NBER WORKING PAPER SERIES

THE IMPACT OF TRADING ON THE COSTS AND BENEFITS OF THE ACID RAIN PROGRAM

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

Working Paper 21383
http://www.nber.org/papers/w21383

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2015, Revised April 2017

Previously circulated as "The Market for Sulfur Dioxide Allowances: What Have We Learned from the Grand Policy Experiment?" Support for this work was provided by Resources for the Future. This paper has benefited from comments made by seminar participants at Carnegie Mellon University, the University of Pennsylvania, RFF, the Association of Environmental and Resource Economists, Georgia State University, the School of Public and Environmental Affairs at Indiana University, the Econometric Society World Congress in Montreal, the Royal Economic Society Conference in Manchester and the UK Applied Environmental Economics Conference. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed a financial relationship of potential relevance for this research. Further information is available online at http://www.nber.org/papers/w21383.ack

The Impact of Trading on the Costs and Benefits of the Acid Rain Program
H. Ron Chan, B. Andrew Chupp, Maureen L. Cropper, and Nicholas Z. Muller
NBER Working Paper No. 21383
July 2015, Revised April 2017
JEL No. Q51,Q52,Q58

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 \$240 million (1995\$). To examine the health effects of trading, we compute the health damages associated with observed sulfur dioxide (SO2) 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 SO2 at a rate equal to its allocation of permits for the year 2002, plus any drawdown of its allowance bank. Damages under the ARP are \$2.4 billion (2000\$) higher than under the No-Trade. This reflects the transfer of allowances from EGUs west of the Mississippi River to units in the eastern US with higher exposed populations.

H. Ron Chan
Arthur Lewis Building-3.078
School of Social Sciences
University of Manchester
Manchester, M13 9PL
ron.chan@manchester.ac.uk

B. Andrew Chupp
School of Public Policy
Georgia Institute of Technology
685 Cherry St.
Atlanta, GA
andrew.chupp@pubpolicy.gatech.edu

Maureen L. Cropper
Department of Economics
University of Maryland
College Park, MD 20742
and NBER
cropper@econ.umd.edu

Nicholas Z. Muller
Department of Economics
Warner Hall, 305D
Middlebury College
303 College Street
Middlebury, VT 05753
and NBER
nicholas.muller74@gmail.com

The Impact of Trading on the Costs and Benefits of the Acid Rain Program

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 of the Acid Rain Program (ARP), enacted under Title IV of the Clean Air Act Amendments of 1990, with the corresponding costs 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 the abatement cost savings of the ARP compared with those of an equally stringent policy that did not allow utilities to trade SO_2 allowances. We estimate a model of compliance behavior for coal-fired electricity generating units (EGUs) covered by the ARP that were not subject to New Source Performance Standards and use the model to compute the compliance cost savings achieved by these units under the ARP in 2002 compared with a uniform performance standard that would have resulted in the same aggregate emissions. We also compute the health damages associated with *all* EGUs covered by the ARP in 2002 and compare it 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. They 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.

[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).

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.

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. We take the initial allocation of allowances to all EGUs covered by the ARP as our no-trade counterfactual and compare damages under this counterfactual to damages under the ARP to estimate the health impacts of trading.

1.2. Our Approach

To measure cost savings from trading under the ARP, we use ex post data to model the compliance behavior of EGUs that were the focus of the ARP: coal-fired units not regulated under New Source Performance Standards (NSPS).[2] We do this for the year 2002. We argue that after 2003, it is difficult to separate the effects of the ARP from other regulations designed to reduce power plant emissions.[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

[2] 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.

[3] 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. Scrubbers were installed at many EGUs between 2006 and 2010. Some were installed in response to signals that EPA intended to drastically reduce SO_2 emissions from power plants below the target under the ARP. About a third of scrubbers installed between 2006 and 2010 were installed either in new sources (as mandated due to NSPS) or retrofitted in boilers subject to New Source Review.

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 $210 million and $240 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). There are at least two reasons for our smaller estimates. Carlson et al. (2000) estimate the gains from trade once the allowance market reaches a steady-state—i.e., when aggregate emissions equal the 8.95 million cap. As noted by Schmalensee and Stavins (2013), the allowance market never reached a steady state—other regulations superseded it. In 2002 aggregate emissions were 10.2 million tons, implying a less stringent cap and a less stringent performance standard than modeled by Carlson et al. (2000). We would therefore expect lower gains from trade. Carlson et al. (2000) also assumed that cost minimization would preclude the installation of scrubbers at non-NSPS plants after 1995. In fact, the number of EGUs with scrubbers at non-NSPS plants increased by 50% between 1995 and 2002.

To compare health damages under the ARP and our no-trade scenario, 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).

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 non-coal 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.

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 for non-NSPS units. In Section 5 we 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 concludes and discusses the policy implications of our results.

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. In Phase I 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.[4] 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. All ARP-regulated units were allocated annual permits in Phase II 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

[4] 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.

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.[5] 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 1977 CAAA in effect required coal-fired power plants built after September 1977 to install scrubbers. NSPS plants were thus required to achieve an emissions rate at least as stringent as was required under the ARP.

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),[6] 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 2013).

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.[7] In 2002, 1,075 coal-fired generating units were regulated under the ARP (see Table 1). These

---

5 Trading under the ARP could not violate the NAAQS.

6 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).

7 After plans for CAIR were announced in 2003, allowance prices rose sharply, signaling the anticipation of a new regulatory regime (Schmalensee and Stavins 2013).

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 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.[8]

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-

[8] 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.

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 to use high-sulfur coal and much less likely to fuel switch than either publicly owned or PUC-regulated units.[9] 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 the compliance behavior of non-NSPS EGUs under the ARP using a discrete choice model of which type of coal to purchase and whether or not to install an FGD. For each EGU the plant manager 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 compliance costs. 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 estimate the cost of delivering coal from each county

[9] 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 compliafootnce choices made by 2002. Only three of the scrubbers installed in divested non-NSPS units were installed after divestiture.

in the 8 coal basins described below to each EGU, as described in the Appendix. Scrubbing costs are also predicted for each unit (see Appendix). 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 removed by scrubbing. The coefficient on this component of costs represents the shadow price of emissions, which we compare to actual allowance prices.

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, or coal from the Uinta basin (Uinta). 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. Because low-sulfur coal from the Uinta basin has higher heat content, we include a separate dummy variable for Uinta basin coal.

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.[10] 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, $\overline{SULFUR_\iota}$,which may limit SO_2 emissions per MMBtu (equation (2)).[11]

$$\min_j C_i(j|\beta, X_i) = \beta^F COALPRICE_i(j) + \beta_i^z SCRUBCOST_i(j) + \beta^P EMISSIONS(j) \\ + \beta^A ASH(j) + PRB(j)(\beta_{0,i}^l + \beta_1^l AGE_i) + \beta_i^U UINTA(j) + \varepsilon_i(j) \qquad (1)$$

[10] When we regress heat input by EGU on year and EGU dummies for the period 1991 to 2005, over 94% of the variation in heat input is described by EGU dummies, suggesting that there is limited variation in unit-specific heat input.

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

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

$C_i(j)$ = unit compliance cost, in cents per MMBtu
$COALPRICE_i$ = delivered coal cost, 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
$EMISSIONS(j) = SULFUR(j) \times (1 - SCRUB(j)\theta)$, $\theta = 0.85$
PRB_j = 1 if coal is from the Powder River Basin
$UINTA_j$ = 1 if coal is from the Uinta Basin
$\varepsilon_i(j)$ = unobserved costs specific to option j

subject to

$$EMISSIONS(j) \leq \overline{SULFUR}_i \tag{2}$$

We allow the coefficients on *SCRUBCOST*, *PRB* and *UINTA* to be random. Although we predict the cost of scrubbing for each EGU, there are factors affecting scrubber cost that we cannot observe. The effective life of the scrubber may be shorter if the plant is preparing to retire the EGU, alternatively, the plant may not have space to retrofit a scrubber on a particular unit. Allowing the coefficient on *SCRUBCOST* to vary across plants controls for the fact that cost-minimizing plants might weight estimated scrubber costs differently, depending on their circumstances. The coefficients on *PRB* and *UINTA* capture the costs of retrofitting boilers; there is no reason to believe that these costs should be the same for all plants.[12]

3.2. Estimation of the Model

We estimate our model of compliance behavior using data for non-NSPS units in 2002.[13] We argue that most units had achieved their optimal compliance strategy under the ARP by this time. 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

[12] We allow the mean cost of retrofitting boilers to burn PRB coal to depend on boiler age. Age was not a statistically significant determinant of the cost of retrofitting to burn coal from the Uinta basin.

[13] 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.

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. 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 *UINTA* to be normally distributed with mean vector *b* and diagonal variance-covariance matrix *V*. Following Fowlie (2010) we treat each plant as a decision maker to allow correlation in unobserved costs for EGUs within each plant. This implies that the random coefficients vary by plant. The likelihood function is therefore given by

$$l(b,V) = \sum_{n=1}^{N} \ln \int_{-\infty}^{\infty} \prod_{i=1}^{I_n} \frac{\exp(-C_i(\mathbf{j}|b,X_i))}{\sum_{j'}^{J} \exp(-C_i(j'|b,X_i))} f(\beta|b,V)d\beta \tag{3}$$

where each plant is denoted as *n* and I_n denotes the set of units within each plant. X_i is the observed vector of covariates for each 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 nest the choice of the county from which coal is purchased within the choice of basin to refine the characteristics of coal purchased. We initially estimate the parameter vector $\boldsymbol{\beta}$ ($\boldsymbol{\beta}^0$), using the average characteristics of coal in each region for all units. Then, conditional on $\boldsymbol{\beta}^0$, we determine for each unit the county within each region that minimizes compliance costs.[15] We

[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.

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 $\boldsymbol{\beta}$ 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, as described in the Appendix. Data on transport costs, together with minemouth prices, are used to estimate the delivered cost of coal for each unit. We average predicted coal prices for the years 2000-2002 to estimate $COALPRICE_i(j)$. Because purchase decisions in 2000-2002 could be based on prices before the year 2000 (e.g., if plants enter into long-term contracts) Appendix Figure 1 plots price trends for the period 1991-2000. We note that the trends across coal regions are approximately parallel over this period.

