The Efficiency and Robustness of Allowance Banking in the U.S. Acid Rain Program*

A. Denny Ellerman and Juan-Pablo Montero

Reprint Series Number 191
*Reprinted from The Energy Journal,
Vol. 28, No. 4, pp. 67-92, (2007),
with kind permission from IAEE.

The Efficiency and Robustness of Allowance Banking in the U.S. Acid Rain Program

A. Denny Ellerman and Juan-Pablo Montero***

This paper provides an empirical evaluation of the efficiency of allowance banking in the nationwide market for sulfur dioxide (SO_2) emission allowances that was created by the U.S. Acid Rain Program. We develop a model of efficient banking, select appropriate parameter values, and evaluate the efficiency of observed temporal pattern of abatement based on aggregate data from the first eight years of the Acid Rain Program. Contrary to the general opinion that banking in this program has been excessive, we find that it has been reasonably efficient. We also identify the erroneous assumptions underlying the earlier view and the conditions required for efficient banking to exist independently of changes in the counterfactual, an attribute we call robustness. These results show that firms use banking provisions in a rational and predictable way and that, at least in the US Acid Rain Program, there is no support for the often expressed concern that banked permits will be used all at once to create emissions spikes.

1. INTRODUCTION

Emissions trading usually refers to trades across space in the same period of time, but it can also refer to trades through time which implies being able to use allowances from one period in earlier or later periods and which thereby changes the temporal pattern of abatement. Several authors have studied the theo-

The Energy Journal, Vol. 28, No. 4. Copyright ©2007 by the IAEE. All rights reserved.

* Center for Energy and Environmental Policy Research (CEEPR) and the Sloan School of Management, Massachusetts Institute of Technology (MIT). Email: ellerman@mit.edu

** Department of Economics, Pontificia Universidad Catolica de Chile, and visiting professor at Sloan School of Management. Email: jmontero@faceapuc.cl

We thank Paul Joskow, Matti Liski, Robert Stavins, Robert Pindyck, Dick Schmalensee, Stephen Smith, Martin Weitzman, seminar participants at Harvard, MIT and University College London and some very helpful anonymous referees for useful comments; Brice Tariel for research assistance; and CEEPR and the U.S. EPA (STAR grant award #R-82863001-0) for financial support. All correspondence should be sent to Ellerman at: MIT/CEEPR (E40-437); 77 Massachusetts Avenue; Cambridge, MA 02139 USA.

retical properties of intertemporal trading,[1] but no work has yet evaluated how firms have actually responded to this feature of tradable permits programs. The U.S. Acid Rain Program (Title IV of the 1990 Clean Air Act Amendments) provides a unique opportunity to do so since it allows banking (but not borrowing) and it is a significant experiment in emissions trading.

One of the most salient features of the Acid Rain Program is that it includes two phases. Generating units larger than 100 MWe capacity and with a 1985 emission rate greater than 2.5 pounds of SO_2 per million Btu (/mmBtu), sometimes described as the "big dirties" became subject in 1995 to a cap based on a benchmark emission rate of 2.5 /mmBtu. The second phase, beginning five years later, differs from the first in both the stringency of the cap and the scope of the program. In Phase II, all fossil-fired generating units greater than 25 MW^e, and regardless of historical emission rate, were brought under the cap, which was based on a benchmark emission rate of 1.2 /mmBtu, approximately half the Phase I rate. This two-phase structure of the program created the incentive to bank since the marginal cost of abatement in Phase II would be expected to be higher than that for Phase I. The design of the program provided the opportunity to take advantage of this incentive by placing no restrictions on the ability of firms to bank. Once an allowance of a particular vintage is activated, it is good until used.

By every measure, banking has been a major form of emissions trading in the U.S. Acid Rain Program. During Phase I, 11.65 million allowances, or 30% of all the allowances distributed in these years, were banked. Equivalently, the reduction in emissions during Phase I was about twice what was required to meet the Phase I cap.[2] Then, 3.7 million, or a little more than a third, of these banked allowances were used to cover SO_2 emissions in excess of the number of new vintage allowances issued for use in 2000-02 under the much tighter Phase II SO_2 emission cap. And, a further 2.5 million banked allowances were used for the same purpose in the next three years, 2003-2005.

