Sulfur Dioxide Control by Electric Utilities: What Are the Gains from Trade?

Curtis Carlson

Department of the Treasury

Dallas Burtraw

Resources for the Future

Maureen Cropper

World Bank and University of Maryland

Karen L. Palmer

Resources for the Future

Title IV of the 1990 Clean Air Act Amendments (CAAA) established a market for transferable sulfur dioxide (SO_2) emission allowances among electric utilities. This market offers firms facing high marginal abatement costs the opportunity to purchase the right to emit SO_2 from firms with lower costs, and this is expected to yield cost savings compared to a command-and-control approach to environmental regulation. This paper uses econometrically estimated marginal abatement cost functions for power plants affected by Title IV of the CAAA to evaluate the performance of the SO_2 allowance market. Specifically,

This paper has benefited from comments and suggestions from Robert Chambers, Denny Ellerman, Suzi Kerr, Richard Newell, Paul Portney, and Byron Swift. We would especially like to thank Don Fullerton for his exceptionally detailed comments, which have significantly improved the paper. Carlson's employment with the Department of the Treasury does not constitute an endorsement by the department of the views expressed in this paper. This paper does not represent the views of the World Bank, its executive directors, or the countries they represent.

[*Journal of Political Economy*, 2000, vol. 108, no. 6]

we investigate whether the much-heralded fall in the cost of abating SO_2, compared to original estimates, can be attributed to allowance trading. We demonstrate that, for plants that use low-sulfur coal to reduce SO_2 emissions, technical change and the fall in prices of low-sulfur coal have lowered marginal abatement cost curves by over 50 percent since 1985. The flexibility to take advantage of these changes is the main source of cost reductions, rather than trading per se. In the long run, allowance trading may achieve cost savings of \$700–\$800 million per year compared to an "enlightened" command-and-control program characterized by a uniform emission rate standard. The cost savings would be twice as great if the alternative to trading were forced scrubbing. However, a comparison of potential cost savings in 1995 and 1996 with modeled costs of actual emissions suggests that most trading gains were unrealized in the first two years of the program.

I. Introduction

For years economists have urged that policy makers use market-based approaches to control pollution (taxes or tradable permits) rather than rely on uniform emission standards or uniform technology mandates (command and control). This advice was largely ignored until the 1990 Clean Air Act Amendments (CAAA) established a market for sulfur dioxide (SO_2) allowances. Along with a cap on overall annual emissions, the SO_2 allowance market gives electric utilities the opportunity to trade rights to emit SO_2 rather than forcing them to install SO_2 abatement technology or emit at a uniform rate. By equalizing marginal abatement costs among power plants, trading should limit SO_2 emissions at a lower cost than the traditional command-and-control approach.

The SO_2 allowance market presents the first real test of the wisdom of economists' advice and therefore merits careful evaluation. Has the allowance market significantly lowered the costs of abating SO_2, as economists claimed it would? An answer in the affirmative would strengthen the case for marketable permits to control other pollutants, such as greenhouse gases. Conversely, if cost savings are small, this would have implications for the design (or even the adoption) of market-based approaches to controlling pollution in the future.

The purpose of this paper is to evaluate the performance of the SO_2 allowance market. Specifically, we ask two questions: (1) How much can the trading of permits reduce the costs of controlling SO_2, compared to command and control; that is, what are the potential gains from trade? (2) Were these trading gains realized in the first years of the allowance market? The answers require that we estimate marginal abatement cost functions for fuel switching at all generating units that do not scrub their emissions, calculate the expected cost of postcombustion abatement (scrubbers), and compute the least-cost solution to achieving the cap on SO_2 emissions. The difference between the least-cost solution and the cost under our counterfactual command-and-control policy represents the potential static efficiency gains from allowance trading. We compute these gains for 1995 and 1996, the first two years of the allowance market, and the expected savings in 2010, when the emissions cap will be stricter and will be applied more broadly and when the allowance market should be functioning as a mature market.

The command-and-control policy against which we measure gains from allowance trading is key to the analysis. A policy that would have imposed end-of-stack abatement technology (scrubbing) would have been significantly more expensive than an emission rate standard applied uniformly to all facilities.[1] A uniform emission rate standard provides firms with considerable flexibility, including the opportunity to take advantage of technical change that is precluded under a more rigid technology standard; hence it is a favorable characterization of a command-and-control approach. In our analysis we evaluate the gains from allowance trading compared to each of two command-and-control alternatives: forced scrubbing and a uniform emission rate standard.[2]

Our approach to evaluating the allowance market is very different from the approach used by other observers to assess market performance. Both the administrator of the Environmental Protection Agency (EPA) and the chair of the Council of Economic Advisers proclaimed the success of the allowance market by comparing allowance prices (circa \$100 per ton in 1997) with estimates of marginal abatement costs produced at the time the CAAA were written (as high as \$1,500).[3] Since the former are much lower than the latter, they concluded that the trading of SO_2 allowances has greatly reduced the cost of curbing SO_2 emissions.

This argument is flawed for two reasons. First, it is inappropriate to judge how well the allowance market is performing simply by comparing current allowance prices with ex ante estimates of marginal abatement

[1] In 1983 the Sikorski/Waxman bill sought to reduce SO_2 emissions by requiring the installation of scrubbers (flue gas desulfurization equipment) at the 50 dirtiest plants. Studies estimate that the annual cost of this proposal would have ranged from \$7.9 billion (Office of Technology Assessment 1983) to \$11.5 billion (Temple, Barker and Sloane 1983), in 1995 dollars.

[2] One justification for the use of an emission rate standard is that it is the approach used to regulate nitrogen oxide emissions under Title IV of the 1990 CAAA.

[3] On March 10, 1997, EPA Administrator Carol Browner argued that "during the 1990 debate on the acid rain program, industry initially projected the cost of an emission allowance to be \$1500 per ton of sulfur dioxide.... Today, those allowances are selling for less than \$100" ("New Initiatives in Environmental Protection," 1997). Likewise, in testimony before the House Commerce Subcommittee on Energy and Power on the economics of the Kyoto protocol (March 1998), Janet Yellen, chair of the Council of Economic Advisers, noted that "emission permit prices, currently at approximately \$100 per ton of SO_2 are well below earlier estimates.... Trading programs may not always bring cost savings as large as those achieved by the SO_2 program."

costs in the least-cost solution. Price can equal marginal abatement cost even if many utilities that might benefit from trading fail to participate in the market. Second, comparing current allowance prices with ex ante estimates of marginal abatement costs shows only that the latter were too high; it does not mean that the allowance market was responsible for the fall in marginal abatement costs.[4]

Our analysis suggests that these claims for the allowance market are misleading, especially the suggestion that formal trading has lowered the cost of SO_2 abatement several-fold. In contrast, we reach the following conclusions.

1. Marginal abatement costs for SO_2 are much lower today than those estimated in 1990. Technical improvements, including advances in the ability to burn low-sulfur coal at existing generators, as well as improvements in overall generating efficiency, lowered the typical unit's marginal abatement cost function by almost \$50 per ton of SO_2 over the decade preceding 1995. The decline in fuel costs lowered marginal abatement costs by about \$200 per ton.

2. This decline in marginal abatement cost, if one assumes that it was not caused by Title IV, has lowered the cost of achieving the SO_2 emission cap under *both* the least-cost solution and enlightened command and control (e.g., under a uniform emission rate standard). This implies that the gains from trade—the cost savings attainable from an allowance trading program—have also fallen over time.[5] We estimate the potential cost savings attributable to formal trading (vs. a uniform emission rate standard) to be \$250 million (1995 dollars) annually during the first phase of the allowance program (1995–2000, which covers the dirtiest power plants). We estimate them to be \$784 million annually during the second phase of the program (beginning in 2000, which covers all plants), about 43 percent of compliance costs under our enlightened command-and-control policy. A comparison of the least-cost solution with a less enlightened command-and-control alternative of forced scrubbing indicates annual savings of almost \$1.6 billion (1995 dollars).

3. Comparing the least-cost solution to achieving actual emission reductions with actual abatement costs indicates that actual compliance costs exceeded the least-cost solution by \$280 million in 1995 and by \$339 million in 1996 (1995 dollars). This suggests that the allowance market did not achieve the least-cost solution, even though marginal

[4] It should also be noted that the ex ante estimates of marginal abatement costs generally pertained to equilibrium in the second phase of the program and therefore cannot be compared with current allowance prices unless they are discounted to the present.

[5] The assumption that the fall in marginal abatement costs was not due to Title IV potentially imparts a downward bias to the estimate of the cost savings under the program. The incentives provided by allowance trading might have accelerated changes in fuel prices, as well as other technological changes (Burtraw 1996).

abatement costs under that solution were approximately equal to allowance prices. The failure to realize potential savings is not surprising. The 1990 CAAA represent a dramatic departure from the pollution regulations to which utilities were previously subject, and taking full advantage of their flexibility may require time. As participants become more familiar with the opportunities presented by the allowance market and ongoing deregulation of the electricity industry provides greater incentives to reduce costs, the volume of trading will no doubt increase and cost savings are more likely to be realized.

The remainder of the paper is organized as follows. Section II provides institutional background on the CAAA. Section III presents the methodology we employ to evaluate the allowance market, including our estimation of marginal abatement cost curves. Section IV estimates potential gains from allowance trading in the long run and explains why these estimates are lower than those predicted when the CAAA were written. Section V evaluates the performance of the allowance market in 1995 and 1996, and Section VI concludes the paper.

II. Institutions

Since 1970, the SO_2 emissions of electric utilities have been regulated in order to achieve federally mandated local air quality standards (the National Ambient Air Quality Standards [NAAQS]). For plants in existence in 1970, these standards, codified in state implementation plans, have typically taken the form of maximum emission rates (pounds of SO_2 per million Btus of heat input). Plants built after 1970 are subject to New Source Performance Standards (NSPS), set at the federal level. Since 1978, NSPS for coal-fired power plants have effectively required the installation of capital-intensive flue gas desulfurization equipment (scrubbers) to reduce SO_2 emissions, an attempt to protect the jobs of coal miners in states with high-sulfur coal. This regulation has significantly raised the costs of SO_2 abatement at new plants in areas in which emissions could have been reduced more cheaply by switching to low-sulfur coal.

During the 1980s, over 70 bills were introduced in Congress to reduce SO_2 emissions from power plants. Some would have forced the scrubbing of emissions by all electric generating units, and others would have provided limited flexibility by imposing uniform emission rate standards, which give firms the opportunity to choose a compliance strategy.

