Journal of Environmental Economics and Management 88 () 180–209

Contents lists available at ScienceDirect

Journal of Environmental Economics and Management

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

The impact of trading on the costs and benefits of the Acid Rain Program[☆]

H. Ron Chan[a], B. Andrew Chupp[b], Maureen L. Cropper[c,*], Nicholas Z. Muller[d]

[a] University of Manchester, UK
[b] Georgia Institute of Technology, USA
[c] Department of Economics, University of Maryland, Resources for the Future, 3114H Tydings Hall, College Park, MD 20742, USA
[d] Carnegie Mellon University, USA

ARTICLE INFO

Article history:
Received 30 March 2017
Available online 8 December 2017

JEL nos:
Q51
Q52
Q58

Keywords:
Sulfur dioxide
Acid rain
Performance standards
Health effects
Pollution permits
Cap and trade

ABSTRACT

We quantify the cost savings from the Acid Rain Program (ARP) by comparing compliance costs for 761 coal-fired generating units under the ARP with compliance costs under a counterfactual uniform performance standard (UPS) that would have achieved the same aggregate emissions in 2002. In 2002, we find compliance costs to be \$200 million (1995\$) lower and health damages to be \$170 million (1995\$) lower under the ARP. We also compare health damages associated with observed SO_2 emissions from all ARP units in 2002 with damages from a no-trade counterfactual. Damages under the ARP are \$2.1 billion (1995\$) higher than under the no-trade scenario, reflecting allowance transfers from units in the western US to units in the eastern US with larger exposed populations.

Introduction

Motivation

The U.S. Acid Rain Program (ARP), enacted under Title IV of the 1990 Clean Air Act Amendments, is widely cited as an example of a successful cap and trade program—one that achieved huge reductions in sulfur dioxide (SO_2) emissions from US coal-fired power plants at a lower cost than a comparable command-and-control regulation (Ellerman et al., 2000; Stavins, 1998; Schmalensee and Stavins, 2013). 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

☆ 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.
We thank Dallas Burtraw, Dick Morgenstern and 3 anonymous referees for their comments.
* Corresponding author.
E-mail address: cropper@econ.umd.edu (M.L. Cropper).

https://doi.org/10.1016/j.jeem.2017.11.004
0095-0696/

in 2000, all electricity generating units (EGUs) greater than 25 MW (MW) were regulated by the ARP. 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 $360 million (1995$) 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]

Were these cost savings realized? There are several studies of Phase I of the ARP which suggest that this was not the case during the early years of the program. Carlson et al. (2000) estimate that actual compliance costs in 1995 and 1996 were slightly higher than they would have been under a uniform performance standard. Swinton (2002, 2004) shows that marginal abatement costs were not equalized across Phase I EGUs, suggesting that the least-cost solution was not achieved. Arimura (2002) finds that uncertainty in regulation led utilities to focus on fuel-switching and blending rather than depending on the allowance market, thus lowering potential cost savings.

An important question is what cost savings were realized once the program was fully operational—i.e., during Phase II of the program. Reductions in compliance costs are a key reason for replacing a command-and-control regulation by a cap-and-trade system. It is, therefore, important, that cost savings be documented. 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 in Phase II are ex ante in nature (Carlson et al., 2000; Ellerman et al., 2000). Carlson et al. (2000) project cost savings based on marginal abatement cost (MAC) functions estimated using pre-ARP (1985–94) data. Their 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. One goal of our paper is to estimate the cost savings associated with allowance trading during Phase II of the ARP compared with an equally stringent uniform performance standard.

There are also concerns that health damages under 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 Fig. 1, in which we estimate the difference in 2002 between $PM_{2.5}$ levels under the ARP and $PM_{2.5}$ levels that 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.

Our approach

The goal of this paper is to compare compliance costs during Phase II of the ARP with the costs of a uniform performance standard that would have achieved the same aggregate emissions as the ARP achieved in Phase II. This is a difficult task. Carlson et al. (2000) project cost savings of the program in the steady state; i.e., once the 8.95 million-ton cap on SO_2 emissions was achieved. In fact, the allowance market never reached a steady state (Schmalensee and Stavins, 2013). Once the health benefits of Title IV were recognized (Burtraw et al., 1998, EPA, 1999), the ARP was replaced by more stringent regulations.[2] In 2003, EPA issued a draft of the Clean Air Interstate Rule (CAIR), which set a cap on SO_2 emissions that was 57% lower than the cap under the ARP. We view compliance decisions made after 2003 as responding to a different regulatory regime. For this reason, our analysis focuses on the year 2002.

In measuring cost savings, we concentrate on the EGUs that were the focal point of the ARP: units that were not already covered by EPA's New Source Performance Standards (NSPS). The NSPS effectively required all EGUs built after 1970 to achieve SO_2 emissions reductions as least as stringent as those enacted by the ARP (see Section *Title IV and Other SO_2 Regulations Facing Coal-Fired Power Plants* below). There were 838 coal-fired EGUs in operation in 2002 not covered by the NSPS. We model the compliance behavior of 761 of these units (hereafter, modeled units).[3]

To estimate the cost savings from the ARP we estimate an econometric model of compliance behavior based on compliance choices observed in 2002. The main methods used to reduce SO_2 emissions are to purchase low-sulfur coal or install a flue-gas desulfurization unit (FGD or "scrubber").[4] Our model is a mixed logit model of the choice of whether to install a scrubber and what type of coal to buy, described by geographic location of coal purchases. This model allows us to predict

[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 $973 million (1990$).

[2] As Schmalensee and Stavins (2013) note, the original motivation for reducing SO_2 under Title IV was to reduce acidic deposition, which can reduce soil quality (through nutrient leaching), impair timber growth, and harm freshwater ecosystems (EPA, 2011). Subsequent research demonstrated the significant health benefits associated with reduced secondary particle formation from lower SO_2 emissions.

[3] As described more fully below, of the 77 EGUs dropped, 36 have no data on coal purchases, 27 purchased coal primarily outside of the United States, and 14 changed to NSPS status shortly after 2002.

[4] Unless otherwise indicated, we use the term low-sulfur coal to refer to coal from the Powder River or Uinta basins.

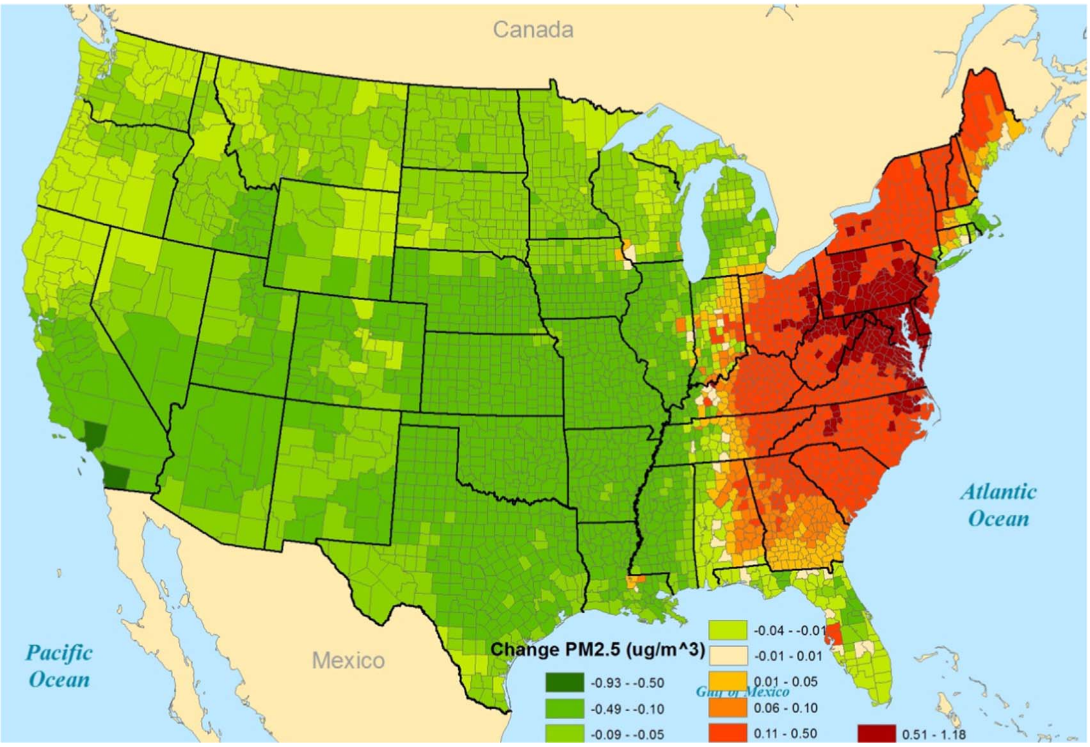

Fig. 1. Difference in $PM_{2.5}$ concentrations in 2002: ARP minus no-trade scenario. This map shows the difference in estimated $PM_{2.5}$ concentrations under the Acid Rain Program and the No-Trade counterfactual. $PM_{2.5}$ concentrations under the ARP are estimated using 2002 emissions from EPA's CEMS database; emissions under the No-Trade Counterfactual assume that each unit emits SO_2 according to its 2002 permit allocation, plus any 2002 drawdown of the allowance bank. Green values indicate lower concentrations under the ARP than under the No-Trade counterfactual, while red values indicate higher concentrations under the ARP. This geographic pattern suggests that trading moved emissions from the West to the East of the US. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

compliance choices under the ARP and under a uniform performance standard (UPS) that achieves the same aggregate emissions as our modeled 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, for the year 2002.

We acknowledge the limitations of our approach. The ARP was a multi-year program. Allowance trading under the ARP gave firms flexibility in achieving emissions reductions: purchasing an allowance is, in effect, an option to install a scrubber or fuel switch at a later date. A fairer comparison to a UPS would involve modeling compliance as a dynamic decision and comparing a multi-year UPS to the ARP, rather than using a static model.

We also acknowledge that a comparison of the ARP and a UPS in 2002 is likely to understate the cost savings that would ultimately have been achieved by the ARP, for two reasons. The first is that the aggregate emissions of our 761 modeled units in 2002 are 40% higher than allowances issued to these units for the year 2002. This implies that our 2002 UPS is much less stringent than the UPS needed to achieve the long-run emissions goal of the ARP. With a less stringent cap, we would expect lower gains from trade than those estimated by Carlson et al. (2000). The second reason is that we ignore the benefits of the banking provisions of the ARP. By letting plants bank allowances for future use, the ARP allowed compliance costs to be postponed. Calculating cost savings under the ARP in a single year fails to capture these benefits.

We also wish to ask whether the health damages associated with SO_2 emissions were lower under the ARP than under a regulatory regime that did not allow trading. We do this using two scenarios. In the first, we compare damages under the ARP and the counterfactual UPS for the 761 coal-fired generating units used to estimate cost savings under the ARP. Although a natural comparison, the counterfactual UPS in fact incorporates trading. As noted above, the emissions of our 761 modeled units in 2002 exceeded their allocated emissions for the year 2002 by 40%; two-thirds of this difference was covered by allowances purchased through trading. Even if damages under the ARP and the UPS counterfactual are approximately equal, trading could have had an impact on health if the allowances obtained by our modeled units came from plants with lower marginal damages.

To examine the full impact of trading on health damages, we construct a no-trade scenario and compare damages under the no-trade scenario with damages under the ARP. The no-trade scenario includes all EGUs covered by the ARP and forces them to emit at the rate prescribed by their initial allocation of allowances, plus any drawdowns of their allowance banks observed in 2002. The impetus of this comparison is to isolate the impact of trading per se, rather than compliance with a uniform performance standard, on health damages.

To compare health damages under the ARP and our counterfactual scenarios, 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 (2000\$) value of a statistical life (VSL).