The occurrence of banking in the Acid Rain Program has not been a surprise. Given the two-phased structure of the program and the opportunity to bank, all analysts expected some amount of banking; however, there was little agreement on the likely size of the bank. Early estimates varied by a factor of five: from two to ten million tons. As Phase I began, one consulting firm created

---

[1] See Rubin (1996) and Cronshaw and Kruse (1996) for general formulations; Schennach (2000) for a formulation specific to the U.S. Acid Rain Program; and Rubin and Kling (1993) for a simulation of a potential banking program for hydrocarbon emission standards imposed on light-duty vehicle manufacturers.

[2] See Ellerman et al. (2000) and Ellerman (2004) for estimates of what emissions would have been without the cap, or counterfactual emissions. Any discussion of the efficiency of abatement, marginal cost, or allowance price implies some level of counterfactual emissions since price and marginal cost reflect the quantity of abatement and abatement is the difference between what is observed and some estimate of what would occur but for the particular policy measure. All subsequent references in this paper to "counterfactual" should be understood to mean counterfactual emissions.

a small sensation by predicting a bank as large as 15 million tons.[3]

Enough years have passed that an evaluation of the temporal efficiency of this aspect of emissions trading can be made. The accumulation phase of the banking period is over, the size of the end-of-Phase-I bank is known, and the rate of draw down in the first six years of Phase II can be observed.[4] In contrast to the earlier papers on banking, which developed theoretical properties or simulated what might occur in a particular program, this paper looks at aggregate behavior in an actual program over a period of time that spans about half of the entire banking period and assesses whether observed behavior is efficient.[5]

Contrary to the common perception of excessive banking during Phase I (Ellerman et al., 2000; Smith et al., 1998), we find that the evolution of the SO_2 allowance bank has been reasonably efficient. We 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 I than seems in retrospect to have been the case. We also show that, while not necessarily true of all programs, efficient banking will generally be robust, that is, little affected by changes in the level of the counterfactual or equivalently by the total abatement required by the cap.

In this paper, we limit attention to the years through 2002 and purposefully ignore the nearly ten-fold increase in SO_2 allowance prices that occurred between late 2003 and late 2005 and the subsequent collapse in early 2006 to a level that was approximately three times the late 2003 level of prices. This increase in price began when it became evident that the proposed implementation of the Clean Air Interstate Rule would effectively lower the SO_2 cap by two-thirds in two steps beginning in 2010 and it was exacerbated by several other convergent factors. [6] In any case, it is too early to evaluate the efficiency of what is effectively a

---

[3] A report from the General Accounting Office (USGAO, 1994) projected a Phase I bank of two million tons. An earlier and more thorough analysis (EPRI, 1993) predicted a bank "between 5 and 10 million tons, with our current projections at the higher end of the range." RDI, a coal and electric utility consulting firm, forecast a 15 million ton bank in mid-1995 as the first emission monitoring reports became available (RDI, 1995). A later EPRI report (EPRI, 1997) written with the benefit of the 1995 compliance data stated: "The bank size by 2000 is surprisingly uncertain—from 10 to 15 million short tons."

[4] We use the term "banking period" to denote the entire multi-year period over which banked, unused allowances are held consisting of an accumulation phase and a draw-down phase. In the U.S. Acid Rain Program, these phases correspond to the entirety of Phase I and some years into Phase II.

[5] Since a permits bank is similar to a non-renewable resource, this paper also adds to the extensive literature on the empirical validity of the Hotelling's rule to predict price and extraction paths (e.g., Farrow, 1985). Our paper provides a simpler test because we do not need to deal with extraction costs (and how they change as the resource is exhausted) and uncertainty regarding the size of the resource.

[6] Since the increase in price since late 2003 is more than can be explained by the new banking period, other factors are also at work, such as changes in coal markets that are affecting the premium for low sulfur coal, the remarkable increase in natural gas prices, and an increase in the construction cost of scrubbers due to higher labor and steel costs,. For an analysis of the effects of a tighter future cap without these confounding effects, see Ellerman (2002) which simulates the Bush Administration's Clear Skies Proposal of February 2002 on banking behavior.

new banking period that was not anticipated in 1995. Accordingly, we restrict our evaluation of the efficiency of banking in the U.S. Acid Rain Program to the years before further, phased-in reductions of the SO_2 cap became a real prospect.