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.

For units that do not install FGDs, the cost of installing and operating scrubbers are estimated as a function of plant and unit characteristics, following Lange and Bellas (2007). As described in the Appendix, we estimate separate models to predict the capital, operating and electricity costs of scrubbers, using data on all EGUs with scrubbers for the period 1991-2005, correcting in each case for selection bias.[17] In general, the costs of retrofitting a unit with a scrubber increase with the age and size of the unit; operating costs (including electricity costs) increase with years since the scrubber has been installed, removal rate, and operating hours. Average installation cost is estimated to be \$346/kW, comparable to values reported by Ellerman et al. (2000).

[16] This procedure is described more fully in the Appendix.

[17] Each model corrects for selection bias by simultaneously estimating the cost equation and the selection equation by full information maximum likelihood. Results are presented in the Appendix.

To estimate the cost model it is necessary to annualize scrubber installation cost. Assuming a discount rate of 11.33% and a 25-year lifetime we annualize predicted installation cost and compute the average cost of scrubbing as the sum of predicted operating cost (including electricity cost) and annualized installation cost, expressed per MMBtu of heat input. Our estimates imply that, on average, a scrubber increases operating costs by about 101 cents per MMBtu, of which 72.9 cents represent capital costs.[18] As Table 5 indicates, there is considerable variation in predicted scrubbing costs across EGUs.

Because the decisions regarding scrubber installation and coal purchase observed in 2002 were in some cases made before 2002, it is important to argue that the coal and scrubbing costs that we estimate are relevant to the choices observed in 2002. As noted above, minemouth coal prices are approximately constant for each coal basin over the period 1991 to 2002. In the case of scrubbers, we note that half of the scrubbers that have been installed in non-NSPS EGUs by 2002 were installed before 1988. Clearly these scrubber were not installed to comply with the ARP.[19] We eliminate these units from the estimation of our model. For the units included in the estimation the average installation date is 1995. The equations used to predict scrubber electricity and operating costs include time dummies to capture trends in these costs. We note that the decadal time dummies in the model used to predict capital costs are not significant at the .05 level.

Our simulation of cost savings under the ARP is based on 761 of the 838 non-NSPS coal-fired generating units in Table 1.[20] 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. Model (1), our base model, contains all of the variables in equation (1). Model (2) allows the price of coal to vary depending

[18] Assuming a removal rate of 85%.

[19] Carlson et al. (2000) attribute their installation to state emissions standards.

[20] Of the 838 non-NSPS units in Table 1, 77 are not used in our simulations: 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.

on whether coal is sourced in-state (Coal Price × In-State). Cheaper in-state coal could reflect PUC policies to encourage the use of in-state coal, given that we are controlling for transportation cost in predicting *COALPRICE*. Model (3) examines the implications of dropping the *UINTA* dummy variable and Model (4) allows some coefficients to vary by regulatory status. These include the coefficient on emissions, which measures the shadow price of emissions, as well as the coefficient on (Coal Price x In-State). To convert the coefficients to monetary units we divide them by the coefficient on *COALPRICE*.[21] Table 7 performs this conversion and adjusts the units in which variables are measured to calculate the shadow price of coal and the mean of the cost of retrofitting boilers to burn low-sulfur coal.

In all models, cost is increasing in coal price, SO_2 emissions, ash, and scrubbing cost. Two important components of unobserved costs—retrofitting costs for PRB coal and for coal from the Uinta basin—both show statistically significant mean effects on compliance costs, with significant dispersion in costs across plants. The cost of retrofitting boilers to burn PRB coal increases with the age of the boiler. Evaluated at the mean of the observations, average annualized cost for using PRB coal is 48 cents per MMBtu in our base model, with a mean of 33 cents for coal from the Uinta basin. The implied mean retrofit costs, in dollars per kW, are $28/kW and $19/kW, respectively.[22] The mean retrofit cost lie within the intervals reported by Ellerman et al. (2000); however, the retrofit cost for coal from the Uinta basin, which has higher heat content than PRB coal, is at the upper end of the interval. We therefore drop it from Model (3) but report simulation results for all four models.

Models (2) through (4) allow for interactions between *COALPRICE* and whether coal is sourced in state. They 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. Model (4) suggests that investor-owned units regulated by PUCs and publicly owned units assign a higher discount to in-state coal than divested units. This result agrees with Cicala (2015), who finds that divested power plants were less likely to purchase in-state coal than non-divested plants (see also Chan et al. 2017).

[21] Only the ratios of coefficients are identified in the mixed logit model. Because the coefficient on *SCRUBCOST* is random, we use the coefficient on *COALPRICE* to convert all other coefficients to monetary terms. We note that the mean of the *SCRUBCOST* coefficient is not significantly different from the *COALPRICE* coefficient at the 10% level in Models (2) – (4). It is significantly different from the coal cost coefficient at the .099 level in Model (1).

[22] In calculating the implied costs per kW, we assume a capacity factor of 0.60 and an operating heat rate of 11 MMBtu/MWh.

In all models the shadow price of SO_2 emissions is roughly equal to observed allowance price.[23] Allowance prices ranged from \$150 to \$200/ton of SO_2 over the period of our study. Model (1) implies that the average shadow price attached to SO_2 emissions was approximately \$162/ton SO_2, with a standard deviation of \$15/ton SO_2. In Model (4) the shadow price on emissions is higher for regulated IOUs (\$185/ton SO_2), than for divested (\$154/ton SO_2) and publicly owned (\$149/ton SO_2) units. This is consistent with the fact that PUC-regulated 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. The differences in shadow prices among the three regulatory statuses are not, however, statistically significant.

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 non-divested 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 9). Cicala (2015) finds that the biggest divergence in methods used by divested versus non-divested plants to reduce SO_2 emissions occurred after the time of our study.

Because we use the models in Table 6 to predict compliance choices in our simulations, it is important to ask how well the models fit. One measure of goodness of fit, percent correctly predicted, suggests that our models predict about 70% of observed choices correctly. To improve model fit, we replace the random coefficients in each model with the mean of the conditional distribution of these coefficients. Specifically, we compute the distribution of each random coefficient for each plant conditional on the plant's observed compliance choices,

$$g_n(\beta|Y_n = \mathbf{J}, X_i, b, V) = \frac{P(Y_n = \mathbf{J}|X_i, \beta) f(\beta|b, V)}{P(Y_n = \mathbf{J}|X_i, b, V)} \tag{4}$$

where Y_n is the vector of decisions made by plant n, $\boldsymbol{J}$ is the observed decision (for all the units owned by plant n), and b and V are the parameters of the Gaussian distributions of the random

[23] 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.

coefficients. We then replace the random coefficients in (1) by the means of these distributions. We also compute the distribution of the error terms $\{\varepsilon_{ij}\}$ conditional on the observed compliance choice and replace the $\{\varepsilon_{ij}\}$ in equation (1) by the means of the conditional distributions of the error terms. The latter capture unobserved components of each compliance choice, analogous to choice-specific fixed effects for each unit. This yields the modified cost function:

$$\begin{aligned}\widehat{C}_i(j|\beta, X_i) = {} & \widehat{\beta}^F COALPRICE_i(j) + \mathbf{E}_i\widehat{\beta}_i^z SCRUBCOST_i(j) + \widehat{\beta}^P EMISSIONS(j) \\ & + \widehat{\beta}^A ASH(j) + PRB(j)(\mathbf{E}_i\widehat{\beta}_{0,i}^l + \widehat{\beta}_1^l AGE_i) + \mathbf{E}_i\widehat{\beta}_i^U UINTA(j) + \mathbf{E}_i\varepsilon_i(j)\end{aligned} \tag{5}$$

Given these modifications, equation (5) perfectly predicts the observed compliance choice for each EGU.[24] When we conduct simulations of the model, we use equation (5).

A second question is how well our model captures observed SO_2 emissions under the ARP.[25] 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. Figure 4 compares the frequency distribution of emissions predicted by our base model with the actual distribution of emissions from the same units observed in 2002. Although the two distributions are not perfectly aligned, they match up well. In the aggregate, the base model predicts ARP emissions from the 761 units to be 7.086 million tons. Measured emissions were 7.094 million tons.

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 using the modified cost function in (5) above. To simulate choices under the uniform performance standard, the permit

[24] That is, it predicts which of the 72 compliance options is chosen. There may still be discrepancies between predicted and observed coal choices at the county level.

[25] It is important to note that observed emissions are not used to estimate our model.

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).

We calculate the cost of compliance under each regime (per MMBtu) as the sum of the coal cost 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. The cost of compliance per MMBtu is given by equation (6):

$$\begin{aligned} COMPC_i(j) = {} & COALPRICE_i(j) + SCRUBCOST_i(j) + \widetilde{\beta}^A ASH(j) + PRB(j)(\mathbf{E}_i\widetilde{\beta}_0^l + \widetilde{\beta}_1^l AGE_i) \\ & + \mathbf{E}_i\widetilde{\beta}_i^U UINTA(j) + \mathbf{E}_i\widetilde{\varepsilon}_i(j) \end{aligned} \tag{6}$$

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

4.2. Simulation Results

We simulate behavior under the ARP and a uniform performance standard using Models (1) through (4) of Table 6. Table 8 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.[26] This does not, however, imply that emissions are predicted perfectly, due to the heterogeneity of the sulfur 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 2.07 to 2.28 pounds of SO_2 per MMBtu. (When

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

weighted by heat input, the UPS is between 1.32 and 1.38 pounds of SO_2 per MMBtu.) This standard is less stringent than the cap implied by 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 5 compares predicted emissions rates under the ARP for Model (1) of Table 6 with the corresponding uniform standard. The 205 units that are above the standard under the ARP must reduce their emissions under the UPS. 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.