Our findings

We estimate the cost savings from emissions trading to be approximately \$200 million (1995\$) in 2002. This is a much smaller estimate than that of Carlson et al. (2000). There are at least two reasons for this. One reason, suggested above, is that the emissions cap in our simulations, which equals actual emissions of our modeled plants in 2002, is much less stringent that the steady-state cap facing these units. With a less stringent cap and a less stringent UPS we would expect lower gains from trade. The second reason is that Carlson et al. (2000) assumed that cost minimization would preclude the installation of scrubbers at coal-fired plants not covered by the NSPS. In fact, the number of EGUs with scrubbers at these plants increased by 50% between 1995 and 2002.

Regarding health impacts, we estimate that in 2002, health damages associated with emissions from our modeled plants would have been slightly (0.18%) higher under the UPS than under the ARP. This is not surprising. EGUs that are below the UPS under the ARP increase their emissions; those above the UPS under the ARP reduce their emissions. Although emissions under the ARP and the counterfactual UPS occur in different places, the exposed populations are high in both cases.

This does not mean, however, that the ARP had no impact on health. In 2002, our modeled units emitted approximately 7 million tons of SO_2, 2 million tons more than their allowance allocations for the year 2002 of 5 million tons. Two-thirds of these allowances were purchased from other EGUs. The health effects of the ARP depend on the location of sellers versus buyers of allowances. To capture the health impacts of trading, we estimate the health damages associated with the observed emissions of *all* units participating in the ARP and compare them with the damages that would have resulted had units emitted SO_2 at a rate determined by the initial distribution of allowances. We find that damages under the ARP exceeded damages under the no-trade counterfactual by \$2.1 billion (1995\$) (1.8% of damages under the ARP). This is because under the ARP, NSPS units and non-coal units transferred or sold allowances to coal-fired units not covered by the NSPS. 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 damages under our no-trade counterfactual should not be compared with the cost savings experienced by our modeled plants: the former are based on all units covered by the ARP; the latter are based on only 761 EGUs. It would therefore be inappropriate to conclude that trading under the ARP yielded negative net benefits. The no-trade counterfactual does, however, suggest the importance of considering differences in marginal damages when designing a cap and trade system, a point made by Muller and Mendelsohn (2009). We emphasize that the difference in compliance costs between the UPS and the ARP for our 761 modeled units indicate significant cost savings from cap and trade. Our estimate of \$200 million (1995\$) is a lower bound to cost savings from trading for the reasons given above.

The paper is organized as follows: Section *Background* 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 *Modeling Compliance Behavior under the ARP*. In Section *Simulation Results* we simulate compliance behavior under a uniform performance standard and compare compliance costs and emissions under the standard and the ARP for modeled units. In Section *The Health Impacts of Trading* we estimate health damages under the ARP and contrast them with damages under a UPS and a scenario in which all units emit SO_2 at a rate determined by the initial distribution of allowances. Section *Conclusions* discusses the policy implications of our results.

Background

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% 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 input in the 1985–87 reference period. Units were also allowed to voluntarily enroll in Phase I, either as substitutes for Table A units or to compensate for reductions in output at Table A units.[5] In Phase II, beginning in 2000, the program was extended

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

to all generating units with a capacity exceeding 25 MW, approximately 1100 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 involved setting emissions limits for existing stationary sources within each state, including power plants. The emissions limits imposed on SO_2 emissions from coal-fired power plants by state and local governments, which we incorporate into our analysis, were sometimes more stringent than the 1.2 pounds of SO_2 per MMBtu of heat input targeted under the ARP.[6] The 1970 CAA also imposed New Source Performance Standards (NSPS) on newly constructed stationary sources, including power plants. Plants built between 1971 and September 1977 were required to reduce their SO_2 emissions to 1.2 pounds per MMBtu. The NSPS enacted under the 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% from ARP levels. Although CAIR was later vacated by the District of Columbia Circuit Court and replaced by the Cross-State Air Pollution Rule (CSAPR),[7] it was clear after December 2003 that EPA aimed to regulate SO_2 emissions from power plants more stringently than under the ARP. We view this as a change in the regulatory regime that effectively signaled the end of the ARP (Schmalensee and Stavins, 2013).

Compliance in Phase II of the Acid Rain Program

Our analysis focuses on the period when the ARP was fully operational—when all coal-fired EGUs were covered by the program—but before plans were announced to more stringently regulate SO_2 emissions. We focus on the year 2002, the third year of Phase II.[8] In 2002, 1,075 coal-fired generating units were regulated under the ARP (see Table 1). These included 237 units that were also regulated under the NSPS and 838 units that were not. We model the compliance decisions of 761 of the 838 units. The 77 units whose behavior we cannot model either have no data on coal purchases (36), buy coal from abroad (27) or are converted to NSPS units shortly after 2002 (14).[9] 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.[10] 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 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, 80 (11%) of our modeled units had installed scrubbers. An additional 28% of these units achieved compliance with the ARP by burning coal exclusively from the Powder River basin (PRB) or the Uinta basin. Remaining units used a combination of blending low-sulfur coal with higher-sulfur coal, using banked allowances, or purchasing additional allowances. The 77 units that we cannot model are, on average, emitting SO_2 at a lower rate than our modeled units (22% of them have installed scrubbers) although they are similar in average heat input. Fig. 2 (a) which shows the location of both modeled and out-of-sample units indicates that out-of-sample units are not clustered geographically.

Table 2 describes compliance choices of the 838 units not covered by the NSPS, according to a unit's status under electricity sector deregulation in 2002. Units may be divested (owned by independent power producers), PUC-regulated (investor-owned utilities whose rates were set by PUCs), or publicly owned. The table indicates the percentage of units that scrubbed and the percentage that used exclusively low-sulfur or high-sulfur coal in 2002. Remaining units blended coal of various sulfur contents. Fig. 2(b) shows the location of modeled units by regulatory status. The percentage of modeled units scrubbing emissions does not differ significantly by regulatory status, although it is slightly higher for publicly-owned units

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

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

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

[9] These 14 units were converted to NSPS units because of New Source Review.

[10] This could be achieved by burning exclusively coal from the Powder River and Uinta basins or by blending this coal with coal from the eastern US.

Table 1
Characteristics of operating coal-fired EGUs in 2002. *Sources:* Data on emissions, heat input, number of scrubber units and emissions rates come from EPA's Continuous Emissions Monitoring System (CEMS). Allocated permits are from the Acid Rain Program compliance records. Coal usage is from EIA form 480.

	NSPS	Not covered by NSPS	
		Sample	Out of sample
Number of Units	237	761	77
Emissions: Total (Tons)	2,310,801	6,881,833	666,753
Emissions: Average (Tons)[a]	9750.2	9043.1	8659.1
Heat Input: Total (M of MMBtu)	8260.0	10770.8	1180.2
Heat Input: Average (M of MMBtu)[a]	34.85	14.15	15.33
Allocated Permits: Total	2,728,357	4,949,239	516,238
Allocated Permits: Average[a]	11512.1	6503.6	6704.4
Number of Scrubbed Units	120	80	17
Number of Units without Scrubbers using PRB or Uinta Coal	85	210	6
Average SO2 Emission Rate (lbs. per MMBtu)[a]	0.5573	1.5101	1.2290

[a] Denotes unweighted average.

(11.9%) and PUC-regulated units (10.5%) than for divested units (9.3%). Most 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.[11] On average, divested and PUC-regulated units were net purchasers of allowances, while publicly owned units were net sellers.[12]

Modeling compliance behavior under the ARP

A model of compliance choice

We model compliance behavior 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.[13] We assume that this choice is made to minimize the costs of fuel purchased plus the costs of compliance with the ARP. Because the type of coal purchased affects SO_2 emissions, we refer to the sum of these as 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 in our econometric model 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 (i.e., the cost of permits); 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 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, which must be covered by permits, are the product of the sulfur content of the coal burned times the fraction of emissions not removed by scrubbing. The coefficient on emissions can be used to calculate the shadow price of emissions, i.e., the perceived price of permits, which we compare to actual allowance prices. In a sensitivity analysis, we allow the shadow price of emissions to vary according to the regulatory status of the unit. PUC-regulated utilities, which could pass compliance costs on to ratepayers, may have attached a different shadow price to emissions than divested utilities (Fullerton et al., 1997; Sotkiewicz and Holt, 2005, Cicala, 2015). We test this empirically.

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

[11] 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 units not covered by the NSPS were installed after divestiture.

[12] We compute permits demanded in 2002 by subtracting 2002 allowances allocated to the unit, plus the allowance bank at the end of 2001, from allowances held at the end of 2002 (before 2002 emissions are deducted). If this number is positive, it indicates that allowances must have been obtained through trades. If it is negative, then allowances were sold. We aggregate permit demand (supply) by regulatory status. The aggregate demand is negative for publicly-owned units, and positive for the other two groups.

[13] As described more fully below, there are 36 coal purchase options which, when crossed with the decision to scrub or not scrub, yields 72 compliance choices.

