. Thus the estimated implicit tax created by the policy may be biased upwards. Alternatively, if the selection into grandfathering based on abatement costs went the other way, then there is a downward bias on the estimate. It is not possible to measure the direction or magnitude of this potential bias.

The results from the MSM estimation are presented in Table 5. Each point estimate is presented with its estimated standard error, and all parameters have small standard errors save σ_A^2 . The fixed cost parameters are normalized to the boiler's profit function $y = Af(v, m, b) - F_c 1(z=1) - F_{nc} 1(z=2) - \tau ue/m - xG - pcm$. The relative magnitudes of the three fixed cost parameters F_c , F_{nc} , and G give information about construction costs. The cost of a new coal boiler is 20 times as large as that of a new non-coal boiler; coal boilers and plants tend to be of much larger capacity, and this is represented both in the profit function of the boiler and in the fixed cost of building one. The cost of adding a scrubber to a coal boiler is about one-thirtieth the cost of the boiler itself, which is consistent with engineering estimates of production costs.

Like the fixed costs, the policy parameter τ can be interpreted in relation to the profit function. The emissions function g gives the emissions intensity in tons of SO_2 per MMBtu fuel input as a function of age and abatement technology. The implicit tax τ is levied on emissions intensity. The amount that is paid by a boiler in the virtual tax (if it is non-grandfathered) is the product of the estimated value of τ and its emissions intensity. For example, a coal boiler aged 20, with no scrubber and not using low-sulfur coal, has an estimated emissions intensity of 0.0016 tons SO_2 per MMBtu fuel input. At the estimated tax value of about 3000, this equates to a tax payment of 4.8. Although this has no units, it can again be compared to the other parameters. Consider the cost of a scrubber, 61.41. This is about 13 times the annual payment in implicit emissions tax for that particular boiler. If that same boiler had a scrubber, its emissions intensity would be 0.000442, and its tax payment would be reduced to 1.3. Thus, utilities are balancing long-run projections of emissions tax payments with scrubber costs to decide when to add a scrubber.

That τ significantly differs from zero implies a rejection of the null hypothesis that grandfathering has no effect on the behavior of plants. Since the implicit tax is paid only by non-grandfathered boilers, a positive τ means that these boilers respond to an incentive to reduce their emissions that the grandfathered boilers do not.

The parameter σ_A^2 gives information about the variance of the random productivity shock. The estimated value of the parameter, about 0.1104, corresponds to a variance of 0.4588.²⁴ Finally, P represents the persistence in the idiosyncratic component of the productivity shock. If $P=0$, then no persistence exists, and if $P=1$ then the idiosyncratic component is constant. Here, P is 0.87. A generator has an 87% chance of staying in its current productivity state and a 13% chance of taking a random draw from the distribution of shocks; the idiosyncratic element of revenue is quite persistent.

4. Counterfactual simulations

The estimated parameters can be used to simulate the effects of policy changes. Specifically, I consider the impact of grandfathering in the CAA by comparing a baseline simulation using the estimated parameters to three counterfactual simulations. In all counterfactuals I present the simulation results for ten periods after the policy change, though the simulations come from an infinite horizon maximization model. In the first counterfactual, I keep all parameters the same, except that I eliminate any grandfathering provisions starting in the first year. Next, I set the implicit tax rate on pollution equal to 90% of its estimated value to simulate what would have happened had the stringency of the grandfathered CAA regulations been reduced by 10% in 1985. Finally, the third counterfactual increases the implicit tax rate on pollution by 10%, to represent the CAA regulations increasing in strength in 1985. The counterfactual that eliminates grandfathering is chosen to quantify the effects of grandfathering itself in the policy, rather than the effects of the stringency of the policy. The two other counterfactuals, which keep grandfathering but alter the stringency of the policy for non-grandfathered units, are chosen to investigate how policy changes that maintain grandfathering will affect outcomes.

²⁴ The variance of a lognormal distribution is $\exp(\sigma^2 - 1) \times \exp(2\mu + \sigma^2)$, and here μ is zero.

