

THE UNITED STATES ACID RAIN PROGRAM:
ARE TRADABLE EMISSION PERMITS WORKING EFFICIENTLY?

by

TODD D ZELLER

B.S., Kansas State University, 2007

A REPORT

submitted in partial fulfillment of the requirements for the degree

MASTER OF ARTS

Department of Economics
College of Arts and Sciences

KANSAS STATE UNIVERSITY
Manhattan, Kansas

2009

Approved by:

Major Professor
Dr. Tracy M. Turner

Abstract

The report examines the extent to which the United States market for tradable pollution permits has been an efficient way to reduce sulfur dioxide emissions from coal-burning power plants. To do so, this report first provides background information on the effects of SO₂ emissions on the environment in the United States. It discusses the initial attempt with the 1970 Clean Air Act to reduce these emissions and its degree of success. The details of the 1990 Amendments are then given. The economic theory behind the different methods of pollution control (quantity regulation, technology mandate, taxation, and emissions restriction through tradable permits) is explained and their efficiency regarding consumer and producer surplus is contrasted. The report then reviews published articles regarding the topic at hand. The U.S. Acid Rain Program has been found to be very efficient in its ability to reduce sulfur dioxide emissions at a low cost to the producers.

Table of Contents

List of Figures	iv
Dedication	v
1. Introduction.....	1
2. Economic Theory of Pollution-Reduction Methods	8
<i>Quantity Regulation: Quantity Restriction vs. Emission Restriction</i>	8
<i>Technology Mandate</i>	11
<i>Taxation: Commodity Tax vs. Emissions Tax</i>	11
<i>Emissions Restrictions with Tradable Permits</i>	13
3. Economic Literature on the Acid Rain Program.....	16
<i>Ellerman and Montero(2007)</i>	16
<i>Schmalensee, Joskow, Ellerman, Montero and Bailey (1998)</i>	21
<i>Ludwig (2004)</i>	23
<i>Tietenberg (2003)</i>	26
4. Conclusion	29
References.....	31
Figures.....	32

List of Figures

Figure 1.1: SO ₂ Allowance Trading Volume and Prices, March 1994-October 2008.....	32
Figure 2.1: Goods Market.....	33
Figure 2.2: Production Restriction.....	34
Figure 2.3: Pollution-Reduction Market with Multiple Firms.....	35
Figure 2.4: Commodity Tax.....	36

Dedication

I dedicate this report to my father, Alan, who has always been there for me. He has continually taken an active interest in my life whether that is attending my extracurricular events, helping me move to Manhattan, advising me with my first car purchase, or aiding me in any other way. He has raised me into the man I am today. Without his love and support, I don't know where I would be today. Thanks Dad!

1. Introduction

Protecting the environment is much more of a concern to society than it has been in the past. This is largely due to the discovery of the harmful effects of pollutants. More and more people are “going green” by recycling, driving more fuel efficient cars, using compact fluorescent light bulbs, for example. Governments around the world have attempted to reduce pollution through a range of means in various industries. This report focuses on the United States’ experience with the sulfur dioxide (SO₂) emissions reduction as part of the U.S. Acid Rain Program.

Acid rain is the result of SO₂ and nitrogen oxides entering the atmosphere; the acids can then travel hundreds of miles from their origins.¹ SO₂ is the primary contributor to acid rain in North America, and two-thirds of those emissions come from coal-fired power plants. Acid rain causes damage to the environment and society in several ways. The first is that it increases the acidity of ponds, lakes, streams, and other bodies of water. This can decrease fish populations, wipe out entire aquatic species, and reduce biodiversity. By the late 1980s, over 650 lakes in the United States could support only acid-tolerant largemouth bass. In one New Hampshire pond, the pH level in 1980 was over 30 times the level it was in 1948 (Gruber, p. 144). Another harmful effect of acid rain is forest erosion; it can injure or even kill many different types of trees, and it degrades the soil, lowering agricultural productivity. Acid rain also damages property. It can ruin paint, corrode metals, and deteriorate stone. In 1985, the U.S. government estimated that acid-rain damage costs about \$5 billion annually (Gruber, p. 146). The

¹ This section draws on Gruber (2005), pp. 144-150.

emitted particulates also reduce visibility. In fact, “sulfate particles account for 50 to 70% of the visibility reduction in the eastern part of the United States” (Gruber, p. 146).

Finally, these same particulates can be inhaled into the lungs and may cause illness or premature death due to heart or lung complications.

Because of these harmful effects, many people may think that the optimal level of pollution is zero. However, the elimination of all pollution is neither possible nor necessarily desirable. The present day economy requires that some pollution exists. For example, air pollution could be reduced by eliminating all automobile usage, but this would lead to longer travel times, higher transportation costs, higher product prices, and even different types of pollution. A change such as this would fundamentally alter our very way of life. Paper mills could also be shut down in order to reduce water pollution, but the lack of paper would also drastically change the way we live our lives. Many other examples like these exist, and thus it can be concluded that, unless we want to make these changes, the optimal level of pollution is not zero. Indeed, as long as the pollution-producing product has value to society, a zero-production level would be economically inefficient.

If not zero, then what is the optimal level of pollution? Using the marginal principle, it can be determined that the pollution level should be at the point where the marginal benefit of abatement equals the marginal cost. The costs of abatement can be determined. Resources such as labor, capital, and land are needed to reduce pollution, and these cost society money. Many benefits exist to cutting pollution as well. Better health, and thus lower health-care costs, can be achieved with cleaner air and water. Some firms, such as farming companies, may have lower production costs when there is clean water.

Also, many people gain utility from a cleaner environment; this can be anything from enjoying a scenic view without smog to being able to swim or boat in an unpolluted lake. However, because firms often have no incentive to reduce their emissions on their own, there will be no abatement, and thus the markets for pollution-producing goods will fail to be efficient.