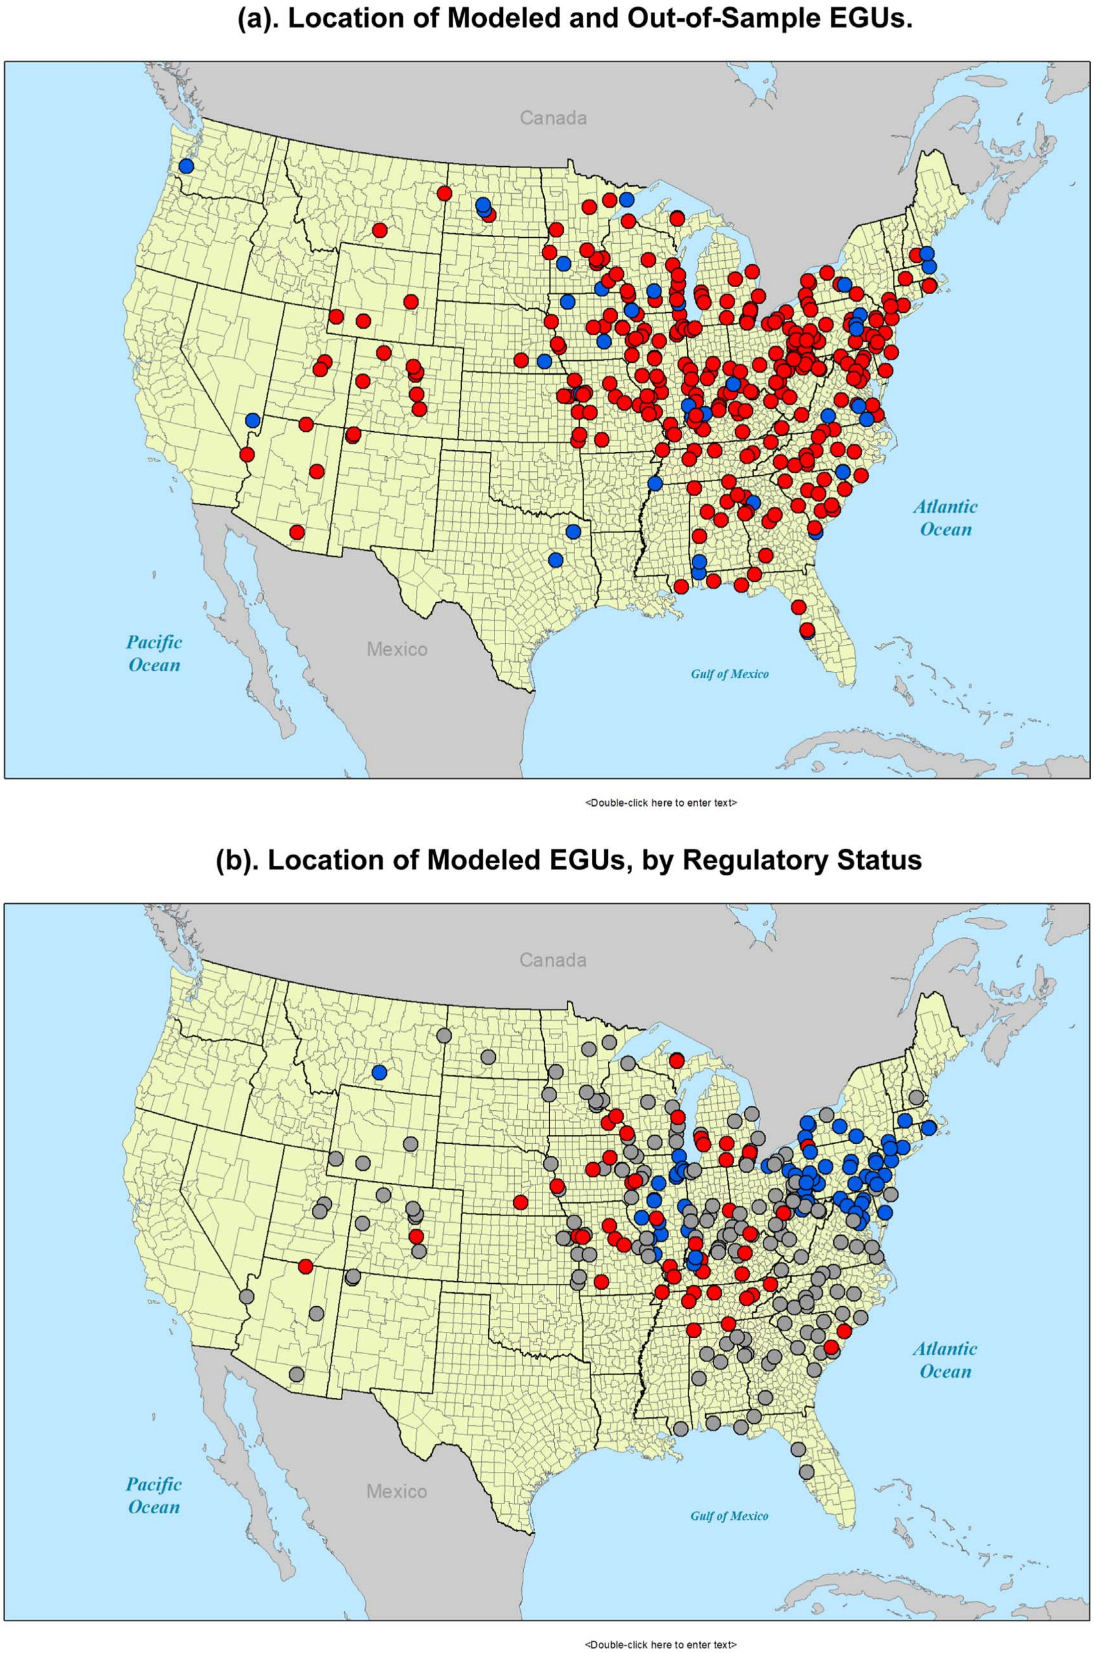

Fig. 2. (a). Location of modeled and out-of-sample EGUs. This map shows the location of 838 coal-fired generating units operating in 2002 that were not subject to New Source Performance Standards. Red dots indicate the 761 EGUs whose behavior we model; blue dots indicate the 77 EGUs dropped from our sample. (b). Location of modeled EGUs, by regulatory status. This map shows the location of our 761 modeled units, distinguished by regulatory status in 2002. Red dots represent publicly-owned plants. Gray dots are IOUs that are PUC regulated. Blue dots indicate divested plants. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 2
Compliance choices in 2002 by regulatory status. *Sources:* Data on regulatory status come from EIA, as do data on coal purchases. Data on emissions and scrubber use come from CEMS.

	Sample (Modeled) Units			**Out of Sample Units**		
	Divested	PUC-Regulated	Publicly Owned	Divested	PUC-Regulated	Publicly Owned
% Scrubbed	9.3	10.5	11.9	26.9	27.8	0.0
% Using low sulfur coal (no scrubber)	19.9	29.4	30.6	15.4	13.9	60.0
% Using high sulfur coal (no scrubber)	44.7	11.6	18.7	7.7	0.0	6.7
Total No. of Units	161	466	134	26	36	15

Notes: "Low (High) sulfur use" refers to units where the majority of purchases originate from the Uinta or Power River Basins (North Appalachian or Illinois Basins). The omitted category is use of coal from other basins or the blending of low and high sulfur coal. For units without coal purchase data, sulfur use is inferred based on the unit's observed emissions rate.

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 plant managers wish to minimize compliance costs (i.e., the costs of fuel plus the costs of abating SO_2 emissions) per MMBtu of heat input for all EGUs. Specifically, we assume that for each EGU i, the compliance option j is chosen that minimizes C_{ij}, the latent (unobserved) compliance cost per MMbtu, subject to the constraint that the EGU not violate state and local emissions standards, $\overline{SULFUR}_i$, which may limit SO_2 emissions per MMBtu. As indicated in Eq. (1), C_{ij} is the sum of $COALPRICE_{ij}$, the delivered coal cost of option j to EGU i in cents per MMBtu, and $SCRUBCOST_{ij}$, the predicted cost in cents per MMBtu of scrubbing unit i if option j is chosen.[14] Both *COALPRICE* and *SCRUBCOST* are predicted for each EGU, as discussed below. The next four terms represent compliance costs whose value we infer from the coefficients of the cost model.[15] The first of these is the cost of SO_2 permits, which is inferred from SO_2 emissions, $EMISSIONS_{ij}$. $EMISSIONS_{ij}$ equals emissions of SO_2 in pounds per MMBtu associated with option j times the average removal rate achieved by scrubbing (assumed to be 85%) if option j entails scrubbing. ASH_{ij} measures the ash content of coal, measured in pounds per MMBtu. PRB_j and $UINTA_j$ are dummy variables equal to 1 if the coal associated with option j is from either the Power River or the Uinta basins.[16] We include these terms in the latent cost function to infer the cost of retrofitting boilers to burn low-sulfur coal. In the case of PRB coal, we allow retrofitting cost to depend on the age of the boiler (AGE_i), calculated using its initial operating date.[17] ε_{ij} is an unobserved cost specific to unit i and option j. Coefficients on *SCRUBCOST*, *PRB* and *UINTA* are subscripted as we allow them to vary across EGUs. In some specifications, the coefficient on *EMISSIONS* is interacted with the regulatory status of the plant.

Eq. (1) is minimized subject to Eq. (2), which states that SO_2 emissions in pounds per MBtu of heat input cannot violate state and local emissions standards.[18]

$$\min_j C_{ij}(\beta) = \beta^F COALPRICE_{ij} + \beta_i^z SCRUBCOST_{ij} + \beta^P EMISSIONS_{ij} + \beta^A ASH_{ij} + PRB_j\left(\beta_{0,i}^w + \beta_1^w AGE_i\right) + \beta_i^U UINTA_j + \varepsilon_{ij} \tag{1}$$

$$EMISSIONS_{ij} \le \overline{SULFUR_i} \tag{2}$$

where $i = 1, 2, \ldots, 761$ (units), $j = 1, 2, \ldots, 72$ (compliance choices).

Estimation of the model

We estimate our model of compliance behavior using data for modeled units in 2002.[19] We argue that most units had achieved their optimal compliance strategy under the ARP by this time.

We estimate the choice of compliance option as a mixed logit model, treating the coefficients on *SCRUBCOST*, *PRB*, and

[14] Note that $SCRUBCOST_{ij} = 0$ for half of the compliance options.

[15] As is customary in discrete choice models (e.g., models of mode choice) we infer the value of physical quantities (e.g., ash) by dividing the coefficient on the physical quantity (e.g., ash) by the coefficient on a variable that is measured in monetary units (e.g., *COALPRICE*). For example, β^A/β^F measures the monetary cost of ash in cents per MMbtu just as dividing the coefficient on the time cost of a trip by the coefficient on the out-of-pocket cost of the trip converts the time cost into cents per minute.

[16] PRB and UINTA are dummy variables that vary only by option *j*. *ASH* and *EMISSIONS* vary by both option and EGU due to the nesting procedure that we describe below. In the nesting procedure, different EGUs may purchase coal from different counties within a basin.

[17] Age was not a statistically significant determinant of the cost of retrofitting to burn coal from the Uinta basin.

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

[19] 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.

UINTA as random; i.e., allowing them to vary by unit. The mixed logit model allows for more flexible substitution patterns than a nested logit model. It is also reasonable that the coefficients on *PRB* and *UINTA*, which capture the costs of retrofitting boilers to burn low-sulfur coal, should vary across units. Allowing the coefficient on *SCRUBCOST* to vary across EGUs controls for the fact that cost-minimizing plants might weight estimated scrubber costs differently, depending on their circumstances, e.g., depending on the age distribution of their generating capital.[20]

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 are identical for all units in the same plant. We treat $\{\varepsilon_{ij}\}$ as independently and identically distributed with a Type I extreme value distribution, and allow the coefficients on *SCRUBCOST*, *PRB*, and *UINTA* to be normally distributed with mean vector $\boldsymbol{b}$ and diagonal variance-covariance matrix $\boldsymbol{V}$. We denote the coefficient vector in the cost function as $\boldsymbol{\beta}$ and the distribution of $\boldsymbol{\beta}$ as $f(\boldsymbol{\beta}|\boldsymbol{b},\boldsymbol{V})$. The log likelihood function is therefore given by

$$l(\boldsymbol{b}, \boldsymbol{V}) = \sum_{n=1}^{N} \ln \int_{-\infty}^{\infty} \prod_{i=1}^{I_n} \frac{\exp\left(-C_{ij}(\beta)\right)}{\sum_{j'}^{J_i} \exp\left(-C_{ij'}(\beta)\right)} f(\beta|\boldsymbol{b}, \boldsymbol{V}) d\beta \tag{3}$$

where each plant is denoted as n and I_n denotes the set of units within each plant.

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% of the unit's coal from one of the eight regions or buying half of the unit's coal from each of two regions.[21] 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.[22] We then replace the *COALPRICE*, *ASH*, and *EMISSIONS* for unit i in region j with the characteristics of the cost-minimizing choice, for all i and j. The likelihood function in Eq. (3) is maximized using the updated coal characteristics, and the procedure is repeated until the parameter vector $\boldsymbol{\beta}$ converges.[23]

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 $COALPRICE_{ij}$.[24] 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 Fig. A.1 plots price trends for the period 1991–2005. We note that the trends across coal regions are approximately parallel over this period.

Imputed delivered coal prices are summarized in Table 4. The entries in Table 4 are the mean of predicted coal prices for 2002 all coal-fired generating units in each Census region, including NSPS units. 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 all units, the cost of installing and operating scrubbers are estimated as a function of plant and unit characteristics, following Lange and Bellas (2005). 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.[25] 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).

[20] One could argue that all coefficients in the cost model should be random; however, allowing the coefficient on *COALPRICE* to be random would complicate the interpretation of the model coefficients. As noted in footnote 15, we divide the coefficient on (e.g.) *ASH*, which is measured in physical units, by the coefficient on *COALPRICE*, which is measured in monetary units to convert the *ASH* coefficient to monetary units. Allowing the coefficient on *COALPRICE* to be random would complicate this.

[21] Only 3% of units buy coal from more than two regions.

[22] In solving this problem, the error term in (1) within the nesting procedure is treated as zero.

[23] This procedure is described more fully in the Appendix. We also report results of estimating our cost model without using the nesting procedure. See Table A.3 of the Appendix.

[24] Fig. A.2 in the Appendix plots predicted and observed average coal prices, by county, based on transactions reported for the 761 modeled EGUs. The correlation coefficient between the two (squared) is 0.67.

[25] 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.

Table 3
Average sulfur content of coal, by coal basin.

Basin	Mean sulfur	Range
North Appalachian, High End	2.8517	(2.0978, 3.5414)
North Appalachian, Low End	1.5608	(0.8756, 2.2840)
Central Appalachian	0.7798	(0.5286, 1.1062)
South Appalachian	1.0483	(0.5318, 1.6121)
Illinois Basin, High End	2.6298	(1.8620, 3.3849)
Illinois Basin, Low End	1.3763	(0.7440, 2.1999)
Uinta Basin	0.4727	(0.3278, 0.7396)
Powder River Basin	0.3831	(0.2288, 0.5306)

Notes: Sulfur is in pounds of S per MMBtu. Range is the 10th and 90th percentiles. Statistics are based on observed transaction data from 1991 to 2005, obtained from EIA form 423.