Compliance costs under the ARP are estimated to be between \$211 million and \$236 million (US\$1995) lower than under the uniform performance standard; significantly smaller estimates than previous studies. There are two reasons for this. Carlson et al. (2000), in comparing the ARP with a uniform performance standard, assume that the uniform emissions standard will reduce all non-scrubbed units to an emissions rate of 1.2 lbs. of SO_2 per MMBtu of heat input, based on 1985-87 heat input. They also assume that no scrubbers would be installed after 1995. Because we are looking at a uniform performance standard that would achieve observed emissions in the year 2002 our standard is a much less stringent standard than 1.2 lbs. of SO_2 per MMBtu of heat input. Indeed, the emissions cap of 7 million tons of SO_2 on the non-NSPS units in our simulation is 40% greater than the cap implied by a standard of 1.2 lbs. of SO_2 per MMBtu of heat input. We would therefore expect the cost savings from trading to be lower than for the more stringent standard.

Secondly, 14 non-NSPS units did install scrubbers after 1995. Carlson et al. (2000) assumed that no scrubbers would be installed after 1995 because their estimated cost per ton of SO_2 reduced via scrubbing exceeded their projected permit price. Our estimates of the cost per ton of SO_2 reduced by installing scrubbers at the 14 units range from \$247 to \$1,702 per ton, a figure much greater than the cost of an SO_2 allowance in 2002. The estimated cost per ton of SO_2 removed for these units, compared with allowance price of \$160 suggests that scrubbing increased compliance costs at these units by \$88 million in 2002.

5. The Health Impacts of Trading

To examine the impact of trading on health damages we expand the scope of our analysis from non-NSPS units to all units covered by the ARP. The emissions of all non-NSPS units in

2002—7.55 million tons—reflect the fact that non-NSPS units were net buyers of allowances from other units covered by the ARP. As Table 1 shows, the emissions of non-NSPS EGUs exceeded allowances allocated for 2002 by over 2 million tons. Approximately two-thirds of these allowances were obtained through trades. The remaining one-third are accounted for by non-NSPS units drawing down their allowance banks (CEMS). At the same time, allowances allocated to NSPS EGUs and non-coal units (primarily natural gas, oil and diesel generating units) in 2002 exceeded their emissions in 2002.[27] To examine the impact of allowance transfers to non-NSPS units, we define a no-trade counterfactual scenario and contrast damages under the no-trade counterfactual with damages under the ARP.

The no-trade counterfactual forces all units covered by the ARP, including oil-fired and gas-fired units, to emit at the rate prescribed by their initial allocation of 2002 allowances, plus any drawdowns of their allowance banks observed in 2002. The ARP version of this scenario uses the actual emissions from each unit under the ARP to calculate damages. We thus compare damages for the no-trade counterfactual with damages from observed emissions, inclusive of permit trading. Under both the ARP and the no-trade counterfactual, aggregate SO_2 emissions are 10.2 million tons. Hence, any difference in damage is due to the geographic distribution of emissions, not the overall amount of discharges.

5.1 Estimating Heath Damages

The health damages due to SO_2 emissions produced by EGUs are estimated using AP2 (Muller 2011, 2012), a stochastic integrated assessment model that links reported and counterfactual emissions to ambient concentrations of air pollutants and ambient concentrations to pollution damages. In order to estimate concentrations, AP2 employs a source-receptor matrix in which each cell ($T_{k,m}$) in the matrix represents the change in ambient concentrations of $PM_{2.5}$ in location (m) due to a one-ton increase of SO_2 emissions from source (k). The source-receptor matrices capture atmospheric processes that link emissions of precursor species (like SO_2) to resulting ambient concentrations of secondary pollutants. Significantly for our study, emissions of SO_2 are connected to concentrations of particulate sulfate, an important constituent of $PM_{2.5}$.

---

[27] NSPS units emitted 410,000 fewer tons of SO_2 and non-coal units 540,000 fewer tons of SO_2 than allowances allocated to them for 2002 (CEMS authors' calculations.)

In studies of the benefits of air pollution policy, such as the ARP, health benefits constitute the majority of monetized benefits, and it is on those that we focus (USEPA 1999). AP2 links ambient concentrations of $PM_{2.5}$ to morbidity and mortality in exposed populations using concentration-response functions from the epidemiological literature. These are combined with county-level population inventories provided by the US Census and baseline incidence rates to calculate health risks. 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.[28]

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).[29] 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 expected life-years remaining.

For each policy simulation, AP2 processes baseline emissions through the source-receptor matrices to estimate baseline $PM_{2.5}$ concentrations, exposures, physical effects, and damages. All baseline emissions (except for SO_2 produced by EGUs) are provided by EPA's National Emission Inventory (NEI) for 2002.[30] Then, for a particular policy scenario, SO_2 emissions from EGUs along with baseline emissions are processed through AP2 to estimate concentrations, exposures, physical effects, and damages. The change in damages due to the change in EGU emissions across policy scenarios is tabulated (1) in total, aggregating across all receptor counties; and (2) by county, to explore spatial patterns in the change in emissions, air quality, and impacts.

[28] 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).

[29] This is approximately equal to EPA's value, $4.8 million (US$1990), adjusted for inflation.

[30] 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.

5.2 Damages under the ARP and under a No-Trade Counterfactual

Table 9 reports the difference between damages under the ARP and the no-trade counterfactual, aggregated across all counties in the US. 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).

Table 9 also displays 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) or with the VSLY reduces aggregate damages under the ARP and the no-trade counterfactual and the difference between them. The difference in damage between the ARP and the no-trade scenario falls to $840 million when either a $2 million VSL is used or when using the VSLY.

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 both cases. This implies that damages under the ARP are almost $4 billion greater than under the no-trade counterfactual. In sum, although the different approaches to mortality damage estimation have a clear impact on the magnitude of damages, trading increased damages in each of the different cases reported in Table 9. The statistical uncertainty associated with the damage estimates is further explored in the Appendix.

Figures 1, 6 and 7 explore the spatial pattern of the difference in damages between the ARP and the no-trade counterfactual. 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 initial allocations, while those incurring lower marginal abatement costs would sell permits. This is suggested by Figure 2(b) 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 SO_2 while remaining in compliance with the ARP. This is especially true of plants in Pennsylvania, Ohio, and West Virginia. 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 damages due to trading exhibit an increase between 1 percent and 5 percent. Figure 6 shows a stark east-west divide extending from Western Indiana down to the Gulf Coast states. Counties east of this line have increasing damages while those to the west have lower damages due to trading. Central and Eastern Pennsylvania, Northern Virginia, and most of Maryland exhibit increases in damages of between 5 and 10 percent. Parts of North Carolina and West Virginia also show increases in damage over 5 percent. And, a few counties in West Virginia show increases greater than 10 percent. Intuitively, these areas incur the largest increases in ambient concentrations (see Figure 1). In contrast, the coastal counties north of New York City, which emitted less than their 2002 allocations (see Figure 2(b)) have lower damages under the ARP.

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 increases in $PM_{2.5}$ that are also the most densely populated: metropolitan areas in the Middle Atlantic states and population centers in Ohio, North Carolina, and South Carolina. Figure 7 shows more variation in damages than Figure 6 due to heterogeneity in population density. For example, damages increase by over \$50 million in fewer than ten counties that are in and near the Pittsburgh, Washington, D.C., and Baltimore metropolitan areas. The vast majority of counties incur less than \$5 million in additional damage. This roughly maps to the counties shown in Figure 1 in which ambient $PM_{2.5}$ increased by less than 0.5 $\mu g/m^3$.[31]

[31] Figure 7 indicates that damages would have been higher under the no-trade counterfactual in southern California than under the ARP. These damages reduce the difference between damages under the ARP and the no-trade counterfactual. We note that, due to state emissions standards, it is unlikely that emissions under the no-trade counterfactual would have been as high as idicated by the initial allocation of allowances, implying a larger difference in damages between the ARP and the no-trade counterfactual.

Figures 6 and 7 quantify the conceptual conclusions of an earlier literature. Mendelsohn (1986) and Stavins (1996) show that if marginal damages and marginal abatement costs are positively correlated, market-based instruments may not increase net benefits relative to command-and-control policies. In the present context, marginal damages are primarily a function of population density: power plants in the (more populous) eastern U.S. tend to have higher marginal damages than facilities in the west. On the cost side, one of the most cost-effective sulfur abatement strategies is the use of low-sulfur coal. Most low sulfur coal is mined in western states. Hence, marginal costs are higher in the east because of the high cost of transporting low-sulfur coal. Putting these patterns together, marginal damage and marginal costs are both higher in the eastern U.S., and, therefore, positively correlated. This implies that facilities that are likely to purchase additional allowances (those with higher than average marginal costs) are also likely to have high marginal damages. Thus, emissions migrate to high damage facilities and, on net, damages increase. Had it been the case that damages and costs were negatively correlated at the margin, trading would have reduced damages, reinforcing abatement cost savings reported above.

6. Conclusions

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 non-coal 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 \$211 million and \$236 million (US\$1995) per year in Phase II of the program, less than half of the savings estimated by Carlson et al. (2000). There are two reasons why our estimates are lower. Carlson et al. (2000) estimate the cost savings from the ARP in the steady-state, when emissions under the program reach the 8.95 million ton cap. We examine the program in 2002 before the program had reached the steady state, when the effective cap—i.e., observed emissions—is higher (10.2 million tons). Indeed, for the non-NSPS units in our simulations, actual emissions in 2002, which form the basis for computing the uniform performance standard, are approximately 40% higher in the aggregate than the 1.2 lb. per MMBtu long-run target under the ARP. A looser target implies lower gains from trade.

It is also the case that Carlson et al. (2000) assumed that plants would achieve the least-cost solution to reducing emissions via fuel switching and that no additional scrubbers would be built after 1995. We 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. 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.

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 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.

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.