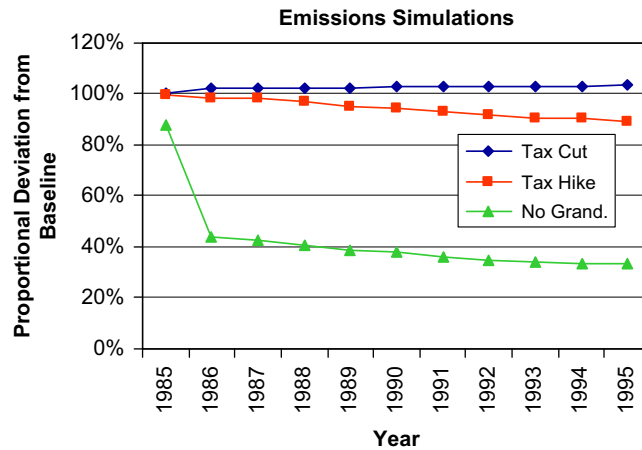


Fig. 2. Results are from counterfactual simulations as described in the text. Values presented are proportional deviations from baseline simulation.

For the eleven periods presented of the three counterfactual simulations, Fig. 2 presents the level of emissions. The line “No Grand.” represents the simulation that eliminates grandfathering. The line “Tax Cut” represents the simulation with the tax rate reduced. The line “Tax Hike” represents the simulation with an implicit tax rate increased by 10%.²⁵ The emissions levels are presented as proportional deviations from the baseline simulation. For example, in 1995 emissions under the “Tax Hike” counterfactual are 90% of emissions under the baseline scenario.

Consider first the “No Grand.” simulation, which shows the largest difference from the baseline; emissions drop sharply. Since older boilers are no longer grandfathered under this counterfactual policy simulation, those boilers now face an implicit emissions tax. Many of these older boilers now choose to adjust their capital to a new vintage to reduce their emissions, especially since they have delayed productivity-enhancing upgrades to avoid that tax. In the next year, many more boilers are newer and cleaner. Once the law is changed, the rate of investment is much higher due to the shock in the policy. Though the number scrapping quickly gets closer to the baseline levels, the initial increase in brand new boilers reduces emissions throughout the simulation period. The magnitude of the change in both emissions and investment is large; in the first year after the policy change 350 additional boilers scrap, and by 1995 emissions are less than 40% of baseline levels. These large outcomes are from a very large policy change: since 87% of boilers are initially grandfathered, eliminating grandfathering amounts to a massive policy shift. The result I find is slightly smaller in magnitude than that of [17], who find that if all plants were subject to NSR standards (that is, if grandfathering were eliminated), then emissions of SO₂ and NO_x would fall by 75%.

The results from the “Tax Cut” and “Tax Hike” counterfactuals follow intuition. When the implicit tax on emissions intensity is increased by 10%, emissions fall, and ten years after the policy change emissions are 10% lower than the baseline. A 10% decrease in the value of the implicit tax sees a rise in emissions, though of a smaller magnitude. Emissions in 1995 are only 4% higher than in the baseline.

The effects of the counterfactual simulations on emissions can be compared with scrapping simulation results. Under the “Tax Cut” scenario, a larger number of boilers initially scrap and replace, while under the “Tax Hike” scenario, a smaller number do. Thus in the “Tax Hike” scenario, boilers are older on average than in the baseline scenario. Since older boilers are on average dirtier, one would expect that in the presence of an increase in the virtual emissions tax, boilers would scrap at a higher rate to avoid paying this increased tax, and likewise after a decrease in the implicit tax rate fewer boilers would scrap. However, here the effects of grandfathering come into play: grandfathered boilers have a valuable status that is lost when they scrap and replace. Increasing the tax rate on emissions increases the value of this status, and makes grandfathered boilers less likely to scrap than in the baseline simulation.