The innovation of Title IV is to move away from these types of uniformly applied regulations. Instead, reductions are to be achieved by setting a cap on emissions while allowing the trading of marketable pollution permits or *allowances*. Each generating unit in the electricity industry is allocated a fixed number of allowances each year and is required to hold one allowance for each ton of SO_2 it emits.[6] Utilities are allowed to transfer allowances among their own facilities, sell them to other firms, or bank them for use in future years.

The eventual goal of Title IV of the CAAA is to cap average annual SO_2 emissions of electric utilities at 8.95 million tons—about half of their 1980 level. This is to be achieved in two phases. In the first phase, which began in 1995, the 110 dirtiest power plants (with 263 generating units) are each allocated allowances sufficient for an emission rate of 2.5 pounds of SO_2 per million Btus of heat input. Firms can voluntarily enroll additional generating units ("compensation and substitution" units) in phase I, subject to the constraint that the average emission rate of all units does not increase. In the second phase, which begins in the year 2000, all fossil-fueled power plants larger than 25 megawatts are annually allocated allowances sufficient for an emission rate of 1.2 pounds of SO_2 per million Btus of heat input. In both phases, heat input is based on the 1985–87 reference period.

Allowance trading takes advantage of the fact that emission control costs vary across generating units and encourages firms with the cheapest control costs to undertake the greatest emission reductions. Unfortunately, firms may not have adequate incentives to minimize SO_2 compliance costs because of decisions made by some state public utility commissions (Bohi and Burtraw 1992; Bohi 1994; Rose 1997). For instance, to protect the jobs of miners in states with high-sulfur coal, some regulators preapproved the recovery of investment in scrubbers, while leaving it uncertain whether the cost of other possible compliance measures would be similarly recoverable. The allowance program itself encouraged scrubbing by allocating 3.5 million "bonus" allowances to firms that installed scrubbers as the means of compliance, for the explicit purpose of protecting jobs in regions with high-sulfur coal. In addition, investments in scrubbers can be depreciated and in some cases expensed (deducted against income tax) as soon as the scrubber is installed. In contrast, in many states the cost of purchased allowances cannot be

[6] Allowances are allocated to units in proportion to emissions during the 1985–87 period. About 2.8 percent of the annual allowance allocations are withheld by the EPA and distributed to buyers through an annual auction run by the Chicago Board of Trade. The revenues are returned to the utilities that were the original owners of the allowances. An emissions cap creates a barrier to entry and commensurate scarcity rents that accrue to owners of existing facilities when allowances are allocated at zero cost to these facilities. These rents would not be present in our command-and-control (performance standard) policy. Fullerton and Metcalf (1997) show that these rents compound economic distortions associated with preexisting taxes, thereby imposing an important source of additional social cost not reflected in firms' compliance costs. Goulder, Parry, and Burtraw (1997) find this additional social cost to have a magnitude similar to that of the compliance cost savings from allowance trading that we identify. They find that the social cost of an emissions cap could be largely alleviated were the program to auction allowances and use the revenue to reduce preexisting labor taxes.

recovered until they are used for compliance (Lile and Burtraw 1998). These facts suggest that—through no fault of its own—the allowance market might not succeed in capturing the potential gains from emission trading, a hypothesis that we investigate below.[7]

III. Methodology

To investigate whether the allowance market has operated efficiently and to estimate the size of potential gains from trading versus other forms of regulation, we estimate marginal abatement cost functions for generating units. These functions can be used to calculate the least-cost solution to achieving an aggregate level of emissions, as well as the expected costs of alternative regulatory approaches.

A. *Calculation of the Gains from Allowance Trading*

The least-cost solution to achieving the SO_2 cap requires minimizing the present discounted value of compliance costs for all generating units over time, subject to constraints on the banking of allowances. Because the SO_2 cap shrinks between phase I and phase II, the banking of allowances will, in general, be optimal (if adjustment costs for reducing pollution are not too great). Thus emissions should be less than allowances in the early years of the program (Rubin 1996). Eventually, however, a steady state is expected in which net contributions to the bank are zero on average and annual emissions equal annual allowances. Rather than solve this intertemporal problem, we sidestep the banking question by taking the banking behavior of firms as given.[8]

Our primary goal is to compute how much more cheaply the chosen level of emissions could be achieved through formal trading within the allowance market than by a uniform emission standard. We calculate the long-run gains from trade by computing the least-cost solution to achieving the emissions cap in the year 2010, when annual allowances should equal annual emissions (Environmental Protection Agency 1995; Electric Power Research Institute 1997). We then contrast this solution with the cost of achieving the cap in 2010 via a uniform emission rate standard (an "enlightened" form of command and control) and with the cost of achieving the cap via forced scrubbing. For 1995 and 1996, the first two years of the allowance market, we compute the potential gains from trade as the difference between the least-cost solution to

[7] Winebrake, Farrell, and Bernstein (1995) and Fullerton, McDermott, and Caulkins (1997) provide estimates of the potential magnitude of inefficiencies that may result, but no author has attempted to estimate actual performance.

[8] We ignore potential future environmental legislation (e.g., for control of particulates, ozone, or greenhouse gases).

achieving *actual* emissions and the cost of achieving these emissions via enlightened command and control. We then calculate the costs *actually incurred* in these two years to learn whether the potential gains from trading have been realized.

The Role of Scrubbing versus Fuel Switching

To calculate the least-cost solution to limiting SO_2 emissions, we must estimate the marginal abatement cost curves of all generating units in the allowance market. In estimating marginal abatement cost functions, we separate plants into those that reduce SO_2 emissions via fuel switching (substituting low-sulfur for high-sulfur coal) and those that have installed scrubbers. As noted above, fuel switching is the chief method of reducing emissions for most power plants. In 1995 only 17 percent of all generating units in the United States used scrubbers. Eighty-six percent of these units were required to do so by law, either to satisfy federal NSPS (61 percent) or state laws (24 percent). The remaining 15 percent of units (28 in number) installed scrubbers specifically to comply with Title IV. We assume that units that were required by law to scrub or chose to do so in phase I continue to do so in the least-cost solution as well as in the command-and-control counterfactual. Hence, the cost of scrubbing at these units is added to the total abatement costs under both scenarios and does not directly affect our estimates of the cost savings from efficient trading.

To see whether additional firms would build scrubbers to minimize compliance costs, we solve for the marginal cost of abatement under fuel switching and compare this with the average cost of abatement via scrubbing. If the marginal cost of abatement under fuel switching is lower than the average cost of abatement under scrubbing, as we find it to be, building additional retrofit scrubbers would not lower costs.[9]

From the perspective of abating SO_2 emissions, the chief difference between units that switch fuel and units that scrub is the shape of their marginal abatement cost (MAC) functions. When electricity output is held constant, plants that switch fuel can reduce the tons of SO_2 they emit by varying the sulfur content of their fuel. If a premium must be paid per million Btus for low-sulfur coal, this implies that the MAC curve

[9] A referee suggests that this comparison should be made using a plant-specific average cost of scrubbing; however, we do not have enough data to calculate an average cost of scrubbing for each coal-fired generating unit in our data set. We compare marginal abatement cost (permit price) in the least-cost solution with average abatement cost under scrubbing on the basis of the 28 units that installed retrofit scrubbers.

slopes down as emissions of SO_2 increase.[10] For plants that scrub, emissions of SO_2 are almost entirely determined by electricity output (heat input). Because scrubbers remove about 95 percent of the sulfur content of coal, emissions are relatively insensitive to the sulfur content of coal burned. Conditional on output, therefore, the MAC curve for scrubbed units is a point. In computing the least-cost solution and the command-and-control alternative, we therefore subtract the emissions of scrubbed units from the emissions cap and solve for the least-cost solution using the estimated MAC curves of units that switch fuel.

Formally, we choose the level of emissions for each fuel-switching unit that minimizes the aggregate cost of achieving the modified overall emissions cap. In general, a generating unit's marginal cost of emissions function depends on output (as well as on emissions and input prices); however, we do not vary electricity output to reduce SO_2. We thus ignore demand-side management as an emissions reduction strategy, as well as the possibility of shifting output from dirty to clean plants to reduce SO_2.[11]

Computation of the Gains from Trade

From a baseline of emissions that would have obtained without the 1990 CAAA, we compute the cost of the least-cost solution for all units that switch fuel as the area under their MAC curves from baseline emissions to emissions under the least-cost solution. For firms whose MAC curves are positive over all relevant emissions levels, the computation is straightforward. For firms whose MAC curves are negative over some range of emissions, we compute the cost of moving from baseline emissions to emissions in the least-cost solution as the area under the portion of the MAC curve that lies above the positive quadrant.[12] To compute total costs under the least-cost solution and in the command-and-control alternatives, the capital and variable costs of retrofit scrubbing are an-

[10] The MAC curve slopes downward because in our long-run cost function capital is fully adjustable and increased capital investments are required to increase the amount of low-sulfur coal burned, in addition to paying a premium per million Btus for low-sulfur (vs. high-sulfur) coal.

[11] In order to switch output among plants, we would have to model the electricity grid, which is beyond the scope of the paper. There is no evidence that utilities have relied on demand-side management to reduce SO_2; indeed, the cost per ton of SO_2 reduced would be much more expensive if achieved through reductions in output than through fuel switching.

[12] Savings at firms with negative abatement costs are not considered cost savings attributable to the trading program. To incorporate units in the least-cost solution for which we have not estimated MAC curves, we allocate allowances, A, to units for which MAC curves are available, solve the least-cost solution, and then multiply total cost by the ratio of total allowances to A. This, in effect, assumes that the aggregate MAC curve for omitted units is identical to that for the units in our data set.

nualized over 20 years using a 6 percent discount rate and are added to the costs of fuel switching. The gains from trade are the difference between total costs under the least-cost solution and total costs under the command-and-control counterfactual.

B. *Estimation of Marginal Abatement Cost Curves*

To estimate MAC functions for plants that switch fuel, we assume that the manager of each power plant minimizes the cost of producing electricity at the generating unit, subject to its production technology and a constraint on SO_2 emissions. This constraint represents the emissions standard facing the plant because of the NAAQS for SO_2.[13] We have chosen the generating unit as the unit of analysis because SO_2 emission standards apply to individual generating units.[14] An alternative approach would be to assume that the manager minimizes the cost of producing a fixed level of output at the plant level, equating the marginal cost of electricity generation across generating units, but this would force us to average emission standards across units faced with different standards. Since the order in which units are brought into service is usually predetermined, we treat output as fixed at the generator level.