We close by noting that our paper should not be interpreted as indicating that the ARP yielded negative benefits compared to a command-and-control alternative. Because of the large benefits of the ARP relative to a no-regulation alternative (USEPA 1999; Burtraw et al. 1998) EPA moved swiftly to tighten the cap on SO_2 emissions beginning in 2003. This forces us to examine the program before it reached the steady-state, in a single year (2002). Our analysis thus fails to capture the cost savings from the banking provisions of the ARP (Ellerman and Montero, 2007) and the impacts of emissions trading on technical improvements in scrubbers (Montero, 2002; Gagelmann and Frondel, 2005). We would, ideally, like to compare the ARP from inception to the steady state to a counterfactual uniform performance standard that was phased in beginning in 1995. This, however, would require a dynamic model of firm behavior, which remains a subject for future research.

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.

Burtraw, D., Krupnick, A., Mansur, E., Austin, D., Farrell, D., 1998. Costs and Benefits of Reducing Air Pollutants Related to Acid Rain. *Contemporary Economic Policy* 16 (4), 379–400.

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. 2017. Efficiency and Environmental Impacts of Electricity Restructuring on Coal-fired Power Plants. *Journal of Environmental Economics and Management* 81 (1): 1-18.

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.

Ellerman, A.D., and J-P. Montero. 2007. The Efficiency and Robustness of Allowance Banking in the U.S. Acid Rain Program. *The Energy Journal* 28 (4): 47-71.

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.

Gagelmann, F., and M. Frondel. 2005. The Impact of Emission Trading on Innovation – Science Fiction or Reality? *Environmental Policy and Governance* 15 (4): 203-211.

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.

Montero, J-P. 2002. Market Structure and Environmental Innovation. *Journal of Applied Economics* 5 (2): 293-325.

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, Richard, and Robert N. Stavins. 2013. The SO_2 Allowance Trading System: The Ironic History of a Grand Policy Experiment. *Journal of Economic Literature* 27 (1): 103-122.

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 non-scrubbed 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 Basins (North Appalachian or Illinois Basins). 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)	101.16	25.04	34.20	224.92
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
Use of PRB coal	0.2352	0.4070	0	1
Use of Uinta coal	0.0769	0.2610	0	1
Use of in-state coal	0.4047	0.4912	0	1

Table 6. Cost Model Estimation Results

	(1)	(2)	(3)	(4)
Mean effects				
Coal Price	0.2869***	0.2840***	0.2915***	0.2722***
	(0.0260)	(0.0303)	(0.0307)	(0.0295)
Emissions	4.6531***	4.6311***	4.0987***	
	(0.4363)	(0.4418)	(0.4297)	
Emissions × PUC-Regulated				5.0220***
				(0.5617)
Emissions × Divested				4.0615***
				(1.0743)
Emissions × Publicly Owned				4.1967***
				(0.8073)
Ash	0.1847***	0.1493***	0.1382***	0.1208***
	(0.0435)	(0.0470)	(0.0457)	(0.0461)
Scrubbing Cost	0.2301***	0.2320***	0.2481***	0.2261***
	(0.0372)	(0.0354)	(0.0456)	(0.0376)
PRB	8.9114***	8.1398***	7.1423***	9.2518***
	(2.1115)	(2.2009)	(2.1414)	(2.2693)
PRB × Age	0.1095***	0.1118***	0.0959**	0.0678
	(0.0393)	(0.0397)	(0.0387)	(0.0439)
Uinta	9.3834***	9.6048***		9.2123***
	(1.2362)	(1.2379)		(1.1049)
Coal Price × In-State		–0.0126***	–0.0140***	
		(0.0047)	(0.0048)	
Coal Price × In-State × PUC-Regulated				–0.0164***
				(0.0056)
Coal Price × In-State × Publicly Owned				–0.0261*
				(0.0141)
Coal Price × In-State × Divested				–0.0026
				(0.0073)
Standard deviations of random coefficients				
Scrubbing cost	0.0985***	0.0992***	0.1143***	0.0962***
	(0.0155)	(0.0153)	(0.0239)	(0.0158)
PRB	6.3967***	6.3457***	6.0128***	6.9948***
	(0.7834)	(0.7361)	(0.8304)	(1.0242)
Uinta	6.3091***	6.6209***		5.4683***
	(0.7885)	(0.7737)		(0.5105)
Log likelihood	–987.1	–941.34	–1048.6	–933.1
Prediction rate (%)	71.48	72.54	68.99	72.40
RMSE				
Emissions rate (lbs/MMBtu)	0.509	0.519	0.530	0.503

Tons SO_2	4777.7	4984.7	4905.8	4602.0

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. All models treat each plant as a decision maker. 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. Interpretation of Model Coefficients

	(1)	(2)	(3)	(4)
Shadow Price for SO_2 ($/ton)	162.21	163.08	140.63	
	(14.86)	(14.89)	(9.90)	
PUC Regulated				184.53
				(22.17)
Divested				149.23
				(38.87)
Publicly-Owned				154.20
				(30.76)
Shadow Price for Ash ($/ton)	12.88	10.52	9.49	8.87
	(2.53)	(2.79)	(2.64)	(2.90)
Retrofitting Cost Associated with				
Uinta (Cents/MMBtu)	32.71	33.82		33.85
	(5.48)	(5.99)		(5.88)
($/kW)	18.91	19.55		19.57
PRB (Cents/MMBtu)	47.72	46.10	38.86	44.85
	(10.43)	(11.11)	(10.16)	(10.19)
($/kW)	27.59	26.65	22.47	25.93

Notes: Dollar values are 1995 US$. Standard errors appear in parentheses. PRB retrofitting costs are calculated based on unit with a mean age (43.631). For calculations in $/kW, we assume an operating heat rate of 11 MMBtu/MWh and a capacity factor of 60%.

Table 8. Simulation Results: ARP and Uniform Standard Counterfactual

Compliance choices

	ARP	Uniform Performance Standard (1)	(2)	(3)	(4)
No scrubber	***681***	***688***	***689***	***689***	***688***
High-sulfur coal	208	195	201	209	204
High end	42	0	0	0	0
Low end	166	195	201	209	204
Medium-sulfur coal	214	205	203	211	203
Low-sulfur coal	189	190	188	183	186
Blend: high & medium	27	55	52	39	50
Blend: high & low	21	21	23	23	23
Blend: medium & low	22	22	22	24	22
Scrubber	***80***	***73***	***72***	***72***	***73***
High-sulfur coal	50	46	45	45	45
Medium-sulfur coal	3	2	2	2	2
Low-sulfur coal	27	25	25	25	26
Predicted emissions (in million tons)					
ARP	7.094[a]	7.086	7.230	7.422	7.164
UPS		7.079	7.233	7.426	7.159
Standard (lbs SO_2 per MMBtu)		2.170	2.070	2.170	2.180
(Weighted)		1.315	1.344	1.380	1.330
Cost savings (in million \$1995)		**236.45**	**223.21**	**223.21**	**211.08**

[a] Denotes actual emissions from CEMS.

Table 9. 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[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

[b] Damages expressed in billions ($2000).
[c] Value in top row for each pair of model parameters corresponds to the ARP.
[d] Value in bottom row for each pair of model parameters corresponds to the 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