This effect of grandfathering on scrapping rates can be seen in Fig. 3. This presents the simulation results for scrapping for the “Tax Cut” and “Tax Hike” counterfactuals, separately for boilers that are initially grandfathered and for those that are initially non-grandfathered. The left-hand axis is the difference between the number of boilers scrapped in the counterfactual and the number scrapped in the baseline. For example, in the first year of the “Tax Cut” counterfactual about 50 more grandfathered boilers are scrapped than in the baseline. The top panel of Fig. 3, which shows the results for boilers that are initially grandfathered, shows the results of changes in the virtual tax rate on scrapping outcomes. When the tax is increased, fewer of these grandfathered boilers scrap, since they have even more incentive to hold on to their valuable grandfathered status. On the other hand, the bottom panel of Fig. 3, which shows the results for boilers that are

²⁵ For clarity, confidence intervals are not presented but can be calculated by taking 100 draws from the asymptotic distribution of the parameters values, simulating the counterfactual for each draw, and calculating the resulting standard errors. This is the same method used to calculate the confidence intervals of the simulated moments in Table 5.

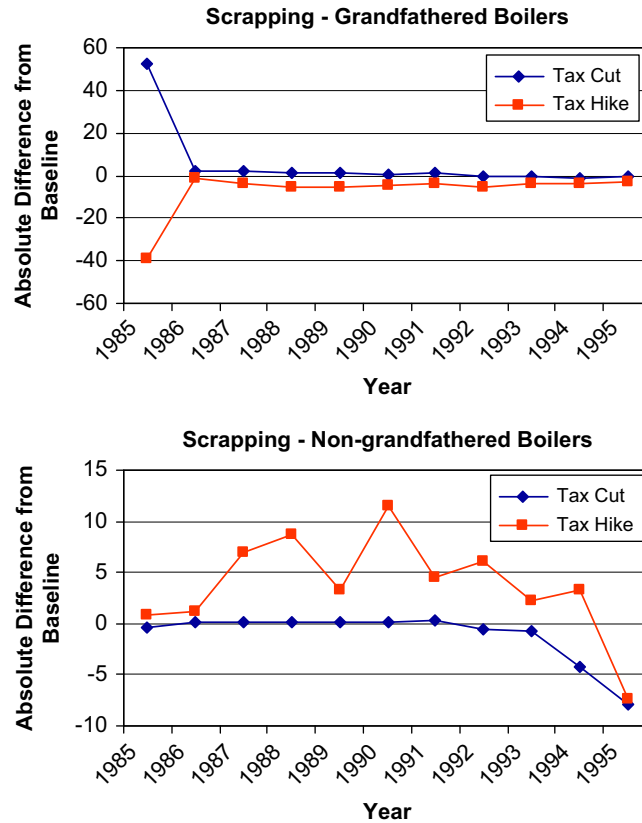


Fig. 3. Results are from counterfactual simulations as described in the text. Values presented are absolute deviations from baseline simulation. The top panel presents simulation results for boilers that are initially grandfathered. The bottom panel presents the results for boilers that are initially non-grandfathered.

initially non-grandfathered, shows no such behavior. After a tax increase, more boilers choose to scrap to newer, cleaner plants, to avoid paying the higher tax on emissions intensity. After a tax decrease the scrapping rate is slightly lower. On the whole, the effect from grandfathered boilers dominates the effect from non-grandfathered boilers, in large part because such a high fraction of boilers (87%) are grandfathered. Fig. 3 also demonstrates how grandfathering creates a schism in the reaction of boilers to policy changes, with grandfathered and non-grandfathered boilers responding in opposite directions. The reduction in new investment comes only from grandfathered boilers, and hence the magnitude and even the existence of a reduction in investment overall depends on the fraction of the fleet in the cross-sectional distribution that is grandfathered at the onset of the policy change.

Comparing the results for emissions in Fig. 2 with those for scrapping in Fig. 3 presents a puzzle. No evidence of a perverse result for emissions is found, yet a perverse result for scrapping is found. These results do not conform to some previous theoretical and empirical findings of a perverse effect of grandfathered policies for automobile emissions standards [1].