Our approach to estimating MAC functions at fuel-switching units is to estimate a cost function and share equations for electricity generation that treat generating capital as variable, using data from the period prior to trading under the CAAA. We treat generating capital as variable to capture capital investments that allow plants to burn low-sulfur coal. Because the firm's desired amount of capital stock is instantaneously achievable in this model, the estimates we obtain are estimates of long-run abatement costs. This is similar to the approach taken by Gollop and Roberts (1983, 1985), who estimated marginal abatement costs at the firm level for 56 coal-fired electric utilities. They examined firms'

[13] Throughout, we assume that electric utilities comply with local permitting constraints set to meet NAAQS. These constraints are distinct from the requirements of the 1990 amendments that established the SO_2 trading program in order to meet regional air quality goals. The assumption that emissions never violate the emissions standard appears justified by EPA data, which show that fewer than 5 percent of the plants in our database were ever in violation of emission regulations during the entire period of our study.

[14] Generating units consist of a generator-boiler pair. For over 85 percent of the generating capacity, there is a one-for-one match between generators and boilers. For the remaining 15 percent, there are multiple generators attached to a boiler or vice versa. Emission standards and allowance allocations apply to the boiler. The continuous emission monitoring system used under Title IV measures emissions at the stack level, and it is often the case that several generating units are attached to one emission stack. For those units that share boilers or stacks or both, we assign emissions on the basis of the percentage of total heat input consumed by each boiler. For generators that share a single boiler, we assign emissions on the basis of the percentage of total electricity output from each generator.

responses to SO_2 regulations between 1973 and 1979 for firms that met emission requirements through fuel switching.[15]

Econometric Model

The manager's problem is to choose labor (l), generating capital (k), and fuel inputs of high- and low-sulfur coal (fhs and fls, respectively) to minimize the cost of producing output q and achieving an emission rate e in time period t, subject to emissions and production constraints. Unit and time indexes are suppressed for convenience:

$$\min_{k,l,fls,fhs} C = p_k k + p_l l + p_{fls} fls + p_{fhs} fhs \tag{1}$$

subject to

$$q(k,\ l,\ fls,\ fhs,\ t) \geq Q,$$

$$e(k,\ l,\ fls,\ fhs,\ t) \leq e^*.$$

In equation (1), e^* represents the emissions standard, typically stated as an emission rate, for example, pounds of SO_2 per million Btus of heat input, averaged over a specified time interval.[16] In the derivation of the cost function to be estimated, one approach would be to replace the chosen values of inputs with the expressions for the optimal input demands as a function of input prices, the level of output, and e^*. For policy purposes, however, we wish to estimate a MAC function that describes the cost of meeting the emission rate actually achieved. For this reason, we write costs as a function of e, the actual emission rate. Because e is an endogenous variable in the cost function, we simultaneously estimate the cost function and an equation to predict e as a function of the emissions standard and other exogenous variables. The cost function to be estimated is thus

$$C = C(p_k,\ p_l,\ p_{fls},\ p_{fhs},\ q,\ e,\ t). \tag{2}$$

The econometric model (eqq. [3]–[6]) consists of the cost function, input share equations, and an equation for the firm's mean annual emission rate. We use a translog form for the cost function, with prices p_i or p_j (for $i,\ j = k,\ l,\ fls,\ fhs$), adding dummy variables (d_m) for each

[15] The econometric estimation of MAC functions is distinct from the approach taken in other analyses of Title IV that rely on engineering estimates of MAC functions (Environmental Protection Agency 1990, 1995; General Accounting Office 1994; Electric Power Research Institute 1995; Kalagnanam and Bokhari 1995; Fullerton et al. 1997; Siegel 1997; Burtraw et al. 1998).

[16] Almost 85 percent of standards are stated as pounds of SO_2 or sulfur per million Btus of heat input. When estimating the cost function, we converted all standards to pounds of SO_2 per million Btus of heat input. Dummy variables were included to distinguish different averaging times.

plant ($m = 1, \ldots, 260$) in the database to measure fixed effects that vary among plants.[17] A quadratic function of time (t) is added to the cost function to capture technical change. Linear time trends enter the input share and emissions rate equations. Dummy variables are included to indicate the type of emission standard the plant faces (w_g, $g = 1, \ldots, 7$) and the time period over which emissions are averaged (v_h, $h = 1, 2, 3$):[18]

$$\begin{aligned}\ln C = {} & \alpha_0 + \sum_m \lambda_m d_m + \sum_j \alpha_j \ln p_j + \alpha_q \ln q + \alpha_e \ln e + \alpha_t t \\ & + \tfrac{1}{2}\sum_i \sum_j \alpha_{ij} \ln p_i \ln p_j + \sum_j \alpha_{jq} \ln p_j \ln q + \sum_j \alpha_{je} \ln p_j \ln e \\ & + \sum_j \alpha_{jt} (\ln p_j) t + \tfrac{1}{2}\gamma_{qq} (\ln q)^2 + \gamma_{qe} \ln q \ln e + \gamma_{qt} (\ln q) t \\ & + \tfrac{1}{2}\gamma_{ee} (\ln e)^2 + \gamma_{et} (\ln e) t + \tfrac{1}{2}\gamma_{tt} t^2 + \epsilon_c,\end{aligned} \tag{3}$$

$$\begin{aligned}s_l = {} & \alpha_l + \alpha_{lk} \ln p_k + \alpha_{ll} \ln p_l + \alpha_{lfls} \ln p_{fls} + \alpha_{lfhs} \ln p_{fhs} \\ & + \alpha_{lq} \ln q + \alpha_{le} \ln e + \alpha_{lt} t + \epsilon_l,\end{aligned} \tag{4}$$

$$\begin{aligned}s_k = {} & \alpha_k + \alpha_{kk} \ln p_k + \alpha_{lk} \ln p_l + \alpha_{kfhs} \ln p_{fls} \\ & + \alpha_{kfhs} \ln p_{fhs} + \alpha_{kq} \ln q + \alpha_{ke} \ln e + \alpha_{kt} t + \epsilon_k,\end{aligned} \tag{5}$$

$$\begin{aligned}\ln e = {} & \beta_0 + \alpha_{e^*} \ln e^* + \sum_g \beta_g w_g + \sum_g \phi_g w_g (\ln e^*) \\ & + \sum_h \delta_h v_h + \sum_j \lambda_j \ln p_j + \delta_q q + \delta_t t + \epsilon_e.\end{aligned} \tag{6}$$

[17] The following conditions are imposed to ensure that the cost function is linearly homogeneous in input prices: $\Sigma_i \alpha_i = 1$ and

$$\sum_j \alpha_{ij} = \sum_j \alpha_{jq} = \sum_j \alpha_{je} = \sum_j \alpha_{jt} = 0, \quad i, j = k, l, fls, fhs.$$

[18] There are seven different types of emission rate standards in our data set. They include pounds of SO_2 emitted per hour, pounds of SO_2 per million Btus in fuel, pounds of sulfur per million Btus in fuel, percentage of sulfur content of fuel by weight, ambient air quality concentration of SO_2, parts per million of SO_2 in stack gas, and "other." The latter three standards, which together make up approximately 1 percent of all observations, could not be directly converted to an emission rate standard in pounds of SO_2 per million Btus in fuel. The standard used for these observations was the highest observed emission rate over the period of observation. Time periods over which emission rates are averaged are divided into periods of less than or equal to 24 hours or greater than 24 hours. A third category is included for units not faced with a known averaging time (e.g., periodic stack testing and not specified).

The estimated model includes input share equations for labor and capital only, and not fuel type. This is necessary because of the large number of zero values for inputs of low-sulfur and high-sulfur coal. At the level of the generating unit, only one type of coal is typically used, implying a zero cost share for the alternative fuel type. To avoid the bias that zero shares would introduce in our estimates, we include only the share equations for generating capital and labor.

The estimation of abatement cost functions is further complicated by the fact that over half of the units in our database exhibit non-cost-minimizing behavior in their choices of fuel at some time during the sample period. As a result of long-term fuel contracts, regulatory incentives, or other unobservable transaction costs associated with fuel switching, these units did not immediately switch to low-sulfur coal when it appeared to be economic for them to do so.[19] In some cases, remaining in long-term contracts may have provided a hedge against price fluctuations. In other cases, utilities may have had little incentive to respond to price changes if fuel prices could be passed on to consumers (Atkinson and Kerkvliet 1989). In any event, these observations violate the assumption of cost minimization implicit in the specification of the model and are therefore excluded when we estimate the cost function.[20]

The cost function, corresponding share equations, and the emission rate equation are estimated by full information maximum likelihood methods using panel data for the period 1985–94. The stochastic disturbances in the estimating equations for any observation are assumed to be correlated across equations.

Our interest centers on the marginal cost of achieving emissions rate e, which can, in turn, be translated into a marginal cost function for tons of SO_2. In general terms, the marginal cost of emissions function, $\partial C/\partial e$, is usually negative over observed ranges of emissions. The negative of this function, $-\partial C/\partial e$, will henceforth be referred to as the marginal abatement cost function. To describe the marginal cost of abating a ton of SO_2, the cost of a given percentage reduction in the

[19] We tested to see whether firms that are apparently violating cost-minimizing behavior are bound by older contracts relative to other firms, on the assumption that the older the contract the less likely it is to reflect current prices facing the firm. Although the contract age of plants that are apparently not cost-minimizing is longer by six months (out of an average contract age of five years), this does not seem to be a large enough difference to account for all of the non-cost-minimizing behavior.

[20] Excluding non-cost-minimizing observations results in eliminating some units for certain years but still enables us to estimate a cost function for these units. They are therefore included in our calculation of the least-cost solution. The number of observations fell from 7,147 to 5,314 after we eliminated observations that violate the assumption of cost minimization. The number of plants included in the cost function estimation fell from 273 to 260, and generating units fell from 761 to 734.

emissions rate can be converted into the equivalent reduction in tons of SO_2.[21]

Data

Our data set consists of virtually all privately and publicly owned phase I coal-fired generating units and all privately owned phase II coal-fired units.[22] These 829 units were responsible for 87 percent of all SO_2 emissions produced by coal-fired power plants in 1985 and 85 percent of all emissions in 1994. For each of the 734 generating units that switch fuel, we compiled data on generating capital, labor, and inputs of high- and low-sulfur coal for 1985–94. The data also include the SO_2 emission rate standard facing the generating unit, its mean annual emission rate in pounds of SO_2 per million Btus of heat input, and output in kilowatt hours. The input prices facing each power plant complete the data set. (See App. A for a more complete description.) For units that scrubbed their emissions over the 1985–94 period, we also obtained information about scrubbing capital in order to compute the average cost of emissions reductions for these units.