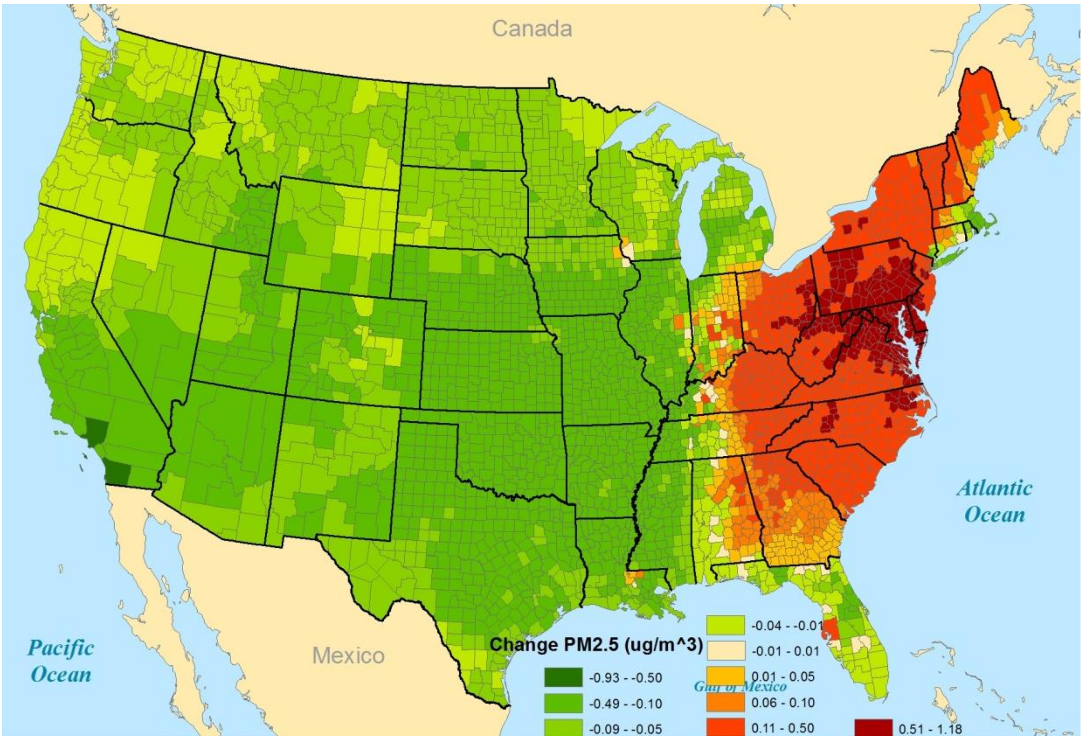

Figure 2(a). Percentage of EGUs Using Low-Sulfur Coal in 2002

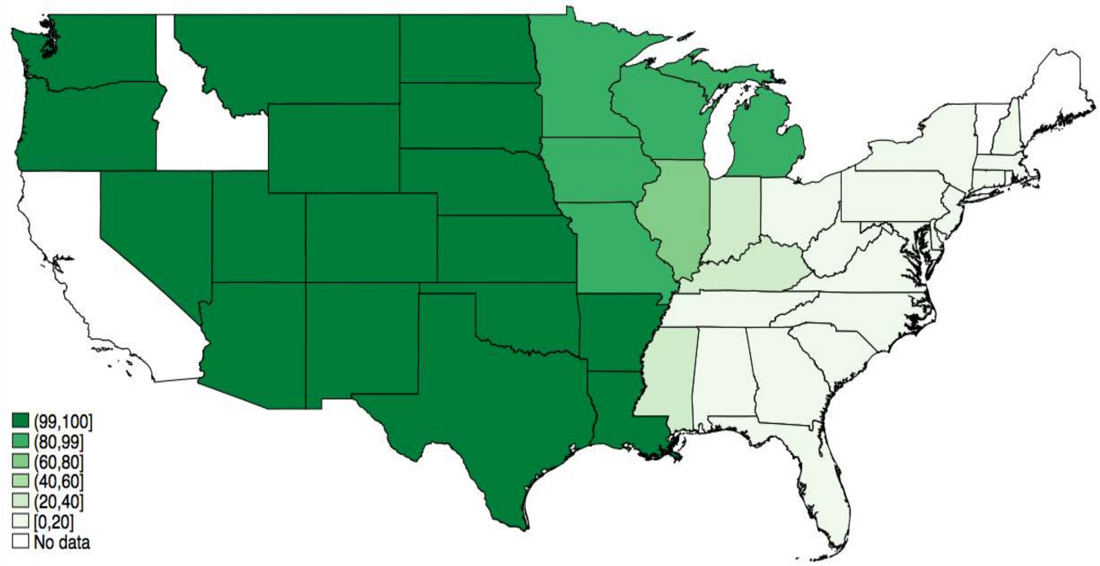

Figure 2(b). Emissions Net of Allocations in 2002

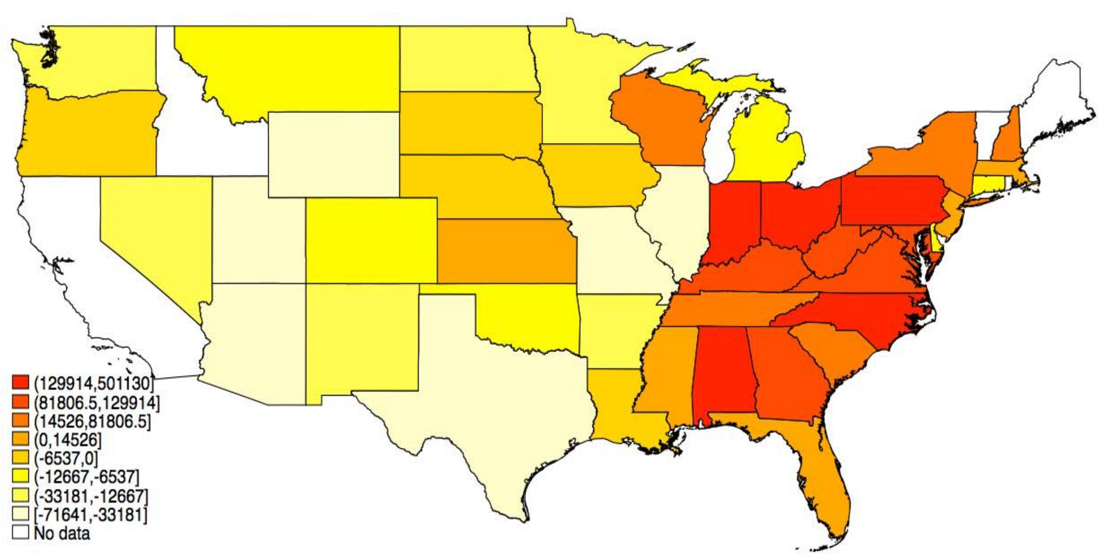

Figure 3(a). Location of Non-NSPS Units by Regulatory

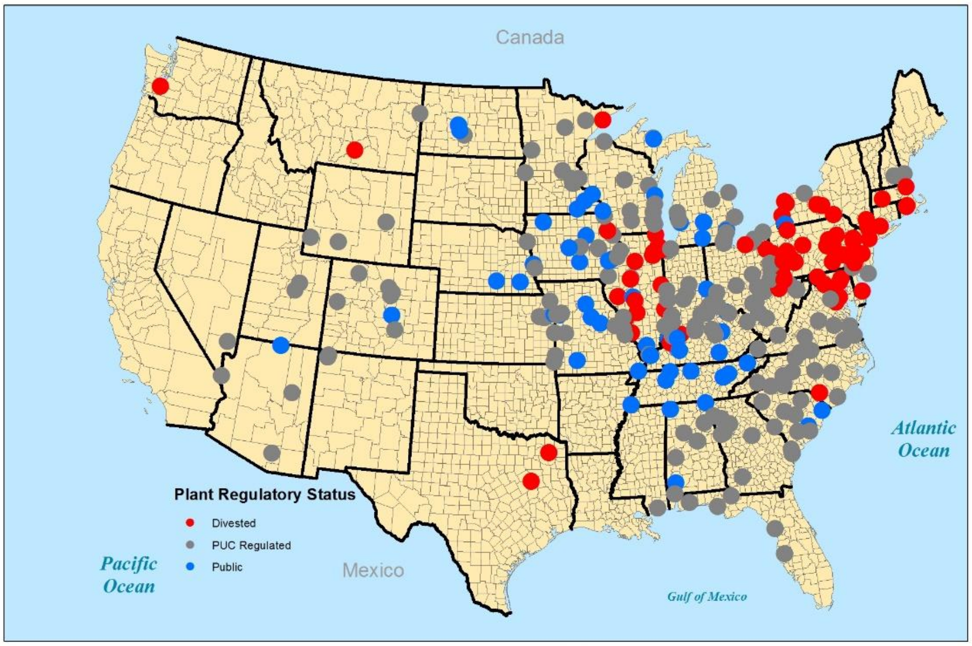

<Double-click here to enter text>

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

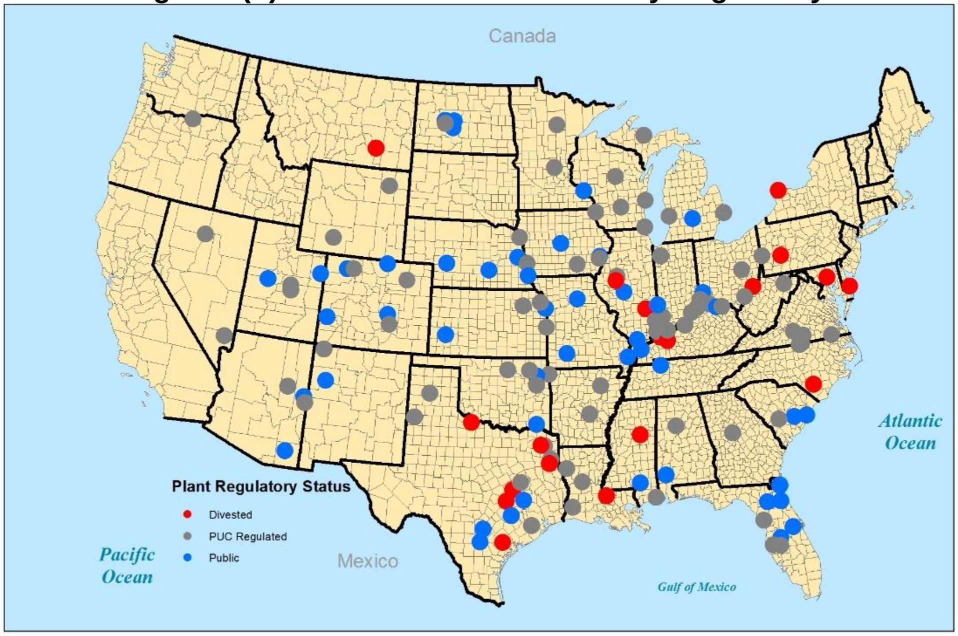

<Double-click here to enter text>

Figure 4. Histogram of Predicted v. Actual Emissions under the ARP

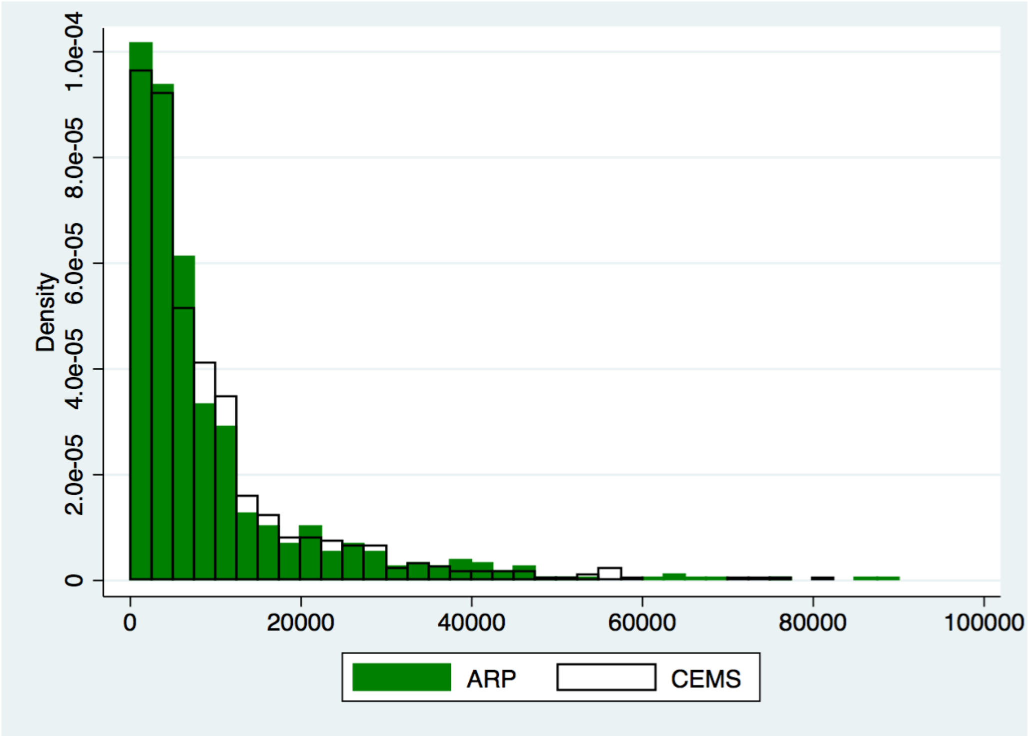

Figure 5. Histogram of Predicted Emission Rates under the ARP

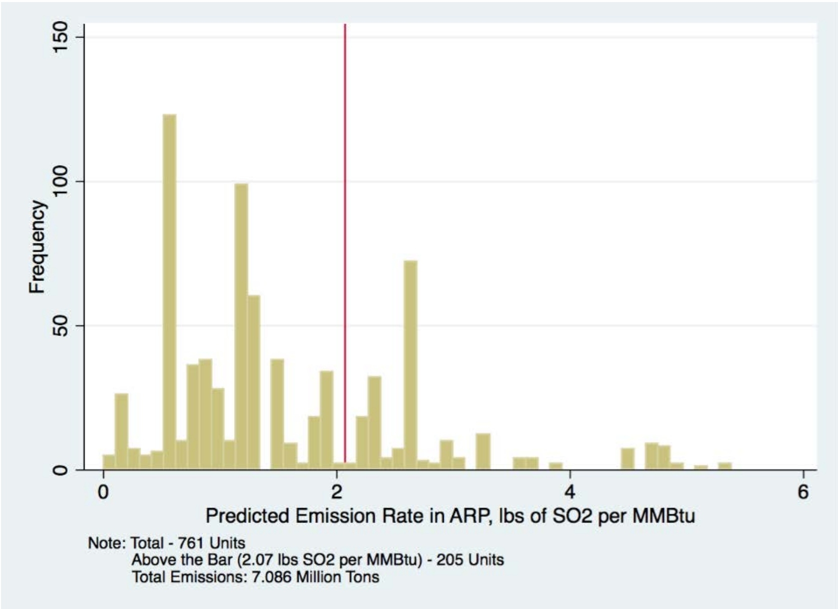

Figure 6. Proportional Difference in Damages: ARP Minus No-Trade Scenario

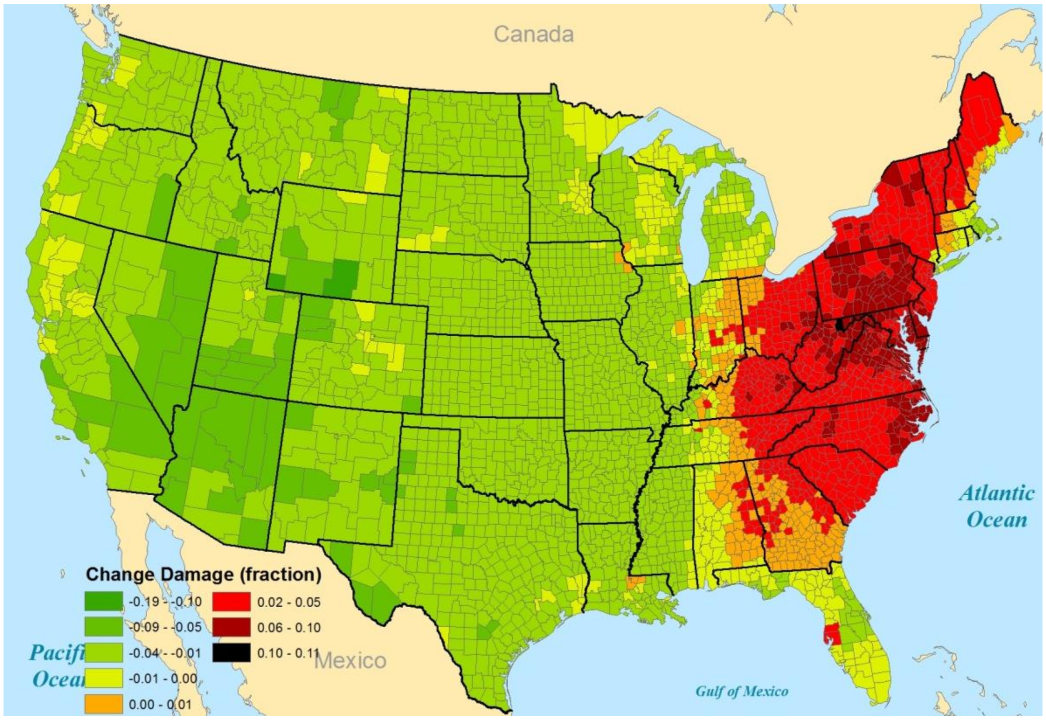

Figure 7. Difference in Health Damages: ARP Minus No-Trade Scenario

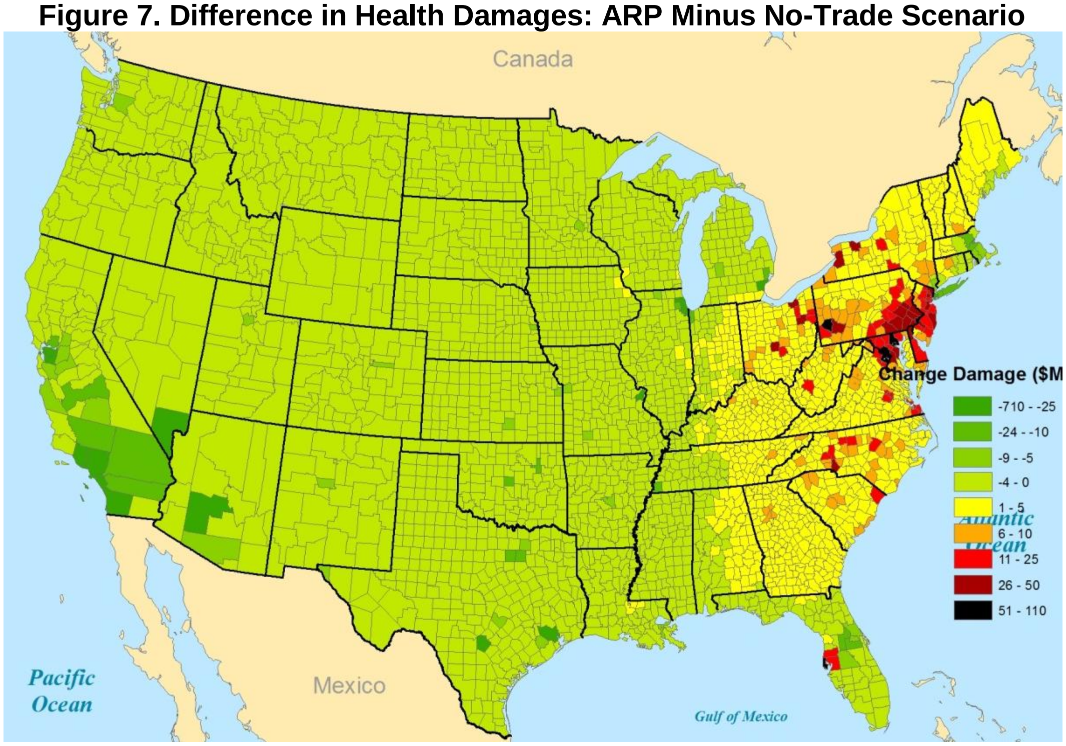

Appendix A. Estimation and Simulation Procedures for the Cost Model

A.1. Estimation of the Mixed Logit Model

We estimate the mixed logit model (equation (3)) treating each plant as a decision maker to allow correlation in unobserved costs for EGUs within each plant, following Fowlie (2010). As noted in the main text, three coefficients are assumed to follow independent Gaussian distributions: the coefficients on scrubbing cost, use of PRB coal (i.e. low sulfur sub-bituminous coal), and the use of Uinta basin coal (i.e. low sulfur bituminous coal). We estimate the model using three thousand Halton draws to simulate the integral in the objective function during maximum likelihood estimation (Train 2009).

Each compliance strategy involves selecting the basin from which to buy coal. Either all coal may be purchased from one basin or 50 percent 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 (2015), to refine the characteristics of the coal purchased within a basin:

1. Start with an estimate of the vector of cost function parameters, $\boldsymbol{\beta}^{(\mathbf{0})}$.
