DISCUSSION PAPER

June 2015 ■ RFF DP 15-25

The Net Benefits of the Acid Rain Program

What Can We Learn from the Grand Policy Experiment?

H. Ron Chan, B. Andrew Chupp, Maureen L. Cropper, and Nicholas Z. Muller

This paper is one in a series of retrospective analyses from RFF's Regulatory Performance Initiative. www.rff.org/regulatoryperformance.

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

RESOURCES FOR THE FUTURE

The Net Benefits of the Acid Rain Program: What Can We Learn from the Grand Policy Experiment?

H. Ron Chan, B. Andrew Chupp, Maureen L. Cropper, and Nicholas Z. Muller

Abstract

This study quantifies the cost savings from the Acid Rain Program (ARP) compared with a command-and-control alternative and also examines the impact of trading under the ARP on health damages. To quantify cost savings, we compare compliance costs for non-NSPS (New Source Performance Standards) coal-fired Electricity Generating Units (EGUs) under the ARP with compliance costs under a uniform performance standard that achieves the same aggregate emissions. We do this for the year 2002, the third year of Phase II of the program. We find annual cost savings of approximately \$250 million (1995\$). To examine the health effects of trading, we compute the health damages associated with observed sulfur dioxide (SO_2) emissions from all units regulated under the ARP in 2002—approximately 10.2 million tons—and compare them with damages from a No-Trade counterfactual in which each unit emits SO_2 at a rate equal to its allocation of permits for the year 2002, plus any drawdown of its allowance bank. Damages under the No-Trade scenario are \$2.4 billion (2000\$) lower than under the ARP. This reflects the transfer of allowances from EGUs west of the Mississippi River to units in the eastern US with higher exposed populations.

Key Words: sulfur dioxide, acid rain, performance standards, health effects, pollution permits, cap and trade

Contents

The Net Benefits of the Acid Rain Program: What Can We Learn from the Grand Policy Experiment?

H. Ron Chan, B. Andrew Chupp, Maureen L. Cropper, and Nicholas Z. Muller*

1. Introduction

Economists have long advocated incentive-based systems of pollution control—in particular, marketable pollution permits—as a more efficient approach to environmental regulation than command and control. In theory, tradable pollution permits should achieve the least-cost solution to achieving a target emissions cap. In a competitive permit market, each source should equate its marginal cost of abatement to the price of a permit, thus guaranteeing that marginal abatement costs are equalized across sources. However, even in theory, pollution permits may not maximize the net benefits of the associated emissions reduction (Mendelsohn 1986; Muller and Mendelsohn 2009). A system of tradable permits may lead to higher damages than a uniform performance standard that achieves the same emissions target if a ton of pollution emitted by buyers of permits has higher marginal damages than a ton of pollution emitted by permit sellers (Mendelsohn 1986). A system of tradable permits may also fail to yield large cost savings relative to a uniform performance standard if other regulations prevent the permit market from reaching the least-cost solution to pollution abatement (Fowlie 2010). For both reasons, the net benefits of a pollution market relative to a uniform standard remain an empirical question.

In this paper, we compare the compliance costs and health damages of the Acid Rain Program (ARP), enacted under Title IV of the Clean Air Act Amendments of 1990, with the corresponding costs and health damages of a uniform performance standard that would have achieved the same aggregate emissions as achieved when the ARP was fully operational (i.e., during Phase II of the program). The Acid Rain Program, which sought to reduce sulfur dioxide (SO_2) emissions from electric utilities to half of their 1980 levels, is often cited as evidence that an emissions trading program can lower the costs of reducing pollution compared with a uniform performance standard (Ellerman et al. 2000; Stavins 1998). Yet there is no comprehensive, ex post evaluation of either the net benefits or the abatement cost savings of the ARP compared

* H. Ron Chan, University of Manchester; B. Andrew Chupp, Georgia Institute of Technology; Maureen L. Cropper, University of Maryland, Resources for the Future, and National Bureau of Economic Research (NBER); Nicholas Z. Muller, Middlebury College and NBER.

with those of an equally stringent policy that did not allow utilities to trade SO_2 allowances. We estimate a structural model of compliance behavior for coal-fired electricity generating units (EGUs) covered by the ARP and use the model to compute the cost savings achieved by the ARP compared with a uniform performance standard that achieves the same aggregate emissions cap. We also compute the health damages associated with the ARP and with a counterfactual no-trade scenario that results in the same aggregate emissions to compute the health impacts of allowance trading.

1.1. Previous Literature

Ex ante studies of the cost savings from allowance trading predicted large cost savings from the program compared with a uniform performance standard, especially in Phase II of the program. Phase I of the ARP, between 1995 and 1999, required the dirtiest 110 coal-fired power plants to reduce their emissions. Beginning in 2000, all EGUs greater than 25 megawatts (MW) were regulated by the program. Ex ante studies of the cost savings from emissions trading predicted much larger cost savings in Phase II of the program, in which all EGUs would participate, than in Phase I. Carlson et al. (2000) predicted cost savings from trading in Phase I of $250 million annually and Ellerman et al. (2000) savings of $360 million (US$1995) annually compared with a uniform performance standard. In contrast, annual Phase II savings were predicted to be $784 million (Carlson et al. 2000) and $1.92 billion (Ellerman et al. 2000).[1]

There is, however, no econometric study of the cost savings achieved by the ARP once the program was fully operational that is based on actual compliance data. Studies of the cost savings delivered by the ARP either are ex ante in nature (Carlson et al. 2000) or focus on Phase I of the program (Arimura 2002; Keohane 2007; Sotkiewicz and Holt 2005; Swinton 2002, 2004). Carlson et al. (2000) project cost savings based on marginal abatement cost (MAC) functions estimated using pre-ARP (1985–94) data. The MAC functions capture the cost of reducing SO_2 emissions only through fuel switching (i.e., substituting low- for high-sulfur coal), not through the installation of flue-gas desulfurization units (scrubbers). In calculating the gains from trade, Carlson et al. assume that no additional scrubbers will be built after 1995.[2] They

[1] EPA (1992) predicted cost savings of $9.6 billion to $13.8 billion over the period 1993–2010, or annualized savings of $689 million to $973 million (US$1990).

[2] In fact, 14 scrubbers were built at EGUs not covered by New Source Performance Standards between 1996 and 2002, the year of our study.

estimate the long-run cost savings from the ARP, compared with a uniform performance standard, by assuming that the ARP will achieve the least-cost solution to the SO_2 cap.

There is, however, no guarantee that allowance trading achieved the least-cost abatement solution. Previous studies suggest that this was not the case during Phase I of the program (Carlson et al. 2000; Sotkiewicz and Holt 2005; Swinton 2002, 2004). Several factors could have prevented electric utilities from reaching the least-cost solution: (1) utilities subject to regulation by Public Utilities Commissions (PUCs) could pass compliance costs on to ratepayers and therefore had no incentive to minimize costs (Sotkiewicz and Holt 2005; Cicala 2015); (2) the fact that PUCs allowed scrubbers to enter the rate base and thus earn a normal rate of return provided incentives to scrub rather than substitute low- for high-sulfur coal (Fullerton et al. 1997; Sotkiewicz and Holt 2005); and (3) uncertainty about the treatment of allowances in the rate base provided incentives to fuel switch rather than purchase allowances (Arimura 2002). The least-cost options for fuel switching were also prevented by regulators who encouraged the purchase of in-state coal (Cicala 2015) or by long-term coal contracts that might, in practice, be difficult to break. We examine the impact of these factors in Phase II once the ARP was fully operational.

There are also concerns that health damages after the ARP were higher than they would have been under a uniform performance standard (Henry et al. 2011). The reason is that, compared with a uniform standard, trading shifted emissions from low marginal abatement cost plants (sellers of permits) located in sparsely populated areas west of the Mississippi River to plants in more densely populated areas east of the Mississippi River (buyers of permits). This is supported by the map in Figure 1, which shows the difference in 2002 between $PM_{2.5}$ levels under the ARP and $PM_{2.5}$ levels that we estimate would have occurred had all EGUs subject to the ARP emitted at a rate equal to their initial allocations of allowances. The map suggests that trading increased $PM_{2.5}$ levels along the Eastern Seaboard, especially in densely populated areas in the Middle Atlantic states.

1.2. Our Approach

To measure cost savings from trading under the ARP, we use ex post data to model the long-run compliance behavior of coal-fired EGUs covered by the program. We focus on all coal-fired generating units not regulated under New Source Performance Standards (NSPS).[3] The main methods used to reduce SO_2 emissions are to purchase low-sulfur coal or install a flue-gas desulfurization unit (FGD). Our model is a mixed logit model of the choice of whether or not to install an FGD and what type of coal to buy, described by geographic location. This model allows us to predict compliance choices under the ARP and under a uniform performance standard (UPS) that achieves the same aggregate emissions as non-NSPS units emitted under the ARP. After estimating the model, compliance choices, compliance costs, and emissions are predicted for each EGU under the ARP and under our counterfactual scenario.

We estimate the cost savings from emissions trading to be between $250 million and $300 million (US$1995) per year, a much smaller estimate than that of Carlson et al. (2000), and a fraction of the cost savings forecast by EPA (1992). We attribute this in part to the failure of regulated units to pursue least-cost compliance options, although we do not find a significant difference between divested, publicly owned, and PUC-regulated EGUs in this regard.

To compare health damages under the ARP and a performance standard, we estimate pollution damages associated with emissions using AP2, an integrated assessment model that links emissions from each power plant to changes in ambient air quality, changes in population exposures to $PM_{2.5}$, and associated health effects. The model (Muller 2011), which is an updated version of the APEEP model (Muller and Mendelsohn 2009; Muller et al. 2011), uses the $PM_{2.5}$ mortality dose-response function estimated by Pope et al. (2002) and values changes in mortality risks using a $6 million (US$2000) value of a statistical life (VSL).

We estimate that in 2002, health damages associated with emissions from non-NSPS plants were approximately the same as they would have been had these plants been subject to a uniform performance standard. This is not surprising. Most non-NSPS plants are located east of the Mississippi River. Although emissions under the ARP and the counterfactual UPS occur in different places, the exposed populations are high in both cases.

This does not mean, however, that the ARP had no impact on health. In 2002, non-NSPS units emitted approximately 7.55 million tons of SO_2, over 2 million tons more than their allowance allocations for the year 2002 of 5.47 million tons. Two-thirds of these allowances

[3] Non-NSPS units were the target of the ARP. Units regulated under the NSPS were required to achieve an emissions rate at least as stringent as the ARP target of 1.2 pounds of SO_2 per million Btu (MMBtu). Non-NSPS units generated over 70% of the SO_2 emissions produced by EGUs in 2002, the year of our study.

were purchased from other EGUs. The health effects of the ARP depend on the location of sellers versus buyers of allowances. To capture the health impacts of trading, we estimate the health damages associated with the observed emissions of *all* units participating in the ARP and compare them with the damages that would have resulted had units emitted SO_2 at a rate determined by the initial distribution of allowances. We find that damages under the ARP exceeded damages under the no-trade counterfactual by $2.4 billion (US$2000) (1.8 percent of damages under the ARP). This is because under the ARP, NSPS units and noncoal units transferred or sold allowances to non-NSPS units. Sellers of allowances were more likely to be located in sparsely populated areas to the west of the Mississippi River, whereas buyers were located in the US Midwest and East.

How do these results compare with EPA's original Regulatory Impact Analysis (RIA) of the ARP? EPA's RIA (1992) estimated the cost savings from trading under the ARP but did not examine the benefits of the program. The RIA projected annual cost savings of $2.1 billion to $2.8 billion (US$1990) between 2000 and 2010.[4] These estimates, which assume that the allowance market would achieve the least-cost solution to the emissions cap, are significantly higher than our estimates. We discuss the reasons for these differences in detail below. EPA did not analyze the health benefits of reducing SO_2 emissions in its 1992 RIA.

The paper is organized as follows: Section 2 discusses the ARP and other regulations affecting SO_2 emissions from coal-fired power plants and describes compliance behavior in Phase II of the ARP. We present our cost model and estimation results in Section 3. In Section 4 we simulate compliance behavior under a uniform performance standard and compare compliance costs and emissions under the standard and the ARP. These emissions estimates are used in Section 5 to estimate health damages under the ARP and under a uniform performance standard affecting all non-NSPS EGUs. We also estimate the damages caused by all units covered by the ARP and contrast them with a scenario in which all units emit SO_2 at a rate determined by the initial distribution of allowances. Section 6 discusses the policy implications of our results.

---

[4] These are annual cost savings for the period 1993 to 2010. They are approximately $800 million to $1.2 billion (US$1995).