To describe sulfur content, we distinguish two classes of coal. Coal that when burned in a standard boiler generates no more than 1.2 pounds of SO_2 per million Btus of heat input is defined as low-sulfur coal; all other is high-sulfur coal. This distinction is not entirely arbitrary. Coal resulting in 1.2 pounds of SO_2 or less is termed "compliance coal" because of its ability to meet the original NSPS, in effect from 1971 to 1978. It will also meet phase II emission standards, on average.[23] To

[21] The marginal cost of a change in the SO_2 emissions rate, e, at a particular value of e is defined as negative one times the product of the elasticity of total cost with respect to the emission rate and the ratio of total cost to the observed emissions rate, or $-\partial C/\partial e = (-\partial \ln C/\partial \ln e)(C/e)$. The marginal cost of abating an additional ton of SO_2 emissions may be derived from the fact that $e = SO_2/\text{mmBtu}$, where mmBtu is millions of Btus of heat input. It follows that

$$-\frac{\partial C}{\partial SO_2} = -\frac{\partial C}{\partial e(1/\text{mmBtu})}$$
$$= \left(-\frac{\partial \ln C}{\partial \ln e}\right)\left(\frac{C}{SO_2/\text{mmBtu}}\right)(1/\text{mmBtu})$$
$$= \left(-\frac{\partial \ln C}{\partial \ln e}\right)\left(\frac{C}{SO_2}\right).$$

[22] The data set excludes all cooperatively owned plants, which are subject to reporting requirements different from those of either privately or publicly owned plants.

[23] An alternative approach to modeling the sulfur content of coal, used by Kolstad and Turnovsky (1998), is to allow plants to select sulfur content as a continuous attribute, given a hedonic price function for coal. We attempted this approach but were unable to obtain reliable estimates of hedonic price functions for each state and year.

approximate the cost of low-sulfur coal for a firm that purchased only high-sulfur coal, we use the average price of low-sulfur coal in the state in which the plant is located. In all cases, we use the contract price rather than the spot price.[24]

Results of the Estimation

To summarize the results of our estimation, we evaluate the MAC function for each fuel-switching unit at 1985 and 1994 emission levels. Table 1 presents the mean and standard deviation of the marginal cost of abating a ton of SO_2, when marginal abatement costs for different units are weighted by SO_2 emissions. In the 1994 time period, 89 percent of all the predicted marginal abatement costs are significantly different from zero at the 5 percent level. (See Appendix table B1 for estimated coefficients.)

Table 1 sheds light on differences in marginal abatement costs between phase I units (including so-called table A units, which are units named in the legislation that must participate in phase I, and compensation and substitution units, which opted into phase I) and phase II units. The table indicates that marginal abatement costs are substantially higher for phase II units ($1,092 per ton of SO_2, on average) than for phase I units ($121 per ton of SO_2, on average). This is not surprising given the much lower emission rates of phase II units. The range of marginal abatement costs is also much higher for phase II than for phase I units. In 1994, marginal abatement costs range from about −$90 per ton for low-cost phase II units to more than $2,700 per ton. The range for phase I units is narrower: from approximately −$260 per ton to $710 per ton.[25]

It is also clear from table 1 that marginal abatement costs have fallen over time for both phase I and phase II units. Indeed, the mean marginal abatement cost has fallen by nearly 50 percent for phase I units and almost 20 percent for phase II units. The fact that average emission rates have fallen over time suggests that the MAC curve for each unit has itself fallen between 1985 and 1994.

There are at least two reasons why MAC curves have fallen. One is that the delivered price of coal, for both high- and low-sulfur coal, has

[24] In 1985, 89 percent of all coal was purchased through long-term contracts rather than on the spot market. Although this percentage has declined through time, 80 percent of all coal was still purchased through long-term contracts in 1995. For this reason we use contract prices throughout the analysis.

[25] Table 1 indicates that marginal abatement costs are negative for at least 10 percent of the units in each category in 1994. As noted above, this failure to take advantage of cost-saving opportunities to switch fuel may be the result of an inability to escape from long-term fuel contracts or insufficient incentives to find the lowest-priced fuel as a result of regulatory fuel adjustment clauses (i.e., non-cost-minimizing behavior).

TABLE 1

1985 AND 1994 WEIGHTED AVERAGE MARGINAL ABATEMENT COSTS AND SO_2 EMISSION RATES FOR COAL-FIRED UNITS WITHOUT SCRUBBERS

	Number of Units	Emission Rate (lbs/mmBtu)	Mean MAC	Standard Deviation	Tenth Percentile	Ninetieth Percentile
			1985			
Phase I	341	3.65	\$250	\$485	\$7	\$582
Table A	241	4.09	\$164	\$308	−\$1	\$401
Compensation and substitution	100	1.97	\$604	\$812	\$57	\$1,108
Phase II	362	1.33	\$1,332	\$1,836	\$183	\$3,924
Total	703	2.45	\$811	\$1,467	\$51	\$2,636
			1994			
Phase I	318	2.82	\$121	\$475	−\$64	\$705
Table A	226	3.15	\$104	\$370	−\$49	\$431
Compensation and substitution	92	1.47	\$190	\$763	−\$93	\$1,499
Phase II	360	1.22	\$1,092	\$2,469	−\$88	\$2,717
Total	678	1.90	\$680	\$1,958	−\$82	\$1,935

NOTE.—Emission rates are weighted by total heat input and marginal abatement costs are weighted by total SO_2 emissions. Figures are based on units that were both included and excluded from the cost function estimation in order to make comparisons between years meaningful.

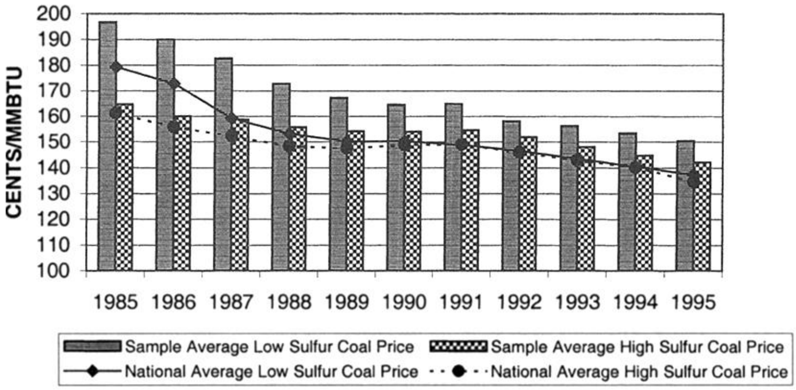

FIG. 1.—Low-sulfur and high-sulfer coal prices

declined over the period. This is illustrated by figure 1, where bars show the nominal (not adjusted for inflation) prices for delivery to utilities of each type of coal, by year, averaged across all units in our sample. In the same figure, lines show U.S. average nominal delivered coal prices, by sulfur content, computed for all utilities in the nation.[26] Figure 1 indicates that the prices of both types of coal fell between 1985 and 1995; however, the price of low-sulfur coal fell faster because of the decline in the cost of transporting low-sulfur coal by rail. What figure 1 does not show is that the price of low-sulfur coal was lower than the price of high-sulfur coal for 20 percent of the units in our sample in 1985 and for 25 percent of the units in our sample in 1994. Over the same period, the quantity of low-sulfur coal delivered to electric utilities rose significantly. The second reason for a fall in the MAC curve is technical progress in abating SO_2 emissions, resulting in part from more general technical progress in electricity generation.

How important are price changes and technical progress in explaining the fall in MAC curves? To answer this question, figure 2 plots estimated

[26] The bars in this chart reflect coal prices as they appear in our data set. In computing the price of low- (high-) sulfur coal, we have weighted the price actually paid by plants that purchase low- (high-) sulfur coal by their heat input and have similarly weighted the *predicted* price of low- (high-) sulfur coal for plants that purchased only high- (low-) sulfur coal. The lines represent national average coal prices, which are computed by averaging the prices paid only by firms that actually purchased each type of coal, including firms excluded from our data set. The resulting prices of low-sulfur coal are slightly lower than our estimates. Our prices of low-sulfur coal reflect the fact that many plants in our data set faced with higher than average prices of low-sulfur coal do not actually purchase low-sulfur coal. Therefore, the average price of low-sulfur coal in our data set is higher than the national average. By symmetric reasoning, the average price of high-sulfur coal in our data set is lower than the national average.

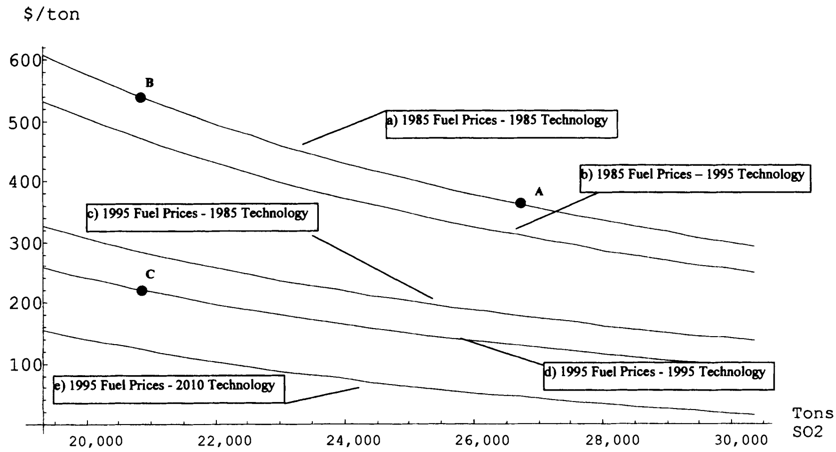

FIG. 2.—Effict of changes in fuel prices and technical change on marginal abatement cost functions.