In order to correct this market failure, the government must step into the picture.² In the United States, regulation concerning SO₂ emissions first began with the 1970 Clean Air Act which limited the atmospheric concentration levels of SO₂ (Gruber, pp. 146-147). It set New Source Performance Standards (NSPS) for all new coal-fired power plants which made them reduce emissions by using low-sulfur coal or by installing scrubbers which remove much of the pollutants from the exhaust fumes. Furthermore, in 1977, scrubbers were required on new plants making the building of new plants even more expensive. This method of regulation resulted in total SO₂ emissions falling by the early 1980s; however, the older and dirtier plants which were not subject to the NSPS continued to pollute. Because these requirements only affected the new plants, there were large incentives to run the older plants past their natural lifetimes to bypass the regulation.

According to the Environmental Protection Agency (EPA, 2008), these loopholes led Congress to pass the 1990 Amendments to the Clean Air Act. Title IV aimed to reduce SO₂ emissions each year by 10 million tons under the 1980 levels. It created two phases of tightened restrictions in order to accomplish this goal. Phase I, which began in

² In some cases, the Coase Theorem may apply whereby the government assigns property rights for the common resource. However, the success of the Coase Theorem requires that bargaining costs are low. Since there are many participants in the energy market, the Coase Theorem is not useful for controlling SO₂ emissions by coal burning power plants.

1995, only affected the 263 dirtiest plants; however, there were 182 additional plants which were used as “compensating plants” that joined the program. Firms were allowed to reduce emissions at these plants instead of at the regulated heavy-polluting plants. Thus, the total number of affected plants was 445. In 2000, Phase II began; the prior restrictions on the dirty plants were tightened, and the program included nearly all fuel-burning power plants. This increased the total number of affected plants to over 2,000. By including the majority of plants, the government was able to prevent them from avoiding pollution regulation.

Furthermore, the EPA indicates that in addition to the goal of reducing SO₂ emissions, the government also wanted to minimize compliance costs and encourage the development of energy efficient technologies. To achieve this, the government included several important features in the program. The introduction of an allowance trading system was crucial. A single allowance lets a plant pollute one ton of SO₂ during the year it was issued or in the future. If a plant holds on to an allowance for future use, this is referred to as “banking” the allowance. There is no limit on the number of allowances or length of time that firms are allowed to bank, and they may use their banked allowances at any time. However, no plant can emit at levels that exceed the limit for maximum atmospheric concentrations of SO₂ set by Title I even if it possess the necessary number of allowances. Once a ton has been emitted, an allowance is retired. The government issues, rather than sells, the allowances to the affected plants based on historic fuel use and a particular emissions rate. The number of allowances is capped at 8.95 million per year which ensures that the emissions reduction will be achieved (EPA). The distribution system tends to favor older plants with regards to the number of allowances they receive.

Once the allowances have been distributed, they can be bought, sold, or banked for future use. All parties are allowed to buy and sell allowances including the general public. In fact, some groups have purchased allowances solely for the purpose of taking them out of the market in order to reduce pollution. The EPA even has an electronic tracking system to keep tally of the allowances. By implementing a trading system, the government allows the affected plants to base their abatement decisions on their marginal costs of abatement making the program more cost effective (which I discuss in greater detail below).

In addition to the trading system, the EPA holds an auction for allowances once each year at the end of March. Anyone is able to participate in the bidding including the power plants, brokers, environmental groups, or anybody else wanting to buy allowances. Only a limited number are auctioned off (2.8% of the year's total allowance allocation), and the winners can use the allowances that year. Also, parties in possession of allowances may offer them to be sold at the auction. Bidders submit forms detailing the number of allowances they want and the price they are willing to pay for them. Thus, the allowances are sold to the highest bidders.³ Having the auctions helps to send a price signal for allowances, and it gives the utilities another opportunity to obtain allowances they may need. The first auction in 1993 resulted in a price of \$131, which was half of the expected price of \$250, but it was actually a good predictor of Phase I prices (Ellerman and Montero, p. 60). The EPA also had a "direct sale" feature in which allowances could be purchased at a fixed price of \$1,500 adjusted for inflation. It was intended to allow the utility companies to build new power plants with the assurance that

³ The auction details were drawn from EPA.

they would be able to obtain the necessary allowances for the new plants, but it was eliminated in 1997 because it turned out to be unnecessary.

Because of the market for allowances and the different ways to obtain them, there is bound to be some variation in their price. Since the supply of allowances is fixed, the price is determined by the demand, and this is determined by the demand for electricity, the relative prices of fuels of differing sulfur content, the cost of scrubbing, and non-Title IV regulatory requirements affecting SO₂ emissions (Ellerman and Montero, 2007, p. 57). For example, if the demand for electricity increases, the price of low-sulfur coal becomes relatively more expensive, scrubbing becomes more costly, or regulatory requirements become stricter, then the demand for allowances will increase, thus increasing the price of allowances. Expectations about future prices can lead firms to bank allowances, reducing the supply of allowances, and thereby increasing the current price of allowances as well.

Over time, there has been a wide variation in the price of permits, which can be seen in Figure 1.1. As mentioned earlier, Phase I prices were fairly close to the first auction price of \$131. The all-time low of under \$70 occurred in 1996. This was largely due to the realization that SO₂ emissions were much lower at the affected plants than had been expected, which reduced the demand for allowances (Ellerman and Montero, pp. 60-61). The record high of just under \$1600 in 1996 was the result of the Clear Air Interstate Rule which would lower the SO₂ cap by two-thirds starting in 2010 (Ellerman and Montero, p. 49). This obviously worried the utility companies that feared they would not be able to reduce their emissions by the necessary amount. Demand for allowances

drastically increased and the price subsequently skyrocketed. As the utilities were able to cope with the new cap expectation, the price fell to the typical sub-\$200 level.