Table 4
Mean values of imputed delivered coal prices, by census region.

	West	South	Midwest	Northeast
North Appalachian, High End	148.6	118.1	116.0	112.0
	(12.40)	(9.70)	(8.87)	(6.91)
North Appalachian, Low End	223.7	150.5	147.3	124.1
	(30.33)	(24.14)	(21.25)	(15.82)
Central Appalachian	228.7	150.4	157.3	155.5
	(28.23)	(20.63)	(20.42)	(14.37)
South Appalachian	185.0	156.4	162.3	166.8
	(19.56)	(16.67)	(17.11)	(15.48)
Illinois Basin, High End	221.2	147.7	131.6	160.5
	(38.19)	(22.05)	(20.76)	(16.29)
Illinois Basin, Low End	210.8	151.2	135.6	159.7
	(32.28)	(18.48)	(17.22)	(14.24)
Uinta Basin	124.1	161.7	145.3	170.8
	(24.44)	(17.44)	(12.81)	(11.89)
Powder River Basin	85.95	129.0	97.26	135.1
	(32.72)	(17.27)	(15.16)	(10.48)

Notes: All coal prices are in 1995 cents per MMBtu; standard deviations in parentheses. Statistics are based on our estimates of the coal price facing each coal-fired EGU in 2002, including NSPS units, using the equation described in the Appendix. The mean and standard deviation are computed using all coal-fired power plants (NSPS as well as modeled units) in each Census region.

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.[26] 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.[27] 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 modeled EGUs by 2002 were installed before 1988. Clearly these scrubbers were not installed to comply with the ARP.[28] 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 the 761 coal-fired generating units in Table 1. We exclude 43 units that installed scrubbers before 1988 from estimation of the model but include them in the simulations reported in Section *Simulation Results*, with the constraint that a scrubber option must be chosen.[29]

Table 5 summarizes the variables entering the compliance cost model. The mean of coal price is the mean of predicted coal price across all 761 modeled units, summarized by coal basin. Scrubbing cost is the mean of predicted scrubbing cost for all 761

[26] A referee correctly notes that scrubber lifetime is likely to be shorter at older EGUs. To investigate this we interact *SCRUBCOST* with EGU age in Appendix Table A.4, column (4). We find no statistically significant effect.
[27] Assuming a removal rate of 85%.
[28] Carlson et al. (2000) attribute their installation to state and local emissions standards.
[29] We also explore alternate cutoff dates for scrubber installation as a sensitivity analysis. Table A.4 of the Appendix shows the impact of using 1985 and 1990 as cutoff dates in estimating the cost model.

Table 5
Summary statistics of model variables.

Variable	Mean	SD	Min	Max	N
Coal price (in cents per MMBtu)					
North Appalachian (High)	114.60	9.074	101.45	157.69	761
North Appalachian (Low)	142.65	22.13	108.82	243.41	761
Illinois Basin (High)	138.10	24.61	99.38	249.79	761
Illinois Basin (Low)	143.10	20.40	111.54	234.95	761
Central Appalachian	152.09	19.89	123.38	254.02	761
South Appalachian	151.47	7.569	135.44	185.69	761
Uinta Basin	153.75	16.32	99.32	181.90	761
Powder River Basin (PRB)	112.34	21.41	43.64	151.13	761
Scrubbing cost (in cents per MMBtu)					
Illinois Basin (High) Coal	105.73	26.62	35.16	227.86	761
PRB Coal	94.78	21.41	33.00	190.98	761
Emissions (in lbs of S per MMBtu)	0.7902	0.7559	0.0542	2.7785	72
Ash (in lbs per MMBtu)	8.8615	1.0723	5.9562	11.0904	72
Unit age	43.631	10.063	11	86	761
Heat input (in thousands MMBtu)	14,144.6	14,263	52.6	87,848.3	761
Divested	0.2116	0.4087	0	1	761
Publicly owned	0.1761	0.3811	0	1	761

Notes: All statistics are based on our 761 modeled units. Coal prices are predicted coal prices, summarized by coal basin.

units, conditional on (a) using high sulfur coal from the Illinois basin and (b) using coal from the PRB. The means of ash and emissions are averages across the 72 choice options. The last four variables in the table describe plant characteristics.

Actual compliance choices of the 761 units are summarized in first column of Table 8. Eighty EGUs are scrubbed: over 60% of these burn high-sulfur coal; one-third burn low-sulfur coal. Of the 681 units that do not scrub, approximately 31% burn high-sulfur coal (coal from the North Appalachian or Illinois basins), 31% burn medium-sulfur coal (coal from Central or South Appalachian basins) and 28% low-sulfur coal (coal from the PRB or Uinta basin). The remainder blend coal of different sulfur contents.

Estimation results

Table 6 presents the parameter estimates for the cost model.[30] Model (1), our base model, contains all variables in Eq. (1). Model (2) allows the weight on the price of coal to vary depending on whether coal is sourced in-state (Coal Price $\times$ In-State). A lower weight on 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 on *EMISSIONS*, *ASH*, *PRB* and *UINTA* to monetary units we divide them by the coefficient on *COALPRICE*.[31] Table 7 performs this conversion and adjusts the units in which variables are measured to calculate the shadow price of SO_2 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.[32] 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 than coal mined out-of-state, although the magnitude of this effect is less than 5%. Model (4) suggests that investor-owned

[30] A referee suggested that we bootstrap the standards errors in the cost model to allow for uncertainty in the estimation of coal prices and scrubbing costs. We did not do this due to the computational burden involved. Calculation of Murphy-Topel (1985) standard errors is not straightforward in the context of the random coefficients logit model with multiple first-stage equations.

[31] As is customary in discrete choice models (see footnote 15) we infer the value of physical quantities by dividing the coefficient on the physical quantity by the coefficient on a variable that is measured in monetary units. 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 *COALPRICE* coefficient at the .099 level in Model (1).

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

Table 6
Cost model estimation results.

	(1)	(2)	(3)	(4)
Mean effects				
Coal price	0.2851***	0.2811***	0.2821***	0.2764***
	(0.0257)	(0.0305)	(0.0300)	(0.0298)
Coal Price × In-State		−0.0129***	−0.0183***	
		(0.0048)	(0.0051)	
Coal Price × In-State × PUC-Regulated				−0.0162***
				(0.0058)
Coal Price × In-State × Publicly Owned				−0.0243*
				(0.0144)
Coal Price × In-State × Divested				−0.0030
				(0.0073)
Emissions	4.6813***	4.6285***	4.0103***	
	(0.4369)	(0.4432)	(0.4240)	
Emissions × PUC-Regulated				4.9917***
				(0.5680)
Emissions × Divested				3.9307***
				(1.0622)
Emissions × Publicly Owned				4.1992***
				(0.7981)
Ash	0.1849***	0.1301***	0.1476***	0.1368***
	(0.0439)	(0.0494)	(0.0473)	(0.0448)
Scrubbing Cost	0.2189***	0.2223***	0.2305***	0.2157***
	(0.0343)	(0.0330)	(0.0424)	(0.0353)
PRB	8.9150***	8.0380***	6.7405***	7.9711***
	(2.1114)	(2.2072)	(2.1414)	(2.172)
PRB × Age	0.1087***	0.1110***	0.0948**	0.1103***
	(0.0395)	(0.0399)	(0.0386)	(0.0411)
Uinta	9.4364***	9.6068***		9.2123***
	(1.2979)	(1.2821)		(1.1049)
Standard deviations of random coefficients				
Scrubbing cost	0.0939***	0.0952***	0.1060***	0.0924***
	(0.0143)	(0.0144)	(0.0222)	(0.0148)
PRB	6.3783***	6.3289***	5.9204***	6.5576***
	(0.8075)	(0.7422)	(0.7907)	(0.8964)
Uinta	6.3128***	6.6197***		6.8085***
	(0.8168)	(0.7905)		(0.9148)
Log likelihood	−987.4	−941.4	−1048.4	−932.6
Prediction rate (%)	71.48	72.40	69.25	71.75
Number of observations	718	718	718	718

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 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.

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).

In all models the shadow price of SO_2 emissions is roughly equal to observed allowance price.[33] 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 \$164/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 (\$181/ton SO_2), than for divested (\$142/ton SO_2) and publicly owned (\$152/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 Fig. 2(b) 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

[33] 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 2000 to convert from pounds to tons and divided by 2 to convert S to SO_2.

Table 7
Interpretation of model coefficients.

	(1)	(2)	(3)	(4)
Permit Price for SO_2 (\$/ton)	164.20	164.63	142.17	
	(14.88)	(15.42)	(9.90)	
PUC Regulated				180.57
				(22.13)
Divested				142.19
				(38.10)
Publicly-Owned				151.90
				(29.72)
Shadow Price for Ash (\$/ton)	12.97	9.25	10.46	9.90
	(2.56)	(3.08)	(2.73)	(2.70)
Retrofitting Cost for				
Uinta (Cents/MMBtu)	33.10	34.17		36.00
	(5.71)	(6.23)		(7.05)
(\$/kW)	19.14	19.76		20.81
PRB (Cents/MMBtu)	47.91	45.81	38.56	46.24
	(4.34)	(5.00)	(4.42)	(4.78)
(\$/kW)	27.70	26.49	22.30	26.73

Notes: All numbers are expressed in 1995\$. Standard errors appear in parentheses. The calculations for PRB retrofitting cost refer to the unit with a mean age (43.631). For calculations associated with \$/kW, we assume an operating heat rate of 11 MMBtu/MWh and a capacity factor of 60%.

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 11). 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,% correctly predicted, suggests that our models predict about 70% of observed choices correctly.[34] 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,

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

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

$$\begin{aligned}\hat{C}_{ij}(\beta) &= \hat{\beta}^F COALPRICE_{ij} + \mathbf{E_i}\hat{\beta}_i^z SCRUBCOST_{ij} + \hat{\beta}^P EMISSIONS_{ij} + \hat{\beta}^A ASH_{ij} \\ &\quad + PRB_j\left(\mathbf{E_i}\hat{\beta}_{0,i}^w + \hat{\beta}_1^w AGE_i\right) + \mathbf{E_i}\hat{\beta}_i^U UINTA_j + \mathbf{E_i}\varepsilon_{ij}\end{aligned} \tag{5}$$

where hats denote estimated values and $\mathbf{E}_i$ denotes the conditional mean of the respective coefficient for unit i. Given these modifications, Eq. (5) perfectly predicts the observed compliance choice for each EGU.[35] When we conduct simulations of the model, we use Eq. (5).

A second question is how well our model captures observed SO_2 emissions under the ARP.[36] 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% 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. Fig. 3 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

[34] For EGU i we calculate the probability of each option being chosen and predict that EGU i will choose the option with the highest predicted probability. If that option is chosen, we say that the model has correctly predicted the choice made by EGU i.

[35] 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.

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

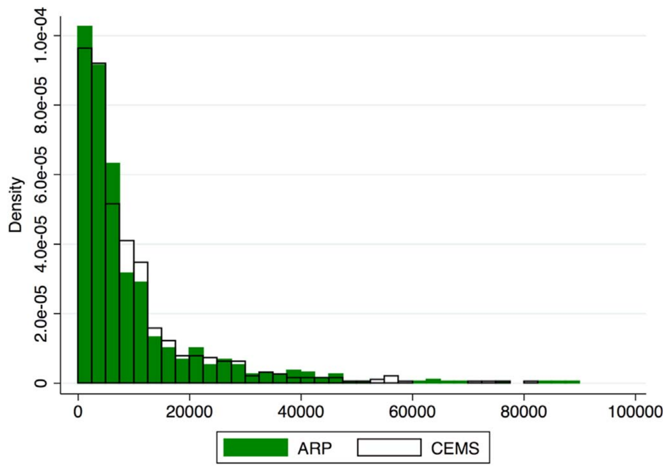

Fig. 3. Histogram of predicted v. actual emissions under the ARP. Green bars represent predicted emissions under the Acid Rain Program, based on Model (1) of Table 6. The clear bars represent the actual emissions distribution, as reported in CEMS. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

base model predicts ARP emissions from the 761 units to be 7.111 million tons. Measured emissions were 7.094 million tons.[37]

Simulation results

Predicting compliance choices, costs, and emissions

To estimate cost savings from the ARP for modeled 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 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).