MAC curves for a generating unit with average phase I input and output characteristics using (*a*) 1985 fuel prices and 1985 technology, (*b*) 1985 fuel prices and 1995 technology, (*c*) 1995 fuel prices and 1985 technology, (*d*) 1995 fuel prices and 1995 technology, and (*e*) 1995 fuel prices and 2010 technology. In all five curves, output as well as all nonfuel input prices are held constant. The effect of technological improvements, represented by the vertical distance between curves *a* and *b*, accounts for about 20 percent of the change of the MAC function, or a decline of about $50 per ton between 1985 and 1995. The effect of changes in fuel prices, represented by the vertical distance between curves *b* and *d*, accounts for the remaining 80 percent of the fall in the MAC function, a decline of about $200 per ton.

This figure also demonstrates why marginal abatement costs computed at the plant's actual level of emissions have fallen even as emissions have themselves declined. Without technological change or changes in fuel prices, an average plant would move in the figure from point *A* in 1985 to point *B* in 1995 (from approximately 27,000 to 21,000 tons of SO_2). Its marginal abatement cost would increase from approximately $360 per ton of SO_2 abated to $540 per ton of SO_2 abated, as emissions were reduced. With changes in fuel prices and technology, however, the unit actually moves from point *A* to point *C*, where its marginal abatement cost is only about half as large as it was originally.[27] If the current trend in technological improvements continues until the year 2010, this *average* unit's marginal abatement cost will fall by an additional $100 per ton (curves *d* to *e*).

IV. The Least-Cost Solution and Potential Gains from Trade in the Long Run

A. *Preferred Estimates of the Least-Cost Solution*

To estimate potential gains from allowance trading in the long run, we use our econometric estimates to compute the least-cost solution to achieving the SO_2 cap in the year 2010. This requires that we make assumptions about parameters that will shift the MAC functions over time: the rate of growth of electricity production (Q), the future path of fuel prices (*pls* and *phs*), and the rate of technical progress. We must also determine the rate at which coal plants in existence in 1995 will be retired from service and what SO_2 emissions would have been in 2010 in the absence of the CAAA. The assumptions we make about these parameters are applied consistently across all scenarios we compare.

Electricity output.—We assume that electricity production averaged over all coal-fired units increases at the rate of 1.49 percent per year.[28] Output is, however, likely to increase more rapidly at scrubbed units, which we assume will be utilized at 80 percent of capacity by 2010.[29] This fixes the emissions of scrubbed units in the long run. We allocate remaining generation and emissions under the cap to fuel-switching units.

Input prices and technical change.—In parameterizing the MAC functions of fuel-switching units, we assume that the real prices of high- and low-sulfur coal remain at 1995 levels and that the rate of technical change experienced between 1985 and 1994 continues through 2010.[30]

[27] It is important to keep in mind that the relative importance of technological change and fuel prices on an individual unit's MAC function depends greatly on where in the United States the unit is located. Generating units located in areas that have had access to relatively inexpensive low-sulfur coal for some time would not see a substantial drop in their MAC functions due to changes in coal prices.

[28] The Electric Power Research Institute (1997) assumes an annual increase in generation from coal-fired facilities of 1 percent per year through 2005 and flat thereafter; the Environmental Protection Agency (1995) assumes an average annual increase of 1.3 percent for 30 years; the 1996 *Annual Energy Outlook* by the Energy Information Administration (EIA) assumes an increase in coal-fired generation of 1.1 percent annually through 2015, but this estimate is revised in 1997 by the EIA to 1.49 percent.

[29] Utilization rates at scrubbed units have been increasing over time. There were 28 generating units with retrofit scrubbers in place by the beginning of 1995. The highest utilization rate in 1995 was 88 percent, and four were above 80 percent utilization.

[30] The fuel price assumption is consistent with the EIA's 1996 *Annual Energy Outlook;* its 1997 edition revised the forecast to indicate that the sulfur premium would shrink slightly further.

Retirement of coal-fired power plants.—We assume that 11 gigawatts of coal-fired capacity in place in 1995 will be retired by the year 2010 and all of that coal-fired capacity will be replaced by natural gas.[31] Economic decisions not directly related to the CAAA are the primary determinant of the timing and decision to switch to natural gas. This will reduce the emissions of these units to negligible levels, thus freeing up allowances and reducing compliance costs for units that remain in the market.

Baseline emissions.—We compute baseline emissions—those that would have prevailed without Title IV—using 1993 emissions *rates* applied to 2010 levels of electricity production. We assume that the declines in emission rates that occurred between 1985 and 1993 were primarily the result of decreases in the price of low-sulfur coal and would have happened in the absence of Title IV.[32]

Continuous emissions monitoring data.—An important feature of the 1990 CAAA is that SO_2 emissions must be measured by a continuous emissions monitoring system (CEMS) rather than estimated on the basis of fuel consumption. Previous studies all use engineering estimates of SO_2 emissions. A comparison of the two measurement techniques reveals that, in 1995, CEMS emissions were about 7 percent higher than estimated emissions, implying that the SO_2 cap is effectively 7 percent below the cap based on engineering formulas. To be consistent with actual practice, we use CEMS data.

Minimum Compliance Costs in the Preferred Case

Under the assumptions above, the minimum annual cost of achieving the SO_2 cap of 8.95 million tons in 2010 is $1.04 billion (1995 dollars). Of this total, $380 million represents the cost incurred by plants that switch fuel, which account for about 60 percent of reductions from baseline emissions.[33] The other 40 percent of reductions come from plants that have built scrubbers. For plants that have installed scrubbers, annualized capital costs are $382 million per year and variable costs $274 million per year. No retrofit scrubbers in addition to those con-

[31] The EIA's 1997 *Annual Energy Outlook* predicts that 22 gigawatts of coal-fired capacity will be retired between 1995 and 2010. Given recent experience with coal plant life extension and developments in monitoring technology that have lowered maintenance costs (Ellerman 1998), we expect substantially fewer coal plants to actually retire over that 15-year horizon.

[32] Emission reductions before 1995 were not bankable, so there is no reason why fuel-switching plants would reduce their emissions in 1990–93 to comply with Title IV. In 1994, however, modifications to existing equipment were made to prepare for compliance in 1995 (Ellerman and Montero 1998).

[33] We have investigated the implied change in the transportation of low-sulfur coal between the least-cost solution and present activity. We find it to be a modest extension of recent trends in the increased use of low-sulfur coal.

structed in phase I, which are assumed to be built under all scenarios, are found to be economic in the least-cost solution.

The marginal cost of emissions reduction, which should approximate the long-run permit price, is $291 per ton of SO_2. This assumes that the marginal ton of SO_2 is reduced via fuel switching, an assumption that is justified if one compares the cost of reducing SO_2 by installing retrofit scrubbers with the cost via fuel switching. Though the useful life of a retrofit scrubber is likely to be close to 20 years (Environmental Protection Agency 1995), the investment decision should reflect current financial and regulatory uncertainties in the industry, which call for a 10-year payback life (Electric Power Research Institute 1997). With this decision rule, the average cost per ton of reducing SO_2 through additional retrofit scrubbers is $360. Since this exceeds the marginal cost of SO_2 reduction via fuel switching, there is no reason why the marginal generating unit would scrub emissions.

B. *Comparisons and Sensitivity Analyses*

Table 2 reports sensitivity analyses for our long-run cost estimates and compares our costs with two previous EPA estimates. Column 1 reports estimates of costs under command and control (a uniform emission rate standard). Column 2 reports costs under the least-cost solution. Columns 3 and 4 report marginal and average costs per ton of abatement in the least-cost solution. Column 5 reports the potential gains from trade (the difference between cols. 1 and 2).

Our preferred estimate of annual CAAA compliance costs of $1.04 billion per year is far lower than the EPA predicted when the 1990 Clean Air Act Amendments were drafted (Environmental Protection Agency 1989, 1990). In fact, it is less than half EPA's estimates of the costs of the trading program, which are reported at the bottom of table 2 in the fourth and fifth rows. This raises two questions: Is our estimate of compliance costs biased downward? If not, why is it so much lower than EPA's original estimates of such costs?

The assumptions made above with regard to electricity generation and fuel prices are likely to overstate rather than understate costs. We assume, for example, the same rate of growth in electricity generated by coal and a slower rate of retirement of coal-fired plants than official predictions (EIA's 1997 *Annual Energy Outlook*). The assumption that high- and low-sulfur coal prices remain at their 1995 level is also conservative: in 1997 the EIA predicts a reduction in the low-sulfur premium below 1995 levels.

The one assumption that might bias our cost estimates downward is our assumption that technical progress will continue from 1995 until 2010 at the same rate as between 1985 and 1994. If we assume, at the

TABLE 2
LONG-RUN (Phase II, Year 2010) COST ESTIMATES (1995 Dollars)

	Total Cost under "Enlightened" Command and Control (Billions) (1)	Total Cost under Efficient Trading (Billions) (2)	Marginal Cost per Ton SO_2 (\$/Ton) (3)	Average Cost per Ton SO_2 (\$/Ton) (4)	Potential Gains from Trade (Billions) (5)
Preferred estimate	1.82	1.04	291	174	.78
1995 technology	2.23	1.51	436	198	.72
1989 prices and 1989 technology	2.67	1.90	560	236	.77
EPA (1990)	...	2.3–5.9	579–760	299–457	...
EPA (1989)	...	2.7–6.2		377–511	...

other extreme, that technical progress stops in 1995, our estimate of compliance costs rises to $1.51 million (1995 dollars), and our estimate of long-run allowance price rises to $436 per ton of SO_2 (see row 2 of table 2). Even this extreme assumption puts our estimate of total compliance costs below the Environmental Protection Agency's (1990) estimate of $2.3–$5.9 billion (table 2) and our estimate of marginal cost ($436) below EPA's estimate of $579–$760 per ton.[34]

It is important for two reasons to understand why these estimates differ. One is to see whether there is a systematic tendency (ex ante) to overestimate the cost of environmental regulations. That costs are systematically overestimated has been alleged by both economists and environmentalists (Goodstein and Hodges 1997; Harrington, Morgenstern, and Nelson 1999) and is an especially timely issue in light of debates over the cost of reducing greenhouse gases.