If an increase in the tax rate leads to less scrapping and hence older boilers than in the baseline, why is not the resulting level of emissions from these older, dirtier boilers higher than in the baseline? The choice between coal and non-coal boilers is affected by the tax rate. Since non-coal boilers emit much less SO_2 , replacing scrapped boilers with non-coal boilers instead of coal boilers reduces emissions, even with no difference in the overall scrapping rate.

The simulated counterfactual policy scenarios affect the choice of scrapped boilers between replacing with coal and non-coal units. The top panel of Fig. 4 presents the number of boilers scrapping and replacing with coal boilers, and the bottom panel presents the number scrapping and replacing with non-coal boilers. Each panel presents the absolute difference of the number scrapping from the baseline for each counterfactual. The first panel shows that the difference between the scrapping rates in the “Tax Cut” and “Tax Hike” counterfactuals is wholly accounted for by the different number of boilers replacing with coal boilers. In the “Tax Cut” counterfactual, 50 more boilers are scrapped and replaced with new coal boilers compared to the baseline, whereas in the “Tax Hike” scenario 40 fewer boilers are scrapped and replaced with new coal boilers. The difference between these two counterfactuals and the baseline in the number of boilers scrapping and replacing with non-coal boilers is almost zero. The apparent contradiction between the outcome for scrapping rates and the outcome for emissions is thus resolved: although a tax decrease results in a higher scrapping rate and hence younger boilers compared to the baseline scenario, the additional boilers that do scrap are replaced by coal

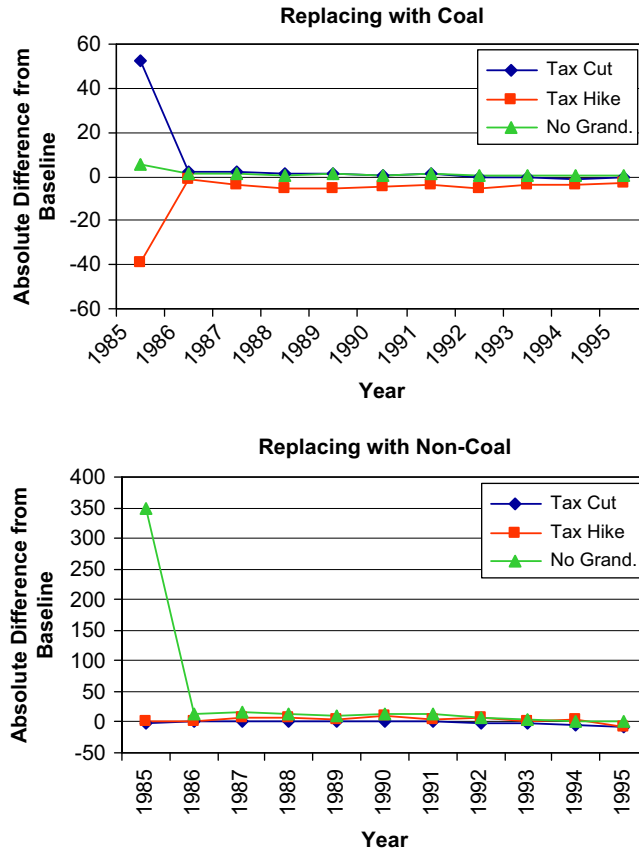


Fig. 4. Results are from counterfactual simulations as described in the text. Values presented are absolute deviations from baseline simulation. Top panel presents the number of boilers scrapping and replacing with coal boilers; the bottom panel presents the number scrapping and replacing with non-coal boilers.