Our simulations assume that the output of each EGU is the same under the UPS as under the ARP. We do not model the level of output associated with each EGU and, hence, do not allow firms to adjust to the UPS by shifting output from high-cost units to low-cost units. According to EIA form 767, few units altered their output as a means of complying with the ARP. 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. This assumption, however, may not be valid for the units pictured in the tail of Fig. 4 which are forced to substantially reduce their emissions under the UPS. The assumption that these units did not adjust their heat input is likely to overstate the costs of the UPS and, hence, the cost savings from trading under the ARP.[38]

We model 761 EGUs in our simulation analysis, including the 43 EGUs we exclude from our estimation, which installed scrubbers before year 1988. We predict their compliance choices based on the models estimated in Section *Estimation Results* using the conditional means of the random coefficients and unobserved costs. In the counterfactual, we allow these EGUs to change their compliance choices, subject to local emission standards, with the additional restriction that a scrubber must be installed (i.e. we remove non-scrubbed options from their choice sets).

We calculate the cost of compliance under each regime (per MMBtu), $COMPC_{ij}$, 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 Eq. (6):

$$COMPC_{ij} = COALPRICE_{ij} + SCRUBCOST_{ij} + \widetilde{\beta^A}ASH_{ij} + PRB_j\left(\mathbf{E_i}\widetilde{\beta^W_{0,i}} + \widetilde{\beta_1^W}AGE_i\right) + \mathbf{E}_i\widetilde{\beta_i^U}UINTA_j + \mathbf{E}_i\tilde{\varepsilon}_{ij} \tag{6}$$

Where $\tilde{\beta} = \hat{\beta}/\hat{\beta}^F$. Note, in this equation, that we use predicted costs of coal and scrubbing rather than the weighted costs in Eq. (5). $COALPRICE_{ij}$ AND $SCRUBCOST_{ij}$ are our estimates of actual fuel and scrubbing costs incurred by EGU i if option j is chosen.

[37] These are measured emissions averaged over the period 2000–2002. Emissions reported in Table 1 are for the year 2002 alone.

[38] Of the 29 EGUs that emitted more than 4 pounds of SO_2 per MMBtu of heat input in 2002, 6 shut down by 2006.

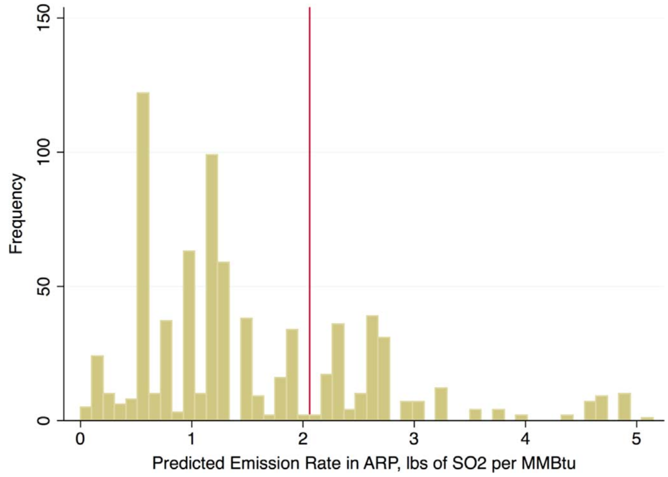

Fig. 4. Histogram of predicted emission rates under the ARP. This graph shows the frequency distribution of predicted emissions under the ARP, in pounds of SO_2 per MMBtu of heat input, based on Model (1) of Table 6.

Although the choice of compliance option is based on weighted costs, actual costs per MMBtu of heat input should be used to compute cost savings, in order to compare our estimates to other authors' estimates. Total compliance costs are calculated using average heat input from 2000 to 2002.[39] The difference between compliance costs under the ARP and the uniform performance standard represent the estimated cost savings from the ARP.

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.[40] 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.[41]

Predicted emissions under the ARP for the 761 modeled units vary from one model to another but are, in the aggregate, within 2% 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 2.06 to 2.29 pounds of SO_2 per MMBtu. (When 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 modeled EGUs in 2002 are 40% higher than allocated permits; hence, the relevant cap should be higher.

Fig. 4 compares predicted emissions rates under the ARP for Model (1) of Table 6 with the corresponding uniform standard. The 204 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 $186 million and $200 million (1995$) lower than under the uniform performance standard; significantly smaller estimates than previous studies.[42] 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 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.

[39] This assumption is valid if the output of the EGU is not changed by the UPS and if the heat rate of the compliance option chosen is the same as the historical heat rate of the EGU. Although scrubbers will affect heat rate, the parasitic load loss associated with a scrubber is small.

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

[41] We acknowledge that it would be desirable to provide bootstrapped standard errors for the estimates of cost savings in Table 8; however, we did not do this due to the computational burden involved.

[42] We estimate that the cost savings under Model (1) of Table 6, $200 million (1995$), are about 20% of compliance costs. If all units were to choose the least-cost source of coal, subject to satisfying state and local emissions standards, total costs would be approximately $985 million (1995$).

Table 8
Simulation results: ARP and uniform standard counterfactual.

Compliance choices		Uniform Performance Standard			
	ARP	(1)	(2)	(3)	(4)
No scrubber	***681***	***691***	***690***	***691***	***692***
High-sulfur coal	208	196	210	216	206
High end	42	0	0	0	0
Low end	166	196	210	216	206
Medium-sulfur coal	214	202	203	212	203
Low-sulfur coal	189	185	177	184	180
Blend: high & medium	27	59	43	35	48
Blend: high & low	21	27	35	22	33
Blend: medium & low	22	22	22	22	22
Scrubber	***80***	***70***	***71***	***70***	***69***
High-sulfur coal	50	45	46	44	45
Medium-sulfur coal	3	2	2	2	2
Low-sulfur coal	27	23	23	24	22
Predicted emissions (in million tons)					
ARP	7.094[a]	7.111	7.244	7.446	7.159
UPS		7.097	7.300	7.445	7.176
Standard level (lbs SO_2 per MMBtu)		2.060	2.120	2.290	2.120
(Weighted)		1.319	1.356	1.383	1.333
Cost savings (in million 1995$)		**199.63**	**199.44**	**185.77**	**190.11**

Note: High sulfur coal refers to coal purchased in North Appalachian or Illinois basins. Medium sulfur coal refers to coal purchased in Central or South Appalachian basins. Low sulfur coal refers to coal purchased in Uinta or Powder River basins.
[a] Denotes actual emissions from CEMS.

Secondly, 14 modeled 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 $1702 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.

The health impacts of trading

We examine the impact of trading on health damages using two scenarios. In the first, we estimate health damages for the 761 generating units used to estimate the cost savings from the ARP. We compare the damages associated with these units under the ARP and under the UPS that yields the same aggregate emissions, using the models from Table 6. A finding that the UPS and the ARP yield equal health damages in this scenario does not, however, indicate that trading had no health impacts. The aggregate emissions of the 761 EGUs we model are much larger in 2002 than their allowance allocations for the year 2002. These emissions were made possible in part by banking, but primarily through the purchase of allowances from other units. This prompts us to examine a second scenario, in which we compare the health damages from *all* units covered by the ARP in 2002 with damages that would have occurred under a no-trade counterfactual.

In the no-trade scenario we expand the scope of our analysis from our modeled units to all units covered by the ARP. Fig. 5, which shows 2002 SO_2 emissions in excess of 2002 allowances for all units covered by the ARP, by state, suggests that units east of the Mississippi River were purchasing allowances from units west of the Mississippi. To examine the impact of allowance transfers to modeled 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.

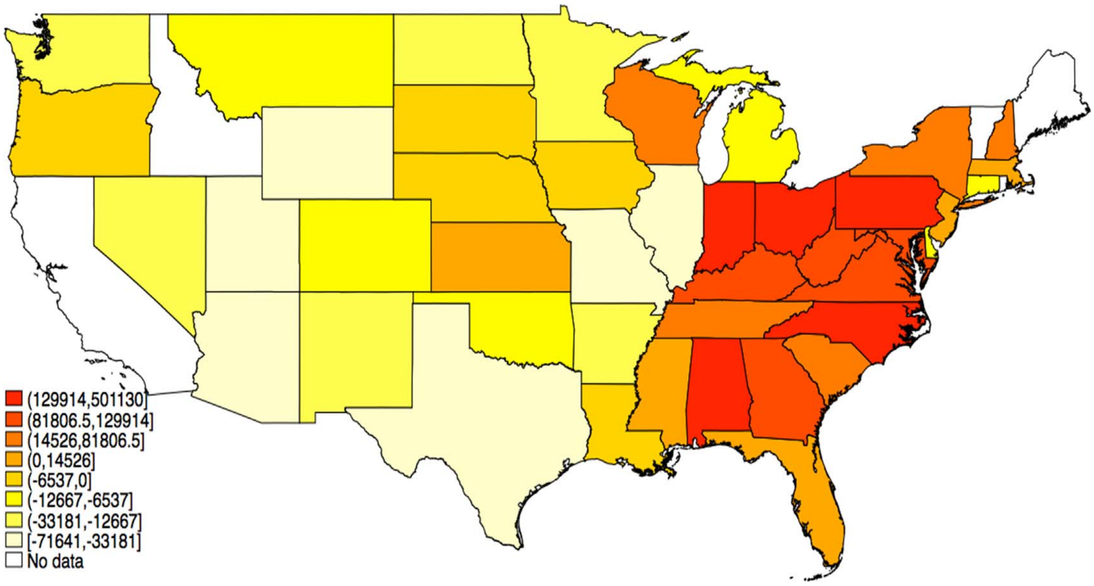

Fig. 5. Emissions of all ARP units, net of allocations, 2002. This map shows, for each state, the difference between tons of SO_2 emitted by all EGUs covered by the ARP in 2002 and the number of 2002 permits allocated to these units. All data come from the CEMS. Positive values suggest EGUs in the state were net purchasers of sulfur dioxide permits, whereas negative values suggest that a state was a net supplier of permits. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

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 $PM_{2.5}$.

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 (EPA, 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.[43]

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 (2000\$).[44] 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.[45] 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,

[43] 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).

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

[45] 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.

Table 9
Comparison of estimated damages from SO_2 emissions under ARP and UPS.

Model	Damage UPS - ARP	Difference UPS - ARP	Difference (UPS − ARP)/ARP	Deaths UPS - ARP	Difference UPS - ARP
Model 1	91.14[a,b,c]			15,965	
	90.97	0.17	0.0018	15,940	25
Model 2	92.91			16,278	
	92.75	0.16	0.0018	16,254	23
Model 3	95.22			16,685	
	94.84	0.38	0.0040	16,624	61
Model 4	91.73			16,070	
	91.52	0.21	0.0023	16,037	33

[a] Damages expressed in billions (1995$).
[b] Value in top row for each pair of model parameters corresponds to the UPS.
[c] Value in bottom row for each pair of model parameters corresponds to the ARP.

aggregating across all receptor counties; and (2) by county, to explore spatial patterns in the change in emissions, air quality, and impacts.

Damages under the ARP and under a UPS

Table 9 reports the difference in damages due to SO_2 emissions under the ARP and the UPS counterfactual simulation for each of the four models of Table 6.[46] In the case of Model (1), mean damages under the ARP amount to $90.97 billion in 2002. For the uniform performance standard, damages are slightly higher, at $91.14 billion. This amounts to a divergence of $170 million (1995$)—0.18% of damages under the ARP—with damages higher under the UPS. In Models (2), (3), and (4) of Table 6, damages are also higher under the UPS. We focus on Model (1) because it fits the data well (see Table 6) and because its predictions of emissions under the ARP are very close to observed emissions from EPA's Continuous Emissions Monitoring System (CEMS) database (see Table 8).

