Estimating the Gains to Emission Trading

Richard Newell, William Pizer and
Jhih-Shyang Shih

March 2004 • Working Paper

Resources for the Future
1616 P Street, NW
Washington, D.C. 20036
Telephone: 202–328–5000
Fax: 202–939–3460
Internet: http://www.rff.org

Estimating the Gains to Emission Trading

Richard Newell, William Pizer, and Jhih-Shyang Shih

Abstract

Over the past twenty years there has been a remarkable trend towards the use of market-based policies to control pollution. That trend has been fueled, in part, by economic arguments that these policies save money—a lot of money. Yet, to date most analyses of the gains to trade have been based on prospective engineering data rather than retrospective empirical cost data, sparking concern that existing estimates ignore actual (versus theorized) command-and-control implementation as well as the practical realities of pollution abatement efforts. This paper addresses such concerns by using data collected from 1979-1985 by the Census Bureau on both pollution abatement costs and abatement levels to estimate control cost functions and the potential gains to emissions trading. Our initial results, focusing on sulfur dioxide controls in the steel industry, find average annual savings of $300,000-$800,000 (1982$) per plant associated with a prospective shift to SO_2 emissions trading, or 5-14% as a share of overall air pollution control costs. The gains as a share of SO_2 control costs would be much higher.

Key Words: policy instruments, market-based, tradable permits

JEL Classification Numbers: Q28, Q38

Contents

Estimating the Gains to Emission Trading

Richard Newell, William Pizer, and Jhih-Shyang Shih*

1. Introduction

Over the past twenty years, there has been a remarkable trend in the United States toward the use of market-based environmental policies, specifically emissions trading. While Title IV of the 1990 Clean Air Act Amendments (the "acid rain trading program") is the most publicized example, there are numerous trading provisions in the regulation of various types of gasoline engines, regional air and water regulations, renewable energy mandates, as well as recent proposals for multipollutant regulation at power plants and economy-wide carbon limits.[1]

In large measure, this trend was initiated and fueled by economic arguments and evidence that emission trading can substantially reduce compliance costs without reducing environmental benefits. The economics literature includes studies based on prospective engineering analysis, some econometric modeling, and more recently global trading models (examining the potential for international climate change mitigation). By far the most common approach, plant-level engineering-economic analyses have estimated the gains from emissions trading to be cost reductions of anywhere from 7% to 95%.[2] In part responding to this startling range of estimates,

* Newell, Pizer, and Shih are Fellows at Resources for the Future, Washington, DC. This paper reports the results of research and analysis undertaken while the authors were research affiliates at the Center for Economic Studies at the U.S. Census Bureau. It has undergone a Census Bureau review more limited in scope than that given to official Census Bureau publications. It has been screened to insure that no confidential information is revealed. We acknowledge financial support from the National Science Foundation Grant SES-9976512. Research results and conclusions expressed are those of the authors and do not necessarily indicate concurrence by the Census Bureau or the National Science Foundation.

[1] See Stavins (2003) for a description of existing trading programs in the United States and elsewhere. Recent multipollutant proposals include S. 366, S. 485, and S. 843, all introduced in the 108th Congress. The Climate Stewardship Act (S. 139, also 108th Congress) proposed an economy-wide carbon cap.

[2] See page 403 of Tietenberg (1992). More recent analyses of policies to address the threat of global climate change

a recent paper by Newell and Stavins (2003) presents a simple "rule-of-thumb" one could use to estimate such gains based on a minimal amount of *ex ante* information; namely, variation in the marginal cost slope and the baseline emission rate. Despite the interest and attention—and the potential magnitude of savings—there has been little effort to complement these *ex ante* engineering studies comparing policy alternatives with econometric estimates of the *ex post* potential cost savings, once regulation in place, using data on historic compliance costs.[3]

This paper uses plant-level Census data on emission abatement and control costs to estimate the gain to emissions trading in four of the most polluting industries (pulp and paper, steel, chemicals, and refining) and for multiple pollutants. Using a panel of data collected between 1979 and 1985 where plants reported both abatement information and pollution control costs, we can estimate a variety of cost models allowing increasing degrees of plant heterogeneity. We can use these estimated models to compute current marginal control costs under existing regulation and to estimate the economic gains associated with equalizing marginal costs through trading.

Our preliminary results consider only sulfur dioxide controls at steel plants. We find that with plausible restrictions on abatement rates, emissions trading for SO_2 alone could reduce total air pollution control costs by 5-14%, or $300,000-$800,000 (1982$) per plant. Expressed as a share of SO_2 control costs alone, the cost reduction would be much higher. Our cost functions indicate a number of counterintuitive results: marginal costs are negative at low levels of

have similarly estimated gains of 50% to 90% for emissions trading among nations (Weyant and Hill 1999).

[3] The singular exception we are aware of is Gollup and Roberts (1985), who find a potential 47% savings from emission trading among electric utilities versus the status quo. Also examining utilities, Carlson et al. (2000) use an econometrically estimated model to conduct a retrospective analysis of the acid rain trading program, but consider a hypothetical performance standard rather than actual regulatory implementation.

abatement and the interaction between SO_2 and particulates is positive—higher levels of abatement of one pollutant make abatement of the other more expensive. These may be real features we need to explain, or may be a result of data problems that we hope to address.

The current draft addresses three key data issues not (or insufficiently) addressed in an earlier version of the paper. First, we have systematically removed implausible values (recoded them as missing) based on a robust measure of sample spread for each variable. Second, we removed observations where the pattern of abatement and costs was sufficiently different from the bulk of the other observations to exceptionally influence the results. Finally, we adopted a multiple imputation approach to address pervasive missing data problems.

In addition to refining the above techniques, we plan to delve more deeply into the observed negative marginal costs and to begin to explore predictors of cost heterogeneity. Despite the preliminary nature of the estimates and the problems that remain to be addressed, our expectation is that these estimates will prove to be a useful benchmark estimate for the potential gains to trade.

2. Background

A variety of criteria can be brought to bear upon the choice of policy instruments to achieve given environmental goals, including the likelihood of achieving the goal, cost-effectiveness, dynamic flexibility, distributional equity, and political and administrative feasibility, to name a few (Bohm and Russell 1985). Over the past three decades, as the stringency of environmental targets has continually increased and experience with market-mechanisms has grown, cost-effectiveness has become an increasingly important criterion. That is, does a particular policy achieve a particular environmental goal at the lowest possible cost?

A key determinant of cost-effectiveness is the heterogeneity of pollution control costs. If marginal control costs were equal and constant at all emission sources, a particular level of abatement would cost the same regardless of how it was implemented. However, there are many reasons why the costs of complying with environmental regulations tend to be heterogeneous, including plant location, size, age, production technology, management, and varying input prices. While it is widely recognized that abatement-cost heterogeneity is a fundamental determinant of the potential cost-savings associated with incentive-based instruments, there is surprisingly little analysis focused on exploring the extent of this heterogeneity from different underlying sources and its relationship to the potential costs or cost savings associated with various policy instruments.