2. Background

2.1. Title IV and Other SO_2 Regulations Facing Coal-Fired Power Plants

The objective of the Acid Rain Program was to reduce sulfur dioxide emissions from fossil-fueled power plants in the United States by 50 percent from 1980 levels. The program was implemented in two phases: In Phase I (1995–99), the most polluting 263 generating units (termed “Table A” units) were required to participate. Table A units were allocated allowances equal to an emissions rate of 2.5 pounds of SO_2 per million Btu (MMBtu) of heat input times the unit’s heat rate in the 1985–87 reference period. Units were also allowed to voluntarily enroll in Phase I, either as substitutes for Table A units or to compensate for reductions in output at Table A units.[5] In Phase II, beginning in 2000, the program was extended to all generating units with a capacity exceeding 25 megawatts, approximately 1,100 coal-fired units. Units were allocated annual permits equal to the product of the target emissions rate—1.2 pounds of SO_2 per MMBtu—and heat input during 1985–87. Under the ARP, units were free to trade permits within and across states. They were also allowed to bank permits for future use but could not borrow permits from future years.

Sulfur dioxide emissions from coal-fired power plants were also regulated under the 1970 Clean Air Act (CAA) and 1977 Clean Air Act Amendments (CAAA). Under the 1970 CAA, states were required to formulate state implementation plans (SIPs) to guarantee that counties within the state did not violate the National Ambient Air Quality Standards (NAAQS). This involved setting emissions limits for existing stationary sources within each state, including power plants. The emissions limits imposed on SO_2 emissions from coal-fired power plants by state and local governments, which we incorporate into our analysis, were sometimes more stringent than the 1.2 pounds of SO_2 per MMBtu of heat input targeted under the ARP.[6] The 1970 CAA also imposed New Source Performance Standards (NSPS) on newly constructed stationary sources, including power plants. Plants built between 1971 and September 1977 were required to reduce their SO_2 emissions to 1.2 pounds per MMBtu. The NSPS enacted under the

[5] As Ellerman et al. (2000) note, “substitution and compensation” units tended to be units with low marginal abatement costs that were enrolled to increase the number of allowances their owners received. Over 150 EGUs were enrolled as “substitution and compensation” units in the first three years of the ARP, with 138 units enrolled in all three years.

[6] Trading under the ARP could not violate the NAAQS.

1977 CAAA in effect required coal-fired power plants built after September 1977 to install scrubbers.

The ARP was followed by attempts to further curb SO_2 emissions from power plants. In December 2003, EPA issued a draft of the Clean Air Interstate Rule (CAIR). Limited to the eastern United States, including 27 states and the District of Columbia, CAIR aimed to mitigate the damages of airborne pollutants that disperse across state borders. CAIR mandated a cap-and-trade system of emissions control for sulfur dioxide and nitrogen oxide emissions, with a goal of reducing SO_2 emissions by 57 percent from ARP levels. Although CAIR was later vacated by the District of Columbia Circuit Court and replaced by the Cross-State Air Pollution Rule (CSAPR),[7] it was clear after December 2003 that EPA aimed to regulate SO_2 emissions from power plants more stringently than under the ARP. We view this as a change in the regulatory regime that effectively signaled the end of the ARP (Schmalensee and Stavins 2012).

2.2. Compliance in Phase II of the Acid Rain Program

Our analysis focuses on the time period when the ARP was fully operational—when all coal-fired EGUs were covered by the program—but before plans were announced to more stringently regulate SO_2 emissions. We focus on the year 2002, the third year of Phase II.[8] In 2002, 1,075 coal-fired generating units were regulated under the ARP (see Table 1). These included 378 units that had participated in Phase I of the program and 697 units that participated only in Phase II of the program. Of the latter, 487 units were not covered by NSPS, while 210 were regulated under the NSPS as well as the ARP. As Table 1 makes clear, units regulated under the NSPS were, on average, emitting at a rate less than half of the target 1.2 pounds of SO_2 per MMBtu. Half of these units had installed scrubbers, and the remainder were burning "compliance coal"—coal that would result in emissions of 1.2 pounds per MMBtu or less. Because the abatement decisions of NSPS units were determined by regulations that preceded the ARP, we exclude them in modeling compliance behavior under the ARP. We also omit the

[7] The DC Circuit Court vacated the Clean Air Interstate Rule, declaring that the system of regional caps was fundamentally flawed. In December 2008, the DC Circuit Court remanded the vacatur, allowing CAIR to remain in place until a new policy consistent with the goals of CAIR could be formulated as a replacement. In July 2011, EPA proposed the Cross-State Air Pollution Rule (CSAPR).

[8] After plans for CAIR were announced in 2003, allowance prices rose sharply, signaling the anticipation of a new regulatory regime (Schmalensee and Stavins 2012).

NSPS units from our simulations of cost savings, effectively assuming that the behavior of the NSPS units was the same under the ARP as under a uniform performance standard.

The compliance choices of remaining units consisted of installing scrubbers, burning low-sulfur coal, or using allowances in excess of those allocated for the year 2002. By 2002, 100 (12 percent) of the non-NSPS units had installed scrubbers. An additional 25 percent of units achieved compliance with the ARP by burning low-sulfur coal. Remaining units used a combination of blending low-sulfur coal with higher-sulfur coal, using banked allowances, or purchasing additional allowances. Banked allowances covered 700,000 tons of emissions. Approximately 38 percent of emissions in 2002 were covered by purchased allowances.[9]

Figure 2 illustrates the geographic pattern of compliance choices. As Figure 2(a) clearly indicates, the percentage of units burning low-sulfur coal is highest in states closest to the Powder River Basin, for which the cost of transporting coal from Wyoming and Utah is much lower than for units east of the Mississippi River. Heterogeneity in the costs of compliance through fuel switching is the main source of cost savings in the allowance market and is reflected in the pattern of allowance trades implied by Figure 2(b). Figure 2(b) shows 2002 SO_2 emissions in excess of 2002 allowances, by state. The map suggests that units east of the Mississippi River were purchasing allowances from units west of the Mississippi.

Table 2 describes compliance according to a unit's status under electricity sector deregulation in 2002. Units may be divested (owned by independent power producers), PUC-regulated (investor-owned utilities whose rates were set by PUCs), or publicly owned. The table indicates the percentage of units that scrubbed and the percentage that used exclusively low-sulfur or high-sulfur coal in 2002. Remaining units blended coal of various sulfur contents. Figures 3(a) and 3(b) show the location of units by regulatory status. We focus on the compliance options chosen by non-NSPS units, which are modeled in Section 3. The percentage of non-NSPS units scrubbing emissions does not differ significantly by regulatory status, although it is slightly higher for divested units (11.2 percent) and PUC-regulated units (12.2 percent) than for publicly owned units (9.4 percent). Most non-NSPS divested units are located east of the Mississippi River, with the majority in the Middle Atlantic states, New England, or Ohio—that is, far from low-sulfur coal. Not surprisingly, divested units were much more likely

[9] We calculate this as the difference between actual emissions and (2002 permits plus banked allowances held at the beginning of 2002), divided by actual emissions.

to use high-sulfur coal and much less likely to fuel switch than either publicly owned or PUC-regulated units.[10] On average, divested and PUC-regulated units were net purchasers of allowances, while publicly owned units were net sellers.

3. Modeling Compliance Behavior under the ARP

3.1. A Model of Compliance Choice

We model long-run compliance behavior of non-NSPS EGUs under the ARP using a static model of the choice of which type of coal to purchase and whether or not to install an FGD. A static model has the virtue of simplicity and allows us to focus on the long-run gains from allowance trading. We model the choice of which type of coal to purchase and whether to scrub as a discrete choice: each EGU must choose which type of coal to buy, indexed by the region from which coal is purchased, crossed with the decision to scrub or not to scrub. We assume that this choice is made to minimize weighted compliance costs, where the weights on different components of compliance costs are a function of plant characteristics, including the plant's regulatory status in the electricity market and whether incentives were provided for the purchase of in-state coal. The choice of coal bought is also subject to state and local emissions standards: types of coal that would violate these standards are eliminated from the choice set.

Compliance costs consist of four components: (1) the direct costs of purchasing coal and scrubbing; (2) the operating costs associated with the ash content of coal; (3) the cost of SO_2 emissions; and (4) the cost of retrofitting the boiler to burn coal with lower sulfur content than the boiler was designed to burn. While the first category of costs can be estimated for each compliance option, the last three are inferred from the coefficients of the cost model. Coal costs are the delivered cost of coal to the unit; we observe this for the chosen option and estimate it for options not chosen, as described below. Scrubbing costs are handled similarly. The operating costs of burning coal will vary with its ash content; hence, we include this characteristic of coal in the cost function and use its coefficient to infer its impact on costs. SO_2 emissions are, by definition, the product of the sulfur content of the coal burned times the fraction of emissions not

[10] This is consistent with results reported by Cicala (2015), who estimates that divested units were 7 percentage points less likely to install additional scrubbers after divestiture than nondivested units. Cicala's analysis covers the period from 1990 through 2009 and indicates that the biggest difference between divested and nondivested units occurred after 2002. We focus on compliance choices made by 2002. Only three of the scrubbers installed in divested non-NSPS units were installed after divestiture.

removed by scrubbing. The coefficient on this component of costs represents the shadow price of emissions, which we compare to actual allowance price.

We include terms in the cost function to indicate whether a particular type of coal requires retrofitting the unit's boiler. The coefficients on these terms capture the cost of retrofitting a boiler to use Powder River Basin (PRB) coal, as well as the retrofitting cost if the unit changes the type of coal burned to comply with the ARP. PRB coal, which is the primary source of low-sulfur coal, has much lower heat content than high-sulfur coal. To burn PRB coal efficiently, boilers designed for high-sulfur coal must be retrofitted. Our choice model estimates this retrofitting cost as a function of boiler age. Changing the type of coal burned to comply with the ARP (e.g., from high- to medium-sulfur coal) may also incur costs associated with boiler retrofitting or with the termination of historic contractual arrangements. Coal procurement data from the early 1980s are used to identify units that have changed their sources of coal to comply with the ARP. We allow adjustment costs to vary with boiler age.

In modeling the compliance decision, we argue that the output of each unit can be treated as fixed: coal-fired units are base-load units, and according to EIA Form 767, few units altered their output as a means of complying with the ARP. Following the literature, we treat electricity production as proportional to heat rate. This allows us to write the cost function as cost per MMBtu of heat input. Specifically, we assume that for each EGU, the compliance option *j* is chosen that minimizes (1) subject to the constraint that the EGU not violate state and local emissions standards $\overline{SULFUR_i}$,which may limit SO_2 emissions per MMBtu (equation (2)).[11]

$$\min\{j\}\, C_i(j) = \beta_i^f COALPRICE_i(j) + \beta_i^z SCRUBCOST_i(j) + \beta^a ASH(j) + \beta^P SO_2EMISSIONS_i + PRB_j\left(\beta_{0i}^l + \beta_1^l AGE_i\right) + MODIFY_{ij}(\beta_{0i}^M + \beta_1^M AGE_i) + \varepsilon_i(j) \tag{1}$$

where i = 1,2,…,N (units), j = 1,2,…,J (compliance choices), and

$C_i(j) =$ unit compliance cost, in cents per MMBtu
$COALPRICE_i$ = delivered coal price, in cents per MMBtu
$SCRUBCOST_i$ = projected scrubbing cost, in cents per MMBtu
$AGE_i =$ age of the unit, calculated using the initial operating date
$ASH(j); SULFUR(j) =$ ash and sulfur content of coal, in pounds per MMBtu

[11] We treat these standards as exogenous to the ARP. Most were imposed in the 1970s and have not been modified since.

$SULFUR\ EMISSIONS_i(j) = SULFUR_i(j) \times (1 - SCRUB(j)\theta)$, $\theta = 0.85$
PRB_j = 1 if coal is from the Powder River Basin
$MODIFY_{ij}$= 1 if coal choice is different from the coal purchased in 1982
$\varepsilon_i(j)$ = unobserved costs specific to option *j*

subject to

$$(1 - \theta(j))SULFUR(j) \leq \overline{SULFUR_\iota}. \tag{2}$$

To incorporate the effect of electricity sector regulations on compliance choices, we allow the coefficients in equation (1) to be functions of the EGU's status under electricity sector regulation. Specifically, we allow these coefficients to vary according to whether the unit was divested by 2002, publicly owned, or an investor-owned utility (IOU) regulated by a PUC.

3.2. Estimation of the Model

We estimate our model of compliance behavior using data for non-NSPS units in 2002.[12] We argue that most units had achieved their optimal compliance strategy under the ARP by this time.[13] It is also the case that at the end of 2003, announcement of the Clean Air Interstate Rule (CAIR) signaled a sharp change in the regulatory regime. This was reflected in the price of allowances, which began to rise sharply in January 2004. Many EGUs installed scrubbers between 2006 and 2010; however, this was in response to signals that EPA intended to drastically reduce the SO_2 emissions from power plants below the target under the ARP.

We estimate choice of compliance option as a mixed logit model, using unit-level data. Specifically, we treat $\{\varepsilon_i(j)\}$ as independently and identically distributed with a Type I extreme value distribution. We allow the coefficients on *SCRUBCOST*, *PRB*, and *MODIFY* to be normally distributed with mean vector B and diagonal variance-covariance matrix Σ. The likelihood function is given by

$$L(B, \Sigma) = \sum_i \sum_j 1(Y_i = \ j) \int_{-\infty}^{\infty} \frac{exp(-C_i(j;\beta,X_i))}{\sum_{j\prime}^{J} exp(-C_i(j\prime;\beta,X_i))} f(\beta|B,\Sigma)d\beta \tag{3}$$

[12] We record whether the unit had a scrubber in operation in 2002. When describing the coal purchasing decision, we average purchases over 2000–2002, since coal purchased in previous years could be burned in 2002.