The rest of the paper is organized as follows. Section 2 presents the model of banking that is used to generate efficient temporal paths of abatement and marginal cost for comparison with observed banking behavior. Section 3 discusses the key assumptions underlying any banking program—program-specific parameters and assumptions about counterfactual emissions, the cost function, and the discount rate—and it provides estimates of the appropriate values for this program. Section 4 compares observed SO_2 allowance banking behavior with simulated efficient paths and draws inferences from that comparison. Section 5 explains the error in the earlier, too facile assumption of excessive banking and the reasons for the robustness of efficient banking in the face of significant errors in expectation concerning counterfactual emissions. Section 6 concludes.

2. A MODEL OF EFFICIENT BANKING[7]

The theory of permits banking follows directly from the theory of nonrenewable resources pioneered by Hotelling (1931). Because the cost of creating the permits that constitute the cap is zero, a banking model with no uncertainty and perfect competition would predict that during the banking period the price $P(t)$ of permits will rise at the risk-free rate of interest r, $\dot{P}(t) / P(t) = r$, where a dot denotes a time derivative. In practice, however, firms will not know with certainty the number of permits they will demand in the near and distant future, and consequently, the market equilibrium price of permits becomes an uncertain variable.[8]

When the price and demand for permits are random variables, investments in abatement and holding permits are no longer risk-free activities, and affected firms can be expected to choose an abatement path that minimizes their expected present value of compliance costs using an appropriate risk-adjusted discount rate ρ. Risk-adverse agents will diversify this risk by holding a portfolio of assets including permits.

Because modeling the efficient path of the SO_2 allowance bank is analogous to modeling the efficient extraction path of an exhaustible resource sold in a competitive market under conditions of uncertainty, we follow the approach put forward by Slade and Thille (1997), who combined the Hotelling model for pricing exhaustible resources with the capital asset pricing model (CAPM) for risky assets. Accordingly, the evolution of allowance prices during the banking period is governed by the arbitrage condition

---

[7] The discussion is limited to banking, although most of what is discussed could apply to borrowing as well. In the Acid Rain Program, borrowing is not allowed.

[8] Since firms will never know for certain the future demand for electricity or the future price of fuels of differing sulfur content, allowance prices will respond to changes in expectation during the banking period, the length of which depends in turn on these expectations. We will return to this complex interaction in our later discussion of the robustness of banking.

$$\frac{\frac{1}{dt}E_t dP(t)}{P(t)} = r + \beta\,(r_m - r) = \rho \qquad (1)$$

where E_t is the expected value operator, r^m is the expected rate of return on a well diversified market portfolio and β, a common financial variable denoting the asset-specific risk premium, is the ratio of the covariance of ρ and r^m to the variance of r^m, that is $\beta = \sigma_{\rho m} / \sigma^2_m$. Note that both r and r^m can change overtime.

In a continuous setting efficient banking also requires instant cost minimization, i.e., at each point in time firms equalize their marginal abatement costs to the current market price. Assuming that there is a sufficiently large number of individual firms so that the aggregate abatement cost function is strictly convex, continuous and twice differentiable, the arbitrage condition can be rewritten as a function of the aggregate marginal abatement costs, $C'(q(t))$, as follows

$$\frac{\frac{1}{dt}E_t C'(q(t))}{C'(q(t))} = \rho \qquad (2)$$

where $q(t)$ is the total amount of abatement at time t.

A further condition that must hold for the entire banking period is that the cumulative number of allowances issued equal cumulative emissions. Therefore, at any time t during the banking period, the time τ at which the bank is expected to end must satisfy this condition, known as the exhaustion condition in exhaustible resource markets.[9]

$$B(t) + \int_t^{\tau} a(t)dt = E_t \left\{ \int_t^{\tau} u(t)dt - \int_t^{\tau} q(t)dt \right\} \qquad (3)$$

where $B(t) \geq 0$ is the size of the allowance bank at t,[10] $a(t)$ is the number of allowances allocated at t that is specified by the legislation, $u(t)$ represents counterfactual emissions, and q(t) signifies abatement. Observed emissions, that is, emissions constrained under the program, are $u(t) - q(t)$.

In addition, a terminal condition must hold at τ. At that point in time and thereafter, multi-year trading caused by the phased-in, discontinuous cap ceases

---

[9] At the end of the banking period and thereafter, agents owning affected units can be expected to maintain some allowance carry-over to deal with the uncertainty about the demand for allowances in any given year or to take advantage of differences between current prices and those expected in the next year, but these amounts will be small by comparison with the carry-over during the initial banking period. Numerical exercises that add a small negative term on the left hand side of equation 3 indicate a minor effect on the efficient banking path. For example, a large carry over of as much as 20% (1.8 million allowances) would increase the size of the bank at the end of Phase I by less than 2%. Accordingly, we ignore this likely post-τ inventory phenomenon in the rest of the paper.