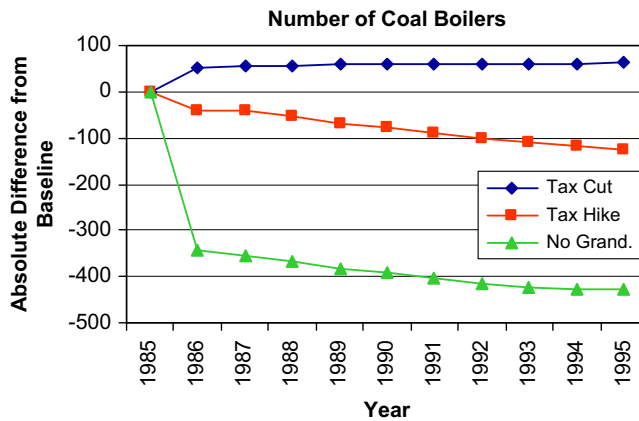


Fig. 5. Results are from counterfactual simulations as described in the text. Values presented are absolute deviations from baseline simulation.

rather than non-coal boilers. These boilers have a higher SO₂ emissions rate, and hence emissions decreases after the tax decrease relative to the baseline. The opposite holds for a tax increase. This explanation also likely accounts for the results from prior literature, here explained through the dynamic model.

The results for the “No Grand.” counterfactual simulation show quite a different pattern. The top panel of Fig. 4 shows that the number of boilers scrapping and being replaced with coal boilers is no different between the baseline and “No Grand.” By contrast, the bottom panel shows that a much larger number of boilers scrap and are replaced with non-coal boilers in this counterfactual. Thus, the change in the overall scrapping rate is almost entirely driven by this different rate

of replacement with non-coal boilers. Because these boilers emit much less SO_2 than coal boilers, emissions under this counterfactual are much lower than the baseline, as shown in Fig. 2. Finally, Fig. 5 summarizes the differences in types of boilers across the three counterfactual simulations by presenting the total number of coal boilers in each simulated year, as compared to the number in the baseline scenario. Because of the different scrapping and replacement rates described above, the number of coal boilers in the economy is highest under the “Tax Cut” simulation and lowest under the “No Grand.” simulation. Comparing this figure to Fig. 2 demonstrates that the simulated changes in emissions are largely driven by boiler type.

How realistic is the assumption that utilities are able to change the boiler type from coal to non-coal or vice versa, and by how much does this assumption affect the results? Plants are often located near a primary fuel source or access point: virtually all coal plants are located near railway lines or at the mouth of a mine, and natural gas plants have laterals that connect to major pipelines. It should be noted, though, that replacement in this model does not require that the boiler be located in the exact same location. A utility could scrap its coal plant located on a railway line and build a gas plant located near a pipeline. However, it may be that utilities have less freedom to change plant type than the model here provides them, for geographic or other constraints (perhaps a utility operates in an area with no access to any natural gas pipelines). Counterfactual results in a model where utilities do not have the option of changing the boiler type (coal boilers can be replaced only with coal) are almost identical to those of the original model: eliminating grandfathering reduces emissions by about 55% compared to the baseline. However, the mechanism by which this emissions reduction occurs is different. In the original model, emissions dropped because many coal boilers were replaced with non-coal boilers. That is not possible here; instead, many more coal boilers add scrubbers, bringing down total emissions. By the tenth year of the simulation, the number of scrubbers increases by a factor of four in the no grandfathering counterfactual compared to the baseline.

Finally, I examine how sensitive the results are to the choice of modeling the policy as an implicit tax on emissions intensity. I estimate an alternative model where the tax is levied on emissions level rather than emissions intensity. In the “Tax Cut” and “Tax Hike” counterfactual simulations of this model, the results are nearly identical to the base model. A tax cut results in more replacement of older plants and higher emissions; a tax hike has the opposite effect. In contrast, while the “No Grand.” results are qualitatively similar, the magnitude of the difference between the counterfactual and the baseline simulation is much less pronounced. Removal of grandfathering results in a drop in emissions compared to the baseline, but only by about 5%, compared to the 60% drop in this counterfactual in the original model. This difference stems from the fact that utilities now have the option of meeting the pollution policy through reduced operating intensity, an option that was not possible when the policy was based on emissions intensity. In the “No Grand.” counterfactual, boilers that were grandfathered and are newly subject to the implicit tax can choose to reduce operating intensity rather than have to scrap and replace. They do so, and therefore there the reduction in emissions is much less pronounced. As argued earlier, this modeling choice is a poorer representation of actual policy than the implicit tax on emissions intensity, since NSPS have always been based on intensity.