As detailed by the EPA, each plant must measure its emissions of various pollutants. A continuous emission monitoring (CEM) system is generally required, and the sources report hourly data to the EPA every quarter. The monitoring is a crucial component of the programs as it ensures that the goals are met and allows the market to properly price the allowances. At the end of the year, if a plant does not have the necessary allowances for its emissions, there is a penalty of \$2,000 adjusted for inflation (base year 1990) for each extra ton of SO₂ it emits. Violating plants must also offset the excess SO₂ emissions with the appropriate number of allowances by having that number of allowances deducted from the coming year's allotment or by setting up some other plan to make up for the violation.

The remainder of this report is organized as follows. Section 2 presents the economic theory pertinent to understanding the relative cost effectiveness of a cap-and-trade system. It discusses different methods of pollution control: quantity regulation, technology mandates, taxation, and emissions restrictions with tradable permits. This is followed by a review of the empirical research on the economic efficiency and impacts of the U.S. Acid Rain Program. The final section concludes this report.

2. Economic Theory of Pollution-Reduction Methods⁴

Governments have several policy options available to them to reduce pollution, and each may be appropriate in different situations. This chapter discusses the economic theory behind each method and its implications for consumer surplus, producer surplus and efficiency. It is also important to note that governments can regulate two different markets. The first and more common is the market for a good; in the case of the Acid Rain Program, it would be the market for electricity. The second, and perhaps more useful, is the market for pollution reduction. The market a government chooses to regulate can have a distinct economic impact.

Quantity Regulation: Quantity Restriction vs. Emission Restriction

The first policy method of pollution control is to directly limit the amount of pollutants each utility is allowed to emit. This can be done with a restriction on the plant's production level of the good (production restriction), electricity, or a restriction on the plant's emissions level (emissions restriction). New technology makes plant-specific emissions restrictions possible. As we will see, emissions restrictions will be more cost effective than production restrictions.

If the government chooses to regulate the goods market, i.e., a production restriction, it will specify the total quantity of production allowed for each of the power plants. Production levels will be restricted in order to reduce emissions. The quantity each plant is allowed to produce may be based on the average emissions produced in the industry. To achieve efficiency, regulators will mandate a production level where

⁴ This section is based on material presented in Rosen and Gayer (2008), O'Sullivan et al. (2008) and Gruber (2005).

$SMC=SMB$.⁵ In general, the use of quantity restrictions in an otherwise competitive market means that the producers are made better off because they are able to receive a higher price per unit when production is reduced. Conversely, consumers of the good are worse off due to the quantity restriction because they are consuming less of the good and have to pay a higher price for it. This is illustrated in Figures 2.1 and 2.2. Figure 2.1 illustrates the market outcome. Figure 2.2 illustrates the quantity restriction which equates SMC with SMB .⁶ We see that the quantity restriction transfers surplus from consumers to producers. The area of triangle ABC is the lost surplus due to the restriction; however, it is more than offset by the value of the pollution reduction due to the regulation.⁷ Even though consumers are made worse off, if the socially optimal level of production is correctly estimated, then society as a whole will be better off. Note that a production restriction does not allow the firms flexibility in choosing how to abate. Moreover, a production restriction does not give firms incentive to find innovative and least-cost ways to abate.

Alternatively, the government can regulate emissions. However, although more cost-effective than a production restriction, an emissions restriction will not be cost effective in general. To see why, consider the following. If the government chooses to establish a market for pollution reduction, then it specifies the total amount of allowable pollution and then divides this amount, possibly equally, between all of the power plants. For this to be possible, it must monitor the emissions at each plant. It also may specify

⁵ The SMC equals the (private) marginal cost (PMC) plus the marginal external damage of the production process. When negative externalities are present the SMC exceeds the PMC and markets generate too much of the good to be socially efficient.

⁶ In this case, the private marginal benefit (PMB) equals SMB because we assume there are no positive externalities associated with electricity production.

⁷ The benefit to society of the regulation can be calculated as the vertical distance AC times the change in Q . The area of this rectangle exceeds the lost $CS + PS$ (area of triangle ABC).

which technologies the plants are allowed to use. To illustrate this case, I will use Gruber's example (pp. 133-138) and Figure 2.3. Total pollution is at 400 units prior to the government intervention. The marginal external damage (MD) of each unit of pollution is \$100 which equals the benefit to society for each unit abated. Suppose that the government wants to reduce pollution by 200 units, and that there are two power plants, A and B, with differing marginal cost of abatement (MC) curves. The government could adopt a policy of equal emissions reduction across plants and require that both plants reduce emissions by 100 units each.

This approach of equal emission reduction across plants will not be the least-cost method to reduce pollution. It would be more efficient for Plant A to reduce pollution more because it can do so more cheaply than Plant B. The socially optimal level for each plant occurs where each plant's MC of abatement intersects the MD curve; this is also where the total marginal cost of abatement curve (MC_T), the horizontal sum of both plants' MC of abatement, intersects the MD curve. Based on differing costs of abatement, Plant A should reduce emissions by 150 units and Plant B by 50 units in order to minimize the total costs of abatement. Thus, the total cost of abating 200 units of pollution will be lower than if each plant has to reduce emissions by 100 units. However, the government cannot know the each plants' MC of abatement. Thus, with a quantity restriction on emissions, the best the government can do is mandate emission reductions based on something other than the plant's MC of abatement.⁸ As we will see, a policy of plant-level emissions restrictions will not be as cost effective as either an emissions tax or a policy that sets a target pollution level for the industry and allows plants flexibility in

⁸ It could use a policy of equal emission reduction across plants or give some plants, such as those with older technology, more leeway on emissions.

how to meet it. A tradable pollution permits system, as I discuss below, allows such flexibility.

Technology Mandate

The 1970 Clean Air Act instituted a technology mandate for newly constructed coal-burning power plants. This type of regulation is known as command-and-control. In general, such a mandate will be a costly way to reduce pollution for two reasons. First, a technology which is efficient for one plant may not be for another. Second, this policy eliminates the incentive to create better abatement technologies because there would be no benefit to using them. Thus, because of the inefficient use of technology, production costs will be higher relative to other policy approaches. Although this method may be simple, it is unlikely to be the best policy choice for a government wanting to reduce pollution at minimal cost to society.