[10] Initially, B (0) = 0 but it would be expected to assume positive values during the banking period.

and the predominant form of emissions trading will be spatial trading. Absent any further reduction of the cap or other changes in conditions that would create a new banking period, and setting aside annual variations that will cause minor changes in annual carry-over, emissions will be expected to be equal to the cap, $a(t)$, for each period of time thereafter; and abatement will be equal to the difference between the cap and counterfactual emissions. Thus, at any time t during the banking period, the time τ at which the bank is expected to end must also satisfy the terminal condition

$$a(\tau) = E_t \{u(\tau) - q(\tau)\} \tag{4}$$

Equations (2), (3) and (4) together with assumptions about both counterfactual emissions, $u(t)$, and the functional form of the aggregate abatement cost function, $C(q)$, can be used to solve numerically for τ and to derive explicit efficient abatement and banking paths during the banking period. Counterfactual emissions are modeled as emissions at $t = 0$ increasing at some rate, $g(t)$, that is, $u(t) = \varepsilon\, u_0\, e^{G(t)}$, where ε represents the share of total uncontrolled emissions subject to the U.S. Acid Rain Program, and $G(t) = \int_0^t g(s)ds$. The value of ε will equal 1.0 in Phase II, when all generating sources are included, and it will be less than 1 (0.57) in the transitional Phase I, when only a subset of sources are included. Since some amount of annual variation around the central expectation for counterfactual emissions can be expected, agents are assumed to adopt a single value of $g(t)$ reflecting the mean expectation at t for the rate of increase in counterfactual emissions during the remainder of the banking period.

Aggregate marginal abatement costs are assumed to depend on aggregate abatement in the following form

$$C'(q(t)) = \alpha_i\, [q(t)]^{\gamma} \tag{5}$$

The scaling parameter, α_i, takes the subscript 1 during Phase I and the subscript 2 thereafter. Two time-differentiated cost functions exist because Phase II expands the scope of the Acid Rain Program to include additional generating units and abatement opportunities. The exponent, γ, reflects the curvature of the relationship and it is assumed to be the same for both the Phase I and Phase II aggregate cost functions and to remain unchanged during the banking period.

Given the functional form in (5), conditions (2) and (4) can be combined to obtain the expected amount of abatement in each period $q(t)$ as function of abatement at τ

$$q(t) = \begin{cases} q(\tau)\left(\dfrac{\alpha_2}{\alpha_1}\right)^{1/\gamma} e^{-\rho(\tau-t)/\gamma} & \text{if } 0 \le t \le T \\ q(\tau)e^{-\rho(\tau-t)/\gamma} & \text{if } T \le t \le \tau \end{cases} \tag{6}$$

where T denotes the end of Phase I and $q(\tau)$ is given by (4). Substituting (6) into (3) yields a single equation that, when fully specified with parameter values, can be expanded to solve for the end of the banking period, τ, and to derive efficient banking paths as is done below for the beginning of the trading program ($t = 0$) using g_0 to denote the expected growth rate of counterfactual emissions at $t = 0$.

$$a_1T + a_2(\tau - T) = \varepsilon u_0 \left(\frac{e^{g_0T} - 1}{g}\right) + u_0 \left[\frac{e^{g_0\tau} - e^{g_0T}}{g}\right] -$$

$$(u_0e^{g_0\tau} - a_2)e^{\frac{-\rho\tau}{\gamma}}\left[\left(\frac{\alpha_2}{\alpha_1}\right)^{\frac{1}{\gamma}}\left(\frac{e^{\frac{\rho T}{\gamma}} - 1}{\rho/\gamma}\right) + \left(\frac{e^{\frac{\rho\tau}{\gamma}} - e^{\frac{\rho T}{\gamma}}}{\rho/\gamma}\right)\right] \qquad (7)$$

The two terms on the left-hand-side of (7) state the number of allowances available in Phase I and during the years of Phase II constituting the drawdown part of the banking period. The first two terms on the right-hand-side give cumulative counterfactual emissions for units affected during Phase I and for all units during Phase II through the end of the banking period. The third term on the right-hand side states the cumulative emission reductions over the entire banking period. The term outside the brackets is the amount of abatement required at $t = \tau$ discounted to $t = 0$ to express the amount of abatement that would occur at $t = 0$. The two terms within the brackets are indices of cumulative abatement, normalized to $q(0)$, for the Phase I units and for all units during the Phase II part of the banking period, respectively. If the two sets of covered units were identical, that is, if $\alpha_2/\alpha_1 = 1$, the two terms would collapse into one.