Most studies of cost effectiveness have used engineering-economic analysis to model prospective control costs.[4] These studies use engineering cost data for "typical" plant types in conjunction with simplified functional forms to represent costs. The models typically compare a uniform performance standard for all sources to a least cost strategy. Such studies have produced a very wide range of estimates of savings—from 7% to 95% with a median of about 80%.[5] It is partially with this variation in mind that Newell and Stavins (2003) sought to develop rules of thumb to allow policymakers to crudely estimate the potential gains from using market-based policies rather than uniform performance standards.

[4] See Atkinson and Lewis (1974), Perl and Dunbar (1982), Seskin et al. (1983), O'Neil et al (1983), Atkinson (1983), Maloney and Yandle (1984), Krupnick (1986), Atkinson and Tietenberg (1982), and O'Ryan (1996), as well as others summarized on page 403 of Tietenberg (1992).

[5] The definition of the environmental goal in some analyses is achieving specified air quality standards, not a given level of abatement. Some studies also consider non-uniform trading, where trading among sources is not one-for-one.

Under their functional form assumptions, Newell and Stavins derive a very simple form for the ratio of the costs of more expensive performance standards to the costs of a market-based policy. Specifically, this ratio equals one plus the coefficient of variation in marginal cost slopes plus the coefficient of variation in baseline (uncontrolled) emission rates.[6] They apply this model to data on nitrogen oxide controls and find a 50% cost savings, comparable to a much more detailed plant-by-plant engineering study.

While the Newell and Stavins approach provides a convenient way to summarize engineering data and estimate the gains to emission trading, it is still fundamentally a prospective analysis most amenable engineering cost data. Even the best studies using this approach are unlikely to capture important elements of firm behavior of which the analyst is unaware. In contrast, Gollup and Roberts (1985) use observed data on utility pollution (SO_2) abatement and production costs to estimate a cost function that explicitly includes emission control rates as a predictor of production costs. They then compute that trading would reduce costs by about one-third. While similar to Gollup and Roberts in that it is based on an econometrically estimated cost function, the strategy of the Carlson et al. (2000) study of the gains from the SO_2 trading program is based on ex ante specification of alternate, ex ante, regulatory frameworks, as in the engineering-economic approaches described above.

Our approach is in the spirit of Gollup and Roberts, using detailed data on reported compliance costs and abatement to estimate cost functions, then computing the gains to trade from these empirical cost functions. However, our study is broader in terms of industries, pollutants, and heterogeneity. While our preliminary focus is sulfur dioxide controls in the steel

[6] See Equation (16) in their paper.

industry, we expect to extend the analysis to chemicals, refining, and pulp and paper, as well as to other air, water, and solid waste pollutants. The panel nature of our data also allows us to estimate plant specific effects using fixed- and random-effect models. Further, the richness of our dataset—the Longitudinal Research Database that is the core of this effort contains considerable data on plant characteristics—allows us to relate observed abatement cost heterogeneity to other plant characteristics. Such relations hold additional value; better information on the sources and extent of cost heterogeneity can help decision makers and analysts focus on cases where potential cost savings are particularly great.

One previous paper has made use of the same Census data on abatement and cost that we utilize to estimate control costs. Hartman et al. (1997) compute average and marginal abatement costs for between four and seven air pollutants over 37 industries. Their focus is on differences across industry, versus the current paper's focus on differences within industries. In particular, their modeling does not consider the interaction among abatement of different pollutants in determining control costs, nor does it take advantage of the panel nature of the data to estimate plant-specific effects (though they do estimate separate models for each industry). Their results show considerable cost heterogeneity *across* industries, viewed either in average or marginal cost terms, and they conclude both that current command-and-control regulation is inefficient and that their cost data can be used to rank (and prioritize) industries for regulation in other countries. While we believe that a more careful approach to modeling costs in each industry is the more valuable approach, their analysis in many ways complements the current effort. Any gains we find within individual industries or across a small set of industries is likely amplified by the heterogeneity they observe across a broad swath of industries.

Regardless of any cost heterogeneity and potential efficiency gains we identify, market-based instruments are by no means a panacea. In some cases, they hold tremendous promise of providing environmental protection cost effectively, but they are not well suited in other cases for a variety of reasons (see, for example, Hahn (1984), Misolek and Elder (1989), Malueg (1990), and Stavins (1995)), and there are, in any event, significant political barriers to their adoption (Keohane et al. 1999). Therefore, it is important for policy makers and policy analysts to understand the magnitude and sources of potential cost savings that may be associated with using an incentive-based instrument for a particular environmental problem. This type of information is also useful for designing more cost-effective standards-based regulations that take account of heterogeneous costs by tailoring regulation according to important sources of heterogeneity.

3. Modeling Abatement Costs

In order to estimate the gains from emissions trading, we must parameterize and estimate a cost function. We can then use that cost function to determine both costs and cost savings compared to a counterfactual with emissions trading. Our choice of functional form is governed in part by available data, such as the existence of data on abatement but not emissions.[7] We are also interested in a cost function that is intuitively appealing; for example, the cost function should describe in a reasonable manner how costs vary across large and small plants. Finally, we seek a degree of parsimony; ideally the gains from trading should be easily computable from a small number of estimable parameters. While we can simulate counterfactual behavior at the

[7] From the outset this excludes the approach in Newell and Stavins, which is based on variation in baseline emission

plant level for an arbitrary cost function, our use of confidential Census data prohibits disclosure of detailed results. For our results to remain transparent, they need to be based on aggregate moments of the data that can be disclosed. Also, such parsimony is helpful in applying this approach to other cases.

Based on this latter criterion of parsimony, we find ourselves with a similar goal as Newell and Stavins (2003) who sought rules-of-thumb for computing cost savings with minimal amounts of information. However, their focus was the relative costs of alternative *prospective* policies—performance standards, uniform reductions, and emission trading—where none are currently in place. Our focus is a counterfactual of trading relative to *existing* policy. More to the point, they required specification of a global abatement cost function that could be used to go from no regulation to the various alternatives. In our calculations, we need only a local cost function that can be used to go from the status quo to the case of emissions trading. For example, while they specify a quadratic cost function that passes through the origin—implying that marginal costs are zero with no abatement—we need not make that assumption. The quadratic model that we develop below need only apply over the region of abatement changes implied by the introduction of trading.

With these concerns in mind, our basic model of abatement costs for observation i (later distinguished by time and plant) is given by

$$C_i/y_i = \gamma_i + \alpha_i(q_i/y_i) + \beta(q_i/y_i)^2 \quad (1)$$

---

rates.

where C_i is total abatement costs, y_i is total output, and q_i is abatement. The parameters γ_i and α_i can vary by observation (plant fixed effects or interactions with other variables), while β is constant across observations.

Several model features are worth noting. First, the model only requires measurement of abatement, abatement cost, and output, consistent with the data available from Census surveys. Second, cost and abatement in the model are scaled by output. In this way, the model already accounts for simple economies of scale (more complicated economies of scale can be incorporated through the specification for γ_i and α_i). At a minimum, we expect the cost of abating one ton to be smaller at a large plant because it is spread out over many more units. This specification accounts for that by starting with a model of *unit* abatement costs (this same approach is taken by Newell and Stavins and others). Finally, it turns out that this specification leads to a simple expression for the gains associated with emission trading. We now turn to developing that expression.