Taxation: Commodity Tax vs. Emissions Tax

The next option for the government to limit pollution is through a corrective tax, also known as a Pigouvian tax. This method forces firms to internalize the cost of the damage caused by the pollution, thus resulting in the socially optimal level of pollution (bringing SMC and SMB into alignment). Figure 2.4 shows the situation in which the government places a per-unit tax on each good produced (commodity tax). This shifts the supply curve upward by the amount of the tax because the marginal cost of each good has increased. The price of the good increases and the quantity produced decreases. The consumer surplus falls and the impact on consumers is equivalent to the case of a production restriction. However, this time producer surplus falls as well because part of

the high product price has to be remitted to the government as tax revenue. For this reason, producers are better off if the government uses a production restriction over a commodity tax since PS is higher with the production restriction. This commodity tax has the same drawbacks as a production restriction, in that it does not encourage firms to seek low-cost abatement methods.

Alternatively, the government could tax emissions directly. Using Gruber's example for a second time, referring to Figure 2.3, the tax will be set equal to the marginal damage at \$100 as before. The plants will abate pollution when the cost of abatement is less than the tax. For Plant A, reducing pollution costs less than \$100 for all units abated up to 150 units. Thus it reduces its emissions by 150 units and pays the tax on the 50 units it does emit. Plant B, on the other hand, will reduce pollution by 50 units and pollute 150 units. This method is likely to be more cost effective than a tax on the production of electricity because it allows firms to choose the amount of abatement based on their marginal cost of abatement. Because it is more cost effective than a tax on the goods market, a tax on emissions may result in a lower price and thus more CS than a tax on the goods produced.

Rather than a tax on plant-specific emissions, heavy polluting plants would rather have a tax on the good which could be based on the industry's average emissions rate; this is because their emissions rate is higher than the average. Low polluting plants, on the other hand, would prefer a tax on plant-specific emissions because their emissions rate is lower than the average, rather than paying the higher "industry-average" rate. Plants that are in the middle ground are indifferent and would pay the same amount in taxes either way.

The greatest advantage of using a corrective tax on emissions over either a production restriction, emissions restriction, or a commodity tax is that it lets the plants choose how much pollution to abate based on their marginal cost of abatement. Moreover, with a tax on emissions, plants begin to switch to cleaner technologies or use low-sulfur coal in order to reduce the amount of taxes they pay. The government also has an additional source of revenue from the tax. A drawback of using an emissions tax rather than the other four methods so far discussed is that the amount plants abate is more uncertain. Thus the actual amount of pollution is uncertain under an emissions tax because it is based on the plants' abatement decisions.

Emissions Restrictions with Tradable Permits

The use of marketable pollution permits or allowances to control emissions is a relatively new innovation. The government sets a target pollution level for the industry, issues enough permits to meet this target, and allows firms to trade these permits. This is also referred to as a cap-and-trade system. Allowing the trading of the permits can be beneficial for both parties. Firms with low abatement costs will sell their permits for a higher price than what it costs them to reduce their emissions while firms with high abatement costs will buy permits for a lower price than what it costs them to reduce their emissions. Using Gruber's example again, referring to Figure 2.3, suppose the government gives each plant 100 permits, each one allowing one unit of pollution. Without trading, each plant would be allowed to pollute 100 units⁹. However, once trading is introduced, both plants can be made better off. The MC of reducing pollution by 100 units is much higher for Plant B, and thus it has an incentive to buy permits from

⁹ This is the outcome of plant-specific emissions restrictions discussed above.

Plant A. As long as the price is above its MC of abatement, Plant A will agree to this trade. Trading will occur until Plant A is reducing its pollution by 150 units, and Plant B by 50 units; at this point the MC of abatement for both plants will be at \$100, the socially optimal level of pollution. Thus, this system is able to take advantage of the differences in abatement costs between the different plants to minimize the total abatement costs.

Tradable permits allow plants to abate according to their MC of abatement. As discussed above, a Pigouvian emissions tax does the same. Because of this, both an allowance system and an emission tax generate lower compliance costs; therefore, the product price will be lower under these policy instruments than if the other methods are used. This would increase consumer surplus under an allowance system or emissions tax relative to the other methods. The producer surplus situation is slightly more complicated and depends on permit prices. For example, heavy-polluting plants prefer a permit system to an emissions tax as long as the permit price times the number of permits needed is less than the taxes that would be owed, which would equal the tax rate times total plant emissions.

Because there are a fixed number of permits, the supply curve is vertical; thus the level of demand for permits determines their price. If the government wants to lower the emissions cap, they can reduce the number of permits over time causing the supply curve to shift to the left. This results in a higher price for the permits; more plants increase their abatement efforts due to the higher price causing total emissions to fall. If plants are able to develop more efficient technologies which allow them to reduce their cost of abatement, then the demand curve for permits will shift to the left resulting in a lower permit price. Thus, unlike with quantity regulation, technology mandates, or commodity

taxes, there are incentives to develop more efficient technology under an allowance system just like there are under a Pigouvian emissions tax.

The greatest advantage of either a cap-and-trade system or an emissions tax is that it “allows sources to select their own compliance strategy. For example, to reduce SO₂ an affected source may repower its plants, use cleaner burning fuel, or reassign some of its energy production capacity from dirtier plants to cleaner ones. Sources also may decide to reduce electricity generation by adopting conservation or efficiency measures. Most options, like fuel switching, require no special prior approval, allowing the source to respond quickly to market conditions without needing government approval” (EPA, Basic Information). Other advantages are that regulators need less information than with quantity restrictions, technology mandates, or commodity taxes, and regulators can be certain of the pollution level because the government sets it itself.