[13] It is also the case that minemouth prices for the major coal basins are stable over the 1994–2004 period (EIA), as are allowance prices.

where Y_i, X_i are the observed choices and vector of covariates for unit *i*.

Estimation of the model requires that we define the choice set for each EGU. We model coal choice as the purchase of coal from one of the six major coal basins (North, Central, and South Appalachian; Illinois; Powder River; and Uinta). The North Appalachian and Illinois basins are each subdivided into two regions based on the sulfur content of coal. The purchase decision is modeled as buying 100 percent of the unit's coal from one of the eight regions or buying half of the unit's coal from each of two regions.[14] These 36 coal purchase options are crossed with the decision to scrub. If a compliance option would violate state or local emissions constraints, the option is dropped from the unit's choice set.

Table 3 describes the sulfur content of coal in each of the eight coal regions. There is clearly considerable variation in sulfur content within each region. To better characterize coal choice, we use a nesting procedure to refine the characteristics of coal purchased by unit *i* in region *j*. We initially estimate the parameter vector β (β^0), using the average characteristics of coal in each region for all units. Then, conditional on β^0, we determine for each unit the county within each region that minimizes compliance costs.[15] We then replace the *COALPRICE*, *ASH*, and *SULFUR* content of coal for unit *i* in region *j* with the characteristics of the cost-minimizing choice, for all *i* and *j*. The likelihood function in equation (3) is maximized using the updated coal characteristics, and the procedure is repeated until the parameter vector β converges.[16]

Implementation of this procedure requires estimating the delivered cost of coal from each county in each coal region to each EGU. Delivered coal prices, together with information on the ash and sulfur content of coal purchased and the distance of the unit from the mine, are used to calculate minemouth prices for all counties. Data on transport costs, together with minemouth prices, are used to estimate the delivered cost of coal for each unit. Imputed delivered coal prices are summarized in Table 4. The table makes clear the cost advantage enjoyed by plants in the West and Midwest: for these plants, low-sulfur coal from the PRB is the cheapest coal to purchase; for plants in the South and Northeast, high-sulfur coal from the North Appalachian basin is cheaper. There is also considerable heterogeneity in coal prices within regions, which aids in identifying the coefficients of the compliance cost function.

---

14 Only 3% of units buy coal from more than two regions.

15 In solving this problem, the error term in (1) is treated as zero.

16 This procedure is described more fully in the Appendix.

For units that do not install FGDs, the cost of installing and operating scrubbers are estimated as a function of plant and unit characteristics (see Chan 2013), following Lange and Bellas (2007). In general, the costs of retrofitting a unit with a scrubber increase with the age and size of the unit; operating costs increase with years since the scrubber has been installed, removal rate, and operating hours.

Our simulation of cost savings under the ARP is based on 761 of the 838 non-NSPS coal-fired generating units in Table 1.[17] We exclude units that installed scrubbers before 1988 from estimation of the model but include them in the simulations reported in Section 4, with the constraint that a scrubber option must be chosen. Table 5 summarizes the variables entering the compliance cost model.

3.3. Estimation Results

Table 6 presents the parameter estimates for the cost model. Models (1) and (2) interact coal price with a dummy variable that indicates whether coal is sourced in-state (Coal Price × In-State) and a dummy variable that indicates coal sourced from nearby mines (Coal Price × Minemouth). Models (3) through (5) add interactions between regulatory status and various components of the cost function: the cost of scrubbing, the in-state discount (Coal Price × In-State), and whether coal comes from the Powder River Basin (PRB). Regulatory status is also interacted with SO_2 emissions to allow the shadow price of emissions to vary by regulatory status. The coefficients on *PRB* and *MODIFY* are random in all models; the coefficient on *SCRUBCOST* is random in Models (2), (4), and (5). A positive coefficient estimate implies that cost is increasing in that argument. Scaling each coefficient by the coefficient on coal price converts it to monetary terms.

In all models, cost is increasing in coal price, SO_2 emissions, ash, and scrubbing cost. There is a 20 percent discount for minemouth coal that is stable across models. Two important components of unobserved costs—retrofitting costs for PRB coal and general modifications—both show statistically significant mean effects on compliance costs, which vary with the age of the boiler. Evaluated at the mean of the observations, average annualized cost for using PRB is about 31 cents per MMBtu, while the general retrofitting cost is about 13 cents per MMBtu.

---

[17] Of the 838 non-NSPS units in Table 1, 77 are not used to estimate the model: 36 have no data on coal purchases, 26 purchase coal primarily outside of the eight coal regions described above, and 15 changed from non-NSPS to NSPS status shortly after 2002.

These cost components also show large variation across units, as implied by their statistically significant standard deviations.

Models (3) through (5) allow for interactions between regulatory status and components of the cost function. Models (1) and (2) suggest that the cost of coal mined in the same state as the EGU (in-state coal) receives a significantly lower weight in the cost function, although the magnitude of this effect is less than 5 percent. Models (4) and (5) indicate that investor-owned units regulated by PUCs assign a higher discount to in-state coal than either publicly owned or divested units. This result agrees with Cicala (2015), who finds that divested power plants were less likely to purchase in-state coal than nondivested plants (see also Chan et al. 2013).

Models (3) through (5) can be used to calculate the shadow price of SO_2 emissions by regulatory status, which can in turn be compared with observed allowance price.[18] Allowance prices ranged from \$150 to \$200/ton of SO_2 over the period of our study. Model (3) implies that the average shadow price attached to SO_2 emissions was lowest for publicly owned units (\$118/ton) and higher for PUC-regulated units (\$174/ton) and divested units (\$161/ton). This is consistent with the fact that PUC-regulated and divested units, many of which are located along the Eastern Seaboard, are far away from low-sulfur coal (see Figure 3(a) and Table 4) and purchased allowances as a method of compliance rather than switching to low-sulfur coal. Publicly owned units were, on average, net sellers of allowances.

Overall, our models do not suggest that divested units behaved significantly differently from IOUs regulated by PUCs. This may seem surprising in view of results obtained by Cicala (2015) and Fowlie (2010), which suggest that divested plants were less likely to install capital equipment as a means of complying with pollution regulations and, in the case of SO_2, more likely to switch to low-sulfur coal than nondivested plants. It should be kept in mind, however, that in most cases, the decision to install a scrubber that was functioning in 2002 at a divested plant was made prior to divestiture: only three scrubbers were installed at divested plants after divestiture (see footnote 10). Cicala (2015) finds that the biggest divergence in methods used by divested versus nondivested plants to reduce SO_2 emissions occurred after the time of our study.

[18]The shadow price of SO_2 is calculated by dividing the coefficient on SO_2 emissions by the coefficient on coal price to scale the parameter to a value in cents. Dividing by 100 gives the price in dollars. This result is multiplied by 2,000 to convert from pounds to tons and divided by 2 to convert S to SO_2.

4. Simulation Results

4.1. Predicting Compliance Choices, Costs, and Emissions

To estimate cost savings from the ARP for non-NSPS units, we predict compliance choices under the ARP and under a uniform performance standard. We calculate the cost of compliance under each regime (per MMBtu) as the sum of the unweighted fuel price and scrubbing cost associated with the option predicted to be chosen, together with the estimated costs of retrofitting the boiler and the estimated operating cost associated with the ash content of the coal burned. Total compliance costs are calculated using average heat input from 2000 to 2002. The difference between compliance costs under the ARP and the uniform performance standard represent the estimated cost savings from the ARP.

Predicted emissions are based on the sulfur content of the coal chosen and the decision whether or not to scrub, as well as the average heat input observed in the data. The sulfur content of the coal type predicted to be chosen yields the emissions rate if no scrubber is installed. If a scrubber is installed, we assume that it removes 85 percent of emissions, which is the average observed removal rate in the data. The emissions rate is multiplied by the heat input used to give predicted emissions in tons.

To predict compliance choices under the ARP and under the uniform performance standard, we use *conditional* distributions of random coefficients and logit error terms. The conditional means of the logit error terms capture idiosyncratic components of costs. We assume these unobserved cost components are permanent and include them in evaluating the counterfactual policy. We estimate the conditional mean of $\varepsilon_i(j)$ for all i and j using 3,000 shuffled Halton draws from the error distribution for each i and j and selecting the draws that lead to the highest predicted probability that the observed compliance choices are chosen under the ARP (see Appendix). Conditional means of the random coefficient on *PRB, SCRUBCOST*, and *MODIFY* are computed similarly. These conditional means are used in the calculation of compliance costs, which are given by equation (4):

$$COMPC_i(j) = COALPRICE_i(j) + SCRUBCOST_i(j) + PRB_j\left(\boldsymbol{E_i}\tilde{\beta}_0^l + \tilde{\beta}_1^l AGE_i\right) + \tilde{\beta}^A ASH(j) + MODIFY_{ij}(\boldsymbol{E_i}\tilde{\beta}_0^M + \beta_1^M AGE_i) + \boldsymbol{E_i}\tilde{\varepsilon}(j) \tag{4}$$

where $\tilde{\beta} = \beta/\beta^f$.

To simulate choices under the uniform performance standard, the permit price component is removed from the cost function (i.e., β^P is set = 0), and a uniform emissions standard is added as an additional constraint to the choice problem. Local emissions standards are still in effect in the counterfactual. The level of the uniform standard is adjusted until aggregate emissions in the counterfactual are equal to those in the ARP (see Appendix for details).

4.2. Simulation Results

In simulating behavior under the ARP and a uniform performance standard, we focus on Models (1) through (4) of Table 6.[19] Table 7 shows predicted compliance choices under the ARP and the uniform performance standard for the 761 units used in our analysis. Because we use conditional means of the error terms and random coefficients for each unit, compliance choices are predicted perfectly under the ARP.[20] This does not, however, imply that emissions are predicted perfectly, due to the heterogeneity of the S content of coal within a basin. Predicted emissions under the ARP vary across the four models in the table. The table also shows predicted emissions under the UPS for each model in Table 6, and associated compliance cost savings, relative to the ARP.

Predicted emissions under the ARP for the 761 non-NSPS units vary from one model to another but are, in the aggregate, within 2 percent of monitored emissions for these units in 2002 (7.094 million tons). The uniform standard needed to achieve the same aggregate emissions as emissions predicted under the ARP ranges from 2.08 to 2.21 pounds of SO_2 per MMBtu. (When weighted by heat input, the UPS is between 1.31 and 1.34 pounds of SO_2 per MMBtu.) This standard is less stringent than the cap implied by the 1.2 pounds of SO_2 per MMBtu. Note from Table 1 that emissions of non-NSPS EGUs in 2002 are 38 percent higher than allocated permits; hence, the relevant cap should be higher.

Figure 4 compares predicted emissions rates under the ARP for Model (2) of Table 6 with the corresponding uniform standard. The 206 units that are above the standard under the ARP must reduce their emissions. Most do so by switching to coal with lower sulfur content than chosen under the ARP, which increases compliance costs. The cost savings achieved by the ARP compared with the UPS reflects the cost of these units moving below the standard.

---

[19] Model (5), which includes a complete set of interactions between regulatory status and components of the cost function, performs no better than Model (4) in predicting compliance choices. A likelihood ratio test fails to reject the null hypothesis that the coefficients of the additional interaction terms in Model (5) are significantly different from zero.

[20] That is, the choice of coal basin and whether a scrubber is installed are predicted perfectly.

Compliance costs under the ARP are estimated to be between $250 million and $300 million (US$1995) lower than under the uniform performance standard, significantly lower than previous estimates. Carlson et al. (2000), in comparing the ARP with a uniform performance standard, assume that the ARP will achieve the least-cost solution to the emissions cap. When we compare for each EGU the cost of chosen compliance option under the ARP with the least-cost method *of achieving their chosen emissions rate*, we find that the least-cost option was not chosen by 23 percent of units.[21] This is illustrated in Figure 5, which plots the cost options available to a particular EGU, using results from Model (2) of Table 6. The options pictured in blue are the least-cost options available. The red dot is the option selected. We approximate the cost of emissions in the neighborhood of the chosen option (the gray dots) and calculate the difference in cents per MMBtu between the chosen option and the least-cost option.

These results are summarized in Table 8. On average, 22.5 percent of units that installed scrubbers, and 22.8 percent of units that did not, spent more than the minimum cost to reach their chosen emissions rate.[22]Approximately 18 percent of divested and publicly owned units were operating in excess of minimum cost per MMBtu, while 22 percent of PUC-regulated and 31 percent of publicly owned EGUs were operating above the minimum cost necessary to achieve their chosen emissions rate. The sum of excess costs in Table 8 totals $206 million (US$1995). We emphasize that we have not solved for the global least-cost solution to the 7.2-million-ton cap; nonetheless, our calculations are suggestive of inefficiencies in compliance, possibly because of long-term coal contracts or lack of incentives to cost-minimize.