We begin by noting that (assuming an interior solution) marginal costs are given by

$$\frac{\partial C_i}{\partial q_i} = \alpha_i + 2\beta\left(q_i / y_i\right) = p_i \qquad (2)$$

where we have introduced the notation p_i to refer to the estimated marginal cost of observation *i*. This is meant to reflect the *emission allowance price* at which level observation *i* would choose the abatement level q_i.

Now suppose emission trading took place. In a competitive equilibrium there would be a single market price, p^*, faced by all plants and they would consequently increase or decrease abatement until their marginal cost was equal to that price. In particular,

$$p^* = \alpha_i + 2\beta\left(q_i^* / y_i\right)$$

or

$$\Delta p_i = \left(p^* - p_i\right) = 2\beta\left(\left(q_i^* - q_i\right)/y_i\right) = 2\beta\left(\Delta q_i / y_i\right), \tag{3}$$

where q_i^* is the new, counterfactual abatement level at plant i where price equals p^* and the Δ before a variable indicates changes when trading is introduced.

How would this equilibrium price be determined? Increases in abatement at one plant must exactly be offset by decreases in abatement at another plant when trading occurs. In this way, we know that in equilibrium, the sum (and average) of the Δq_i's must equal zero.[8] Rearranging (3) we have

$$\Delta q_i = \frac{y_i}{2\beta}\Delta p_i . \tag{4}$$

Therefore,

$$\begin{aligned} \sum\nolimits_i \Delta q_i &= 0 \\ \sum\nolimits_i \frac{y_i}{2\beta}\Delta p_i &= 0 \\ \sum\nolimits_i \frac{y_i}{2\beta}\left(p^* - p_i\right) &= 0 \\ p^* &= \sum\nolimits_i \left(\frac{y_i}{\sum\nolimits_j y_j}\right) p_i \end{aligned} \tag{5}$$

In other words, the equilibrium price when emission trading is allowed will equal the output weighted average of pre-existing marginal costs. We explore the intuition for this expression further below.

---

[8] We eventually simulate trading across years as well as plants. This may seem odd. However, this simply serves to average behavior across multiple observations of the same plant, and provides a more representative result.

We can now compute the cost savings. At each plant, the net cost savings including any transfer associated with permit buying and selling will equal a Harberger triangle: the change in price, times the change in quantity, times one-half. Equation (4) allows further simplification:

$$\begin{aligned}\Delta C_i &= \frac{1}{2}\Delta p_i \Delta q_i \\ &= \frac{1}{2}\frac{y_i}{2\beta}\left(\Delta p_i\right)^2,\end{aligned}$$

where ΔC_i is the cost savings from trade at plant i.[9] This allows us to write the average gain across all observations as

$$\begin{aligned}N^{-1}\sum\nolimits_i \Delta C_i &= N^{-1}\sum\nolimits_i \frac{1}{2}\frac{y_i}{2\beta}\left(\Delta p_i\right)^2 \\ &= N^{-1}\frac{1}{4\beta}\sum\nolimits_i y_i\left(p^* - p_i\right)^2 \\ &= \frac{1}{4\left(\beta/\bar{y}\right)}\operatorname{var}_y p_i \qquad (6)\end{aligned}$$

where var_y refers to the output (y) weighted variance, $\bar{y}$ is the average plant size, and N is the number of observations. In other words, the average gain from emissions trading in this model equals the output-weighted squared variation in marginal costs, divided by four times the marginal cost slope for the average plant. Recall from (2) that the slope of the marginal cost schedule is β/y_i.

Consider the intuition for Equation (5) and (6). If plants were the same size, our marginal cost model would be linear in abatement, rather than linear in the abatement *rate*, and the equilibrium price would be the average price. However, with our model of marginal costs linear

---

[9] This also follows from change in costs associated with the change in abatement given by (1), plus the change in abatement times the market value p^*.

in the abatement rate, larger plants have more abatement opportunities—hence the output weighting. The same holds true for the cost savings expression $N^{-1}\Sigma\, \Delta C_i$. If plants were all of size one, the average cost savings would simply be the average squared price change, divided by twice the marginal cost slope. With different sized plants, we have to weight larger plants more to account for greater abatement opportunities in our model.

By estimating values for the elements of (6), we can derive an econometrically-based estimate of cost savings associated with emissions trading relative to the status quo. As we explore further below, however, while this formula is intuitive and crisp, it assumes an interior solution. This may be unrealistic in situations where that solution would imply abatement rates that are beyond physically plausible values. We therefore adopt a modified approach to estimating cost savings that takes into account this and other concerns.

4. Data and Model Estimation

Our data come from the Annual Survey of Manufactures and the Census of Manufactures, collectively referred to as the Longitudinal Research Database (LRD). This database contains information on the value of shipments by commodity code at the plant level. This data is combined with information on abatement and control costs from the Pollution Abatement and Control Expenditure Survey (PACE). The PACE Survey collected information on abatement quantity only in the years 1979, 1980, 1981, 1982, 1984, and 1985. These six years make up the time dimension of our panel. The unit dimension of our panel is comprised of those observations that report non-zero values of both air pollution control costs and air pollution abatement (for at least one air pollutant).

While the PACE survey includes both capital costs and operating and maintenance (O&M) expense (including depreciation on previous capital), we only use data on O&M. Note that PACE does not break down air pollution O&M costs by individual pollutant. Abatement is broken down by pollutant category, although the breakdown varies from year to year. We can consistently identify particulates, sulfur dioxide, nitrogen oxides+, and "other."[10] In these preliminary results, we focus on sulfur dioxide because its role in fine particulate formation and long-range transport makes it a regional or national pollution problem, so that large-scale trading would be appropriate. Particulates, in contrast, are typically a local pollution problem and broad geographic trading would fail to address local concerns. The other two pollution categories are combinations of multiple pollutants that are difficult to directly interpret. Our preliminary results also focus on the steel industry based on what we believe is a more straightforward pollution control problem (end-of-pipe treatment on a boiler), versus chemicals, refining, or pulp and paper, where pollution control is more complex.[11]

Table 1 provides descriptive statistics for the different samples used in this analysis; the samples are described further below (note that there is not much variation in first and second moments across the data sets). Output is the implied quantity index associated with a chain-weighted divisia price index computed over time for each plant, normalized to 1982 dollars.[12] We convert nominal abatement costs into 1982 constant dollars using the producer price index

---

[10] Nitrogen oxides+ (or NO+) refers to nitrogen oxides, carbon monoxide, and volatile organic compounds. The "other" category includes lead, toxic, and radioactive pollution.

[11] See Hartman et al. (1997) for a description of control technologies in various industries.

[12] The weights in the price index come from the detailed value-of-shipments data in the LRD; these are combined with national-level producer prices at the six-digit level from the Bureau of Labor Statistics. Note that we do not try to establish any cross-plant price variation, but rather adjust the output measure—which is really a value not a quantity—over time for price changes specific to the output bundle for each plant.

for all manufactured goods (note that year 2003 dollars are 39% higher than 1982 dollars based on inflation in the PPI). All other data elements are directly reported values. We see that the typical plant in the sample shipped roughly $350 million in products each year and abated about 1.5 tons of sulfur dioxide per thousand dollars of output. Total annual air pollution abatement costs were almost $6 million per plant, or about 1.4% of annual revenue.