Finally, one possible disadvantage of the tradable-permit system is that it can create pockets of pollution. This would occur if heavy-polluting plants, whose abatement costs are high, purchase allowances from newer, cleaner plants, whose abatement costs are low. As a result, the heavy-polluting plants continue to pollute at high levels creating areas with large amounts of pollution, known as hotspots.¹⁰ Note that hotspots similarly can arise with emissions taxes. Heavy polluting plants with high abatement costs will pay the tax and continue to pollute; firms with lower MC of abatement will reduce pollution to avoid the tax. This simplified analysis of the way in which a tradeable permit system works ignores the role that banking plays in impacting emissions abatement. I discuss the

¹⁰ A hotspot can arise from the presence of one large, heavy polluting plant and does not require a geographic concentration of plants. However, Swift (2000) finds that the amount of emissions exceed the number of permits in several regions, suggesting that heavy polluters are geographically concentrated.

implications of banking for the cost effectiveness of permit markets below in the review of Ellerman and Montero (2007).

3. Economic Literature on the Acid Rain Program

Ellerman and Montero(2007)

Ellerman and Montero examine the efficiency of banking allowances in the U.S. Acid Rain Program by developing a model of efficient banking and using the first eight years of data. They also identify erroneous assumptions in the earlier views as well as the necessary conditions for efficient banking to exist independently of changes in the counterfactual (estimates for the emissions that would have occurred without the cap).

The authors claim that the two-phase feature in the Acid Rain Program is very significant. In Phase I, the “big dirties” were subject to an emission cap whereas in Phase II, nearly all coal-burning power plants were subject to the cap. This cap became stricter in the second phase as well, but compliance was about twice the required level. Because the marginal cost of abatement in Phase II was higher, there was incentive for the plants to bank their allowances; in fact, 30% of the allotted allowances were banked during Phase I. Over 1/3 of these banked allowances were used in the first three years of Phase II. Thus, many analysts believed that there was too much banking in Phase I.

However, Ellerman and Montero argue that "this misperception has been due to 1) a failure to understand how an initial error in expectation concerning counterfactual emissions affects banking behavior and 2) an assumption of greater irreversibility in abatement at the beginning of Phase 1 than seems in retrospect to have been the case" (p. 49). They are also able to show that, in general, efficient banking will be robust. In other

words, it will not be affected much by changes in the level of the counterfactual or the total abatement required by the cap. They use the data through 2002 ignoring 2003-2006 because of the disruption in the market caused by the announcement of the Clean Air Interstate Rule which would lower the SO₂ cap by two-thirds.

A permit-banking model with no uncertainty and perfect competition results in permit prices rising by the risk-free interest rate. However, firms do not know how many permits they will need in the future, so the equilibrium price of permits becomes an uncertain variable. Because of this uncertainty, holding permits and investing in abatement are no longer risk-free activities. In order to minimize their compliance costs, risk-averse firms will diversify their portfolios, which will include permits. Efficient banking requires that firms set their marginal abatement cost to the price of permits at each point in time. Another condition is that the total number of allowances equals total emissions throughout the banking period. At the end of this period, the amount of abatement will equal the difference between the cap and the counterfactual. Combining these conditions, the authors are able to obtain the expected abatement in each period as a function of abatement at the end of the banking period, and then they can derive efficient banking paths. The model can be adapted to changes in agents' expectations and changes in parameter values during the banking period.

The authors next list the parameter values they used in their model. The allowance cap for Phase I was 7.62 million allowances while the initial counterfactual emissions was 9.07 million tons of SO₂. The cap for Phase II was 9.39 million allowances while the

counterfactual emissions was 15.79 million tons of SO₂.¹¹ Even though the demand for electricity was expected to grow between 1.5-2.5% per year, the expectations for the counterfactual emissions varied widely. High estimates predicted emissions growth of 1.25% per year, and low estimates predicted constant emissions or a decline of 0.5-1.0% per year. The authors split the high and low estimates and use 0.65% for the annual growth in the counterfactual emissions. They assume a linear relationship between the quantity and the marginal cost of abatement. Because allowance prices will be equal to the marginal cost, the price path will depend on the discount rate. High rates lead to shorter banking periods, lower initial prices, and greater increases in marginal abatement cost during the banking period. After the banking period, prices for allowances increase at a lower rate. The authors find that there is no undiversifiable risk due to holding allowances so the risk-free rate is the most appropriate discount rate; this is estimated to be about 3.5% during Phase I.

The actual allowance banking closely follows the path predicted by Ellerman and Montero. It follows a path according to a discount rate of 4% and a growth rate of counterfactual emissions of 0.65%. The banking follows an optimal path given these two values. However, in 2000 the actual banking path diverges from the path using constant parameters. The real discount rate fell from 4% in 2000 to 2.75% in 2002. This partially explains the change but cannot solely justify the decline in the draw-down rate, the rate at which banked allowances were used. The most likely explanation is that the expected growth in counterfactual emissions changed. This is likely to be the case because the

¹¹Phase II counterfactual emissions includes the projected emissions of all regulated plants whereas Phase I counterfactual emissions includes the initially affected plants. This is the reason for the large increase in counterfactual emissions from Phase I to Phase II.

growth in electricity demand in the 1990s was higher than expected, natural gas prices were high in 2000-2001, and Congress was considering lowering the SO₂ cap by two-thirds. All of these would cause future allowance prices to be higher.

As mentioned earlier in the introduction of this report, the first EPA auction in March 1993 resulted in an allowance price of \$131; this was about half the price of the trades occurring at the time (\$250-\$300). The expected Phase I price was \$250; however, this auction price was a good predictor of actual Phase I prices. SO₂ emissions in the Midwest fell in part as a result of the deregulation of the railroads in the 1980s.¹² Because emissions were so much lower than expected, allowance prices fell even more to an all-time low just under \$70 in 1996. Nearly half the abatement in early Phase I was due to the installation of scrubbers, which is capital-intensive, requires significant lead-time, and is irreversible for long periods. Also, multi-year contracts for low-sulfur coal were signed to get around the expected price increase of this coal once the U.S. Acid Rain Program started; these contracts also contributed to irreversibility by limiting power plants to switch back to high-sulfur coal when the higher prices did not occur.