5. Estimating Health Damages

To compare net benefits under the ARP and the uniform performance standard, we must compute damages under each program. The damages due to SO_2 emissions produced by EGUs are estimated using AP2 (Muller 2011, 2012), a stochastic integrated assessment model that links emissions to ambient concentrations of air pollutants and ambient concentrations to pollution damages. AP2 is an updated version of APEEP, which has been used extensively in prior

[21] It is also the case that we are comparing compliance costs under the two policies only for non-NSPS units; that is, we assume that NSPS units would behave the same under both policies. Carlson et al. (2000) base their estimate of cost savings on all nonscrubbed units.

[22] This does not imply that all units that installed scrubbers under the ARP should have done so. We do not solve for the global least-cost solution to the cap.

research.[23] The air quality model in AP2 employs a source-receptor matrix in which each cell ($T_{i,j}$) in the matrix represents the change in ambient concentrations of $PM_{2.5}$ in location (j) due to a one-ton increase from source (i). The source-receptor matrices capture atmospheric processes that link emissions of precursor species (NO_x, SO_2) to resulting ambient concentrations of secondary pollutants. Significantly for our study, emissions of SO_2 are connected to concentrations of ammonium sulfate, an important constituent of $PM_{2.5}$.

In studies of the benefits of the ARP, health benefits constitute the majority of monetized benefits, and it is on those that we focus. AP2 links ambient concentrations of $PM_{2.5}$ to morbidity and mortality using concentration-response functions from the literature. These are combined with county-level population inventories provided by the US Census and baseline incidence rates to calculate health risks.[24] As in previous studies, adult mortality constitutes the most important health risk associated with $PM_{2.5}$ exposure. This study uses results from Pope et al. (2002) to link $PM_{2.5}$ to adult mortality. A recent meta-analysis (Roman et al. 2008) is used in a sensitivity analysis.[25]

Concentration response functions translate exposures, by county and age group, into changes in mortality risk. We value these risks using a VSL of $6 million (US$2000).[26] In the default modeling setup, the $6 million VSL is applied uniformly to all exposed populations. In a sensitivity analysis, the value of a statistical life-year (VSLY) approach is used. This strategy relies on detailed life-expectancy information to tabulate the number of expected life-years remaining for each population age cohort. Changes in life-years remaining due to $PM_{2.5}$ exposure are valued at $200,000 per life-year. This approach places a higher value on mortality risks faced by younger populations, since such age groups have more life-years remaining, in expectation.

[23] APEEP has been used in Muller and Mendelsohn (2007, 2009, 2012); NRC (2010); Muller et al. (2011); and Henry et al. (2011).

[24] Populations are coded according to 19 age groups, since baseline incidence rates (especially death rates) vary considerably according to age of the population.

[25] Pope et al. (2002) forms the basis for benefit estimates in the first prospective study of the 1990 CAAA (EPA 1999). Roman et al. (2008) was used in the second prospective study (EPA 2011).

[26] This is approximately equal to EPA's value, $4.8 million (US$1990), adjusted for inflation. In a sensitivity analysis, the $2 million (US$2000) VSL reported by Mrozek and Taylor (2002) is used.

5.1. Calculation of Damages from Coal-Fired Power Plants

For each policy simulation, AP2 processes baseline emissions through the air quality model source-receptor matrices to estimate baseline $PM_{2.5}$ concentrations, exposures, physical effects, and damages. All baseline emissions (except for SO_2 produced by electric generating units) are provided by EPA's National Emission Inventory (NEI) for 2002.[27] Then, for a particular policy scenario, SO_2 emissions from EGUs generated by the cost model are matched by EGU to AP2, and concentrations, exposures, physical effects, and damages are reestimated.[28] The change in damage due to the simultaneous change in EGU emissions is tabulated (1) in total, across all receptor counties; and (2) by county, to explore spatial patterns in the change in emissions, air quality, and impacts due to the different policy scenarios modeled in this paper.

5.2. Damages under the ARP and under Two Alternative Scenarios

We examine damages under the ARP and under two alternative scenarios. In Scenario 1, we estimate damages for the 761 non-NSPS coal-fired generating units used to estimate the cost savings from the ARP. As noted above, these units generated 70 percent of SO_2 emissions from coal-fired power plants in 2002. We compare the damages associated with these units under the ARP and under the uniform performance standard that yields the same aggregate emissions, using Model (2) from Table 6. As noted above, the cap in Scenario 1 reflects trading under the ARP: for non-NSPS units, the emissions cap exceeds allocated emissions for the year 2002 by 38 percent; two-thirds of this difference was covered by allowances purchased through trading.

To examine the impact of trading on health damages, we construct Scenario 2, which we term the No-Trade Scenario. Scenario 2 includes all EGUs covered by the ARP and forces them to emit at the rate prescribed by their initial allocation of allowances, plus any drawdowns of their allowance banks observed in 2002. The impetus of this comparison is to isolate the impact of trading per se, rather than compliance with a uniform performance standard, on health damages.[29] In order to construct the no-trading counterfactual, we begin by determining the quantity of allowances allocated to each unit for the year 2002. To this we add the difference

---

[27] These emissions are allocated by county of location and height of release into AP2. All non-EGU emissions for the coterminous United States are included in AP2.

[28] Note that this approach differs from the way AP2 has been applied in most prior research. That is, the model is *not* used to compute marginal ($/ton) damage. Rather, EGU emissions are changed simultaneously in order to reflect the concurrent emissions change that would have occurred with policy implementation.

[29] We do not compute cost information for this scenario. Our cost model applies only to non-NSPS units.

between each unit's permit bank in 2002 and in 2003, only if the unit was actually drawing down the bank. This difference represents permits drawn from the bank in 2002 in order to allow additional emissions in 2002. Thus, total emissions from a particular EGU, under a no-trade policy, are calculated as the allocation plus permits used from the bank. The ARP version of this scenario uses the actual emissions from each unit under the ARP to calculate damages. Under both the ARP and the no-trade counterfactual, aggregate SO_2 emissions are 10.2 million tons.

5.2.1. Scenario 1: Damages at Non-NSPS Units under the ARP versus a Uniform Performance Standard

Table 9 reports the difference in damages due to SO_2 emissions under the ARP and the uniform performance standard counterfactual simulation (Scenario 1) using Model (2) of Table 6.[30] Damages under the ARP amount to $103.44 billion in 2002. For the uniform performance standard, damages are slightly higher, at $103.45 billion. This amounts to a divergence of only $9 million (US$2000)—0.01 percent of damages under the ARP—with damages higher under the UPS. In Models (1), (3), and (4) of Table 6, damages are also higher under the UPS (see Appendix Table A.1), but the differences are small. We focus on Model (2) because it fits the data well (see Table 6) and because its predictions of damages under the ARP are very close to damages using observed emissions from the CEMS database. Models (3) and (4) of Table 6 underpredict damages based on observed emissions by about 2 percent; Model (2) comes within 0.2 percent of damages estimated using observed emissions.

We conclude based on our simulations that for non-NSPS plants, health damages under the ARP were no greater in the aggregate than they would have been under a counterfactual UPS, implemented in 2002. This suggests that for these plants, the ARP delivered a small cost savings with no increase in health damages.

5.2.2. Scenario 2: A No-Trade Counterfactual

Table 10 reports the difference between damages under the ARP and the No-Trade counterfactual. The table indicates that trading facilitated by the ARP increased adverse impacts by approximately 1.8 percent, or in absolute terms, by $2.44 billion (US$2000). We would expect units facing relatively high marginal abatement costs, such as those in the eastern United States farther from low-sulfur coal, to purchase permits under the ARP and emit more than their

[30] Scenario 1 damages for Models (1) through (4) of Table 6 are reported in Appendix Table A.1.

initial allocations, while those incurring lower marginal abatement costs would sell permits. This is suggested by Figure 2(a) and borne out by Figure 1, which shows modeled $PM_{2.5}$ concentrations attributable to actual CEMS emissions minus $PM_{2.5}$ concentrations attributed to the No-Trade counterfactual. Figure 1 clearly indicates that firms and facilities in the eastern United States increased emissions relative to their initial allocations: firms in these areas purchased permits in order to emit more and remain in compliance with the ARP. Firms in the western half of the country were clearly net sellers, abating more and enabling higher emissions east of the Mississippi River.

Figure 6 shows the proportional difference in damages under the ARP minus the No-Trade Scenario. The percentage change in damages is roughly proportional to the difference in $PM_{2.5}$ concentrations shown in Figure 1. Most counties showing an increase in damage due to trading exhibit an increase between 1 percent and 5 percent. Counties adjacent to the large metropolitan areas tend to exhibit a higher proportional change, on the order of 5 percent to 10 percent. And a few counties in West Virginia show increases greater than 10 percent.

Figure 7 expresses the difference in health damages in dollar terms. The absolute difference in damages reflects differences in the exposed population as well as differences in $PM_{2.5}$ levels. Thus, the biggest dollar differences in damages occur in the areas in Figure 1 with the greatest increase in $PM_{2.5}$ that are the most densely populated: metropolitan areas in the Middle Atlantic states and population centers in Ohio, North Carolina, and South Carolina.

5.2.3. Sensitivity Analyses

Tables 9 and 10 also display the results from a sensitivity analysis exploring alternative approaches to modeling damages from SO_2 emissions. The sensitivity analysis focuses on different ways to model the mortality impacts from $PM_{2.5}$ exposure because prior research has shown that the largest single contributor to air pollution damage is premature mortality risk (EPA 1999; NRC 2010; Muller et al. 2011). The sensitivity analyses include (a) using a lower (2 million) VSL applied to persons of all ages; (b) using a VSLY based on a value of $200,000 per life-year; and (c) using Roman et al. (2008) to model $PM_{2.5}$ mortality risks rather than Pope et al. (2002).

Replacing the $6 million VSL with the $2 million VSL reported in Mrozek and Taylor (2002) reduces aggregate damages in both scenarios, as well as the difference between damages under the ARP and the counterfactual. In Scenario 1, the difference between damages under the ARP and the UPS falls from –$9 million to –$2 million. In Scenario 2, the use of a lower VSL reduces monetary damages under the ARP to $52.3 billion and damages under the No-Trade counterfactual to $51.4 billion. The difference in damage between the ARP and the uniform performance standard falls to $870 million, compared with $2.44 billion in the base case.

The use of a VSLY also reduces aggregate damages in both scenarios; however, while it narrows the difference between damages under the ARP and the counterfactual in Scenario 2, which fall to $880 million from $2.44 billion, the use of a VSLY has the opposite effect in Scenario 1. The difference in damages between the ARP and the UPS increases in absolute value from –$9 million to –$106 million. This is the result of differences in the age structure of the exposed populations under the ARP and the UPS.

Replacing the dose-response function relating $PM_{2.5}$ exposure to mortality risk in Pope et al. (2002) with the relationship reported in Roman et al. (2008), which suggests that $PM_{2.5}$ has a 60 percent larger effect on mortality rates, raises damages estimates in all cases. In Scenario 2, this implies that damages under the ARP are almost $4 billion greater than under the No-Trade counterfactual. In Scenario 1, the ARP has marginally higher (rather than lower) damages than the UPS.

In sum, although the different approaches to mortality damage measurement or estimation have a clear impact on the magnitude of damages, the central findings of our base cases hold: For non-NSPS units, damages are similar under the ARP and the UPS. When we compare damages under the ARP to damages under the No-Trade counterfactual, trading increased damages in each of the different cases reported in Table 10.

6. Conclusions and Implication for Retrospective Policy Analysis

6.1. Summary of Results

In this study, we quantify the cost savings from the ARP compared with a command-and-control alternative and also examine the impact of trading under the ARP on health damages from SO_2. To quantify cost savings, we compare compliance costs for non-NSPS coal-fired EGUs under the ARP with compliance costs under a uniform performance standard that achieves the same aggregate emissions. We do this for the year 2002, the third year of Phase II of the program. The emissions of non-NSPS units in 2002 were approximately 7.55 million tons of SO_2, over 2 million tons more than allowances allocated to these units for the year 2002 under the ARP. The difference represents the effects of allowance purchases from NSPS and noncoal units regulated under the ARP and the drawing down of allowance banks. To examine the health effects of trading, we compute the health damages associated with observed SO_2 emissions from *all* units regulated under the ARP in 2002—approximately 10.2 million tons—and compare them with damages from a No-Trade counterfactual. In the No-Trade counterfactual, each unit emits SO_2 at a rate equal to its allocation of permits for the year 2002, plus any drawdown of its allowance bank.

We find the cost savings from the cap-and-trade system—the difference between the costs of coal purchase and scrubbing under the ARP and the uniform performance standard—to be positive, but lower than findings of previous studies. Specifically, we estimate this difference to be between \$250 million and \$300 million (US\$1995) per year in Phase II of the program, less than half of the savings estimated by Carlson et al. (2000). Carlson et al. (2000) assume that firms will achieve the least-cost solution to reducing emissions via fuel switching and that no additional scrubbers will be built after 1995. Our analysis suggests that a least-cost solution was not achieved.