Fig. 6 shows the spatial patterns of the difference in damages between the ARP and the UPS for Model (1). Damages are lower under the UPS in the Northeast and Middle Atlantic states but higher in parts of the East North Central, East South Central and South Atlantic Census divisions. Fig. 4 helps to explain these results. Under the UPS, all units that are above the standard in Fig. 4 must reduce their emissions under the UPS; however, units below the standard may increase their emissions, to the extent permitted by state and local emissions standards. These standards are, on average, less stringent in the East North Central, East South Central and South Atlantic Census divisions than in the Middle Atlantic division. On net, locations with higher damages under the UPS have monetary damages that are $170 million higher than under the ARP, corresponding to an additional 25 lives lost.

We conclude based on our simulations that for modeled plants, health damages under the ARP were no greater in the aggregate than they would have been under a counterfactual UPS, implemented in 2002; in fact, they are $170 million less. This suggests that for these plants, the ARP delivered a $200 cost savings and a $170 million reduction in health damages.

Damages under the ARP and under a no-trade scenario

Table 10 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%, or in absolute terms, by $2.10 billion (1995$).

Table 10 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 $6 million VSL (2000$) with the $2 million VSL (2000$) 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 $750 million (1995$) 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

[46] Appendix Table B.1 reports a sensitivity analysis of the damages associated with Model (1) of Table 6.

Table 10
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	119.89[a,b,c]			21,003	
	117.79	2.10	0.0175	20,618	386
Alternative Dose-response[d]	189.83			34,304	
	186.40	3.43	0.0181	33,666	638
VSLY	51.88			21,003	
	51.13	0.75	0.0144	20,618	386
$2M VSL	46.16			21,003	
	45.43	0.75	0.0162	20,618	386

[a] Damages expressed in billions (1995$).
[b] Value in top row for each pair of model parameters corresponds to the ARP, using observed emissions.
[c] Value in bottom row for each pair of model parameters corresponds to the No-Trade Scenario.
[d] Uses dose-response function for $PM_{2.5}$ mortality from Roman et al. (2008).

reported in Roman et al. (2008), which suggests that $PM_{2.5}$ has a 60% larger effect on mortality rates, raises damages estimates in both cases. This implies that damages under the ARP are approximately $3.4 billion (1995$) 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 10. The statistical uncertainty associated with the damage estimates is further explored in the Appendix.

Figs. 1, 7 and 8 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 Fig. 5 and borne out by Fig. 1, which shows modeled $PM_{2.5}$ concentrations attributable to actual emissions, as measured by EPA's Continuous Emissions Monitoring System (CEMS) minus $PM_{2.5}$ concentrations attributed to the no-trade counterfactual. Fig. 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.

Fig. 7 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 Fig. 1. Most counties showing an increase in damages due to trading exhibit an increase between 1% and 5%. Fig. 7 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%. Parts of North Carolina and West Virginia also show increases in damage over 5%. And, a few counties in West Virginia show increases greater than 10%. Intuitively, these areas incur the largest increases in ambient concentrations (see Fig. 1). In contrast, the coastal counties north of New York City, which emitted less than their 2002 allocations (see Fig. 5) have lower damages under the ARP.

Fig. 8 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 Fig. 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. Fig. 8 shows more variation in damages than Fig. 7 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 Fig. 1 in which ambient $PM_{2.5}$ increased by less than 0.5 $\mu g/m^3$.[47]

Our no-trade counterfactual reinforces the importance of considering marginal damages in designing a trading scheme, a point made many years ago by Mendelsohn (1986) and Savins (1996) and reiterated more recently by Holland and Yates (2015) and Fowlie and Muller (2017). 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-

[47] Fig. 8 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 indicated by the initial allocation of allowances, implying a larger difference in damages between the ARP and the no-trade counterfactual.

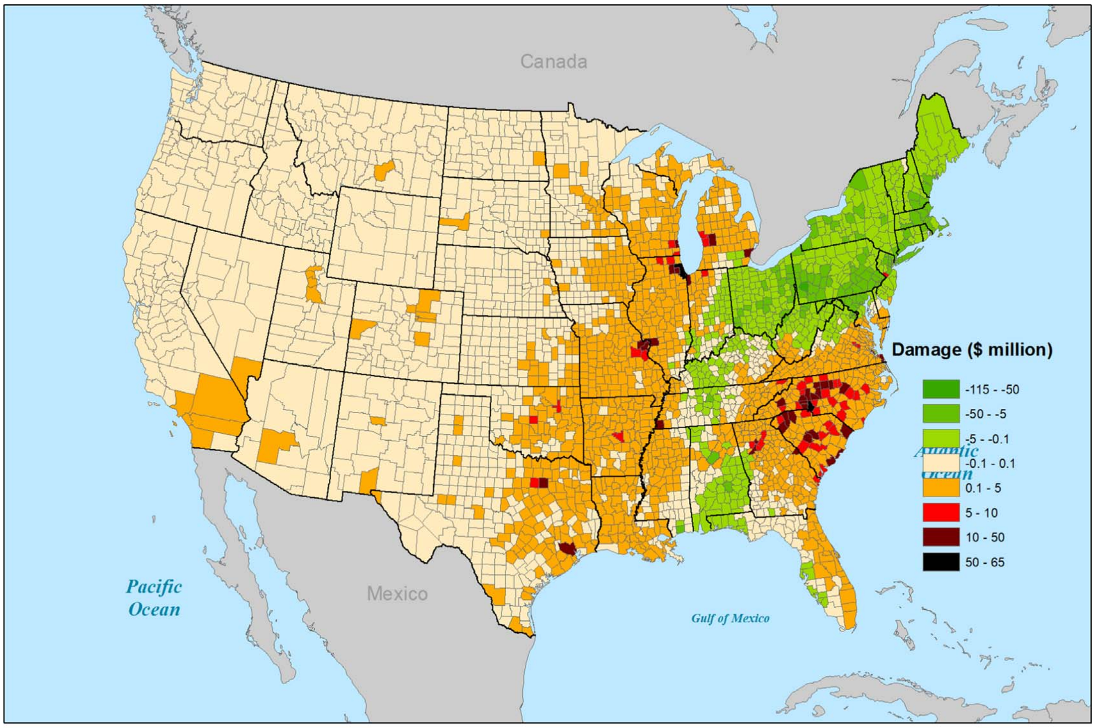

<Double-click here to enter text>

Fig. 6. Difference in damages: ARP minus UPS. This map shows the difference in monetized health damages from our 761 modeled EGUs in 2002 under the ARP minus the UPS. Monetized damages in both cases are based on Model (1) of Table 6. Green areas are locations in which damages are lower under the UPS than the ARP. In locations colored orange or red, damages are higher under the UPS. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

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 emphasizes the importance of allowing trading ratios in a cap-and-trade scheme to reflect differences in marginal damages. The possibility of doing this in the context of the ARP is considered by Muller and Mendelsohn (2009).

Conclusions

In this study we quantify the cost savings from the ARP compared with a command-and-control alternative and examine the impact of trading under the ARP on health damages from SO_2. To quantify cost savings, we compare compliance costs for 761 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. We also examine the difference in health damages between the ARP and the counterfactual UPS.

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 approximately $200 million in 2002. This represents about 20% of the compliance costs incurred by our 761 EGUs, but is lower than the cost 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 modeled 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 our modeled 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

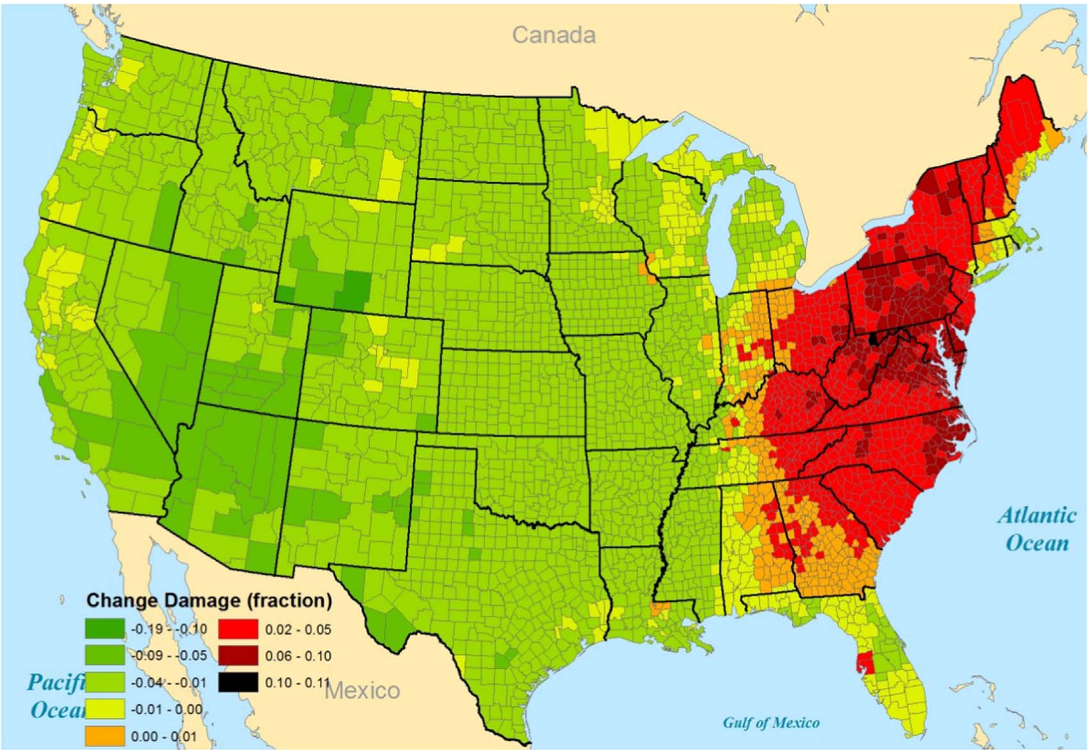

Fig. 7. Proportional difference in damages: ARP minus no-trade scenario. This map shows monetized health damages in 2002 associated with the ARP minus damages associated with the No-Trade scenario, expressed as a proportion of damages under the ARP. Emissions are based on all units covered by the ARP, including NSPS units and non-coal units. Damages under the ARP are based on actual 2002 emissions. Emissions under the No-Trade scenario equal 2002 allowances plus observed drawdowns of the emissions bank. Red values indicate increased damages under the ARP relative to the No-Trade scenario. Green values indicate lower damages under the ARP. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

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 find that health damages associated with emissions from our 761 modeled plants would have been slightly (0.18%) higher under a UPS than under the ARP: health damages from the 761 EGUs would have been $170 million (1995$) lower under the ARP. This difference does not, however, imply that trading under the ARP necessarily resulted in lower health damages than a command-and-control counterfactual.