The second reason is that the factors that explain why estimates of compliance costs have fallen also explain why the costs of command-and-control approaches to reducing SO_2 have fallen and why the potential gains from allowance trading are also lower than originally anticipated. One reason for EPA's high estimates of compliance costs could be failure to foresee the continued fall in the low-sulfur coal premium, as well as continuing technical progress in fuel switching.[35] To estimate the potential magnitude of these effects, we recompute the least-cost solution using 1989 prices and technology.[36] Columns 2 and 3 of the third row of table 2 show that both total costs and marginal abatement costs rise by about 90 percent (relative to our "preferred" estimates in the first row). Total costs rise from $1.04 billion to $1.90 billion, and marginal abatement costs rise from $291 to $560. When fuel switching determines the marginal cost of compliance, using 1989 fuel prices and technology can produce marginal cost estimates approximately as large as those that were predicted when Title IV was written ($579–$760; see col. 3 of table 2). Total costs also increase, in part because a higher sulfur premium lowers the percentage of emissions reductions that can be obtained for free. Under our preferred scenario, 57 percent of emissions reductions from plants that switch fuel are obtained by realizing negative marginal abatement costs (switching to cheaper low-sulfur

[34] Recent engineering studies that acknowledge the use of low-sulfur coal for compliance have also identified the declining trend of marginal and annual total costs of compliance (Electric Power Research Institute 1995, 1997; Environmental Protection Agency 1995).

[35] We cannot assert this with certainty since assumptions regarding coal prices and changes in technology are not transparent in EPA's reports (1989, 1990).

[36] We also assume that emissions are estimated on the basis of fuel consumption, as they were in studies prior to the passage of Title IV.

coal). This figure, however, falls to 21 percent when 1989 prices and technology are used.[37]

Failure to foresee changes in prices and technical progress, however, does not explain all of the difference in total cost estimates. Also important are differences in the baseline from which emissions reductions are measured. In all our calculations, we assume that the emission *rates* (pounds of SO_2 per million Btus) that would have prevailed in the absence of the 1990 CAAA are those that prevailed in 1993. They are much lower than 1989 emission rates; hence the reductions in emissions necessary to achieve the 8.95 million ton cap, by our calculations, are much lower than imagined in 1989 (specifically, about 2 million tons lower). When MAC curves are held constant, lowering the necessary reduction in emissions will lower total compliance costs.

Finally, EPA's estimates of total compliance costs are higher than ours because it assumed that more retrofit scrubbers would be built (37) than were actually constructed (28). It also failed to foresee a 50 percent fall in the cost of scrubbing that we identify.

C. *Potential Gains from Trade*

We now consider the cost of meeting the SO_2 cap using a command-and-control approach and compute the potential gains from trade as the difference between this cost and the cost of compliance under the least-cost solution. We first model command and control as a uniform performance standard that is designed to achieve the same level of emissions as the trading program, consistent with the goal of Title IV to achieve an average emissions rate of 1.2 pounds of SO_2 per million Btus of heat input.[38]

For our preferred case, we estimate the potential gains from trade compared to the "enlightened" command-and-control scenario to be \$784 million (43 percent of the cost of command and control). The potential gains from trade are estimated by subtracting the total costs

[37] While this 57 percent estimate may seem high, we have evidence from 1995 and 1996 that utilities are realizing such economic cost savings. In 1995, one-quarter of a potential \$443 million in savings from fuel switching was realized. In 1996, half of \$644 million in potential savings was realized. We believe that increased competition in the electric utility industry will motivate generators to take advantage of these savings.

[38] The uniform emission rate standard does not take into account the fact that some units may face unrealized "economic" emission reductions beyond those mandated by the standard. Therefore, emissions are lower under the uniform standard than they are under a trading program, which provides firms with higher abatement costs the flexibility to capture the slack in the effective emission constraint at other firms (Oates, Portney, and McGartland 1989). It should be noted that if a uniform emission rate is set so that total emissions under a trading program equal those under a uniform emission rate standard, it will still be the case that costs under the trading program will be less than under a uniform emission standard. Potential gains from trade will be lower, however.

under efficient trading in column 2 of table 2 ($1.04 billion) from the total costs under a command-and-control program in column 1 ($1.82 billion). While these gains constitute 43 percent of the cost of command and control, they are not as large as those originally predicted. The General Accounting Office (1994), for example, estimated that a command-and-control cap on emissions at each generating unit would cost approximately $5.3 billion annually and that the reduction in costs from efficient trading would be $3.1 billion (about 60 percent of the command-and-control figure).

The explanation for our more modest estimates of trading gains is clear: the factors that have caused marginal abatement costs to fall also would have lowered the costs of achieving the SO_2 emissions cap via command and control. These factors include the fall in the price of low-sulfur coal and, to some extent, technical improvements that have facilitated fuel switching.[39] It should also be noted that, in addition to *lowering* MAC curves (see fig. 2), the fall in low-sulfur coal prices has made MAC curves *more homogeneous*. The reason is that the cost of transporting low-sulfur coal to more distant locations, for example, the East and Southeast, has fallen, rendering differences in transportation cost a less important component of the overall cost of fuel switching. Since a major source of trading gains is differences in MAC curves among units in the market, this increased homogeneity is also partly responsible for low gains from trade.

One alternative to our preferred assumptions is the possibility that, in the absence of the cap and trading program, a different command-and-control policy would have been adopted. In place of our enlightened command-and-control policy of an emission rate performance standard, we considered the possibility of forced retrofit scrubbing to achieve an equivalent level of aggregate emissions. We assume the cost per ton of additional retrofit scrubbing to be the same as that observed for retrofit scrubbers built in phase I.[40] The aggregate cost of this approach would have been $2.6 billion per year. Compared to this alter-

[39] Some of the technological developments in fuel blending may not have occurred under a uniform emissions standard since blending of coals with different sulfur contents by itself (i.e., without the option of purchasing allowances) generally would not be sufficient to achieve the required emission reductions. Similarly, there would have been less incentive to improve the performance of scrubbing equipment under a uniform command-and-control emission rate standard, and the witnessed improvements may not have been realized. To the extent that the effects of the allowance trading program on technological change in emissions reduction are reflected in our data, our estimates of the costs of a uniform emission rate standard and the potential gains from trade are likely understated.

[40] That is, we assume a 20-year depreciation schedule with a historic cost of capital, which was applicable to retrofit scrubbers initiated before the start of phase I. These assumptions result in a lower cost of scrubbing than if we were to assume that some of these scrubbers would be built after 1995, when a higher cost of capital for the industry would be applicable.

native, the trading program generates potential cost savings of $1.6 billion ($2.6 − $1.04 billion). About half of these savings are captured by moving from forced scrubbing to a more flexible uniform performance standard ($2.6 − $1.82 billion = $780 million).

The relatively modest potential gains from trading relative to a more flexible uniform performance standard ($784 million) should not be interpreted as a criticism of the allowance market, but they are likely to have an impact on market performance. If potential gains from trade are small and transaction costs of using the market are substantial, utilities will be less eager to trade allowances. In the next section, we analyze the performance of the SO_2 allowance market in 1995 and 1996 to determine the potential gains from trade under a perfectly functioning market and how much of these gains actually have been realized.

V. The Performance of the Allowance Market in 1995 and 1996

In 1995, the aggregate emissions of phase I units were approximately 5.3 million tons, rising to 5.44 million tons in 1996. To compute the least-cost method of achieving these emissions levels, we parameterize MAC functions for each unit using actual output levels and input prices. Technical progress is assumed to occur at the same annual rate observed between 1984 and 1995. We take as our baseline 1993 emission *rates,* which we apply to 1995 and 1996 electricity generation to predict emissions in the absence of Title IV.

The least-cost solution yields a common marginal abatement cost (for the last ton emitted by all units that switch fuel) and a set of efficient emission levels for all generating units. To compute total costs, we integrate each unit's MAC curve from baseline emissions to emissions under the least-cost solution. The total cost of *actual* emissions in each year is computed analogously, except that integration under the MAC curve occurs from baseline emissions to actual 1995 and 1996 emissions.[41]

Column 1 of table 3 suggests that the allowance market did not realize potential gains from trade in 1995. The estimated annual command-and-control cost of achieving 1995 emissions is $802 million (1995 dollars) as shown in row 1 of column 1. The estimated cost with efficient trading is $552 million as shown in row 2; hence, the potential gains from trading are estimated to be $250 million. Row 5 reports the esti-

[41] Emission rates are based on Department of Energy–Energy Information Administration engineering estimates, not CEMS data. Because both the heat input and SO_2 emissions estimated by the DOE-EIA are lower than the CEMS measurements, the estimated emission rates under DOE-EIA and CEMS are equal to each other on average.

TABLE 3
PHASE I (1995 and 1996) COST ESTIMATES (1995 Dollars)

	1995 (1)	1996 (2)
Total cost under "enlightened" command and control (millions)	802	777
Total cost under efficient trading (millions)	552	571
MAC under efficient trading ($/ton)	101	71
Potential gains from trade (millions)	250	206
Actual total compliance cost (millions)	832	910

mated *actual* cost of achieving 1995 emissions to be $832 million.[42] The difference between the costs under efficient trading and actual compliance costs is $280 million, suggesting that potential cost savings were unrealized in the first year of the allowance market. Indeed, the fact that our estimate of actual compliance costs exceeds our estimate of the cost of command and control suggests that the uniform performance standard would have been no less efficient than the actual pattern of emissions chosen by utilities. Our confidence in these results is strengthened by noting that our estimate of the actual compliance costs is close to estimates obtained by Ellerman et al. (1997) ($728 million) based on a survey of the industry.

Performance under the program did not change dramatically in 1996. Our estimate of the command-and-control cost of achieving 1996 emissions is shown in column 2 of table 3 to be $777 million (1995 dollars). The least-cost solution under efficient trading is $571 million, and the potential gains from trade are estimated to be $206 million ($777 − $571 million). The estimated *actual* cost of achieving 1996 emissions increased slightly from the previous year, to $910 million, partly because of increased utilization of scrubbed units. This suggests that $339 million of potential cost savings were unrealized in the second year of the allowance market.

The failure of the allowance market to achieve the least-cost solution in 1995 and 1996 is neither surprising nor alarming. One reason why our estimates of actual compliance costs exceed estimated compliance costs under the least-cost solution is that our model does not account for short-term adjustment costs that may be faced by firms in the first two years of the program. Adjustment costs associated with changing fuel contracts and capital expenditures as well as regulatory policies may make it appear that firms have failed to minimize costs when they have actually done so. Indeed, this fact may explain why our estimates of

[42] All the estimates in table 3 for 1995 include the capital and variable costs of scrubbing ($496 million).

actual compliance costs exceed estimates of compliance costs under command and control in the short run.