4.1 Estimable Model

To estimate the model (1), we have to specify the further parameterization of γ_i and α_i as well as append a stochastic term. We actually go a bit further, however, and consider the possibility of a higher order, cubic term in abatement. If a cubic term is included, this implies that the formula given in (6) can only be approximated because the slope is not constant, a point we will return to in the next section when we calculate the gains to emission trading.

This leaves us with our most general estimable model for a single pollutant as:

$$C_{i,t}/y_{i,t} = \gamma_i + \boldsymbol{\theta}\cdot\mathbf{z}_{i,t} + (\alpha + \boldsymbol{\delta}\cdot\mathbf{x}_{i,t})(q_{i,t}/y_{i,t}) + \beta(q_{i,t}/y_{i,t})^2 + \tau(q_{i,t}/y_{i,t})^3 + \varepsilon_{i,t} \tag{7}$$

where now *i* indexes over plants and *t* indexes over time, τ is the additional cubic term, and ε is a stochastic disturbance. The specification includes a fixed-effect for each plant γ_i, additional controls given by $\boldsymbol{\theta}\cdot\mathbf{z}_{i,t}$, and observation-specific marginal cost effects given by $\boldsymbol{\delta}\cdot\mathbf{x}_{i,t}$. The **z**'s include linear, quadratic, and cubic terms in other pollutants, interactions among other pollutants, and output and output squared. The **x**'s include other pollutants.

Written generally for all four pollutants, indexed by *j*, Equation (7) becomes:

$$\begin{aligned} C_{i,t}/y_{i,t} = \gamma_i + \theta_1 y_{i,t} + \theta_2 y_{i,t}^2 + \textstyle\sum_j \alpha_j (q_{j,i,t}/y_{i,t}) + \sum_j \sum_{j'>j} \delta_{j,j'} (q_{j,i,t}/y_{i,t})(q_{j',i,t}/y_{i,t}) \\ + \textstyle\sum_j \beta_j (q_{j,i,t}/y_{i,t})^2 + \sum_j \tau_j (q_{j,i,t}/y_{i,t})^3 + \varepsilon_{i,t} \end{aligned}$$

We assume the disturbance ε is iid and normal and estimate the model using least squares.[13] Note therefore that under this specification the intercept of the marginal abatement cost function varies across observations while its slope is constant.

4.2 Implausible, Missing, and Influential Data

Before estimating the model, we first need to address rather pervasive data quality problems—problems that have discouraged many researchers from working with this data. Some observations reveal implausible features in comparison to the rest of the sample. Others exert a disproportionate influence on the estimated parameters. And finally, there is a complex pattern of missing data, which varies by survey year. Before even addressing these problems, we first eliminated observations that had zero or missing data for all four pollutants and/or output—this shrunk our dataset from an initial size of around 850 to 775.[14] We then adopted a three stage approach to refine the data before estimating our models in order to arrive at relatively stable parameter and welfare gain estimates that are not sensitive to individual observations and are robust to missing data concerns. We note that our approach is based on our best judgement and that the literature provides little systematic guidance on procedures for "cleaning" data prior to estimation. Rather than "sweep this under the rug" as happens with many analyses, we make our approach as explicit as possible given the unusual degree of data problems in our sample.

First, we recoded variable observations where the values were strongly suspect based on a robust measure of variation for each variable. Specifically, we recoded values to missing if

[13] A more reasonable assumption, especially given our data problems, is that the errors are mean zero but not iid, and therefore use robust standard error estimates. This is something we intend to do.

[14] Because we scale all our variables by output, a zero for output makes it impossible to create the variables we use in the rest of our approach. Zeros or missing values for all four pollutants suggests little information is contained in the observation; based on our imputation approach (below) we would at best predict mean values conditional on

they were greater than the median plus 20 times the interquartile range (75th-25th percentile, equivalent to about 1.3 standard deviations in a normal distribution). This affected only a few percent of the observations for each variable. Because we initially scaled our data by output, this kind of variation appeared implausible to us and more likely reflected miscoding. Given a key element of our approach involves imputing missing data, there is also no loss of sample size or information from the rest of an observation with one implausible value.

Second, we sequentially dropped observations that exerted undue influence on the cubic abatement terms. By treating missing values as zeros and estimating the full cubic model with fixed effects as well as indicator variables for zero/missing values, using the DFBETA statistic[15] we identified the observations whose deletion caused the largest normalized (i.e., change scaled by standard error) change in any of the four cubic terms.[16] We did this sequentially until the largest effect was no more than ½ a standard error. This resulted in the deletion of about 25 observations.

We focused on the cubic coefficients because, among the linear, quadratic, and cubic parameters defining the marginal cost schedule, they should be the most sensitive to deletion of a single observation. In addition, because our primary interest revolves around simulating

---

output.

[15] The DFBETA test statistic, according to Belsley et al. (1980), is one of the three regression diagnostics highlighted by Chatterjee and Hadi (1986). It provides the most succinct measure of influence when the point of concern is a particular regression coefficient.

[16] An interesting question is how to simultaneously deal with missing data and influence. We did not want to model the missing data without addressing observations that seemed to be distorting the model; yet, we were concerned that the identification of influential observations might be disturbed by the missing data. In the end, we decided that influence could be detected with missing values coded as zero, and attempted to control somewhat for this problem by using indicator variables for zero/missing values. Note that we also included observations with the dependent variable coded as zero at this stage.

movement along the marginal cost schedule, it made sense to focus on the sensitivity of the marginal cost schedule (versus precision or prediction).

The last step of our data preparation involves multiply imputing missing data.[17] This was complicated by the fact that in the last two years of the sample, 1984 and 1985, the data appear to have already been "cleaned" in that there is no missing data.[18] In contrast, over 1979-82, we find that there is a mixture of missing values (61%) and reported zero values (6%) scattered across observations. Therefore, the first step in our imputation is to estimate a model of missingness across those zero and missing observations in 1979-82, and then randomly impute whether zeros in 1984-85 are truly zero or are missing values miscoded as zeros.

We then sequentially consider each of the four abatement explanatory variables (abatement of particulates, sulfur dioxide, nitrogen dioxides+, and other) and the dependent variable, air pollution abatement costs, all scaled by output. For each variable, we impute using all of the right- and left-hand side variables, including quadratic and interaction terms, as well as dummy variables that equal one for missing and zero values, and deleting any terms that include the variable being imputed.[19] After imputing one variable, we do not use those imputed values to impute later variables and continue to impute with missing values as zero. An interesting

[17] See Rubin (1987). The basic idea of imputation is to replace missing values with values that are consistent with observed relations among the data where values are not missing. Partially observed observations can thereby contribute information without biasing the results. Simple use of predicted values causes a problem with estimated standard errors; however, because the imputed values perfectly follow the estimated relationship and make the estimated values appear more precise than they are. The solution is to multiply impute with a stochastic disturbance, so that the imputed data has variation similar to observed data. Variation in parameter estimates across imputations can be added to the average within-imputation variation to produce standard errors that reflect both the variation in the observed data and the additional variation introduced by imputation.