Our cost data suggest that many generating units were not using the cheapest method of complying with the program. Comparing the costs of achieving the emissions rate selected by each EGU under the ARP with the least-cost method of achieving this emissions rate suggests that 23 percent of units could have reduced the cost of achieving their chosen emissions rate, at a cost savings of \$206 million (US\$1995). We also note that the number of scrubbers installed between 1996 and 2002 at non-NSPS plants suggests that the *global* least-cost solution was not achieved. The cost per ton of SO_2 removed by the 14 scrubbers installed between 1996 and 2002 is much higher than the cost of purchasing an SO_2 allowance.[31]

We also find that health damages under the ARP were greater than under the No-Trade counterfactual. These damages primarily represent adult premature mortality, as estimated by Pope et al. (2002) and valued using a \$6 million VSL (US\$2000). The mean difference in damages is \$2.44 billion (US\$2000), or about 1.8 percent of damages under the ARP. Health damages were greater under the ARP than in the No-Trade Scenario in densely populated areas in the Northeast and Middle Atlantic states (see Figure 6). This reflects the trading of allowances from units west of the Mississippi River to units east of the Mississippi River. As Henry et al. (2011) note, there is a positive correlation between marginal abatement costs for SO_2 and

[31] Fourteen scrubbers were installed at non-NSPS plants between 1996 and 2002. Our estimates of the cost per ton of SO_2 reduced range from \$247 to \$1,702 per ton, a figure much greater than the cost of an SO_2 allowance in 2002.

marginal damages from SO_2 emissions. When allowances are traded one-for-one, it is not surprising that emissions would increase in areas with higher marginal damages.

6.2. Comparison with Results of the Original RIA

EPA's RIA for ARP (EPA 1992) estimated the cost savings from allowance trading versus a command-and-control approach to achieving the same SO_2 cap. The RIA did not discuss the benefits of capping SO_2 emissions. At the time of the original RIA, the health benefits of reducing SO_2 emissions were not fully understood; the RIA predated the publication of Dockery et al. (1993) and Pope et al. (1995), which related fine particles to premature mortality. The original motivation for Title IV was to reduce acid deposition in the eastern United States. Since that time, it has become clear that the majority of benefits from the ARP in monetary terms come from human health improvements (EPA 1999, 2011). At the time of the RIA, however, the links among sulfur dioxide emissions, fine particles, and human health had not reliably been quantified.

Cost estimates in the RIA are derived from the Coal and Electric Utilities Model (CEUM). For each plant, the model estimates the compliance costs by simulating different strategies and their costs, and then assumes that the plant will choose the least-cost solution. The model is a combination of engineering assumptions and economic models, so that one plant's choices affect the national coal market. The RIA calculates that the benefits from trading—that is, the difference in compliance costs between the ARP and a uniform performance standard of 1.2 pounds of SO_2 per MMBtu—would be \$0.4 billion to \$0.6 billion (US\$1990) annually during Phase I, \$2.1 billion to \$2.8 billion (US\$1990) annually between 2000 and 2010, and \$1.3 billion to \$1.4 billion (US\$1990) thereafter. The high benefits from trading are due in part to assumptions about the inability of EGUs in the eastern United States (those burning bituminous coal) to switch to low-sulfur coal. The RIA assumes that retrofitting of boilers burning bituminous coal to burn PRB coal is impossible, and hence, all of these EGUs are forced to scrub under the UPS scenario.

There are several reason for the large difference between our estimates of cost savings and the original RIA. Because we use actual emissions by non-NSPS EGUs as the cap in our counterfactual UPS, our performance standard is less stringent than the one assumed in the RIA. The main reason for cost differences, however, lies in the assumption that fuel-switching is impossible for EGUs burning bituminous coal. This effectively forces EGUs in the eastern United States to install FGDs under the UPS, which greatly increases counterfactual compliance costs and, hence, the estimated gains from trade. As shown in Table 7, EGUs move toward lower-sulfur coal in our UPS, and there is no more scrubbing than under the ARP.

We believe that effectively assuming scrubbing in the command-and-control counterfactual biased upward EPA's estimates of the gains from trade. A more reasonable estimate of trading gains in Phase II of the program is provided by Carlson et al. (2000); however, as noted above, that estimate is also high, given its assumption of cost-minimizing behavior by power plants.

6.3. Implications of Our Results for Policy

What are the implications of our analysis for the design of environmental policy? One implication is that conventional estimates of the gains from cap and trade (versus command and control), which assume that cap and trade will achieve the least-cost solution to the emissions cap, must be viewed as an upper bound to realized trading gains. Our analysis suggests that the least-cost solution to observed emissions in 2002 was not achieved. As explained in the preceding section, this helps explain most of the difference between Carlson et al.'s (2000) estimates of trading gains and ours.

A second implication is that environmental policy should consider the impact of trading on the spatial distribution of emissions. The impact of trading under a cap-and-trade program on health damages depends on how the program is structured (e.g., are permits traded one-for-one or at the ratio of marginal damages) and on the correlation between marginal abatement costs and marginal damages. In a program like the ARP, emissions tend to flow from facilities with low marginal abatement costs to those facing higher costs of abatement. If damages are also higher at high-cost plants, total damage may rise.

In the context of SO_2, and other local air pollutants, damages per ton are higher for plants in, or upwind from, population centers (Fann et al. 2009; Muller and Mendelsohn 2009; Levy et al. 2009). In addition, the costs for SO_2 tend to be lower for firms in the West because of access to low-sulfur (Powder River Basin) coal. Broadly, abatement costs rise from West to East. Since population densities (and marginal damages) also follow this pattern, damages and costs are positively correlated. Ton-for-ton trading increases damages, as the early theoretical models predicted (Mendelsohn 1986). This, of course, need not be the case for all cap-and-trade programs, but the issue needs to be examined when selecting among policy options.

References

Arimura, T. H. 2002. An Empirical Study of the SO_2 Allowance Market: Effects of PUC Regulations. *Journal of Environmental Economics and Management* 44 (2): 271–89.

Bhat, C. R. 2001. Quasi-random Maximum Simulated Likelihood Estimation of the Mixed Multinomial Logit Model. *Transportation Research Part B* 35: 677–93.

Carlson, C. P., D. Burtraw, M. Cropper, and K. Palmer. 2000. SO_2 Control by Electric Utilities: What Are the Gains from Trade? *Journal of Political Economy* 108 (6): 1292–1326.

Chan, H. S. 2013. How Large Are the Cost Savings from Emissions Trading? An Ex Post Evaluation of the Acid Rain Program. Working paper.

Chan, H. S., H. Fell, I. Lange, and S. Li. 2013. Efficiency and Environmental Impacts of Electricity Restructuring on Coal-fired Power Plants. Working paper.

Cicala, S. 2015. When Does Regulation Distort Costs? Lessons from Fuel Procurement in U.S. Electricity Generation. *American Economic Review* 105 (1): 411–44.

Dockery, D. W., C. A. Pope, X. Xu, J. D. Spengler, J. H. Ware, M. E. Fay, B. G. Ferris Jr., and F. E. Speizer. 1993. An Association between Air Pollution and Mortality in Six U.S. Cities. *New England Journal of Medicine* 329: 1753–59.

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

EPA (US Environmental Protection Agency). 1992. Regulatory Impact Analysis of the Final Acid Rain Implementation Regulations. Washington, DC: EPA.

———. 1999. The Benefits and Costs of the Clean Air Act from 1990 to 2010. Washington, DC: EPA.

———. 2011. The Benefits and Costs of the Clean Air Act from 1990 to 2020. Washington, DC: EPA.

Fann, N., C. M. Fulcher, and B. J. Hubbell. 2009. The Influence of Location, Source and Emission Type in Estimates of the Human Health Benefits of Reducing a Ton of Air Pollution. *Air Quality Atmosphere and Health* 2: 169–76.

Fowlie, M. 2010. Emissions Trading, Electricity Restructuring, and Investment in Pollution Abatement. *American Economic Review* 100 (3): 837–69.

Fullerton, D., S. Dermott, and J. Caulkins. 1997. Sulfur Dioxide Compliance of a Regulated Utility. *Journal of Environmental Economics and Management* 34 (1): 32–53.

Henry, D. D., N. Z. Muller, and R. Mendelsohn. 2011. The Social Cost of Trading: Measuring the Increased Damages from Sulfur Dioxide Trading in the United States. *Journal of Policy Analysis and Management* 30 (3): 598–612.

Keohane, N. O. 2007. Cost Savings from Allowance Trading in the 1990 Clean Air Act: Estimates from a Choice-Based Model." In *Moving to Markets in Environmental Regulation: Lessons from Twenty Years of Experience*, edited by C. D. Kolstad and J. Freeman. New York: Oxford University Press.

Lange, I. A., and A. S. Bellas. 2007. The 1990 Clean Air Act and the Implicit Price of Sulfur in Coal. *B.E. Journal of Economic Analysis & Policy* 7 (1): 41.

Levy, J. I., L. K. Baxter, and J. Schwartz. 2009. Uncertainty and Variability in Health-Related Damages from Coal-Fired Power Plants in the United States. *Risk Analysis* 29 (7): 1000–1014.

Mendelsohn, R. 1986. Regulating Heterogeneous Emissions. *Journal of Environmental Economics and Management* 13: 301–12.

Mrozek, J. R., and L. O. Taylor. 2002. What Determines the Value of Life? A Meta-analysis. *Journal of Policy Analysis and Management* 21 (2): 253–70.

Muller, N. Z. 2011. Linking Policy to Statistical Uncertainty in Air Pollution Damages. *B.E. Press Journal of Economic Analysis & Policy* 11 (1).

———. 2012. The Design of Optimal Climate Policy with Air Pollution Co-benefits. *Resource and Energy Economics* 34 (4): 696–722.

Muller N. Z., and R. O. Mendelsohn. 2007. Measuring the Damages due to Air Pollution in the United States. *Journal of Environmental Economics and Management* 54: 1–14.

———. 2009. Efficient Pollution Regulation: Getting the Prices Right. *American Economic Review* 99 (5): 1714–39.

———. 2012. Efficient Pollution Regulation: Getting the Prices Right: Mortality Rate Update (Corrigendum). *American Economic Review* 102 (1): 613–16.

Muller, N. Z., R. O. Mendelsohn, and W. D. Nordhaus. 2011. Environmental Accounting for Pollution in the U.S. Economy. *American Economic Review* 101 (5): 1649–75.

NRC (National Research Council). 2010. Hidden Costs of Energy: Unpriced Consequences of Energy Production and Use.

Pope, C. A., R. T. Burnett, M. J. Thun, E. E. Calle, D. Krewski, K. Ito, and G. D. Thurston. 2002. Lung Cancer, Cardiopulmonary Mortality, and Long-Term Exposure to Fine Particulate Air Pollution. *Journal of the American Medical Association* 287 (9): 1132–41.

Pope, C. A, M. J. Thun, M. M. Namboodiri, D. W. Dockery, J. S. Evans, F. E. Speizer, and C. W. Heath Jr. 1995. Particulate Air Pollution as a Predictor of Mortality in a Prospective Study of U.S. Adults. *American Journal of Respiratory and Critical Care Medicine* 151 (3): 669–74.

Revelt, D., and K. Train. 2000. Customer-Specific Taste Parameters and Mixed Logit: Households' Choice of Electricity Supplier. Department of Economics, UCB.

Roman, H. A., K. D. Walker, T. L. Walsh, L. Conner, H. M. Richmond, B. J. Hubbell, and P. L. Kinney. 2008. Expert Judgement Assessment of the Mortality Impact of Changes in Ambient Fine Particulate Matter in the U.S. *Environmental Science and Technology* 42: 2268–74.

Schmalensee, R., and R. N. Stavins. 2012. The SO_2 Allowance Trading System: The Ironic History of a Grand Policy Experiment. Working paper 2012-012. MIT Center for Energy and Environmental Policy Research.

Sotkiewicz, P. M., and L. Holt. 2005. Public Utility Commission Regulation and Cost-Effectiveness of Title IV: Lessons for CAIR. *Electricity Journal* 18 (8): 68–80.

Stavins, R. N. 1998. What Can We Learn from the Grand Policy Experiment? Lessons from SO_2 Allowance Trading. *Journal of Economic Perspectives* 12 (3): 69–88.

Swinton, J. R. 2002. The Potential for Cost Savings in the Sulfur Dioxide Allowance Market: Empirical Evidence from Florida. *Land Economics* 78 (3): 390–404.

———. 2004. Phase I Completed: An Empirical Assessment of the 1990 CAAA. *Environmental and Resource Economics* 27 (3): 227–46.

Train, K. E. 2009. *Discrete Choice Methods with Simulation*. Cambridge: Cambridge University Press.