2. For each generating unit i and each alternative j, a coal type k (e.g., county) is selected within alternative j. The coal type k is associated with attributes *COALPRICE(k;j)*, *SULFUR(k;j)* and *ASH(k;j)*. For each unit i and each basin j the k is chosen that minimizes the deterministic version of the compliance cost function in equation (1). Call this $k^*(i,j)$.
3. 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 a new parameter vector $\boldsymbol{\beta}^*$.
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 20-40 iterations depending on the number of control variables.

A.2. Simulation of Compliance Costs and Emissions under a Uniform Performance Standard

For each of the cost models in Table 6, we compute compliance choices under the ARP using equation (5) and compliance costs under the ARP using equation (6). This section describes how a uniform performance standard (UPS) is constructed for each model in Table 6.

For each model in Table 6, compliance choices under the ARP are computed using the modified cost function presented in equation (5) of the text. The emissions corresponding to these predicted compliance choices are calculated by multiplying the emissions rate for each EGU by the average heat input used in 2000–02. This yields aggregate emissions for each model, as shown in Table 8. To determine the UPS corresponding to aggregate emissions we proceed as follows:

1. Set the shadow price of permits in equation (5) equal to zero and start with a uniform emissions standard $\bar{s}^{(0)}$. For each EGU, solve for the compliance strategy that minimizes compliance costs, excluding coal types that violate the uniform emissions standard $\bar{s}^{(0)}$.
2. Compute the emissions associated with the compliance strategy in Step 1. If aggregate emissions exceed predicted emissions under the ARP, repeat Step 1 with $\bar{s}^{(t)} = \bar{s}^{(t-1)} - 0.01$ until emissions in the counterfactual are approximately equal to the emissions under the ARP.
3. Calculate the cost of the UPS for each EGU using equation (6).

A.3. Data Sources and Prediction of Coal Prices and Scrubbing Costs

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.