[18] Except for a number of observations with all variables coded as missing.

[19] The use of dummy variables for zero and missing values is a crude approach to control for the fact that missing and zero values may not behave in the way as observed, non-zero values. Generally, these coefficients are significant in the imputation equations.

question is whether the approach could be improved by repeating the process until some kind of convergence was achieved. In any case, we repeat the entire imputation process (including prediction of missing/zero in 1984-85) five times to create a set of multiple imputations. When we report values for the imputed data, these represent averages over all five sets of results.

4.3 Estimation Results

Table 2 shows results for the full cubic model as well as versions with either the interactions among pollutants or the cubic terms excluded. We report results for the multiply imputed data, as well as the data with missing values simply coded as zero.[20] (We have also included results reported in an earlier version of the paper for reference.) Each of the models has its advantages.

The full cubic model is the most flexible of the models and encompasses the other two as special cases. This suggests the cubic model is to be preferred. The earlier theoretical model for computing welfare gains, based on a constant marginal cost slope, is best represented by the quadratic model. Finally, if we want to visualize the marginal cost schedule from the estimated values, the task is complicated by interactions that require information about covariates and shift the marginal cost schedule among observations. This favors the no interaction model.

Looking at the results, we first note that there is little apparent difference between the imputed and the non-imputed results. The statistically significant changes are related to coefficients on nitrogen oxides+ and other pollutants. There are differences that show up when we calculate welfare gains, however. The slightly more positive coefficients on the sulfur

[20] For this data set, observations with missing values for total abatement costs are dropped; that is why the reported sample size is smaller in Table 1.

dioxide terms for the quadratic and full cubic models imply an upward sloping region of the marginal cost schedule for the unimputed data over the observed range of values, while the imputed data reveal a consistently downward sloping marginal cost schedule over the same region. Thus, we cannot compute gains for the quadratic and full cubic model using the imputed data.

This leads us to important distinctions across models. Namely, the coefficients that concern us both make more sense and are more significant in the cubic model without interactions. As noted above, the other two models imply downward sloping marginal cost schedules for sulfur dioxide. Further, among the eight relevant coefficients for the marginal cost slopes (quadratic and cubic terms), five are significant in the no-interaction model—more than the other two models. Noting that the quadratic particulates term is almost significant at the 5% level, and that the mixture of pollutants in the "other" abatement category makes that schedule suspect, the results are even stronger—the slopes of all of the marginal cost schedules are sensible for this model. These schedules are shown over the relevant range of abatement in Figure 1.[21]

The one feature that we should note before we move past the pollutant interactions is the significant and positive interaction between particulates and sulfur dioxide. This result has appeared in all of our models and goes against our intuition that there should be co-benefits associated with reductions in one pollutant in terms of inadvertent reductions in other pollutants. The two explanations we have are that either there really is a penalty—reductions in one raise the cost of reducing the other, perhaps because the control devices do not work well together. Or,

[21] Note that the range for particulates in this figure is 0.05 tons per $100 while the range for sulfur dioxide and

respondents do not include any co-benefits in their reported abatement. For example, even if an SO_2 scrubber removes particulates, only the particulate abatement from the particulate-specific control device (say, a baghouse) is reported. Both of these effects could explain a positive interaction term.

Viewing the marginal cost schedules, the most disturbing feature is that marginal costs are negative over a farely wide range of values for both sulfure dioxide and nitrogent oxides+ (compare the figure's range to the values in Table 1). Is it possible that control costs actually decline as abatement increases at low levels? Or, is this reflecting fixed costs, an artifact of functional form restrictions, unobserved heterogeneity, influential observations, or simply bad underlying data? This is something we will need to investigate further. Interestingly, Carlson et al. (2000) (see Table I, p. 1307) also find a sizable range of negative marginal costs in their estimation of the marginal abatement costs of SO_2 in electric power plants.

5. The Gains from Emission Trading

We can use the models estimated in the previous section to compute the gains from emission trading among plants. In particular, we use the estimated parameter values to compute marginal costs for each observation using (2). With these observation-specific marginal cost predictions in hand, we can make a simple calculation of the output-weighted marginal cost variance as suggested by Equation (6) in order to compute the cost savings.

---

nitrogen oxides+ is 0.05 tons per $1000.

5.1 Beyond the Simple Formula for Cost Savings

For quadratic models, we can simply insert the estimated quadratic coefficient into Equation (6) to calculate the gain. In the case of our cubic model, we will need to approximate the slope over the relevant range, either explicitly averaging or reading an approximation off the graph in Figure 1. A second concern that arises in this simple calculation is that it does not impose any constraints on the control rates implied by emissions trading—plants move up and down their marginal cost schedules until all have the same marginal cost and the aggregate change in abatement is zero. When plants are operating on the *same* marginal cost schedule, as shown in left panel of Figure 2, movement toward the average (or output-weighted average) price implies that control rates are moving toward a similar central value. However, when plants operate on *different* marginal cost schedules, as shown in the right panel of Figure 2, movement toward an average price can imply control rates moving away from central values. In this case, control rates tending to plus or minus infinity could arise.

One way to address both the cubic cost function and this concern about plausible control rates is to simulate the behavior of each plant using the actual cubic cost function and imposing the constraint that control rates lie within the empirically observed range. We can also handle downward sloping ranges of the marginal cost schedules by assuming flat schedules before the minimum (for particulates) and after the maximum (for SO_2).[22] With shallower marginal cost slopes and/or movement near the extrema, many plants will tend to the imposed limits, zero and some maximal rate, when trading occurs (e.g., in the right panel of Figure 2, only one plant appears to have an interior solution). For the extreme of perfectly flat but heterogeneous

[22] This does raise the question, however, of exactly where the data lie—do the downward sloping portions of the schedules reflect actual behavior or are they extrapolations of functional form?

marginal costs, this more careful approach amounts to lining up plants in order of marginal costs and finding the price level where plants below the threshold, controlling at the maximal rate, abate the same amount as observed in the sample.

A final concern (and one that remains unaddressed) is whether facilities that do not abate in the sample are substantially different than those that do abate. In the extreme, we might imagine that those who do not abate actually have no emissions to abate—making their control costs infinite. Unfortunately, we have been unable to find any information to help us confront this question. We explored matching abatement to fuel use data which is available for four of the six years with the thought that abatement and emissions could be related to fuel.[23] Perhaps surprisingly, the fraction of plants reporting any sulfur dioxide abatement was not much higher among plants reporting the use of coal and/or oil (the two fuels that contain sulfur)—in part because sulfur is also contained in the raw iron used to make steel at some facilities.[24]

5.2 Estimates of Cost Savings

The results using both the formulaic and simulation approach are presented in Table 3 for sulfur dioxide. We report results for all data sets and models where marginal costs are (partially) upward sloping. The formulaic approach uses Equation (6), coupled with the reported variance, mean plant size from Table 1, and marginal cost slope. This slope parameter is equal to the quadratic parameter from Table 2 plus 3 times the cubic parameter (if non-zero) times 0.002, the

[23] We thank Kerry Smith for this suggestion.