The second reason why cost savings were not realized is that little trading occurred during the first two years of the market. That trading began slowly is to be expected. Title IV represents a dramatic departure from traditional environmental regulation. It requires utilities to manage a financial asset—emission allowances—for which there is no precedent. It also requires a well-functioning market in allowances, which takes time to establish. Allowance trades are growing in volume (Environmental Protection Agency 1999). Economically significant trades between separate utility holding companies have doubled every year since the inception of the program through 1997, which suggests that utilities are increasingly taking advantage of the allowance market as a means to reduce compliance costs (Kruger and Dean 1997). In addition, the number of allowances used for compliance that were obtained through interfirm transactions increased by 50 percent between 1995 and 1996 (Environmental Protection Agency 1999).

We note, in closing, that the failure of firms to realize cost savings through trading cannot be inferred simply by comparing the marginal cost of the last ton emitted in the least-cost solution with the price of an allowance. Table 3 suggests that the marginal cost of abatement in the least-cost solution ($101 in 1995 and $71 in 1996) is close to the price at which allowances were trading (about $90). The similarity of these numbers does not, however, demonstrate that the market was operating efficiently. The two could be similar even if many participants opted out of the market, which was in fact the case. The allowance price was set by the subset of utilities that entered the market; those that did not failed to capture potential gains from trade.

VI. Conclusions

When the market for sulfur dioxide allowances was envisioned in the late 1980s, the cost of complying with the proposed SO_2 cap was thought to be much higher than it has, in fact, turned out to be. Likewise, the potential trading gains associated with the market were predicted to be much higher than the estimates presented above. The relatively lower trading gains that we predict for the allowance market in the long run are largely the result of two factors: declines in the price of low-sulfur coal and improvements in technology that have lowered the cost of fuel switching. These factors have lowered the gains from trade in two ways. First, they have lowered marginal abatement cost curves for most generating units, which has lowered the cost of achieving the cap either through a uniform emission rate standard or through allowance trading. Second, because spatial differences in coal prices (which include transportation costs) have been reduced, MAC curves have become more homogeneous. This has also lowered the gains from trade.

Our results have several important lessons for policy makers as they consider adopting an allowance trading approach to regulating other utility emissions such as nitrogen oxides and greenhouse gases. First, our findings lend support to the theory that the costs of compliance with incentive-based regulation are often overestimated ex ante. We show that estimates of the costs of compliance with the SO_2 reduction goals under Title IV have fallen substantially over time as a result of a combination of unanticipated declines in coal prices and technical change. This suggests that attempts to estimate the future costs of other pollution control programs may be similarly flawed, especially given the difficulty in forecasting future trends in technological change. This technology forecasting task is made more complicated by the introduction of greater competition in electricity markets, which is expected to accelerate the pace of technical change.

Second, our results suggest that, when policy makers design an allowance market, it is important for them to consider the source of trading gains and how these gains might change over time. The source of trading gains in the SO_2 allowance market is spatial differences in the price of high- versus low-sulfur coal. As these price differences have diminished, so have potential trading gains. The market for carbon dioxide is initially likely to generate large trading gains because coal-fired power plants, by converting to natural gas, can reduce their carbon dioxide emissions at a lower cost than oil- and gas-fired plants. Once this conversion is completed, however, trading gains within the electric utility industry will diminish.

Finally, our results suggest that it will take time for allowance markets to mature and, therefore, for the potential gains from trade to be realized. We suggest that, on the whole, the market failed to realize potential gains from trade in 1995 or 1996. The reluctance of many firms in the utility industry to take advantage of the allowance market may be a result of features of utility regulation that have limited incentives to participate in the market. As competition increases within the generation segment of the industry and as we enter the second phase of the allowance program, we expect to see greater use of the market to reduce the costs of environmental compliance. Yet formal trading in the SO_2 allowance market may still not achieve large cost savings compared to a uniform performance standard. The flexibility of the trading program has encouraged utilities to capitalize on advantageous trends, such as changing fuel prices and technological innovation that might have been delayed or discouraged by traditional regulatory approaches. The SO_2 program shows that a market in tradable emission rights is, indeed, feasible. As the electric utilities industry becomes more competitive, one would expect the advantages of emission trading programs for other pollutants to become more evident.

Appendix A

Data for estimating the generating unit cost functions and input share equations come from the Energy Information Administration's Form EIA-767 and the Federal Energy Regulatory Commission's (FERC's) Form 1 for the period 1985–96. Electric utility plant capital stock comes from Form 1 for the period 1982–96. Prior to this, capital stock data come from the EIA's annual report *Electric Plant Cost and Power Production Expenses* (*Expenses*) and precursors to this report. Coal prices come from the *Monthly Report of Cost and Quality of Fuels for Electric Plants* (*Monthly Report*) for 1985–96. The following list describes each of the variables that enter the cost function and input share equations.

Q, output: Electrical generation (kilowatt hours) by generating unit. Source: EIA-767.

p_{hs}, price of high-sulfur fuel: As in Gollop and Roberts (1983, 1985), the price of high-sulfur fuel is the weighted average price, in cents per million Btus, of high-sulfur fuel bought by the utility that owns the generator. An emission boundary of 1.2 pounds of SO_2 per million Btus of heat input is used to differentiate low- and high-sulfur coal. If the utility bought no high-sulfur fuel, then the price is equal to the price of low-sulfur fuel bought by the utility multiplied by the ratio of high- to low-sulfur coal prices in the state in which the plant is located. Source: *Monthly Report.*

p_{ls}, price of low-sulfur fuel: Measured in the same manner as the price of high-sulfur fuel. Source: *Monthly Report.*

p_l, wage rate: The utility's total labor expenditures are divided by the sum of the number of full-time employees and one-half of the number of part-time and temporary employees working for the utility. Publicly owned plants' wage rates are equal to the average wage rate of privately owned plants in the state in which the plant is located, or the region surrounding the state if no privately owned plants are located in the same state. Source: Form 1.

p_k, rental price of generation capital: The rental price of generation capital is equal to the utility's cost of capital plus the depreciation rate, adjusted for changes in the cost of construction (Cowing, Small, and Stevenson 1981). That is,

$$p_k = (R_u + DE) \times HW_{r,t},$$

where R_u is the utility's cost of capital, DE represents the depreciation rate, and $HW_{r,t}$ is the Handy-Whitman index of electric utility construction costs, which varies by region of the country and year, adjusted to reflect a base year of 1990. The financial cost of capital for privately owned plants is estimated as the sum of the long-term debt interest rate, the preferred stock dividend rate, and the required return on equity capital, where each factor is weighted by its respective capital structure proportion. The financial cost of capital for publicly owned plants is equal to the long-term debt interest rate reported in Moody's Municipal and Government Manual. Data for jointly owned plants come from the utility indicated as the operator by EIA-767. The depreciation rate is assumed to be 5 percent and is applied to the undepreciated value of capital stock remaining in each year. This is based on a decay pattern defined by the 1.5 declining balance method and a 30-year asset life. Source: Form 1 and *The Handy-Whitman Index of Public Utility Construction Costs* (calculated by Whitman, Requardt and Assoc., January 1995).

e^*, emission standard: The emission standard, in pounds of SO_2 per million Btus of heat input. Source: EIA-767.

e, average emission rate: The annual average emission rate for each utility plant. Source: Calculated by the EIA from information in EIA-767.

k, generation capital stock: The capital stock for each plant is calculated as follows:

$$CS_t = CS_{t-1} + \frac{NI_t}{HW_{r,t}}, \quad t = 1951, \ldots, 1995,$$

where CS_t is the adjusted capital stock for year t, NI_t is the net investment for year t, and $HW_{r,t}$ is the Handy-Whitman index, which varies by region of the country, for year t, adjusted to reflect a base year of 1990. The plant's net capital stock is equal to the initial investment in buildings and equipment plus the costs of additions minus the value of retirements. Each generator's capital stock is the product of the plant's capital stock, the generator unit's share of the plant's total generation capacity, and the percentage of time that the boiler was under load. Source: *Expenses,* EIA-767, and Form 1.

Generation capital expenditure: The product of the deflated generation capital stock and the rental price of generation capital.

Labor expenditure: The product of the wage rate and the total number of employees working at the plant multiplied by the generator's share of the plant's total generation capacity.

High-sulfur fuel expenditure: The product of heat input from high-sulfur coal, in millions of Btus, and the price of high-sulfur fuel. The type of fuel burned by the generating unit is determined from the unit's SO_2 emission rate before SO_2 removal.

Low-sulfur fuel expenditure: The product of heat input from low-sulfur coal, in millions of Btus, and the price of low-sulfur coal.

Scrubbing capital: The quantity of scrubbing capital at plants that currently scrub is determined by a lump-sum investment made when the scrubber is put in place. The annual cost of scrubbing capital is determined using a method analogous to that employed for generation capital with two exceptions. First, for generating capital we use firm-specific measures of the cost of capital, and for scrubbers we use an industry average cost of capital because only a few firms built scrubbers and we want scrubber costs to reflect the cost to the industry as a whole. Second, for generating capital the data on depreciation of in-place capital come from the FERC, whereas we assume the depreciation schedule for scrubbers. Capital investments in retrofit scrubbers built in phase I are amortized using a levelized capital recovery factor of 11.3 percent, which reflects an industry average nominal cost of capital of 9.5 percent (a real cost of 6.5 percent) and a 20-year cost recovery period. For additional investments, we assume a nominal cost of capital of 11.5 percent (a real cost of 8.5 percent) and a 10-year cost recovery period, yielding a levelized capital recovery factor of 17.3 percent.