Tables and Figures

Table 1. Characteristics of Operating Coal-Fired EGUs in 2002

	Phase I		Phase II	
	Non-NSPS	NSPS	Non-NSPS	NSPS
Number of units	351	27	487	210
Total emissions (tons)	4,070,639	292,649.7	3,477,947	2,018,152
Total heat input (M of MMBtu)	6,089.2	940.87	5,861.8	7,319.1
Total allocated 2002 permits	2,886,593	371,357	2,578,884	2,357,000
Number of scrubbed units	41	15	59	109
Number of nonscrubbed units burning western coal	86	11	125	74
Average SO_2 emissions rate (lbs. per MMBtu)	1.6061	0.6389	1.3964	0.5468

Table 2. Compliance Choices in 2002 by Regulatory Status

	Divested units		PUC-regulated units		Publicly owned units	
	Non-NSPS	NSPS	Non-NSPS	NSPS	Non-NSPS	NSPS
% scrubbed	11.2	54.5	12.2	39.1	9.4	68.4
% using low-sulfur coal (no scrubber)	17.6	21.2	26.5	46.1	30.2	25.0
% using high-sulfur coal (no scrubber)	39.0	0.0	10.8	0.0	16.8	0.0
Total no. of units	187	33	502	128	149	76

Notes: Low (high) sulfur use refers to units where the majority of purchases originate from the Uinta or Powder River Basin (North Appalachian or Illinois Basin). For units without coal purchase data, sulfur use is inferred based on the unit's observed emissions rate.

Table 3. Average Sulfur Content of Coal, by Coal Basin

Basin	Mean sulfur	Range
North Appalachian, High End	2.7785	(2.0646,3.4062)
North Appalachian, Low End	1.5685	(0.8979,2.2406)
Central Appalachian	0.7636	(0.5376,1.0376)
South Appalachian	1.0789	(0.5802,1.4730)
Illinois Basin, High End	2.7700	(1.9804,3.4998)
Illinois Basin, Low End	1.2233	(0.7264,1.6833)
Uinta Basin	0.4792	(0.3072,0.8182)
Powder River Basin	0.3611	(0.2269,0.4816)

Notes: Unit is in pounds of S per MMBtu. All summary statistics are based on observed transaction data from 1991 to 2010. Range is based on the observed 10th to 90th percentile.

Table 4. Mean Values of Imputed Delivered Coal Prices, by Census Region, in 1995 Cents

	West	South	Midwest	Northeast
North Appalachian, High End	150.6	120.3	118.9	113.2
North Appalachian, Low End	216.9	146.6	143.6	121.9
Central Appalachian	228.1	148.0	155.4	152.0
South Appalachian	177.7	149.5	155.3	160.1
Illinois Basin, High End	217.7	144.0	130.6	158.6
Illinois Basin, Low End	208.7	150.6	135.6	159.0
Uinta Basin	122.3	161.7	144.4	169.9
Powder River Basin	83.78	126.3	95.52	133.1

Table 5. Summary Statistics of Model Variables

Variable	Mean	SD	Min	Max
Coal price (in cents per MMBtu)				
North Appalachian (High)	109.24	9.0742	96.08	152.32
North Appalachian (Low)	149.85	22.134	116.02	250.62
Illinois Basin (High)	141.17	24.608	102.45	252.86
Illinois Basin (Low)	144.54	20.397	112.97	236.38
Central Appalachian	157.97	19.892	129.26	259.89
South Appalachian	148.51	7.569	132.48	182.73
Uinta Basin	153.75	16.321	99.32	181.90
Powder River Basin (PRB)	112.72	21.407	44.01	151.50
Scrubbing cost (in cents per MMBtu)	41.954	25.159	14.46	531.99
Unit age	43.631	10.063	11	86
Heat input (in thousands MMBtu)	14,144.6	14,263	52.6	87,848.3
Phase I designation	0.4205	0.4940	0	1
Divested	0.2116	0.4087	0	1
Publicly owned	0.1761	0.3811	0	1
Modified boiler post-ARP	0.4271	0.4950	0	1
Use of PRB coal	0.2352	0.4070	0	1
Use of in-state coal	0.4047	0.4912	0	1

Table 6. Cost Model Estimation Results

	(1)	(2)	(3)	(4)	(5)
Mean effects					
Coal price	0.1848***	0.2274***	0.2008***	0.2405***	0.2420***
	(0.0132)	(0.0185)	(0.0139)	(0.0184)	(0.0192)
Emissions	3.0599***	3.7886***			
	(0.2778)	(0.2986)			
Emissions × PUC-Regulated			3.4990***	4.3963***	4.3300***
			(0.2531)	(0.3364)	(0.3381)
Emissions × Divested			3.2365***	3.7141***	3.8668***
			(0.4326)	(0.4770)	(0.5052)
Emissions × Publicly Owned			2.3723***	2.6818***	2.6610***
			(0.3230)	(0.4126)	(0.4111)
Ash	0.1223***	0.1689***	0.1621***	0.1885***	0.1793***
	(0.0219)	(0.0309)	(0.0275)	(0.0313)	(0.0329)
Scrubbing Cost	0.2018***	0.5125***	0.2013***	0.5461***	0.5651***
	(0.0151)	(0.1132)	(0.0150)	(0.1445)	(0.2169)
Scrubbing Cost × Divested			0.0418**	0.0496	0.0432
			(0.0211)	(0.0534)	(0.0561)
Scrubbing Cost × Publicly Owned					–0.0480
					(0.0537)
Modification	1.1855*	0.7936	1.0621	0.5749	0.5811
	(0.6430)	(0.7628)	(0.6650)	(0.7969)	(0.8063)
Modification × Age	0.0314**	0.0424**	0.0378**	0.0510***	0.0501***
	(0.0145)	(0.0173)	(0.0151)	(0.0182)	(0.0183)
PRB	3.1042***	4.1451***	3.7545***	4.5285***	4.3138***
	(0.9316)	(1.1280)	(0.9759)	(1.1466)	(1.1678)
PRB × Age	0.0573***	0.0700***	0.0532***	0.0674***	0.0708***
	(0.0194)	(0.0228)	(0.0206)	(0.0237)	(0.0240)
Coal Price × Minemouth	–0.0400***	–0.0410***	–0.0403***	–0.0476***	–0.0462***
	(0.0122)	(0.0129)	(0.0128)	(0.0138)	(0.0137)
Coal Price × In-State	–0.0071***	–0.0111***			
	(0.0019)	(0.0025)			
Coal Price × In-State × PUC-Regulated			–0.0078***	–0.0101***	–0.0099***
			(0.0024)	(0.0031)	(0.0030)
Coal Price × In-State × Publicly Owned			0.0075*	–0.0086	–0.0087
			(0.0038)	(0.0061)	(0.0061)
Coal Price × In-State × Divested					–0.0034
					(0.0051)
PRB × Publicly Owned			1.1246*	0.9719	1.0437
			(0.6196)	(0.6965)	(0.7119)
PRB × Divested					0.4083
					(0.6693)

Standard deviations of random coefficients

Scrubbing cost		0.2194***		0.2416***	0.2457*
		(0.0627)		(0.0848)	(0.1276)
Modification	0.6880	1.5591***	0.8745*	1.7877***	1.8102***
	(0.6158)	(0.4241)	(0.4794)	(0.3940)	(0.4094)
PRB	2.0491***	2.5016***	2.1984***	2.4585***	2.4930***
	(0.3039)	(0.3631)	(0.2942)	(0.3336)	(0.3411)
Log likelihood	–1026.7	–986.1	–1009.0	–975.6	–974.5
Prediction rate (%)	70.57	74.38	71.22	75.03	75.03
RMSE					
Emissions rate (lbs/MMBtu)	0.492	0.477	0.494	0.477	0.477
Tons SO_2	4823.3	4796.1	4434.6	4368.2	4494.0

Notes: All standard errors are robust standard errors, outputs from a random coefficient logit model. *, **, and *** indicate statistical significance at the 10, 5, and 1% levels. A positive coefficient implies that the cost is increasing in that component. In all specifications, NSPS units are dropped. All models are estimated based on observed choices for generating units that have not installed a scrubber or that installed a scrubber after 1988. Prediction rates are the percentage of sample units that actually used the choice with the highest predicted probability from the mixed logit model. Errors in predicting emissions are computed by comparing emissions, based on each model, with monitored emissions from EPA's Continuous Emission Monitoring System (CEMS).

Table 7. Simulation Results: ARP and Uniform Standard Counterfactual

Compliance choices		Uniform Performance Standard			
	ARP	(1)	(2)	(3)	(4)
No scrubber	***681***	***689***	***682***	***692***	***684***
High-sulfur coal	208	202	197	201	200
High end	42	0	0	0	0
Low end	166	202	197	201	200
Medium-sulfur coal	214	220	214	219	217
Low-sulfur coal	189	188	187	189	185
Blend: high & medium	27	32	39	33	33
Blend: high & low	21	25	22	28	26
Blend: medium & low	22	22	23	22	23
Scrubber	***80***	***72***	***79***	***69***	***77***
High-sulfur coal	50	44	49	43	47
Medium-sulfur coal	3	2	3	1	3
Low-sulfur coal	27	26	27	25	27
Predicted emissions (in million tons)					
ARP	7.094[a]	7.191	7.204	7.090	7.044
UPS		7.187	7.188	7.059	7.047
Standard level (lbs SO_2 per MMBtu)		2.100	2.210	2.090	2.080
(Weighted)		1.335	1.336	1.312	1.309
Cost savings (in million $1995)		296.43	253.18	262.64	244.05

[a] Denotes actual emissions from CEMS.

Table 8. Cost of Reaching Chosen Emissions Rate under the ARP in Excess of Minimum Cost

	Total units	Not reached least-cost		Average excess cost (cents/MMBtu)		Average excess cost (million $)	
	761	173	(22.73%)	Mean	Std. dev.	Mean	Std. dev.
Scrubbed	80	18	(22.50%)	20.03	15.46	2.155	2.645
Nonscrubbed	681	155	(22.76%)	7.884	6.338	1.083	1.873
Regulated	466	102	(21.89%)	8.981	8.410	1.221	1.488
Divested	161	29	(18.01%)	7.441	4.891	1.112	1.102
Public	134	42	(31.34%)	10.73	10.62	1.187	3.200

Table 9. Scenario 1: Comparison of Estimated Damages from SO_2 Emissions under ARP and Uniform Performance Standard

IAM Model parameters	**Damage ARP UPS**	**Difference (ARP – STD)**	**Difference (ARP – STD)/ARP**	**Deaths**	**Difference (ARP – STD)**
Default	103.4[a,b,c]			16,133	
	103.5[d]	–0.009	–0.0001	16,133	0
Alternative	163.8			26,350	
Dose-response[e]	163.8	0.002	0.0000	26,349	1
VSLY	44.6			16,133	
	44.7	–0.106	–0.0024	16,133	0
$2M VSL	39.8			16,133	
	39.8	–0.006	–0.0001	16,133	0

[a] Under scenario 1, only non-NSPS plants are included in both ARP and UPS simulations.
[b] Damages expressed in billions ($2000).
[c] Value in top row for each pair of model parameters corresponds to ARP.
[d] Value in bottom row for each pair of model parameters corresponds to UPS.
[e] Uses dose-response function for $PM_{2.5}$ mortality from Roman et al. (2008).

Table 10. Scenario 2: Comparison of Estimated Damages from SO_2 Emissions under ARP and No-Trade Counterfactual

IAM Model parameters	Damage ARP No-Trade	Difference (ARP – No Trade)	Difference (ARP – No Trade)/ARP	Deaths ARP No Trade	Difference (ARP – No Trade)
Default	135.8[a,b,c,d]			16,296	
	133.3	2.44	0.0184	16,165	130
Alternative	214.9			26,616	
Dose-response[e]	211.0	3.98	0.0182	26,402	214
VSLY	58.7			16,296	
	57.9	0.88	0.0136	16,165	130
$2M VSL	52.3			16,296	
	51.4	0.87	0.0172	16,165	130

[a] Under scenario 2, both NSPS and non-NSPS plants are included in both ARP and Allocation simulations.
[b] Damages expressed in billions ($2000).
[c] Value in top row for each pair of model parameters corresponds to ARP.
[d] Value in bottom row for each pair of model parameters corresponds to No-Trade Scenario.
[e] Uses dose-response function for $PM_{2.5}$ mortality from Roman et al. (2008).