[24] It was true that coal users had significantly higher abatement when they did abate. This suggests that coal could be a useful control variable and/or used to scale abatement. Scaling abatement would not work, however, because of the many cases where abatement is non-zero even though coal is not used. The use of coal as a control variable is also problematic because it is not observed in 1982 and 1984—roughly 1/3 of our sample.

mean of the observed sulfur abatement rates. The simulation results are presented both in dollar terms and as a percentage of total air pollution control costs.

First we would note that the formulaic estimates for the quadratic models tend to be dramatically over estimated compared to the simulation results. This suggests that the heterogeniety in marginal costs would otherwise lead to extreme and implausible control rates outside the observed range. Second, we would note that the formulaic approach overestimates gains for convex marginal costs and underestimates gains for concave marginal costs. This follows from the basic geometry of the area under linear (quadratic) versus convex or concave marginal cost schedules.

In any case, we focus our attention on the simulation results. These estimates are remarkably consistent—about 5-14% or $300,000-$800,000 annually per plant (1982$). Our preferred model, the cubic model without interaction based on the imputed data, is at the low end of that range. However, using capital expenditure data to estimate the share of air pollution control costs attributable to sulfur,[25] the cost savings for sulfur dioxide alone would appear to be 50-100+%.

In Table 3 we also report the output-weighted mean of predicted marginal cost for the sample (as well as the simulated market clearing price). This reveals another intriguing feature of the data: not only are there negative marginal costs, but the market clearing price is almost always negative based on our estimates. This suggests that while marginal costs are positive over a wide range of the non-zero values, the negative values dominate the market equilibrium.

[25] Sulfur dioxide capital expenditures were $43.6M, $67.3M, $0.8M, and $1.3M in 1979, 1980, 1982, and 1983, respectively (other years were not reported due to disclosure concerns). Total air pollution capital expenditures were $403.9M, $382.1M, $175.8M, and $60.7M in those same years, leaving sulfur dioxide expenditures at 11% of the total. Note that there are many reasons to believe that the share of SO_2 O&M costs in total air pollution costs are

6. Next Steps

The work so far has focused on developing a parsimonious, estimable model of abatement costs and using the model to estimate the gains from emissions trading. We have also tackled many data problems, including implausible or missing values, and unduly influential observations. While our original expectation was to use a Bayesian modeling approach to tackle all these problems in a unified framework, using a data augmentation approach to address missing data and a hierarchical model to deal with heterogeneity and measurement error, we believe that our current approach is sufficient.[26]

Aside from refining these techniques and extending the results both to other pollutants and other sectors, our main concern at this point are the negative marginal costs we consistently find at low abatement rates. This was not solved by more flexible (i.e., cubic) functional forms. One hypothesis is that fixed costs exist, suggesting the inclusion of indicator variables for whether any abatement is occuring for a individual pollutants, perhaps interacted with scale. Other than that, we can continue to explore more flexible forms, including higher order terms or spline functions.

We also considered more flexible models of heterogeneity, including esimation of a random-coefficient model. In a random coefficient model, observations for each plant are assumed to have separate coefficients based on a specified random distribution across all plants (similar to a Bayesian hierarchical model). From an estimation perspective this is an efficiency issue, as the random elements can be viewed as part of the error. However, we can also derive

---

proportionally higher than their share in capital costs, e.g., higher energy and disposal costs.

[26] Chapter 5 of Gelman et al. (1995) describes hierarchical model estimation; see Tanner and Wong (1990) for a description of data augmentation. Our decision not to follow that route was influenced primarily by the logistic

the plant-specific coefficients and use them to estimate the gains to trade.[27] At this point, we do not believe further heterogeniety along these lines is useful.

The last part of the original research proposal was to consider whether heterogeniety in marginal costs could be predicted by other, more easily observed, covariates. The motivation is that if there are reasonably good predictors of cost heterogeniety, they might provide regulators with guidance about when market-based policies would offer siginificant gains. We still intend to explore this issue once we feel the cost function modeling is working well.

7. Conclusions

The potential cost savings from emissions trading, relative to uniform standards-based approaches, has been a primary impetus to environmental policy reform over the past two decades. However, most analyses outside the electricity sector have relied on *ex ante* engineering-economic studies. Such studies are inherently limited in their ability to capture plant-level heterogeneity as well as the practical realities of emissions control. This paper presents the first empirical estimates of the potential gains to emissions trading based on *ex post* data on compliance costs.

The preliminary estimates looking at sulfur dioxide control costs in the steel industry suggest savings of 5-14% of total air pollution control costs and a much higher share of sulfur dioxide control costs. The estimates are also higher if control rates are not restricted to lie within the range observed in the sample.

---

difficulties of bringing in, and programming with, new software at the secure Census research facilities.

[27] For example, see p. 478 of Greene (1990),

This preliminary work has focused on a number of significant data problems that need to be addressed. These include the treatment of implausible reported abatement rates and control costs, combinations of abatement rates and control costs that lead a few observations to strongly influence the results, reported zero-values of abatement that may in fact be missing, and more general missing values that are may or may not be zero. This project has yet to examine whether cost heterogeneity can be traced to more easily observed facility features, which could be an important result for policy design. Finally, we have yet to fully address the meaning and sensibility of negative marginal costs for low values of abatement.

Nonetheless, we find these preliminary results encouraging and believe this work could eventually provide a useful benchmark in understanding the pattern of existing pollution control costs and the potential gains to emission trading.

Tables

Table 1: Descriptive Statistics

Variable	*Old Data*	*Missing as Zero*	*Imputed*
Output ($1000)	349,000	444,000	429,000
	(450,000)	(586,000)	(573,000)
Abatement (tons/$1000)			
Sulfur dioxide	0.0015	0.0016	0.0021
	(0.0050)	(0.0062)	(0.0068)
Particulates	0.056	0.044	0.041
	(0.072)	(0.062)	(0.061)
Nitrogen oxides+	0.0033	0.0026	0.0029
	(0.017)	(0.0126)	(0.0124)
Others	0.00175	0.0008	0.0019
	(0.0131)	(0.0038)	(0.0043)
Abatement costs ($)	5,890,000	5,502,000	5,127,000
	(11,100,000)	(10,600,000)	(10,300,000)
Abatement costs/output	0.0141	0.0107	0.0100
	(0.0179)	(0.0134)	(0.0133)
Number of obs: ~600	~600	~700	~750
Number of plants: ~150	~150	~170	~190

Note: Values have been rounded and sample size has been blurred to prevent complementary disclosure if sample changes in the future. Nitrogen oxides+ includes carbon monoxide and volatile organic compounds. Others includes lead, hazardous, and radioactive materials.