It was the combination of the low levels of counterfactual emissions and the appearance of significant irreversibility that led many people to believe that there was too much banking of allowances. However, using their model, the authors are able to make predictions with varying levels of initial counterfactual emissions. They use a range of +20% to -20% and assume complete reversibility of abatement techniques. The initial price and quantity of abatement fall sharply with lower levels of initial counterfactual

¹² Because demand for low-sulfur coal was growing, Union Pacific built tracks in the Powder River Basin, which was previously monopolized by Burlington Northern Sante Fe, thus creating competition for the transportation of this coal. Without the deregulation, this would not have been possible.

emissions, but the amount of banking and the predicted initial emissions barely change at all. Errors in expectation of the counterfactual lead to changes in current abatement rather than banking of allowances.

Because the authors have found banking to be efficient, they conclude that sufficient reversibility must have existed. Switching to low-sulfur coal, the abatement technique having the most reversibility and requiring the least lead time, rose by about 63% while scrubbing increased just slightly. Some of the abatement by switching fuel types is presumed to have been canceled or reversed as prices fell, but as prices began to rise again, the cancelled abatement was restored.

Ellerman and Montero concluded that the banking of allowances was rather efficient. The program shows that properly constructed markets produce favorable results. This efficiency was surprising because many experts, including the authors themselves in their previous work, believed banking to be excessive. Banking was relatively robust, and the irreversibility in certain abatement decisions was not enough to prevent abatement reduction approximately equal to the error in expected initial counterfactual emissions. However, this did have a cost that showed up in the choice of abatement technique; there was less switching to low sulfur coal during Phase I than would have been optimal.

The authors discovered several implications. The correlation of the returns from holding allowances with those from having a diversified portfolio of equities is low or even zero. Thus the discount rate will be low and the role of banking is important compared to a case in which the discount rate is high. Next, changes in counterfactual emissions have a stronger impact on abatement and allowance prices than banking

practices. Finally, in cap-and-trade programs, banking will increase abatement and allowance prices during the early phases because of voluntary action. Although there were some agents whose behavior was inefficient, they did not have a big impact on the system. The U.S. Acid Rain Program cannot be said to be efficient in every sense but has shown that its participants most often act in rational ways. There was also fear that emissions would spike as plants use all of their banked allowances at once, thus creating unpredictable variations in emission levels. Some years could have low emissions as many allowances are banked while other years would have high levels of emissions as those allowances are used. However, the evidence does not support this fear.

Schmalensee, Joskow, Ellerman, Montero and Bailey (1998)

Schmalensee et al. examine the compliance costs and allowance market performance under the U.S. Acid Rain Program and the early lessons learned from it. After the authors describe how the program works, which I have already done in Chapter 1, they then explain what happened to SO₂ emissions. The data used comes from the EPA's National Allowance Data Base, Supplemental Data File, Allowance Tracking System, and Emissions Monitoring System. Total SO₂ emissions fell considerably between 1990 and 1994; about 80% of this decline was due to the larger market area of the Powder River Basin. The deregulation of the railroads and entry of United Pacific resulted in a decline of rail rates thus expanding this market's size. Emissions also fell sharply in 1995 and were much lower than the allowances issued. Emissions remained lower than the allowances in 1996 as well. There are several possible reasons for the overcompliance. First, allowance prices are expected to be higher in Phase II than in Phase I, thus it may be efficient to save some allowances for future use. The other

possibility is that the market did not foresee the rail rate declines, and thus they overinvested in scrubbers and made long-term contracts for more low-sulfur coal than needed. Another reason for the dramatic decline in emissions was that a large number of volunteer plants joined the program and reduced their emission levels. Some did this to substitute for the higher cost, dirty plants (known as Table A plants), while others did it in order to be compensated with allowances or to be exempt from NO_x regulation.

Next Schmalensee et al. discuss how the emissions were reduced. They first compute the counterfactual emissions which were calculated by taking the product of 1993 emission rates times the fuel heat input, measured in lb/MMBtu (pounds per million British thermal units), in 1996. The counterfactual measures what emissions would have been in 1996 if the fuel usage was the same in 1996 as in 1993. Because the estimated counterfactual emissions were above the 1993 emissions rate, it is clear that compliance was not obtained by a decline in the plants' use, i.e., cutting production. The utilization of scrubbers by 27 plants accounted for 45% of the emissions reduction while switching, mostly to lower-sulfur coal, accounted for the other 55%. Table A plants accounted for at least 95% of the emissions reduction in both years.

The authors then examine what occurred in the allowance market. The pricing data comes from trade press reports, price indices from the Emissions Exchange, Fieldston, and Cantor Fitzgerald, and the EPA's annual auction. The EPA reported the allocation of allowances in 1992 so that early trading could occur. This created information on what the price of an allowance should be. Early trading occurred at \$300 and \$265 per allowance which was expected. However, the 1993 EPA auction resulted in a price of \$131; this was thought to be an abnormality, but the 1994 auction had a similar

result. Shortly after, the rest of the market moved towards that price. The market also appeared to develop rather efficiently in that one price was emerging without the presence of speculative bubbles: millions of allowances were privately traded, and the average winning bids in the auctions converged towards the lowest winning bids as time progressed.

Schmalensee et al. then briefly discuss the costs of compliance. The cost of reducing emissions by 3.9 million tons in 1995 was approximately \$726 million or about \$187 per ton. This is on the lower end of the early cost estimates. They estimate a savings of 25-34% compared to a system with the same number of allowances but with no trading. Finally, the authors explain the lessons learned. The Acid Rain Program has achieved the emissions goal on time, without much litigation, and has cost less than expected. This is partially due to the lower rail rates, but this still remains a big accomplishment for an environmental program. It shows that tradable-permit programs can work the way theory predicts, and that they can adapt to other economic forces. Also, efficient markets for the allowances take time to mature. However, despite the short-term success of the U.S. Acid Rain Program, tradable-permit programs are not suited for every situation, and many factors must be properly executed for them to work.