Prediction of Coal Prices

Coal procurement data were obtained from EIA Form 423, which records coal transactions between mines and plants in the United States. Each transaction contains delivered prices, mine locations, coal quality (heat, sulfur, and ash contents) and contractual arrangements. These data were used to estimate regression models to predict plant-specific coal prices for our sample plants. We run the following non-linear regression for each coal-producing region using transaction-level data from 1991 to 2005:

$$COALPRICE_{ijt} = \exp\begin{bmatrix} \alpha_1 \ln SULFUR_{ijt} + \alpha_2 \ln ASH_{ijt} \\ +\alpha_3\left(\ln SULFUR_{ijt}\right)^2 + \alpha_4\left(\ln ASH_{ijt}\right)^2 \\ +\alpha_5\left(\ln SULFUR_{ijt}\right) \times \left(\ln ASH_{ijt}\right) + \alpha_6 SPOT_{ijt} + \delta_t \end{bmatrix}$$
$$+ \tau DISTANCE_{ij} + e^{\varepsilon_{ijt}} \ldots\ldots\ldots\ldots\ldots\ldots (A.1)$$

where $COALPRICE_{ijt}$ is the real delivered price of coal (in cents per MMBtu), $SULFUR_{ijt}$ and ASH_{ijt} are the sulfur and ash content of the coal, $SPOT_{ijt}$ is a dummy which equals 1 if the coal transaction is executed on the spot market (or via a contract of a less than 12 months duration), and δ_t is year dummy. $DISTANCE_{ij}$ is the railroad distance (county to county) between coal mine *i* and coal-fired power plant *j* gathered from CTA Transportation Networks, scaled by the heat content of coal. Therefore, τ measures transportation costs per ton-mile. Table A.1 tabulates regression results. Across all models there is a price premium for lower sulfur content, lower cost for coal sold on a spot market or via a short-term contract. Aligning with EIA estimates, the transportation rates from each basin also vary. Equation (A.1) is used to predict coal prices by basin, as shown in Figure A.1

To match the plant-level coal-procurement 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 markedly 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.

Prediction of Scrubbing Costs

Scrubber costs (both capital and operating costs) and attributes, including removal rate, hours of operation and age of the scrubbers, were obtained from EIA Forms 767 and 860. Because we need to predict scrubbing costs for all units, predicting scrubber costs using only those units that chose to install a scrubber may induce selection-bias in the estimates (Keohane 2004). We therefore estimate Heckman selection models for capital and for operating costs, using regulation status, sulfur premium on coal and geographic variables as variables excluded from the cost models but included in the selection equations.

Operating costs for scrubbers, as recorded in the EIA surveys, do not include energy costs associated with operating the scrubber. This cost is often substantial (constituting about 5-10% of the total cost of a scrubber) therefore it is important for us to account for it (Bellas and Lange, 2008). We therefore estimate a Heckman selection model similar to those above, with scrubber energy consumption as the dependent variable.

Table A.2 shows regression estimates for operating costs, capital costs and energy consumption. Across all specifications Log(Coal Consumption), a measure of the size of the EGU, is less than one, suggesting economies of scale in scrubber operation. The selection equation indicates that a scrubber is more likely to be installed on an EGU if the plant faces a higher sulfur premium or the unit is regulated under New Source Performance Standard (NSPS).

Using the estimates, we predict energy consumption, operating cost and capital cost for each unit in our sample.[32] To convert energy consumption (in MWh) into energy costs for scrubbers, we assume that the operating heat rate for all units in our sample is 11 MMBtu/MWh, which is the mean operating heat rate found in Linn et al. (2014). We multiply the energy consumption by the assumed operating heat rate and coal price and divide by the total fuel consumption to compute scrubber energy cost per MMBtu (of the unit). We find an energy cost of about 4 cents per MMBtu (in 1995$), approximately 5% of the total operating plus capital costs of a scrubber.

[32] Since we estimate costs based on a logarithmic regression, the consistent estimate of the predicted cost takes the form exp $(\hat{\mu} + 0.5\,\hat{\sigma}^2)$, where $\hat{\mu}$ and $\hat{\sigma}^2$are the estimated mean and variance of log(cost).

Figure A1. Trends in Predicted Coal Prices by Basin.

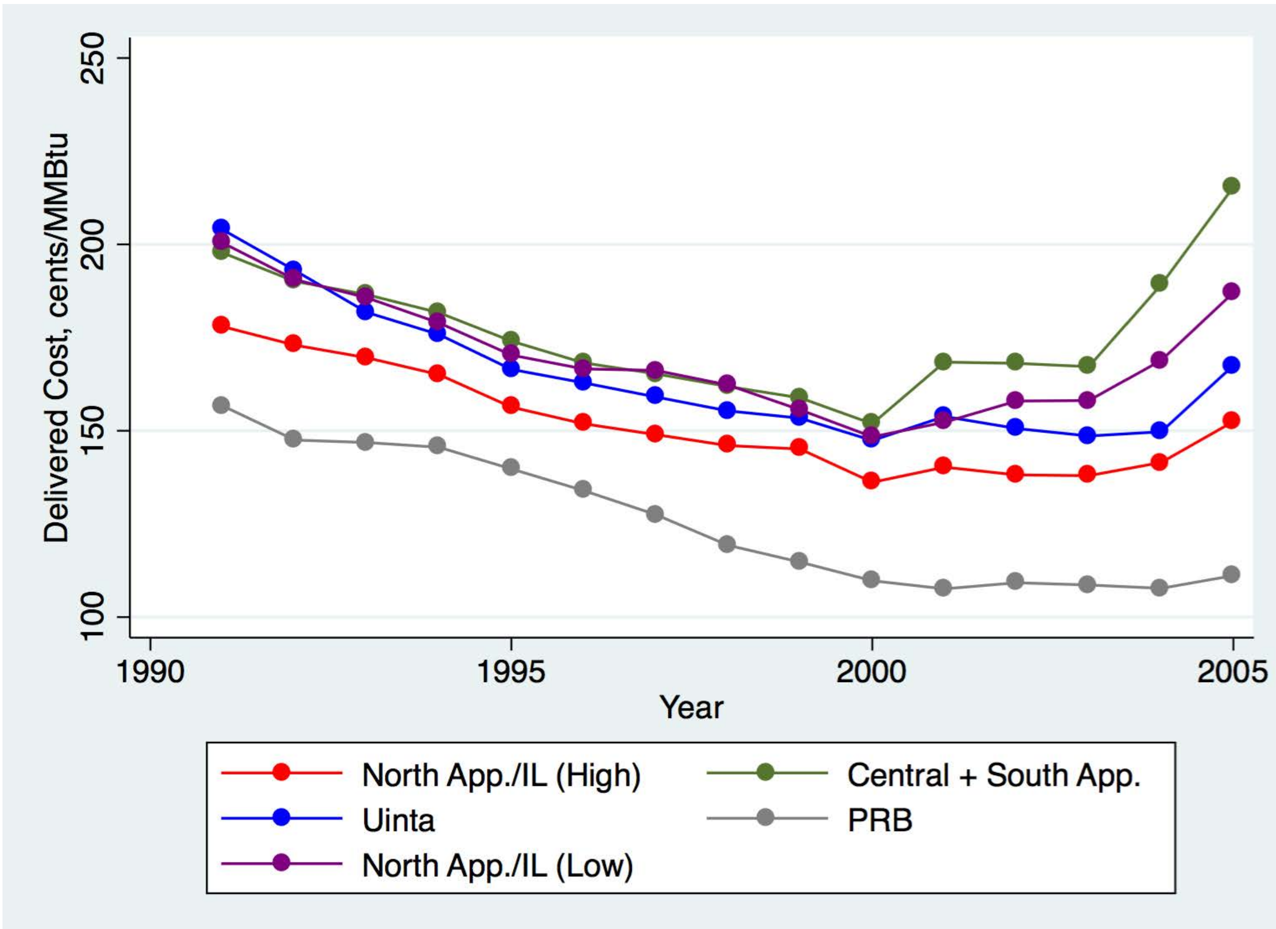

Table A.1: Coal Price Estimates

	North App. (High)	Central App.	South App.	Illinois (High)	Uinta	PRB	North App. (Low)	Illinois (Low)
ln SULFUR	-0.724***	-0.231***	-0.212**	-0.040	-0.084	-1.172***	-0.193***	-0.317***
	(0.103)	(0.015)	(0.107)	(0.036)	(0.078)	(0.303)	(0.016)	(0.044)
(ln SULFUR)^2	0.144***	-0.037***	0.055***	0.065***	0.035*	-0.023	-0.017***	-0.039***
	(0.019)	(0.004)	(0.015)	(0.008)	(0.018)	(0.035)	(0.005)	(0.008)
ln ASH	-0.181	0.334***	1.507***	0.441***	-0.507***	1.578**	0.824***	0.966***
	(0.147)	(0.017)	(0.236)	(0.046)	(0.115)	(0.636)	(0.023)	(0.071)
(ln ASH)^2	0.003	-0.109***	-0.374***	-0.125***	0.124***	-0.270*	-0.190***	-0.254***
	(0.037)	(0.004)	(0.052)	(0.011)	(0.024)	(0.141)	(0.005)	(0.017)
ln SULFUR x ln ASH	0.181***	0.065***	-0.175***	-0.048**	0.039	0.623***	0.014*	0.087***
	(0.047)	(0.007)	(0.046)	(0.019)	(0.034)	(0.134)	(0.007)	(0.022)
Spot	-0.273***	-0.056***	-0.230***	-0.106***	-0.192***	-0.339***	-0.080***	-0.096***