Table 2: Regression Results

	old data	*missing as zero data*			*imputed data*		
Sulfur dioxide	–785**	–580**	–580**	–430	–370**	**–520****	–290
	(308)	(170)	(270)	(270)	(160)	**(250)**	(260)
(Sulfur dioxide)^2	47.8	11	38**	4.8	–8.2**	**36****	–16
(x 1000)	(28.1)	(7.6)	(14)	(15)	(2.8)	**(12)**	(15)
(Sulfur dioxide)^3			–450**	77		**–460****	73
(x 1000)			(160)	(180)		**(140)**	(170)
Particulates	14.1	6	74**	98**	11	**74****	110**
	(19.4)	(21)	(35)	(34)	(21)	**(33)**	(34)
(Particulates)^2	111**	100	–360	–600**	79	**–400**	–720**
	(43.4)	(65)	(230)	(230)	(64)	**(230)**	(230)
(Particulates)^3			770	1200**		**860****	1400**
			(400)	(400)		**(390)**	(390)
Nitrogen oxides+	40.6	–11	–1000**	–640**	–71	**–620****	–550**
	(188)	(180)	(310)	(320)	(180)	**(230)**	(260)
(Nitrogen oxides+)^2	–0.7	0.3	25**	16**	0.690	**17****	15**
(x 1000)	(1.0)	(1.4)	(7.5)	(7.7)	(1.5)	**(6)**	(6.3)
(Nitrogen oxides+)^3			–140**	–98**		**–100****	–91**
(x 1000)			(45)	(46)		**(38)**	(39)
Other	36.5	12	–860	–110	15	**–230**	–110
	(80.5)	(340)	(510)	(520)	(200)	**(250)**	(260)
(Other)^2	0.50	2.4	4.6	140	–0.43	**10**	16
(x 1000)	(1.5)	(9.5)	(37)	(38)	(11)	**(24)**	(25)
(Other)^3			–600	–230		**–170**	–500
(x 1000)			(670)	(700)		**(490)**	(540)
Particulates x SO_2	4500**	2800**		3400**	3700**		4600**
	(1900)	(720)		(810)	(740)		(890)
Particulates x NO+	1590	110		190	180		–280
	(977)	(1400)		(1400)	(1400)		(1400)
Particulates x other	–490	–1800		–1900	–5500**		–5000
	(1470)	(2800)		(2900)	(2600)		(2700)
SO_2 x NO+	–38.8	1.7		–3.1	–1.6		–1.5
(x 1000)	(25.7)	(8.4)		(8.5)	(6.4)		(6.4)
SO_2 x other	–20.4	–160**		–160**	22**		29**
(x 1000)	(19.3)	(71)		(76)	(10)		(12)
NO+ x other	–9.8	–14		–2.5	1.4		1.9
(x 1000)	(13.8)	(25)		(26)	(7.5)		(7.6)
Output (x 10^{-6})	–23**	–10**	–20**	–20**	–10**	**–10****	–10**
	(5.2)	(4)	(4)	(4)	(3)	**(3)**	(3)
(Output)^2 (x 10^{-12})	5.8**	3**	3**	3**	2**	**3****	2**
	(2.2)	(1)	(1)	(1)	(1)	**(1)**	(1)

Note: Standard errors shown in parentheses; double asterisks indicate significance at the 5% level; boldface indicates preferred results.

Table 3: Estimated Gains to Sulfur Dioxide Emissions Trading

	Quadratic Approximation Using Equation (6)				*Simulation*		
Model	*Avg. MC ($/ton)*	*Variance (x1000)*	*MC Slope (x1000)**	*Savings ($1000/obs)*[†]	*Eq. Price ($/ton)*	*Savings ($1000/obs)*	*% of total costs*
Quadratic (old data)	–575	1700	47.8	3,030	–2300	1,450	25
Quadratic (missing as zero)	–450	320	11	3,200	–570	536	9.7
Cubic, no interactions (missing as zero)	–500	40	33	134	–440	296	5.4
Cubic (missing as zero)	–300	330	5.2	7,000	–510	776	14
Cubic, no interactions (imputed)	–440	42	33	140	**–370**	**302**	**5.9**

Note: Marginal cost slopes* for cubic models are based on the slope at 0.002 tons/$1000 (the mean of the observed values). Other results are simulated based on actual cost function and constraining abatement rates to observed data range. Boldface indicates preferred results.