5. Conclusion

This paper investigates how environmental regulations impact plant investment decisions and plant emissions. Regulations that are grandfathered can have short-term perverse effects. When grandfathered regulations are strengthened, grandfathered units may avoid being subject to the regulation by withholding investment in newer capital. If newer capital is less polluting, the reduction in new investment may increase emissions. Also, when grandfathered regulations are weakened, emissions can decrease in the short run, as grandfathered plants lose their valuable status that keeps them from investing. For electric power plants subject to the Clean Air Act, I find that removing grandfathering from the regulations in 1985 leads to a 60% decrease in emissions by 1995. A marginal change in the stringency of grandfathered regulations leads to a perverse effect on investment but not on emissions. This is because utilities can choose boilers of different emissions intensities.

Future research could extend the model to answer other questions. Here, investment only occurs on the extensive margin: plants make a discrete choice decision about whether or not to scrap. One could allow for a continuous level of investment. The policy here is constant and without uncertainty. Adding uncertainty allows for a policy that changes over time. Alternative policies can also be considered within this model. For example, to counteract the investment disincentive created by grandfathering, regulators could provide direct subsidies to investment. Finally, the fixed cost parameters F_c , F_{nc} , and G are constant over time and firms in the model. Capital investment costs are likely to change and react to macroeconomics conditions [18], so these adjustment cost variables may be modeled as endogenous.

When modeling the impacts of policy that contain grandfathering provisions, whether in environmental regulations or elsewhere, it is important to consider how these provisions affect the behavior of regulated firms or individuals and the potential for unintended consequences. A dynamic structural model as presented here can capture the regulatory impacts and the incentives created by grandfathering. The cross-sectional distribution of plant ages and grandfathering statuses, the relative profitability and emissions intensity of capital of different vintage, and the ability of plants to invest in abatement measures or choose among different types of capital all can affect how regulated entities respond to these policies.

Acknowledgments

I thank the National Science Foundation Graduate Research Fellowship program for financial assistance, the National Center for Supercomputing Applications for computer use and Natalie Ko and Nat Keohane for data. I also thank R.J. Briggs, Russ Cooper, Dean Corbae, Don Fullerton, Arik Levinson, Rob Stavins, two anonymous referees, and seminar participants at Texas, Maryland, Stanford, Princeton, Brown, Washington, Northwestern, Harvard, Georgetown, the University of California Energy Institute, the Econometric Society meetings in Minneapolis, and Camp Resources XIII in Wilmington, NC, for helpful comments.