Figure 1. Difference in $PM_{2.5}$ Concentrations in 2002: ARP Minus No-Trade Scenario

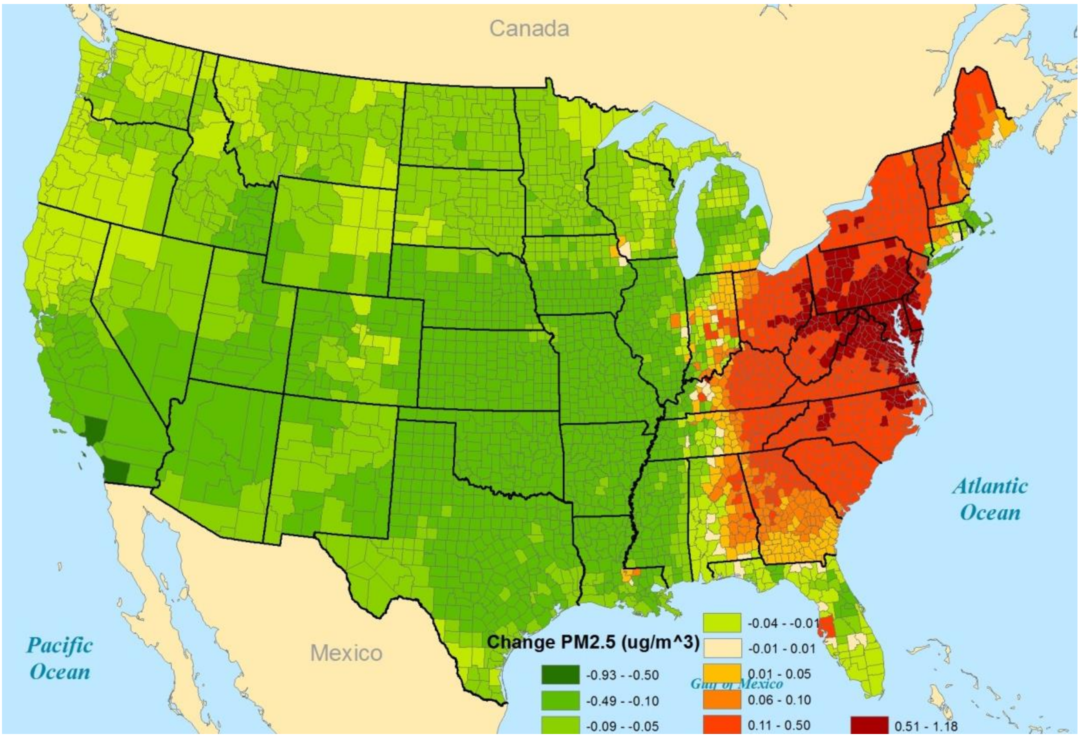

Figure 2(a). Percentage of EGUs Using Low-Sulfur Coal in 2002

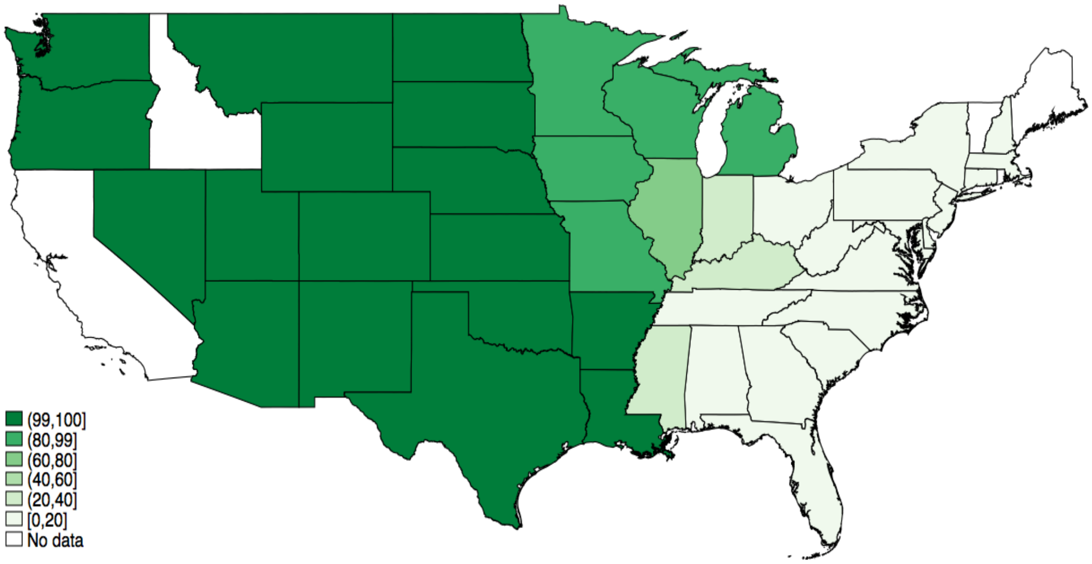

Figure 2(b). Emissions Net of Allocations in 2002

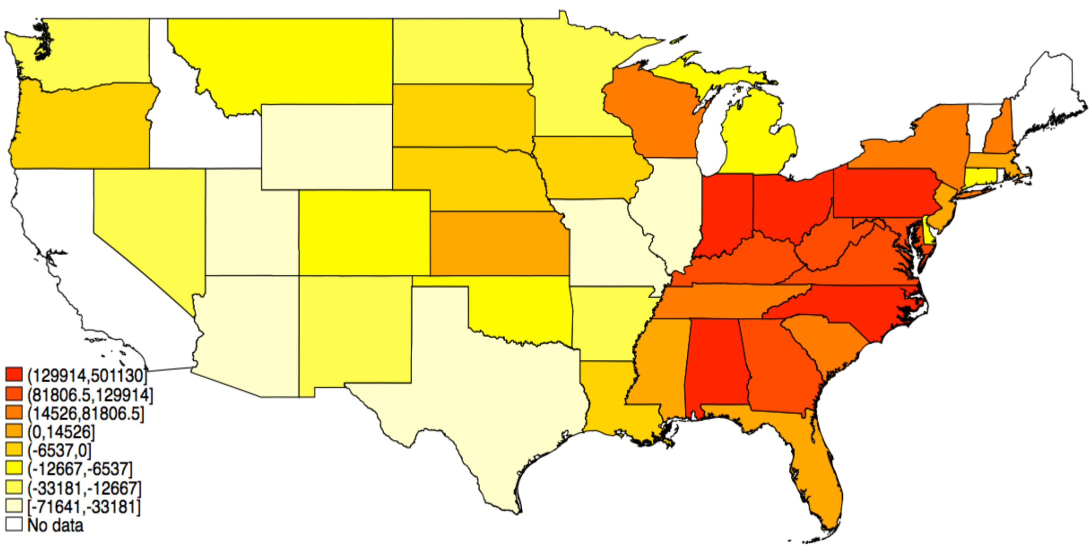

Figure 3(a). Location of Non-NSPS Units by Regulatory Status

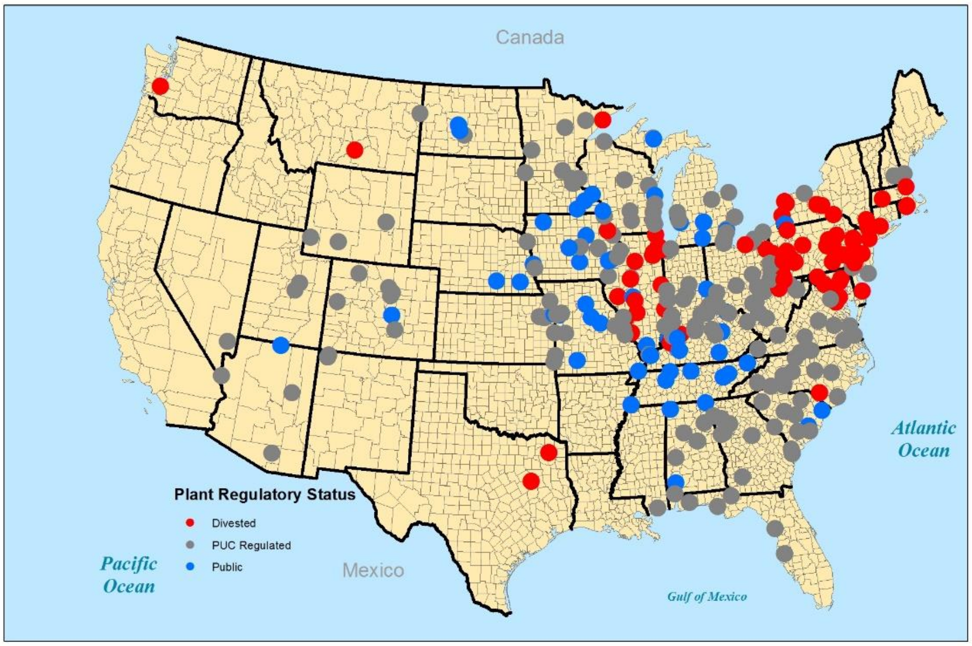

<Double-click here to enter text>

Figure 3(b). Location of NSPS Units by Regulatory Status

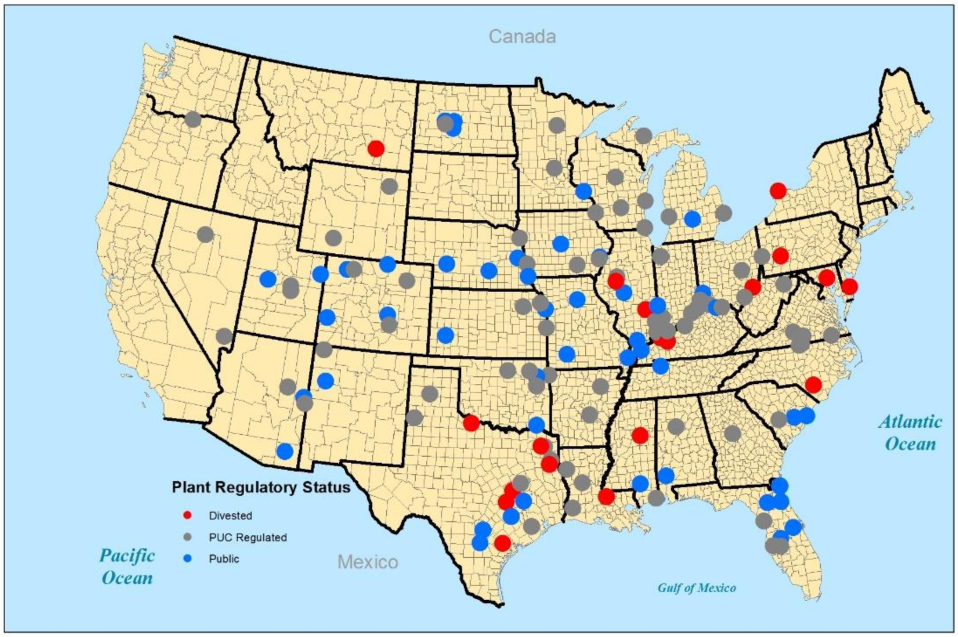

<Double-click here to enter text>

Figure 4. Histogram of Predicted Emissions Rates under the ARP, Model (2)

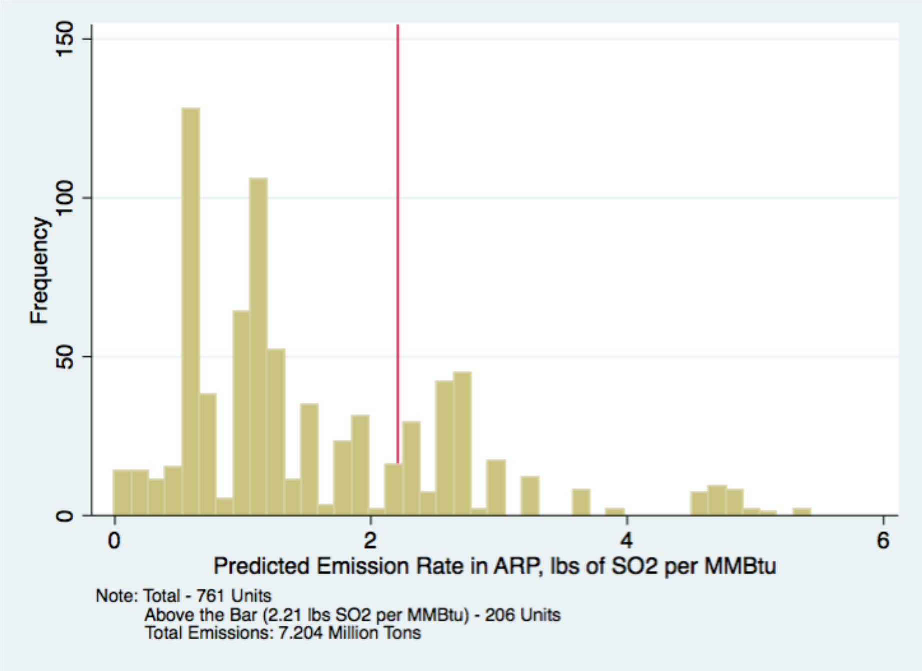

Figure 5. Cost of Achieving of Various Emissions Rates

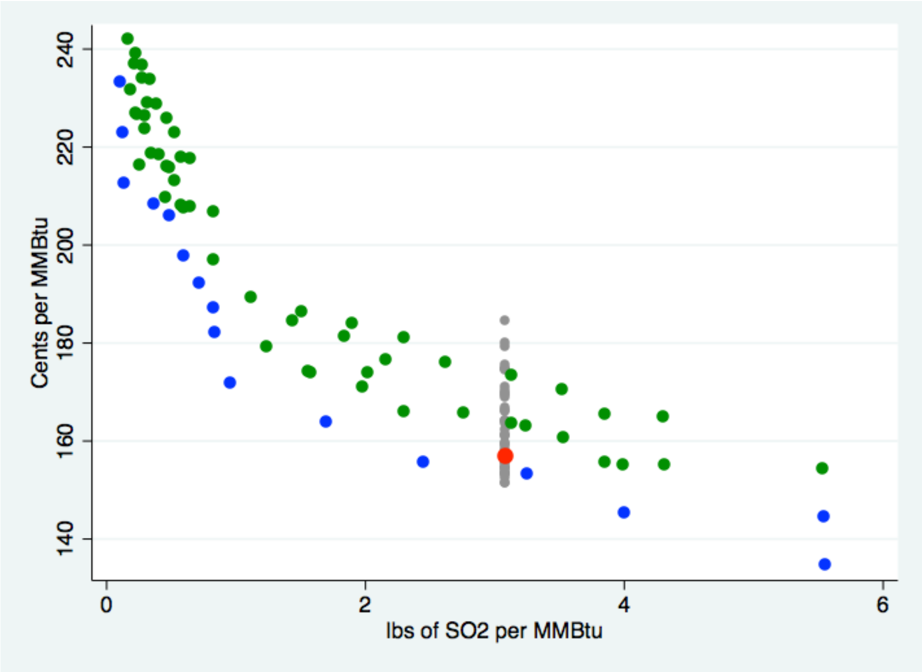

Notes: Calculations are based on Model (2) of Table 6. Blue dots identify the least-cost options and the red dot the chosen option.