To further 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 that health damages under the ARP are 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 $6 million VSL (2000$). The mean difference in damages is $2.10 billion (1995$), or about 1.8% 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 Fig. 7). 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 provides evidence that trading under the ARP yielded cost savings. The $200 million (1995$) that we document for 2002 is, indeed, a lower bound to cost savings, even for the year 2002. Our analysis fails to capture the cost savings from the banking provisions of the ARP (Ellerman and Montero, 2007) and the impacts of emissions trading on

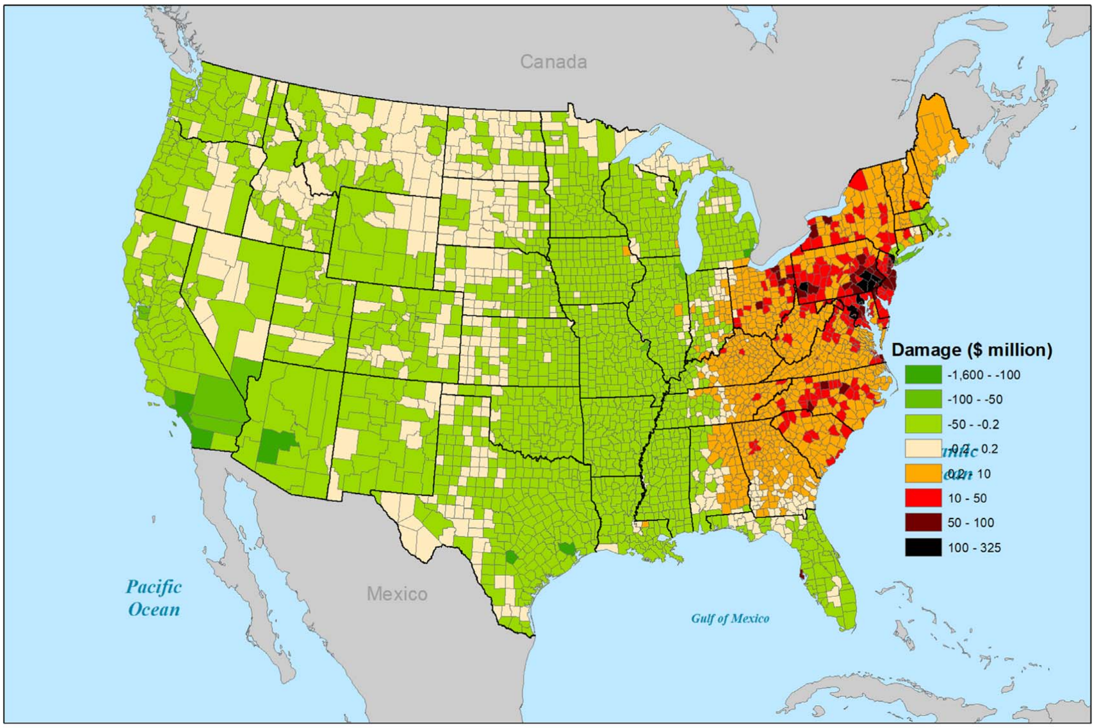

<Double-click here to enter text>

Fig. 8. Difference in health damages: ARP minus no-trade scenario. This map shows monetized damages (1995$) associated with the ARP in 2002 minus damages associated with the No-Trade scenario in 2002. Emissions are based on all units covered by the ARP, including NSPS units and non-coal units. Damages under the ARP are based on actual 2002 emissions. Emissions under the No-Trade scenario equal 2002 allowances plus observed drawdowns of the emissions bank. Red values indicate higher damages under the ARP compared to the No-Trade scenario. Green values indicate lower damages under the ARP. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

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.

Appendix A. Estimation and simulation procedures for the cost model

Prediction of coal prices and scrubbing costs

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:

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

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 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, a lower price for coal sold on a spot market or via a short-term contract. Aligning with EIA

Table A.1
Coal price equations.

	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: The table presents results of estimating Eq. (A.1) for each coal basin. 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.

estimates, the transportation rates from each basin also vary. Eq. (A.1) is used to predict coal prices by basin, as shown in Fig. A.1.

We also use equation to predict delivered coal price to each plant. To indicate how well our model predicts delivered coal prices, Fig. A.2 plots the average of predicted coal prices from each coal-mining county to our modeled plants again observed coal prices. The correlation coefficient between average predicted and observed prices (squared) equals 0.67.

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).

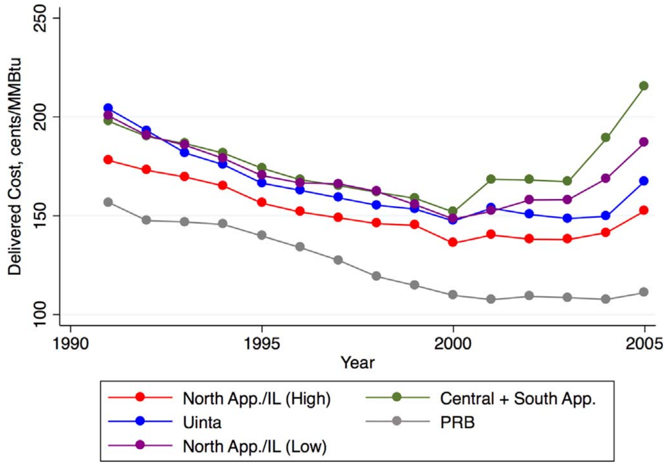

Fig. A.1. Trends in predicted coal prices by basin. This graph shows predicted delivered coal prices, by basin, for the period 1991–2005. Predicted prices are based on estimates of Eq. (A.1) and are averaged across all units for which data are available. Because coal purchases over the 2000–2002 period could be based on earlier prices, it is reassuring that the price trends are parallel from 1991 through 2000.

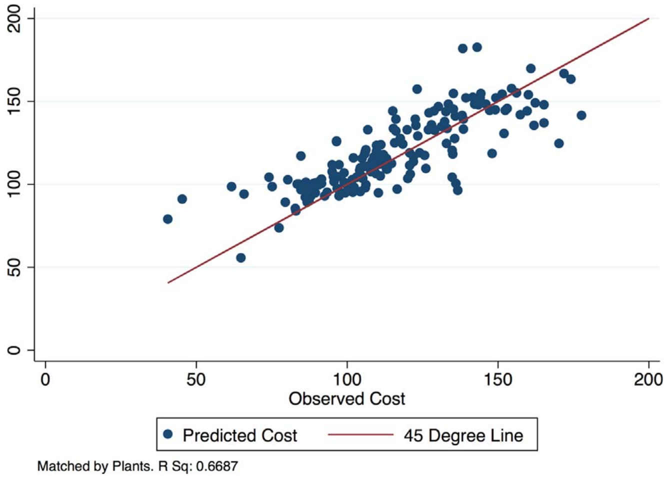

Fig. A.2. Comparison of actual and predicted average coal prices, 2002. This graph plots predicted and observed average coal prices, by county, based on transactions reported for the 761 modeled EGUs. The correlation coefficient between the two (squared) is 0.67.

Using the estimates, we predict energy consumption, operating cost and capital cost for each unit in our sample.[48] 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.

Estimation of the mixed logit model

We estimate the mixed logit model (Eq. (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% may be purchased from each of two basins. We split the two high-sulfur coal basins, the North Appalachian and

[48] 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).

Table A.2
Scrubber cost equations.

	(1) Energy Cons.	(2)	(3) Operating Cost	(4)	(5) Capital Cost	(6)
	OLS	Heckman	OLS	Heckman	OLS	Heckman
Main Estimation Equation						
Log(Scrub Age)	0.199*	0.243**	0.214**	0.263***		
	(0.115)	(0.108)	(0.095)	(0.102)		
Log(Unit Age)	0.753***	0.645***	0.002	0.042	0.420**	0.796***
	(0.180)	(0.210)	(0.137)	(0.168)	(0.162)	(0.216)
Log(Coal Consumption)	0.709***	0.787***	0.607***	0.606***	0.855***	0.888***
	(0.089)	(0.122)	(0.073)	(0.116)	(0.104)	(0.148)
Log(Removal Rate)	0.901***	0.890***	0.425***	0.515***	0.936***	0.991***
	(0.215)	(0.218)	(0.136)	(0.134)	(0.232)	(0.238)
Log(Op. Hours)	0.710***	0.748***	0.255	0.285	−0.497**	−0.443*
	(0.198)	(0.194)	(0.200)	(0.176)	(0.234)	(0.234)
PRB	−0.049	−0.050	−0.381***	−0.475***		
	(0.260)	(0.258)	(0.145)	(0.131)		
Flue Gas Entering Rate			1.021***	1.178***		
			(0.274)	(0.297)		
Selection Equation						
Log(Sulfur Premium)		0.514		0.389		1.620***
		(0.431)		(0.396)		(0.427)
NSPS		0.810***		0.812***		0.516*
		(0.121)		(0.117)		(0.309)
Log(Coal Consumption)		0.354***		0.363***		0.447***
		(0.047)		(0.046)		(0.064)
Northeast		−0.740**		−0.764**		−0.848**
		(0.350)		(0.348)		(0.380)
South		−1.149***		−1.215***		−0.812***
		(0.275)		(0.249)		(0.310)
Midwest		−1.158***		−1.223***		−1.443***
		(0.311)		(0.287)		(0.319)
Year = 1992	−0.370***	−0.375***	−0.111***	−0.107***		
	(0.132)	(0.131)	(0.038)	(0.037)		
Year = 1994	−0.605***	−0.603***	−0.208***	−0.219***		
	(0.172)	(0.169)	(0.074)	(0.076)		
Year = 1996	−0.756***	−0.740***	−0.436***	−0.442***		
	(0.211)	(0.210)	(0.119)	(0.131)		
Year = 1998	−0.761***	−0.738***	−0.483***	−0.508***		
	(0.237)	(0.237)	(0.135)	(0.149)		
Year = 2000	−0.864***	−0.838***	−0.614***	−0.656***		
	(0.193)	(0.194)	(0.149)	(0.172)		
Year = 2002	−1.046***	−0.974***	−0.588***	−0.627***		
	(0.216)	(0.226)	(0.187)	(0.234)		
Year = 2004	−1.304***	−1.224***	−0.648***	−0.684***		
	(0.258)	(0.267)	(0.149)	(0.202)		
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.

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 (2014), to refine the characteristics of the coal purchased within a basin:

1. Start with an estimate of the vector of cost function parameters, $\beta^{(0)}$.
2. For each 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 Eq. (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, $\left|\beta^{(t)} - \beta^{(t-1)}\right|<1\times10^{-6}$.

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.

Sensitivity analyses

We present two sets of sensitivity analyses. As suggested by a referee, we estimate the models in Table 6 without using the nesting procedure described above. These results, displayed in Table A.3, indicate that the coefficients of the cost models estimated without the nesting procedure are similar to coefficients estimated using the nesting procedure (cf. Table 6 of the text). The nesting procedure, however, refines the choice of coal within a basin and, hence, the sulfur content of the coal burned. As shown at the bottom of Table A.3, predicted SO_2 emissions are more than 13% larger than actual emissions (7.094 million tons) when models of compliance behavior are estimated without the nesting procedure. In contrast, Model 1, estimated using the nesting procedure, produces aggregate emissions (7.080 million tons) that are almost identical to those observed under the ARP.

Table A.3
Estimation results without nesting procedure.

	(1)	(2)	(3)	(4)
Mean effects				
Coal price	0.2377***	0.2137***	0.2251***	0.2029***
	(0.0257)	(0.0216)	(0.0216)	(0.0224)
Coal Price × In-State		−0.0123***	−0.0128***	
		(0.0036)	(0.0034)	
Coal Price × In-State × PUC-Regulated				−0.0120***
				(0.0046)
Coal Price × In-State × Publicly Owned				−0.0073
				(0.0067)
Coal Price × In-State × Divested				−0.0271***
				(0.0069)
Emissions	4.8678***	4.3931***	4.2920***	
	(0.5597)	(0.5815)	(0.5548)	
Emissions × PUC-Regulated				4.7517***
				(0.6697)
Emissions × Divested				3.5445***
				(1.1905)
Emissions × Publicly Owned				3.6495***
				(0.9338)
Ash	0.2254	0.4738**	0.4722**	0.4402**
	(0.1903)	(0.1914)	(0.1958)	(0.1893)
Scrubbing Cost	0.2213***	0.2046***	0.2251***	0.2078*
	(0.0406)	(0.0370)	(0.0469)	(0.1182)
PRB	10.2715***	9.3974***	8.4418***	9.4280***
	(1.9587)	(2.1417)	(1.8557)	(2.6436)
PRB × Age	0.0624	0.0869**	0.0875**	0.0820***
	(0.0405)	(0.0382)	(0.0358)	(0.0463)
Uinta	7.6219***	8.2866***		8.8145***
	(0.9593)	(1.2427)		(1.4137)
Standard deviations of random coefficients				
Scrubbing cost	0.0929***	0.0899***	0.1005***	0.0964
	(0.0160)	(0.0195)	(0.0230)	(0.0791)
PRB	5.7835***	6.3973***	6.0120***	7.2331***
	(0.8403)	(1.0733)	(1.1747)	(1.1208)
Uinta	5.2096***	6.7419***		7.2464***
	(0.5014)	(0.9612)		(1.1476)
Log likelihood	−1154.4	−1124.7	−1189.6	−1110.9
Prediction rate (%)	66.36	67.02	64.26	67.94

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). Predicted emissions in all models = 8.050 million tons.