Figures

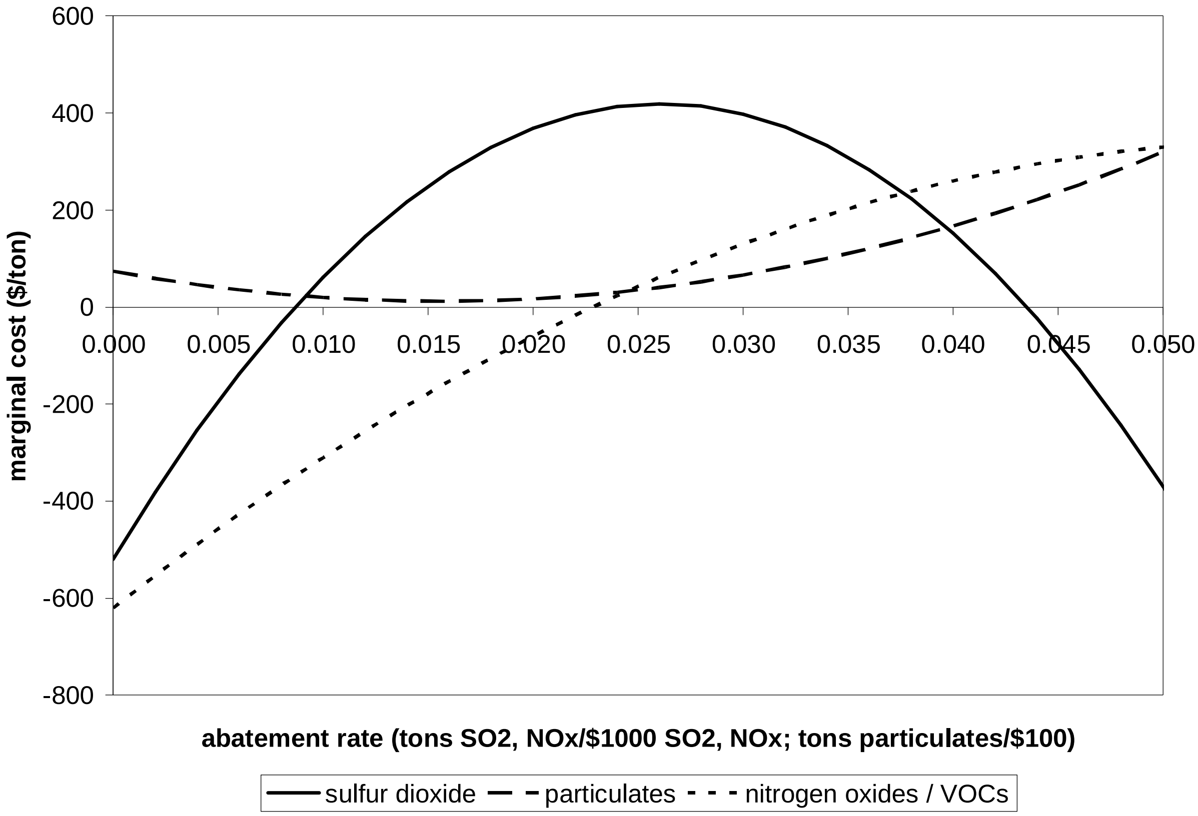

Figure 1: Marginal Cost Schedule for Preferred Results

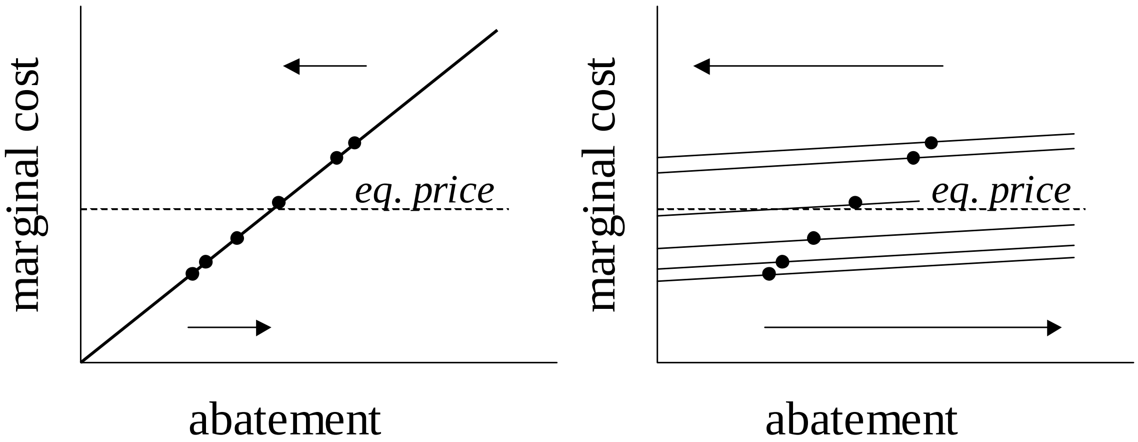

Figure 2: Movement to Emission Market Price Under Trading

References

Atkinson, Scott E. 1983. Marketable pollution permits and acid rain externalities. *Canadian Journal of Economics* XVI (4):704-722.

Atkinson, Scott E., and Donald H. Lewis. 1974. A Cost-Effectiveness Analysis of Alternative Air Quality Control Strategies. *Journal of Environmental Economics and Management* 1:237-250.

Atkinson, Scott E., and Tom Tietenberg. 1982. The Empirical Properties of Two Classes of Designs for Transferable Discharge Permit Markets. *Journal of Environmental Economics and Management* 9:101-121.

Belsley, D.A., E. Kuh, and R.E. Welsch. 1980. *Regression Diagnostics: Identifying Influential Data and Sources of Collinearity*. New York: Wiley.

Bohm, Peter, and Clifford S. Russell. 1985. Comparative Analysis of Alternative Policy Instruments. In *Handbook of Natural Resource and Energy Economics*, edited by A. V. Kneese and J. L. Sweeney. Amsterdam: North-Holland.

Carlson, Curtis, Dallas Burtraw, Maureen Cropper, and Karen L. Palmer. 2000. Sulfure Dioxide Control by Electric Utilities: What Are the Gains from Trade? *Journal of Political Economy* 108 (61):1292-1326.

Chatterjee, Samprit, and Ali S. Hadi. 1986. Influential Observations, High Leverage Points, and Outliers in Linear Regression. *Statistitical Science* 1 (3):379-393.

Gelman, A., J. B. Carlin, H. S. Stern, and D. B. Rubin. 1995. *Bayesian Data Analysis*. New York: Chapman and Hall.

Gollup, Frank M., and Mark J. Roberts. 1985. Cost-minimizing Regulation of Sulfur Emissions: Regional Gains in Electric Power. *Review of Economics and Statistics* 67 (1):81-90.

Greene, William H. 1990. *Econometric Analysis*. New York: Macmillan.

Hahn, Robert W. 1984. Market Power and Transferable Property Rights. *Quarterly Journal of Economics* 99:753-765.

Hartman, Raymond S., David Wheeler, and Manjula Singh. 1997. The cost of air pollution abatement. *Applied Economics* 29:759-774.

Keohane, Nathaniel O., Richard L. Revesz, and Robert N. Stavins. 1999. The Positive Political Economy of Instrument Choice in Environmental Policy. In *Environmental and Public Economics, Essays in Honor of Wallace E. Oates*, edited by A. Panagariya, P. Portney and R. Schwab. London: Edward Elgar.

Krupnick, Alan J. 1986. Costs of Alternative Policies for the Control of Nitrogen Dioxide in Baltimore. *Journal of Environmental Economics and Management* 13:189-197.

Maloney, Michael T., and Bruce Yandle. 1984. Estimation of the Cost of Air Pollution Control Regulation. *Journal of Environmental Economics and Management* 11:244-263.

Malueg, David A. 1990. Emission Credit Trading and the Incentive to Adopt New Pollution Abatement Technology. *Journal of Environmental Economics and Management* 16:52-57.

Misolek, W. S., and H. W. Elder. 1989. Exclusionary Manipulation of Markets for Pollution Rights. *Journal of Environmental Economics and Management* 16:156-166.

Newell, Richard G., and Robert N. Stavins. 2003. Cost Heterogeneity and Potential Savings from Market-Based Policies. *Journal of Regulatory Economics* 23 (1):43-59.

O'Neil, William, Martin David, Christina Moore, and Erhard Joeres. 1983. Transferable Discharge Permits and Economic Efficiency: The Fox River. *Journal of Environmental Economics and Management* 10:346-355.

O'Ryan, Raúl E. 1996. Cost-Effective Policies to Improve Urban Air Quality in Santiago, Chile. *Journal of Environmental Economics and Management* 31:302-313.

Perl, Lewis J., and Frederick C. Dunbar. 1982. Cost Effectiveness and Cost-Benefit Analysis of Air Quality Regulations. *American Economic Review, Papers and Proceedings* 72 (2):208-213.

Rubin, Donald B. 1987. *Multiple Imputation for Nonresponse in Surveys*. New York: Wiley.

Seskin, Eugene P., Jr. Robert J. Anderson, and Robert O. Reid. 1983. An Empirical Analysis of Economic Strategies for Controlling Air Pollution. *Journal of Environmental Economics and Management* 10:112-124.

Stavins, Robert N. 1995. Transaction Costs and Tradeable Permits. *Journal of Environmental Economics and Management* 29:133-147.

———. 2003. Experience with Market-Based Environmental Policy Instruments. In *Handbook of Environmental Economics*, edited by K.-G. Mäler and J. Vincent. Amsterdam: Elsevier Science.

Tanner, M., and W. Wong. 1990. The Calculation of Posterior Distributions by Data Augmentation. *Journal of the American Statistical Association* 82:528-550.

Tietenberg, Tom. 1992. *Environmental and Natural Resource Economics*. 3rd ed. New York: HarperCollins.

Weyant, John P., and Jennifer Hill. 1999. Introduction and Overview: The Costs of the Kyoto Protocol. *Energy Journal* (special issue):vii-xliv.