Ludwig (2004)

Ludwig examines whether the permit-trading feature of the Acid Rain Program has a strong impact on SO₂ emissions and thus whether this feature is a cause of pollution hotspots. After discussing how permit-trading programs function and the details of the Acid Rain Program, she briefly describes the findings of Swift (2000) on hotspots. He gives several reasons why trading will have little effect on hotspots. First, the program

has air quality standards that would prevent hot spots from occurring. Second, other factors like plant location, plant size, and fuel utilization are more important contributors to hotspots than trading. Swift also found that regions had total emissions between 64-77% of their permits issued, meaning that some of the permits were banked, and that “81 percent of permits used to offset emissions come from within the state.” This means that there were not large numbers of permits moving to just a few destinations and causing hotspots; rather, his findings suggest that trading may actually prevent hotspots altogether.

Ludwig investigates the impact of net permit flows to a plant on the plant’s emissions level using cross-sectional plant-level data in 2002. Net permit flows (PNET) in 2002 are calculated as net permit purchases: the total number purchased minus the total number sold in 2002. Plants with a high PNET ought to have a high emissions level, unless the plant is buying the permits to bank them for future use.

Ludwig uses a traditional linear regression and Ordinary Least Squares analysis. Her model uses SO₂ emission levels as the dependent variable and plant size, primary fuel type, SO₂ controls, and net flow of permits as the explanatory variables. SO₂ emissions were only taken from active participants in the program, and they are measured in tons per year. Heat input is used as a proxy for plant size and is the product of the quantity of fuel used and the fuel’s heat content. Its coefficient is expected to be positive because larger plants emit more SO₂. A dummy variable is used for fuel type where the primary use of coal is assigned a value of 1; all other fuel types are assigned a value of 0. This coefficient is also expected to be positive because coal has higher sulfur content than the other fuel types. A dummy is also used for SO₂ controls. Plants that have some sort of

control have a value of 1, and ones that do not have a value of 0. Controls reduce emissions so this coefficient is expected to be negative. She uses data from 2002, the most recent year at the time, which came from the EPA's website.

In her regression results, Ludwig finds all of the coefficients to have the expected signs, and all of the variables were significant at the 99% confidence level. She finds the coefficient for net permit inflow to be 0.06 meaning that, *ceteris paribus*, each additional permit above a balance of zero increases SO₂ emissions by 0.06 tons per year. Note that this coefficient is significantly less than one, so that one permit bought on net increases SO₂ emissions by far less than one ton. Rather, the other explanatory variables such as plant size and fuel type appear to be more important in explaining emissions levels. However, this result may be an artifact of the year 2002: impending legislation proposed lowering the SO₂ cap by two-thirds, leading firms to bank. Running this model for a year in the 1990s prior to the announcement of this legislation may have different results.

Ludwig then tests for potential econometric problems. She finds that multicollinearity is not a problem within her model, but heteroskedasticity is. The coefficients remain unbiased, but the t-statistics and F-statistics cannot be trusted because the variances and standard errors are underestimated. Although she does not do so, the heteroskedasticity problem can be eliminated by using robust standard errors. Note that a correction could change the level of significance for PNET from significant to insignificant, which would further support her conclusion that PNET has little impact on SO₂ emissions. In light of her results, Ludwig concludes that hotspots should no longer prevent policymakers from using permits trading systems. She suggests that incentive

based programs, like the Acid Rain Program, be used not only for environmental policy, but for a multitude of policies.

Tietenberg (2003)

Tietenberg examines the experiences of programs using tradable permits, primarily ones involving air-pollution control, water supply, and fisheries management. His goal is to draw some conclusions from these experiences to provide insight on how to implement tradable-permit programs resulting from the Kyoto Protocol. The author reports that the use of tradable-permit programs is becoming more prevalent when it comes to limiting access to the “commons” by which he means public goods such as the atmosphere, fish populations, and the wetlands. The reason for this is that typical command-and-control regulations do not sufficiently protect the value of the resources being used. The Kyoto Protocol will expand their use even further using three mechanisms: emissions trading (allows trading of assigned amounts), joint implementation (gives emissions-reduction credit to an Annex I¹³ country when they help finance projects that reduce emissions in another Annex I country), and the Clean Development Mechanism (CDM) in which Annex I parties finance projects in non-Annex I countries to receive emission reductions. It also appears that starting new tradable-permit programs becomes easier with more experience.

Tietenberg found both successes and failures with tradable-permit programs. Often times, they are more cost effective than traditional policies. There have been many lessons learned from past experiences that can be applied to the climate change case. The

¹³ Annex I countries are those listed in Annex B of the Kyoto Protocol and primarily include developed countries.

Kyoto Protocol aims to control six greenhouse gases. Experience with multi-pollutant programs is limited, so while past experience can be useful and informative, it is by no means a sure thing. Cap-and-trade programs performed both economically and environmentally better than credit-trading systems¹⁴ because the latter have higher transaction costs and regulatory barriers. Monitoring and enforcement are vital to the success of these programs. Carbon emissions may be easy to estimate, but other greenhouse gases are not quite as easy to monitor. Some countries cannot monitor emissions as easily as other countries. There are differing degrees of enforcement capability among nations as well which can hinder the progress and efficiency of international programs. Countries with little or no enforcement could just end up selling their permits to countries where enforcement is much stricter. Thus monitoring and enforcement are much easier said than done. Other problems may arise as well. Spatial externalities may not be a problem for CO₂ emissions, but they have been an issue for fisheries and water allocation and may be problematic for other greenhouse gases. Leakages, or the relocation of firms, either within a country or between a country that has signed the protocol and one that has not, could occur if the firms want to avoid regulation.