References

- [1] Howard Gruenspecht, Differentiated regulation: the case of auto emissions standards, *American Economic Review* 81 (3) (1982) 328–331.
- [2] Randy Nelson, Tom Tietenberg, Michael Donihue, Differential environmental regulation: effects on electric utility capital turnover and emissions, *Review of Economics and Statistics* 75 (2) (1993) 368–373.
- [3] Arik Levinson, Grandfather regulations, new source bias, and state air toxics regulations, *Ecological Economics* 28 (2) (1999) 299–311.
- [4] Michael Maloney, Gordon Brady, Capital turnover and marketable pollution rights, *Journal of Law and Economics* 31 (1) (1988) 203–226.
- [5] John List, Daniel Millimet, Warren McHone, The unintended disincentive in the Clean Air Act, *Advances in Economic Analysis and Policy* 4 (2) (2004) (Article 2).
- [6] James Bushnell, Catherine Wolfram, Enforcement of Vintage-Differentiated Regulations: The Case of New Source Review (2008).
- [7] John Rust, Optimal replacement of GMC bus engines: an empirical model of Harold Zurcher, *Econometrica* 55 (5) (1987) 999–1033.
- [8] Don Fullerton, Garth Heutel, The general equilibrium incidence of environmental taxes, *Journal of Public Economics* 91 (3–4) (2007) 571–591.
- [9] Gilbert Metcalf, Environmental levies and distortionary taxation: pigou, taxation and pollution, *Journal of Public Economics* 87 (2) (2003) 313–322.
- [10] Lans Bovenberg, Ruud de Mooij, Environmental levies and distortionary taxation, *American Economic Review* 84 (4) (1994) 1085–1089.
- [11] Russell Cooper, John Haltiwanger, Laura Power, Machine replacement and the business cycle: lumps and bumps, *American Economic Review* 89 (4) (1999) 921–946.
- [12] Curtis Carlson, Dallas Burtraw, Maureen Cropper, Karen Palmer, Sulfur dioxide control by electric utilities: what are the gains from trade?, *Journal of Political Economy* 108 (6) (2000) 1292–1326.
- [13] Jerome Adda, Russell Cooper, *Dynamic Economics: Quantitative Methods and Applications*, MIT Press, Cambridge, 2003.
- [14] Christian Gourieroux, Alain Monfort, *Simulation-Based Econometric Models*, Oxford University Press, Oxford, 1996.
- [15] Bong-Soo Lee, Beth Fisher Ingram, Simulation estimation of time-series models, *Journal of Econometrics* 47 (2–3) (1991) 197–205.
- [16] Daniel McFadden, A method of simulated moments for estimation of discrete response models without numerical integration, *Econometrica* 57 (5) (1989) 995–1026.
- [17] Biewald, Bruce, White, David, Woolf, Tim, Ackerman, Frank, Moomaw, William. Grandfathering and Environmental Comparability: An Economic Analysis of Air Emission Regulations and Electricity Market Distortions. National Association of Regulatory Utility Commissioners, 1998.
- [18] Wald, Matthew, 2007. Costs surge for building power plants. *The New York Times*, July 10, 2007.
- [19] Jerome Adda, Russell Cooper, Balladurette and Juppettee: a discrete analysis of scrapping subsidies, *Journal of Political Economy* 108 (4) (2000) 778–806.
- [20] Mark Doms, Timothy Dunne, Capital adjustment patterns in manufacturing plants, *Review of Economic Dynamics* 1 (2) (1998) 409–429.
- [21] Geman, Ben, 2008. Energy law's emissions mandate largely exempts corn ethanol. *Greenwire*, March 28, 2008.
- [22] Paul Joskow, Contract duration and relationship-specific investments: empirical evidence from coal markets, *American Economic Review* 77 (1) (1987) 168–185.
- [23] Nathaniel Keohane, Erin Mansur, Andrey Voynuv, Averting regulatory enforcement: evidence from new source review, *Journal of Economics and Management Strategy* 18 (1) (2009) 75–104.
- [24] Oivind Nilsen, Fabio Schiantarelli, Zeros and lumps in investment: empirical evidence on irreversibilities and nonconvexities, *Review of Economics and Statistics* 85 (4) (2003) 1021–1037.
- [25] Office of Technology Assessment, *Acid Rain and Transported Air Pollution: Implications for Public Policy*, Office of Technology Assessment, Washington, DC, 1984.
- [26] Steven Shavell, On optimal legal change, past behavior, and grandfathering, *Journal of Legal Studies* 37 (1) (2008) 37–85.
- [27] Robert Stavins, Vintage differentiated environmental regulation, *Stanford Environmental Law Journal* 25 (1) (2006) 29–66.
- [28] Stephen. Ryan, *The Costs of Environmental Regulation in a Concentrated Industry*, Massachusetts Institute of Technology, Department of Economics, 2009.
- [29] Ian Lange, Allen Bellas, The 1990 clean air act and the implicit price of sulfur in coal, *The B.E. Journal of Economic Analysis & Policy (Topics)* 7 (1) (2007).