Table A.4
Sensitivity analysis of Model 1.

Model	Baseline	Scrubber cutoff date = 1990	Scrubber cutoff date = 1985	Interact Scrubcost and Age
	(1)	(2)	(3)	(4)
Mean effects				
Coal price	0.2851***	0.2845***	0.2853***	0.2843***
	(0.0257)	(0.0259)	(0.0256)	(0.0255)
Emissions	4.6813***	4.6751***	4.6388***	4.6476***
	(0.4369)	(0.4374)	(0.4319)	(0.4313)
Ash	0.1849***	0.1844***	0.1825***	0.1861***
	(0.0439)	(0.0439)	(0.0434)	(0.0432)
Scrubbing Cost	0.2189***	0.2140***	0.2538***	0.1714***
	(0.0343)	(0.0317)	(0.0328)	(0.0446)
Scrubbing Cost x Age				0.0014
				(0.0010)
PRB	8.9150***	8.9265***	10.5150***	8.8059***
	(2.1114)	(2.1289)	(2.2026)	(2.1873)
PRB × Age	0.1087***	0.1094***	0.0748*	0.1138***
	(0.0395)	(0.0396)	(0.0400)	(0.0408)
Uinta	9.4364***	9.4983***	9.2683***	9.3320***
	(1.2979)	(1.3202)	(1.0382)	(1.1239)
Standard deviations of random coefficients				
Scrubbing cost	0.0939***	0.0897***	0.1083***	0.0981***
	(0.0143)	(0.0147)	(0.0150)	(0.0181)
PRB	6.3783***	6.4270***	7.2273***	6.5604***
	(0.8075)	(0.8312)	(0.8016)	(0.7848)
Uinta	6.3128***	6.3354***	6.2181***	6.2744***
	(0.8168)	(0.8169)	(0.6205)	(0.7432)

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).

Our second set of sensitivity analyses, presented in Table A.4, varies the set of EGUs used to estimate the compliance cost model and alters the specification of the model to allow scrubbing costs to be affected by EGU age. In the text, we estimate models of compliance costs omitting EGUs that installed scrubbers before 1988 on the grounds that these scrubbers were not installed to comply with the ARP.[49] As a sensitivity analysis, we estimate Model 1 using, alternately, 1985 and 1990 as cutoff years. Using 1990 as the cutoff year has little impact on model coefficients (cf. columns (1) and (2)). Using 1985 as a cutoff year (column (3)) alters some of the coefficients, but does not change our estimates of cost savings. The last column of Table A.4 interacts SCRUBCOST with EGU age. We use a scrubber life of 25 years in annualizing scrubbing costs, regardless of the age of the EGU. To allow for the fact that annual scrubber costs would be higher for older EGUs (with shorter expected lifetimes) we include a term interacting SCRUBCOST with age. The coefficient on the interaction term is not significantly different from zero.

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 Eq. (5) and compliance costs under the ARP using Eq. (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 Eq. (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 Eq. (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

[49] When Title IV was written, the period 1985–87 was chosen as the basis for assigning SO_2 allowances—it was considered a period when plants were unlikely to have anticipated Title IV. This is the primary basis for our choice of the year 1988 as a cut-off year.

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 Eq. (6).

Appendix B. Further analysis of health damages

In this appendix we provide additional detail on the health impacts of imposing a uniform performance standard and, for the no-trade counterfactual, 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 Table B.1 we show the impact of sensitivity analyses in estimating health damages under the ARP and the UPS using Model (1) of Table 6 in the text.

Table B.2. presents Monte Carlo simulations of damages for the no-trade scenario, treating 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 1000 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.2 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
Comparison of estimated damages from SO_2 emissions under ARP and UPS.

IAM Model Parameters	Damage UPS - ARP	Difference UPS - ARP	Difference (UPS – ARP)/ARP	Deaths UPS - ARP	Difference UPS - ARP
Default	91.14[a,b,c]			15,965	
	90.97	0.17	0.0018	15,940	25
Alternative Dose-response[d]	144.31			26,078	
	144.06	0.26	0.0018	26,036	42
VSLY	39.41			15,965	
	39.26	0.15	0.0038	15,940	25
\$2M VSL	35.10			15,965	
	35.02	0.07	0.0022	15,940	25

[a] Damages expressed in billions (1995\$).
[b] Value in top row for each pair of model parameters corresponds to the UPS based on Model (1) of Table 6.
[c] Value in bottom row for each pair of model parameters corresponds to the ARP based on Model (1) of Table 6.
[d] Uses dose-response function for $PM_{2.5}$ mortality from Roman et al. (2008).

Table B.2
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

ARP – NTC = Difference in damage between observed emissions and the no-trade counterfactual.
[a] =95% confidence intervals in parenthesis.

Appendix C. Data appendix

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 and the banking of allowances, permitting us to compute the no-trade counterfactual. Data on heat input come from CEMS.

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

References

Arimura, T.H., 2002. An empirical study of the SO_2 allowance market: effects of PUC regulations. J. Environ. Econ. Manag. 44 (2), 271–289.
Burtraw, D., Krupnick, A., Mansur, E., Austin, D., Farrell, D., 1998. Costs and benefits of reducing air pollutants related to acid rain. Contemp. Econ. Policy 16 (4), 379–400.
Carlson, C.P., Burtraw, D., Cropper, M., Palmer, K., 2000. SO_2 control by electric utilities: what are the gains from trade? J. Political Econ. 108 (6), 1292–1326.
Chan, H.S., 2014. An Ex Post Evaluation of the US Acid Rain Program (Dissertation). University of Maryland.
Chan, H.S., Fell, H., Lange, I., Li, S., 2017. Efficiency and environmental impacts of electricity restructuring on coal-fired power plants. J. Environ. Econ. Manag. 81 (1), 1–18.
Cicala, S., 2015. When does regulation distort costs? Lessons from fuel procurement in U.S. electricity generation. Am. Econ. Rev. 105 (1), 411–444.
Ellerman, A.D., Joskow, P.L., Schmalensee, R.M., Montero, J.-P., Bailey, E.M., 2000. Markets for Clean Air. Cambridge University Press, Cambridge.
Ellerman, A.D., Montero, J.-P., 2007. The efficiency and robustness of allowance banking in the U.S. Acid Rain Program. Energy J. 28 (4), 47–71.
EPA (US Environmental Protection Agency), 1992. Regulatory Impact Analysis of the Final Acid Rain Implementation Regulations. EPA, Washington, DC.
EPA (US Environmental Protection Agency), 1999. The Benefits and Costs of the Clean Air Act from 1990 to 2010. EPA, Washington, DC.
EPA (US Environmental Protection Agency), 2011. The Benefits and Costs of the Clean Air Act from 1990 to 2020. EPA, Washington, DC.
Fann, N., Fulcher, C.M., Hubbell, B.J., 2009. The influence of location, source and emission type in estimates of the human health benefits of reducing a ton of air pollution. Air Qual. Atmos. Health 2, 169–176.
Fowlie, M., 2010. Emissions trading, electricity restructuring, and investment in pollution abatement. Am. Econ. Rev. 100 (3), 837–869.
Fowlie, M. and Muller, N., 2017. Market-based Emissions Regulation When Damages Vary across Sources: What Are the Gains from Differentiation? Working Paper.
Fullerton, D., Dermott, S., Caulkins, J., 1997. Sulfur dioxide compliance of a regulated utility. J. Environ. Econ. Manag. 34 (1), 32–53.
Gagelmann, F., Frondel, M., 2005. The impact of emission trading on innovation – science fiction or reality? Environ. Policy Gov. 15 (4), 203–211.
Henry, D.D., Muller, N.Z., Mendelsohn, R., 2011. The social cost of trading: measuring the increased damages from sulfur dioxide trading in the United States. J. Policy Anal. Manag. 30 (3), 598–612.
Holland, S.P., Yates, A.J., 2015. Optimal trading ratios for pollution permit markets. J. Public Econ. 125, 16–27.
Lange, I.A., Bellas, A.S., 2005. Technological change for sulfur dioxide scrubber under market-based regulation. Land Econ. 81 (4), 546–556.
Levy, J.I., Baxter, L.K., Schwartz, J., 2009. Uncertainty and variability in health-related damages from coal-fired power plants in the United States. Risk Anal. 29 (7), 1000–1014.
Mendelsohn, R., 1986. Regulating heterogeneous emissions. J. Environ. Econ. Manag. 13, 301–312.
Montero, J.-P., 2002. Market structure and environmental innovation. J. Appl. Econ. 5 (2), 293–325.
Mrozek, J.R., Taylor, L.O., 2002. What determines the value of life? A meta-analysis. J. Policy Anal. Manag. 21 (2), 253–270.
Muller, N.Z., 2011. Linking policy to statistical uncertainty in air pollution damages. B. E. Press J. Econ. Anal. Policy 11 (1).
Muller, N.Z., 2012. The design of optimal climate policy with air pollution co-benefits. Resour. Energy Econ. 34 (4), 696–722.
Muller, N.Z., Mendelsohn, R.O., 2009. Efficient pollution regulation: getting the prices right. Am. Econ. Rev. 99 (5), 1714–1739.
Muller, N.Z., Mendelsohn, R.O., Nordhaus, W.D., 2011. Environmental accounting for pollution in the U.S. economy. Am. Econ. Rev. 101 (5), 1649–1675.
Murphy, K.M., Topel, R.H., 1985. Estimation and inference in two-step econometric models. J. Bus. Econ. Stat. 3 (4), 370–379.
NRC (National Research Council), 2010. Hidden Costs of Energy: Unpriced Consequences of Energy Production and Use.
Pope, C.A., Burnett, R.T., Thun, M.J., Calle, E.E., Krewski, D., Ito, K., Thurston, G.D., 2002. Lung cancer, cardiopulmonary mortality, and long-term exposure to fine particulate air pollution. J. Am. Med. Assoc. 287 (9), 1132–1141.
Roman, H.A., Walker, K.D., Walsh, T.L., Conner, L., Richmond, H.M., Hubbell, B.J., Kinney, P.L., 2008. Expert judgement assessment of the mortality impact of changes in ambient fine particulate matter in the U.S. Environ. Sci. Technol. 42, 2268–2274.
Schmalensee, Richard, Stavins, Robert N., 2013. The SO_2 allowance trading system: the ironic history of a grand policy experiment. J. Econ. Lit. 27 (1), 103–122.

Sotkiewicz, P.M., Holt, L., 2005. Public utility commission regulation and cost-effectiveness of title IV: lessons for CAIR. Electr. J. 18 (8), 68–80.
Stavins, R.N., 1998. What can we learn from the grand policy experiment? Lessons from SO_2 allowance trading. J. Econ. Perspect. 12 (3), 69–88.
Stavins, R.N., 1996. Correlated uncertainty and policy instrument choice, J. Environ. Econ. Manag. 30, 218–232.
Swinton, J.R., 2002. The potential for cost savings in the sulfur dioxide allowance market: empirical evidence from Florida. Land Econ. 78 (3), 390–404.
Swinton, J.R., 2004. Phase I completed: an empirical assessment of the 1990 CAAA. Environ. Resour. Econ. 27 (3), 227–246.
Train, K.E., 2009. Discrete Choice Methods with Simulation. Cambridge University Press, Cambridge.
Bellas, Allen, Lange, Ian, 2008. Impacts of market-based environmental and generation policy on scrubber electricity usage. Energy J. 29 (2), 151–164.
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, N., 2004. "Environmental Policy and the Choice of Abatement Technique: Evidence from Coal-fired Power Plants." Unpublished.
Linn, Joshua, Mastrangelo, Erin, Burtraw, Dallas, 2014. Regulating greenhouse gases from coal power plants under the Clean Air Act. J. Assoc. Environ. Resour. Econ. 1 (1/2), 97–134.