Tietenberg also claims that setting the cap is a vital step, but if there is no consensus on what that cap should be, the goal of reducing emissions may not be reached. Another important requirement for success is emitters' ability to find cost-effective ways to reduce their emissions, but some emitters, especially those in

¹⁴ In a credit trading system, a non-regulated plant can cut emissions below a specified level and earn credit for making the emissions reduction. It can then sell that credit to a regulated plant, which can use that credit instead of cutting its own emissions.

developing countries, may not be able to do this. The CDM is intended to counter this problem. It has been observed that over time, programs move towards fewer restrictions as parties become more familiar with them, but too many restrictions from the start may doom them to failure. Also, flexibility (i.e. allowing banking and borrowing) in the system is important to reduce the costs of compliance.

The author also lists lessons for program design. The costs of monitoring, enforcing, and administering the system can be paid for by fees placed on each permit. Because monitoring in the Kyoto Protocol will involve some self-reporting, there may be some deception. He suggests creating “layers of veracity checks” in order to reduce this. Also, trading should only be allowed between eligible parties. Eligibility is defined as having an approved domestic enforcement system and having been in compliance in the previous time period. Price information and other relevant data should also be easily accessible in order to make the market more transparent. Reduction obligations for future periods should also be made known. This would make clearer the value of investments. The program should be designed to account for extreme circumstances such as price spikes. Certain measures to prevent the price-spike problem could prevent extreme economic damage as a result of a severe change in circumstances. Also, joint implementation and the CDM must be carefully utilized as it will likely take some time before parties involved have enough experience with them to use them efficiently.

Finally, Tietenberg reports that there are two expectations of economic theory that do not hold in reality. The first is that the main purpose of a tradable-permit system is to protect the value of the resource, not the resources itself. However, the strictness of the cap and the degree of compliance with the cap can be affected by policy. The second is

that there is a tradeoff between efficiency and equity regardless of the initial distribution of the permits. Theory says that trade allows the permits to move towards those who value them the most. However, in practice, the permits are typically given to historic users, and thus other means of ensuring equity are required. Therefore, these systems must still deal with the efficiency-equity issue. Tradable-permit programs are not right for every situation, but they will likely be very important when it comes to climate change.

4. Conclusion

The U.S. Acid Rain Program was a revolutionary departure from traditional pollution control methods. It was able to reduce SO₂ emissions by more than the required level on time and at a lower-than-expected cost. This was an extraordinary feat that proved that economic theory has a place in the world other than just the classroom. Methods such as quantity regulation and commodity taxes have their strengths and their weaknesses, but in theory tradable-permit programs and taxes on emissions have the most efficient outcome for society. After seeing the success of the Acid Rain Program, this appears to be the case.

Review of the literature seems to agree with this conclusion. Ellerman and Montero concluded that the banking of allowances was rather efficient and relatively robust. Schmalensee et al. found the compliance costs to be on the lower end of the expectations, and that this program was significantly less expensive than a similar program without trading. Ludwig found that permit trading had little effect on SO₂ emissions, and thus was not a major cause of pollution hotspots. Finally, Tietenberg


examined the experiences of cap-and-trade systems around the world and found a number of reasons to implement them. However, all of the authors agree that a cap-and-trade system may not be appropriate in all markets, and that certain measures must be taken in order for them to be as successful as the U.S. Acid Rain Program. By learning from the experiences and lessons of this program and those of similar programs around the world, legislators will be better prepared to implement policies that are able to reduce or even reverse the damage caused by pollution and obtain the environmental quality that today's society so dearly values.

References

- Ellerman, A. Denny and Juan Pablo Montero. "The Efficiency and Robustness of Allowance Banking in the U.S. Acid Rain Program." *Energy Journal*, vol. 28, no. 4, 2007, pp. 47-71.
- Environmental Protection Agency. *Acid Rain Program*. August 25, 2008.
<<http://www.epa.gov/airmarkt/progsregs/arp/index.html>>. April 1, 2009.
- Gruber, Jonathan. Public Finance and Public Policy. New York: Worth Publishers, 2005.
- Ludwig, Lindsay C. "The U.S. Acid Rain Program and Its Effect on SO₂ Emissions." *Issues in Political Economy*, vol. 13, August 2004.
- O'Sullivan, Arthur, Steven Sheffrin and Stephen Perez. Microeconomics: Principles and Tools. Upper Saddle River: Prentice Hall, 5th Edition, 2008.
- Rosen, Harvey S. and Ted Gayer. Public Finance. New York: McGraw-Hill Irwin, 8th Edition, 2008.
- Schmalensee, Richard, Paul L. Joskow, A. Denny Ellerman, Juan Pablo Montero and Elizabeth M. Bailey. "An Interim Evaluation of Sulfur Dioxide Emissions Trading." *The Journal of Economic Perspectives*, vol. 12, no. 3, Summer, 1998, pp. 53-68.
- Swift, Bryon. "Allowance Trading and SO₂ Hot Spots - Good News from the Acid Rain Program." *Environment Reporter* vol. 31, no. 19, 2000, pp. 954-959.
- Tietenberg, Tom. "The Tradable-Permits Approach to Protecting the Commons: Lessons for Climate Change." *Oxford Review of Economic Policy*, vol. 19, no. 3, Autumn 2003, pp.400-419.

Figures


Figure 1.1: SO₂ Allowance Trading Volume and Prices, March 1994-October 2008



Note: Transfer volume data provided only through December 2007.


Source EPA (transfer volume); Cantor CO₂e (price), 2008.

Figure 2.1: Goods Market, Generates a Product Price Where $PMC=PMB$




PMC – Private Marginal Cost
PMB – Private Marginal Benefit
SMB – Social Marginal Benefit
CS – Consumer Surplus
PS – Producer Surplus

Figure 2.2: Production Restriction.




The social marginal cost (SMC) exceeds the PMC by the amount of the external damage. Q^* is found where $SMC = SMB$.

Figure 2.3: Pollution-Reduction Market with Multiple Firms, Firms' MC and Society's MB of Abatement (Not Production).



Source: Gruber, p. 136.

Figure 2.4: Commodity Tax



The tax brings SMC and SMB into alignment.