Figure 6. Proportional Difference in Damages: ARP Minus No-Trade Scenario

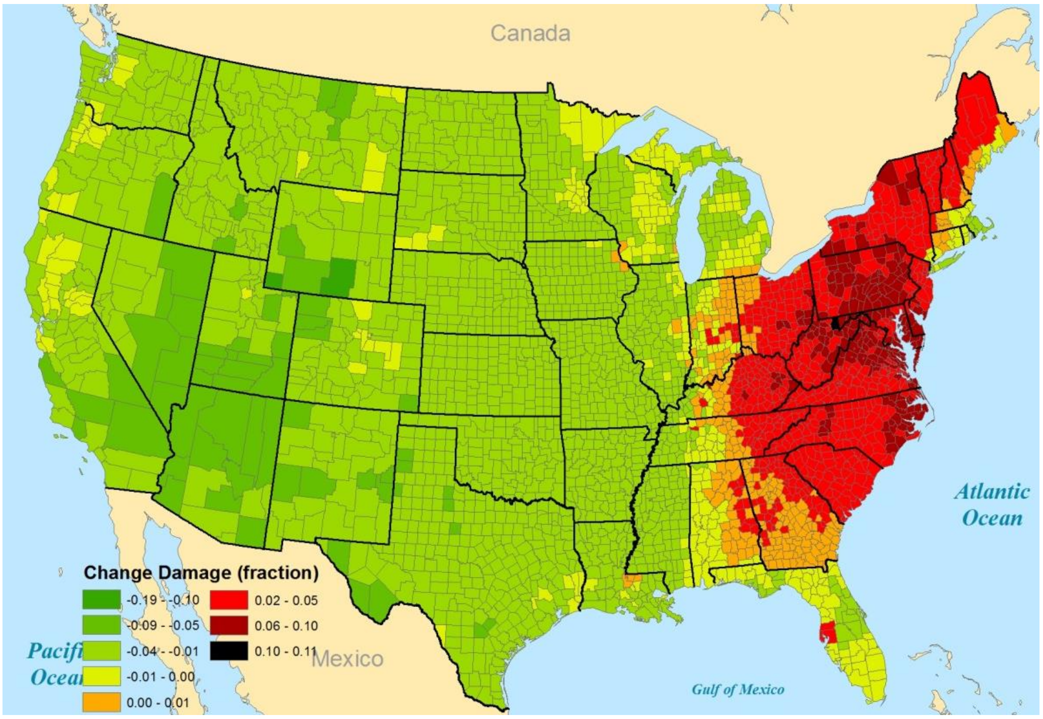

Figure 7. Difference in Health Damages: ARP Minus No-Trade Scenario

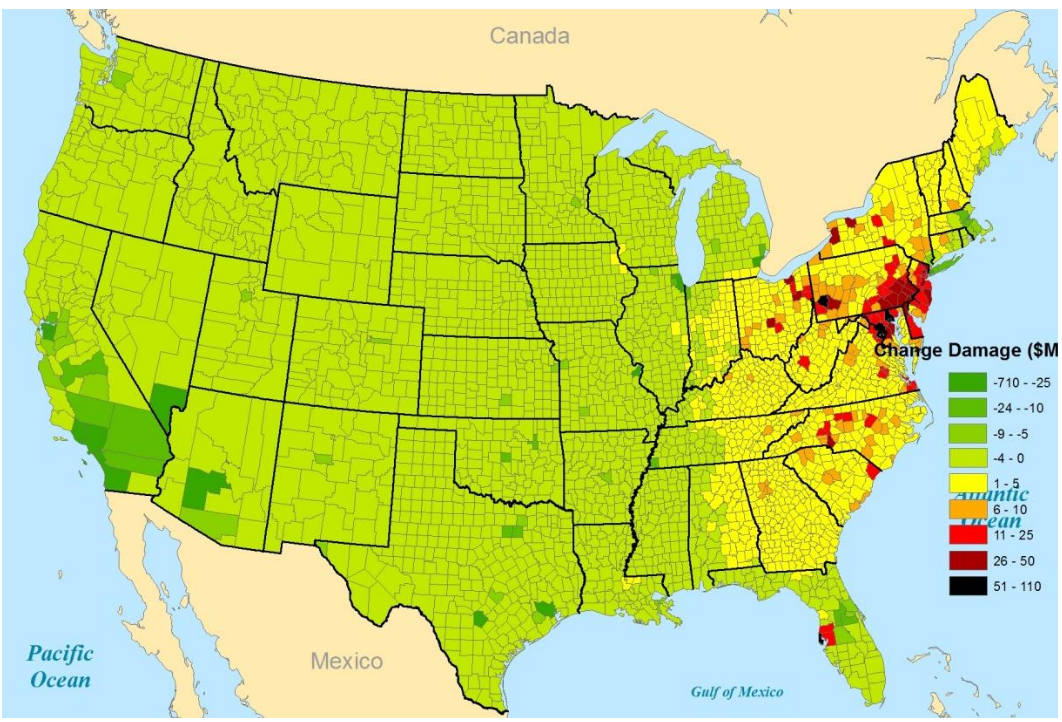

Appendix. Estimation and Simulation Procedures for the Cost Model

A.1. Estimation of the Mixed Logit Model

We choose to estimate a mixed logit model rather than a conditional logit model for two reasons. First, the weights placed on each cost component may vary across units. Second, some of the coefficients capture the cost of retrofitting boilers (e.g., to burn PRB coal), and there is no reason to believe that these costs should be the same for all units. Three thousand Halton draws are used to simulate the integral in the objective function during maximum likelihood estimation (Train 2009). As noted in the main text, three coefficients are assumed to follow independent Gaussian distributions: the coefficients on scrubbing cost, use of PRB coal, and whether the source of the unit's coal has changed since 1982.

Each compliance strategy involves selecting the basin from which to buy coal. Either all coal may be purchased from one basin or 50% may be purchased from each of two basins. We split the two high-sulfur coal basins, the North Appalachian and Illinois basins, into two basins based on the observed sulfur content in each county to make them more homogenous. However, because the variance in sulfur content within each coal basin remains large, it is extremely difficult to accurately estimate the unit's emissions rate without further refining the attributes of the coal purchased. We therefore use the following iterative procedure, as documented in Chan (2013), to refine the characteristics of the coal purchased within a basin:

1. Start with an estimate of the vector of cost function parameters, $\beta^{(0)}$.
2. For each alternative *j*, each generating unit *i* picks a coal type *k* within each alternative j. The coal type k is associated with attributes *COALPRICE(k;j)*, *SULFUR(k;j)* and *ASH(k;j)*. Unit i picks k, for each j, to minimize a deterministic version of the compliance cost function in equation (1).
3. After determining the optimal $k^*(i,j)$ for each *i* and *j*, unit *i* is assumed to choose $k^*(i,j)$ if it chooses alternative *j*. Substitute the attributes of coal type $k^*(i,j)$ into the matrix X_i in the mixed logit model.
4. Rerun the maximum simulated likelihood procedure on the mixed logit model based on these new attributes to obtain β^*.
5. Update $\beta^{(t)} = 0.8\beta^{(t-1)} + 0.2\beta^*$and repeat Steps 2 to 4 until $\beta^{(t)}$ is sufficiently close to $\beta^{(t-1)}$, that is, $\left|\beta^{(t)} - \beta^{(t-1)}\right| < 1 \times 10^{-6}$.

Each coal type *k* is defined as a mine-producing county or a 50–50 blend between two counties. We chose the county as the level of disaggregation given that it is the smallest geographic unit we observe in the data. The procedure generally reaches convergence in 20 iterations.

A.2. Simulation of Compliance Costs and Emissions under the ARP and the Uniform Emissions Standard Cases

Compliance choices and emissions for the case of a uniform performance standard are computed using the following procedures:

1. Estimate unit-specific conditional distributions for the random coefficients (Revelt and Train 2000)

$$g_i(\beta|D_i = Y, X_i, B, \Sigma) = \frac{P(D_i = Y|X_i, \beta) f(\beta|B, \Sigma)}{P(D_i = Y|X_i, B, \Sigma)}$$

   where D_i is the decision made by i, Y is the observed decision, and B and Σ are the parameters of the Gaussian distributions for the random coefficients.

2. Estimate the conditional means of the logit error terms, which represent unobserved compliance costs, for each unit i and compliance option j using shuffled Halton draws (Bhat 2001). Treat them as separate unit- and alternative-specific constant terms.

3. Compute the total compliance cost, as well as predicted emissions, based on the predicted choice for each unit. We multiply the emissions rate and average cost by the average heat input used in 2000–02 to calculate aggregate emissions and total costs. These are the predicted costs and emissions under the ARP.

4. Set the shadow price of permits to be zero and start with a uniform emissions standard $\bar{s}^{(0)}$. Repeat the iterative procedure above but excluding coal types that violate the uniform emissions standard $\bar{s}^{(0)}$. Predict the optimal compliance strategy $j*$ that minimizes the new compliance cost function, using the conditional distributions in Steps 1 and 2.

5. Compute the aggregate compliance cost and emissions as in Step 3, using the same observed heat input in MMBtu. If aggregate emissions exceed the predicted emissions in the ARP, repeat Step 4 with $\bar{s}^{(t)} = \bar{s}^{(t-1)} - 0.01$ until emissions in the counterfactual are approximately equal to the emissions in the ARP.

A.3. Data Sources

Our data come from the US Energy Information Administration (EIA) and the US Environmental Protection Agency (EPA). Emissions at the generating unit level come from the Continuous Emission Monitoring System (CEMS), made available by EPA. CEMS monitors power plants at hourly intervals to measure compliance and tracks sulfur dioxide emissions, total heat input (in MMBtu), and gross generation, allowing us to calculate actual emissions rates. The ARP compliance records from EPA provide information regarding allowance allocations, trades, and the banking of allowances, permitting us to compute the no-trade counterfactual.

Coal procurement data were obtained from EIA Form 423, which records coal transactions between mines and plants in the United States. Detailed information includes mine locations, coal quality (heat, sulfur, and ash contents), contractual arrangements, and transaction prices, in the form of delivered prices. These data were used to estimate regression models to predict region-plant-specific coal prices for our sample plants. Scrubber costs were obtained from EIA Form 860. Capital and operating costs are predicted by estimating regression models using observed costs and attributes (see Chan 2013 for more details).

To match the plant-level coal-purchase data to our analysis at the generating unit level, we use the following algorithm. For plants with similar emissions rates across EGUs, we assume all units use the average type of coal that the plant purchased. For plants with scrubbers installed in some but not all EGUs, we assign the cheaper coal (i.e., coal with higher sulfur content) that the plant purchased to the units with scrubbers and cleaner coal to units without scrubbers. For plants with considerably different emissions rates, we record the two types of coal that were used most intensively and match coal with higher sulfur content to the EGUs with higher emissions rates.

Table A.1. Scenario 1: Comparison of Estimated Damages from SO_2 Emissions under the ARP and a Uniform Performance Standard for Models (1) – (4) in Table 6

Econometric model	Damage ARP UPS	Difference (ARP – STD)	Difference (ARP – STD)/ARP	Difference (ARP – CEMS)	Difference (ARP – CEMS)/ CEMS
Model (2)	103.4[a,b,c]				
	103.5[d]	–0.009	–0.0001	0.236[e]	0.0023
Model (1)	103.2				
	103.8	–0.596	–0.0058	0.010	0.0001
Model (3)	101.5				
	101.8	–0.312	–0.0031	–1.685	–0.0163
Model (4)	100.9				
	101.8	–0.884	–0.0088	–2.338	–0.0227

[a] Under scenario 1, only non-NSPS plants are included in both ARP and UPS simulations.
[b] Damages expressed in billions ($2000).
[c] Value in top row for each pair of model parameters corresponds to ARP.
[d] Value in bottom row for each pair of model parameters corresponds to uniform performance standard.
[e] CEMS corresponds to observed emissions as reported by EPA.