	(0.008)	(0.001)	(0.009)	(0.003)	(0.008)	(0.010)	(0.002)	(0.004)
Distance	0.575***	1.368***	0.534***	1.717***	0.829***	0.903***	1.392***	1.449***
	(0.057)	(0.009)	(0.131)	(0.020)	(0.013)	(0.008)	(0.017)	(0.024)
Year = 1992	5.544***	4.725***	3.737***	4.518***	5.398***	2.316***	4.126***	4.080***
	(0.158)	(0.019)	(0.273)	(0.052)	(0.139)	(0.717)	(0.026)	(0.073)
Year = 1994	5.487***	4.668***	3.653***	4.451***	5.260***	2.289***	4.047***	3.982***
	(0.158)	(0.019)	(0.272)	(0.052)	(0.140)	(0.717)	(0.026)	(0.073)
Year = 1996	5.391***	4.562***	3.569***	4.322***	5.142***	2.115***	3.959***	3.862***
	(0.159)	(0.019)	(0.272)	(0.052)	(0.140)	(0.717)	(0.026)	(0.073)
Year = 1998	5.300***	4.501***	3.605***	4.292***	5.067***	1.839**	3.923***	3.822***
	(0.159)	(0.019)	(0.272)	(0.052)	(0.140)	(0.717)	(0.026)	(0.074)
Year = 2000	5.197***	4.416***	3.464***	4.192***	4.981***	1.607**	3.762***	3.743***
	(0.159)	(0.019)	(0.273)	(0.052)	(0.140)	(0.718)	(0.027)	(0.073)
Year = 2002	5.106***	4.572***	3.438***	4.284***	5.016***	1.592**	3.885***	3.781***
	(0.158)	(0.019)	(0.273)	(0.052)	(0.139)	(0.718)	(0.026)	(0.073)

Year = 2004	5.203***	4.737***	3.511***	4.281***	5.007***	1.550**	4.014***	3.814***
	(0.158)	(0.019)	(0.272)	(0.052)	(0.140)	(0.717)	(0.026)	(0.073)
Observations	11520	129598	4944	25178	11399	40881	49999	11679
Adj. R Squared	0.147	0.438	0.377	0.442	0.345	0.329	0.396	0.482

Note: Robust standard errors are reported in parentheses. ***, **, * denote statistical significance at 99%, 95%, 90% level. All year fixed effects are included in all models however only the even numbered year dummies are displayed. A constant is excluded in all models.

Table A.2: Scrubber Cost Estimates

	(1)	(2)	(3)	(4)	(5)	(6)
	Energy Cons.		Operating Cost		Capital Cost	
	OLS	Heckman	OLS	Heckman	OLS	Heckman
Main Estimation Equation						
Log(Scrub Age)	0.373**	0.382***	0.174*	0.243**		
	(0.145)	(0.122)	(0.097)	(0.115)		
Log(Coal Consumption)	0.722***	0.800***	0.520***	0.576***	0.855***	0.888***
	(0.113)	(0.131)	(0.065)	(0.123)	(0.104)	(0.148)
Log(Removal Rate)	0.771***	0.826***	0.438***	0.562***	0.936***	0.991***
	(0.216)	(0.212)	(0.133)	(0.127)	(0.232)	(0.238)
Log(Op. Hours)	0.807***	0.749***	0.348*	0.351*	-0.497**	-0.443*
	(0.285)	(0.225)	(0.181)	(0.195)	(0.234)	(0.234)
Flue Gas Entering Rate			0.961***	1.079***		
			(0.290)	(0.327)		
Log(Install. Age)					0.420**	0.796***
					(0.162)	(0.216)
Selection Equaton						
Log(Sulfur Premium)		0.458		0.365		1.620***
		(0.421)		(0.400)		(0.427)
NSPS		0.837***		0.808***		0.516*
		(0.116)		(0.117)		(0.309)
Log(Coal Consumption)		0.349***		0.362***		0.447***
		(0.047)		(0.046)		(0.064)
Northeast		-0.713**		-0.767**		-0.848**
		(0.338)		(0.339)		(0.380)
South		-1.093***		-1.209***		-0.812***
		(0.274)		(0.248)		(0.310)
Midwest		-1.102***		-1.212***		-1.443***
		(0.308)		(0.287)		(0.319)

Year = 1992	-0.369***	-0.348***	-0.085**	-0.105***		
	(0.139)	(0.133)	(0.037)	(0.034)		
Year = 1994	-0.491***	-0.486***	-0.182***	-0.204***		
	(0.179)	(0.166)	(0.067)	(0.069)		
Year = 1996	-0.543***	-0.539***	-0.415***	-0.418***		
	(0.199)	(0.191)	(0.116)	(0.117)		
Year = 1998	-0.462**	-0.459**	-0.449***	-0.469***		
	(0.222)	(0.211)	(0.103)	(0.106)		
Year = 2000	-0.521***	-0.512***	-0.570***	-0.606***		
	(0.165)	(0.158)	(0.115)	(0.128)		
Year = 2002	-0.649***	-0.616***	-0.540***	-0.563***		
	(0.186)	(0.179)	(0.164)	(0.186)		
Year = 2004	-1.003***	-0.835***	-0.618***	-0.626***		
	(0.232)	(0.206)	(0.107)	(0.133)		
Install Year <= 1979					0.521	0.598
					(0.399)	(0.414)
Install Year = [1980,1989]					0.334	0.578*
					(0.262)	(0.348)
Install Year = [1990,1994]					-0.350	0.035
					(0.398)	(0.547)
No. of Observations	2624	18858	2620	18833	373	1615

Note: Robust standard errors are reported in parentheses. ***, **, * denote statistical significance at 99%, 95%, 90% level. All year and decade fixed effects are included in models (1)-(4) and (5)-(6) respectively. We do not report all year fixed effects in models (1)-(4) for exposition purposes. The excluded categories for time/year fixed effects are: year 1991 for models (1)-(4); installation year > 1995 for models (5)-(6) The dependent variables in models (1)-(2), (3)-(4), and (5)-(6) are Log(Energy Consumption), Log(Operating Costs) and Log(Capital Costs) respectively.

Appendix B. Monte Carlo Analysis of Health Damages

In this appendix we evaluate the impact of uncertainty in air quality modeling, dose-response functions linking $PM_{2.5}$ exposure to human health outcomes, and the VSL on health damages. In Monte Carlo simulations each of the major input parameters to AP2 is treated as a random variable. From each of these distributions a realization is drawn and damages are computed under both the ARP and the no-trade counterfactual. The damage estimates are stored and the process is repeated 1,000 times; each time damages are computed conditional on a different draw from the input distributions. This procedure yields an empirical distribution of damage estimates for the observed ARP, the no-trade counterfactual, and for the difference between the two.

Air quality model uncertainty is represented by estimating the standard errors associated with the ($T_{k,m}$) in the source-receptor matrices, for different bearing and distance bands between each source (k) and receptor pair (m), (see Muller, 2011). These standard errors are then used to construct empirical distributions for each transfer coefficient.

The concentration-response functions that govern the $PM_{2.5}$-mortality link are empirical functions estimated in the epidemiological literature. The reported standard errors of parameters in these functions are used to construct empirical distributions for the concentration-response functions. Finally, the VSL distribution is built using the mean and standard deviation reported by USEPA (EPA 2012).

In order to estimate the contributions of uncertainty from each input parameter distribution to the cumulative uncertainty in damage estimates, we run several additional simulations. We begin with the case in which the parameters of the air quality model, VSL and dose-response function are all treated as random. The mean and standard deviation of the difference in damages between the ARP and the no-trade scenario are calculated as well as the coefficient of variation (standard deviation divided by the mean). Next, one of the input parameters, the VSL for example, is set to its deterministic value and the Monte Carlo simulations are re-run. The coefficient of variation is re-computed and compared to the all-stochastic case. A large drop in the coefficient of variation indicates that uncertainty in the VSL (in this example) contributes a significant share of the total uncertainty in damages. This process is then repeated for the air quality model and the dose-response parameter.

Table B.1 shows the results from these simulations. With all parameters modeled as stochastic, the mean difference in damage between the ARP and the no-trade counterfactual is \$2.35 billion with a 95% confidence interval of \$2.25 billion to \$2.45 billion. The standard deviation is $1.61 billion, yielding a coefficient of variation of 0.69. The next simulation treats the source-receptor matrix as deterministic while the VSL and mortality dose-response inputs are stochastic. This does not affect the mean difference in damage, but the standard deviation falls from 0.69 to 0.55. Fixing the mortality dose-response parameter to its deterministic value has a very similar effect to the air quality model. However, when the VSL is modeled deterministically, the coefficient of variation falls to 0.52. This suggests that the uncertainty in the VSL parameter contributes the greatest share of statistical uncertainty to the damage estimates.

Table B.1. Monte Carlo Simulation Results: Difference in Damage between ARP and No-Trade Counterfactual.

	All Inputs Stochastic	Air Quality Model	Mortality Dose-Response	VSL
Policy Scenario	**(ARP – NTC)**	**(ARP – NTC)**	**(ARP – NTC)**	**(ARP – STD)**
Mean	2.35	2.36	2.37	2.34
	(2.25, 2.45)[A]	(2.28, 2.44)	(2.28, 2.45)	(2.26, 2.42)
Standard deviation	1.61	1.29	1.30	1.22
Coefficient of Variation	0.69	0.55	0.55	0.52

[A] = 95% confidence intervals in parenthesis.

ARP – NTC = Difference in damage between observed emissions and the no-trade counterfactual.

References

Bellas, Allen, and Ian Lange. 2008. "Impacts of Market-based Environmental and Generation Policy on Scrubber Electricity Usage." *The Energy Journal*, 29 (2): 151-164.

Chan, H. S. Ron 2015. "How Large Are the Cost Savings from Emissions Trading? An Evaluation of the U.S. Acid Rain Program." Unpublished.

EPA (US Environmental Protection Agency). 2012. Environmental Benefits Mapping and Analysis Program 2015. http://www2.epa.gov/sites/production/files/2015-04/documents/benmap-ce_user_manual_march_2015.pdf

Keohane, Nathaniel O. 2004. "Environmental Policy and the Choice of Abatement Technique: Evidence from Coal-fired Power Plants." Unpublished.

Linn, Joshua, Erin Mastrangelo, and Dallas Burtraw. 2014. "Regulating Greenhouse Gases from Coal Power Plants under the Clean Air Act." *Journal of the Association of Environmental and Resource Economists*, 1 (1/2): 97-134.

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