Appendix B

TABLE B1
EMISSION EQUATION AND COST FUNCTION PARAMETER ESTIMATES

β_0	−1.18189	δ_q	.025434	α_{lq}	−.038168
	(.283785)		(.444724E-02)		(.471865E-03)
α_{e*}	.824970	δ_t	.020624	α_{flsq}	.029624
	(.090186)		(.276917E-02)		(.668690E-02)
β_{DH}	−1.08051	α_0	8.93057	α_{fhsq}	−.129491E-03
	(.389955)		(.226447)		(.676841E-02)
β_{DM}	−.125737	α_1	.534165	α_{kq}	.867337E-02
	(.087204)		(.030666)		(.101037E-02)
β_{DP}	−.131133	α_{fls}	−.845765	α_{le}	−.538269E-02
	(1.55036)		(.105390)		(.148797E-02)
β_{OT}	−.419885	α_{fhs}	−.404111	α_{flse}	−.454225
	(.076088)		(.109651)		(.022820)
β_{SB}	.751207	α_k	1.71571	α_{fhse}	.451805
	(.104390)		(.054526)		(.022838)
β_{SU}	−.340344	α_q	−.254062	α_{ke}	.780243E-02
	(.092253)		(.013426)		(.222931E-02)
Φ_{DH}	1.32542	α_e	−.340125	α_{lt}	.011471
	(.570489)		(.033449)		(.778806E-02)
Φ_{DM}	−.333184	α_t	.050289	α_{flst}	−.010615
	(.092656)		(.747385E-02)		(.387403E-02)
Φ_{DP}	.106607	α_{ll}	.038505	α_{fhst}	.587411E-02
	(1.87445)		(.095939E-02)		(.389566E-02)
Φ_{OT}	.090034	α_{lfls}	−.198530E-02	α_{kt}	.501848E-02
	(.090867)		(.395130E-02)		(.533437E-03)
Φ_{SB}	−.763329	α_{lfhs}	−.019384	γ_{qq}	.084244
	(.091229)		(.442843E-02)		(.100102E-02)
Φ_{SU}	−.019927	α_{lk}	−.017135	γ_{qe}	.019087
	(.100186)		(.486692E-02)		(.185753E-02)
δ_{24}	.099986	α_{flsfls}	−.015430	γ_{qt}	−.104989E-02
	(.012408)		(.078950)		(.390694E-03)
δ_{oth}	.059288	α_{flsfhs}	.149895	γ_{ee}	.072446
	(.020120)		(.078656)		(.016756)
λ_l	−.394573	α_{flsk}	−.132480	γ_{et}	.259888E-02
	(.031436)		(.667833E-02)		(.120259E-02)
λ_{fls}	.354107	α_{fhsfhs}	−.080841	γ_{tt}	.531355E-03
	(.018044)		(.079025)		(.581379E-03)
λ_{fhs}	−.215433	α_{fhsk}	−.049670		
	(.029490)		(.693520E-02)		
λ_k	−.721643	α_{kk}	.199285		
	(.074711)		(.788783E-02)		

NOTE.— Parameter labels: l = labor, fls = low-sulfur coal, fhs = high-sulfur coal, k= capital, t = time, q = electricity output, e = emission rate, e^* = emission standard, DH = pounds of SO_2 emitted per hour, DM = parts per million of SO_2 in stack gas, DP = pounds of SO_2 per million Btus of fuel, OH = other, SB = pounds of sulfur per million Btus of fuel, SU = percentage of sulfur content of fuel by weight, 24 = greater than 24 hours, oth = other or nonspecified averaging times.

References

Atkinson, Scott E., and Kerkvliet, Joe. "Dual Measures of Monopoly and Monopsony Power: An Application to Regulated Electric Utilities." *Rev. Econ. and Statis.* 71 (May 1989): 250–57.

Bohi, Douglas R. "Utilities and State Regulators Are Failing to Take Advantage of Emission Allowance Trading." *Electricity J.* 7 (March 1994): 20–27.

Bohi, Douglas R., and Burtraw, Dallas. "Utility Investment Behavior and the Emission Trading Market." *Resources and Energy* 14 (April 1992): 129–53.

Burtraw, Dallas. "The SO2 Emissions Trading Program: Cost Savings without Allowance Trades." *Contemporary Econ. Policy* 14 (April 1996): 79–94.

Burtraw, Dallas; Krupnick, Alan J.; Mansur, Erin; Austin, David; and Farrell, Deirdre. "Costs and Benefits of Reducing Air Pollutants Related to Acid Rain." *Contemporary Econ. Policy* 16 (October 1998): 379–400.

Cowing, Thomas G.; Small, Jeffrey; and Stevenson, Rodney E. "Comparative Measures of Total Factor Productivity in the Regulated Sector: The Electric Utility Industry." In *Productivity Measurement in Regulated Industries,* edited by Thomas G. Cowing and Rodney E. Stevenson. New York: Academic Press, 1981.

Electric Power Research Institute. "The Emission Allowance Market and Electric Utility SO_2 Compliance in a Competitive and Uncertain Future." TR-105490. Palo Alto, Calif.: Electric Power Res. Inst., September 1995.

———. "SO_2 Compliance and Allowance Trading: Developments and Outlook." TR-107891. Palo Alto, Calif.: Electric Power Res. Inst., September 1997.

Ellerman, A. Denny. "Note on the Seemingly Indefinite Extension of Power Plant Lives: A Panel Contribution." *Energy J.* 19, no. 2 (1998): 129–32.

Ellerman, A. Denny, and Montero, Juan Pablo. "The Declining Trend in Sulfur Dioxide Emissions: Implications for Allowance Prices." *J. Environmental Econ. and Management* 36 (July 1998): 26–45.

Ellerman, A. Denny; Schmalensee, Richard; Joskow, Paul L.; Montero, Juan Pablo; and Bailey, Elizabeth M. "Emissions Trading under the U.S. Acid Rain Program: Evaluation of Compliance Costs and Allowance Market Performance." Manuscript. Cambridge: Massachusetts Inst. Tech., Center Energy and Environmental Policy Res., October 1997.

Energy Information Administration. *Monthly Report of Cost and Quality of Fuels for Electric Plants.* Washington: Energy Information Admin., 1985–96.

———. *Annual Energy Outlook.* Washington: Energy Information Admin., 1996–97.

Environmental Protection Agency. "Economic Analysis of Title V [sic] (Acid Rain Provisions) of the Administration's Proposed Clean Air Act Amendments." Report prepared by ICF Resources Inc. Washington: Environmental Protection Agency, September 1989.

———. "Comparison of the Economic Impacts of the Acid Rain Provisions of the Senate Bill (S. 1630) and the House Bill (S. 1630)." Report prepared by ICF Resources Inc. Washington: Environmental Protection Agency, July 1990.

———. "Economic Analysis of the Title IV Requirements of the 1990 Clean Air Act Amendments." Report prepared by ICF Resources Inc. Washington: Environmental Protection Agency, September 1995.

———. "Cumulative Summary of Allowance Trading." www.epa.gov/acidrain/trading.html, 1999.

Fullerton, Don; McDermott, Shaun P.; and Caulkins, Jonathan P. "Sulfur Dioxide Compliance of a Regulated Utility." *J. Environmental Econ. and Management* 34 (September 1997): 32–53.

Fullerton, Don, and Metcalf, Gilbert E. "Environmental Controls, Scarcity Rents, and Pre-existing Distortions." Working Paper no. 6091. Cambridge, Mass.: NBER, July 1997.

Gollop, Frank M., and Roberts, Mark J. "Environmental Regulations and Productivity Growth: The Case of Fossil-Fueled Electric Power Generation." *J.P.E.* 91 (August 1983): 654–74.

———. "Cost-Minimizing Regulation of Sulfur Emissions: Regional Gains in Electric Power." *Rev. Econ. and Statis.* 67 (February 1985): 81–90.

Goodstein, Eban, and Hodges, Hart. "Polluted Data." *American Prospect* 8 (November–December 1997): 64–69.

Goulder, Lawrence H.; Parry, Ian W. H.; and Burtraw, Dallas. "Revenue-Raising versus Other Approaches to Environmental Protection: The Critical Significance of Preexisting Tax Distortions." *Rand J. Econ.* 28 (Winter 1997): 708–31.

General Accounting Office. "Air Pollution: Allowance Trading Offers an Opportunity to Reduce Emissions at Less Cost." GAO/RCED-95-30. Washington: Gen. Accounting Off., 1994.

Harrington, Winston; Morgenstern, Richard D.; and Nelson, Peter. "On the Accuracy of Regulatory Cost Estimates." Discussion Paper no. 99-18. Washington: Resources for the Future, January 1999.

Kalagnanam, Jayant, and Bokhari, Farasat. "A Market Simulation-Based Cost Module." Manuscript. Pittsburgh: Carnegie Mellon Univ., Dept. Engineering and Public Policy, 1995.

Kolstad, Charles D., and Turnovsky, Michelle H. L. "Cost Functions and Nonlinear Prices: Estimating a Technology with Quality-Differentiated Inputs." *Rev. Econ. and Statis.* 80 (August 1998): 444–53.

Kruger, Joseph, and Dean, Melanie. "Looking Back on SO_2 Trading: What's Good for the Environment Is Good for the Market." *Public Utilities Fortnightly* 135 (August 1997): 30–37.

Lile, Ron, and Burtraw, Dallas. "State-Level Policies and Regulatory Guidance for Compliance in the Early Years of the SO_2 Emission Allowance Trading Program." Discussion Paper no. 98-35. Washington: Resources for the Future, May 1998.

"New Initiatives in Environmental Protection." *Commonwealth* (March 31, 1997).

Oates, Wallace E.; Portney, Paul R.; and McGartland, Albert M. "The *Net* Benefits of Incentive-Based Regulation: A Case Study of Environmental Standard Setting." *A.E.R.* 79 (December 1989): 1233–42.

Office of Technology Assessment. "An Analysis of the 'Sikorski/Waxman' Acid Rain Control Proposal: H.R. 3400, the National Acid Deposition Control Act of 1983." Staff memorandum. Washington: Off. Technology Assessment, rev. July 12, 1983.

Rose, Kenneth. "Implementing an Emissions Trading Program in an Economically Regulated Industry: Lessons from the SO_2 Trading Program." In *Market Based Approaches to Environmental Policy: Regulatory Innovations to the Fore,* edited by Richard F. Kosobud and Jennifer M. Zimmerman. New York: Van Nostrand Reinhold, 1997.

Rubin, Jonathan D. "A Model of Intertemporal Emission Trading, Banking, and Borrowing." *J. Environmental Econ. and Management* 31 (November 1996): 269–86.

Siegel, Stuart A. "Evaluating the Cost Effectiveness of the Title IV Acid Rain Provisions of the 1990 Clean Air Act Amendments." Ph.D. dissertation, Carnegie Mellon Univ., 1997.

Temple, Barker and Sloane, Inc. "Evaluation of H.R. 3400: The 'Sikorski/Waxman' Bill for Acid Rain Abatement." Report. Washington: Edison Electric Inst., September 20, 1983.

Winebrake, James J.; Farrell, Alexander E.; and Bernstein, Mark A. "The Clean

Air Act's Sulfur Dioxide Emissions Market: Estimating the Costs of Regulatory and Legislative Intervention." *Resource and Energy Econ.* 17 (November 1995): 239–60.