Equation (7) implicitly defines the efficient banking program at $t = 0$, when $B(0) = 0$, on the assumption that parameter values will not change during the banking period; however, it can be easily adapted to incorporate subsequent changes in agents' expectations or in other parameter values during the banking period. The observed banking path will then reflect segments of differing efficient paths each reflecting successive starting points, the accumulated bank as of each starting point, and the changed parameter values.

3. PARAMETER VALUES FOR THE U.S. ACID RAIN PROGRAM

3.1 Allowances and Counterfactual Emissions

The allowance cap and the assumptions concerning counterfactual emissions define what would be the annual required reduction of emissions absent any banking, as is depicted in Figure 1. The annual allowance caps are specified in the legislation and implementing regulations. The average annual, aggregate allowance cap for Phase I, a_1, is 7.62 million allowances, the cap for Phase II, a_2, is 9.39 million allowances, and the transition from Phase I to Phase II, T, occurs at $t = 5$.[11]

Figure 1. Title IV Caps and Counterfactual Emissions

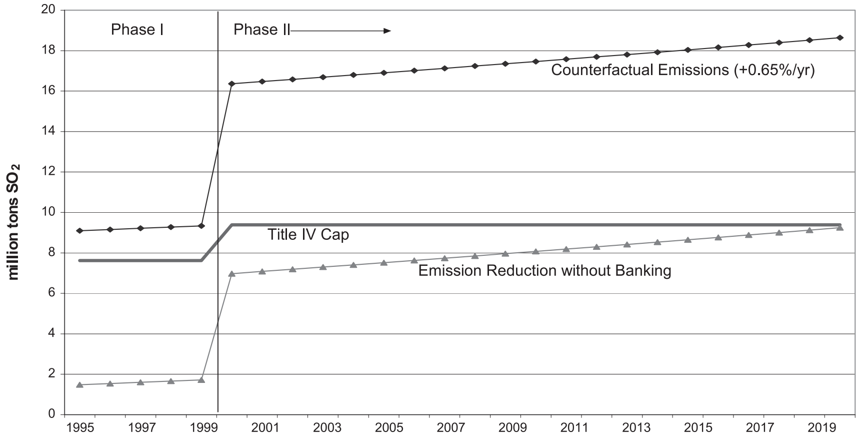

Initial counterfactual emissions can be estimated with considerable accuracy. For the generating units first affected in Phase I and remaining so since then, we assume that initial counterfactual emissions, $u_1(0)$, are 9.07 million tons, as determined by a simple technique for calculating the counterfactual.[12] The initial counterfactual for the much larger universe of units affected in 2000, $u_2(0)$, is the sum of the counterfactual for the Phase I units and observed 1995 emissions of 6.72 million tons for the units first affected in 2000, or a total of 15.79 million tons.

Estimation of the expected growth rate in counterfactual emissions, g, is subject to much greater uncertainty. Both EPRI and EPA's contractor, ICF, conducted careful early studies for the purpose of analyzing the effect and cost of the Acid Rain Program and they contained estimates, as of the early 1990s, of what

---

[11] A total of 38.09 allowances were distributed for the five Phase I years and the annual cap cited here is this cumulative sum divided by five. In fact, more allowances were allocated in 1995 and 1996 than in 1997-99; however, the distribution of the total five-year amount among years in Phase I is without importance from the standpoint of an efficient banking program at reasonable discount rates because of the short duration of Phase I. The Phase II cap is an average for the period 2000-09 including the following components: 8.9 million allowances from the basic allocation distributed annually to units and through the EPA auction, 0.10 million allowances to ¶ 410 opt-in units, and an assumed annual average of 0.39 million bonus and extra allowances over this period. Various bonus and extra allowances amounted to 0.96, 0.56, and 0.53 million allowances in 2000, 2001, and 2002, respectively, but they will be fewer in subsequent years.

[12] A simple method for computing a counterfactual assumes that, in the absence of the SO_2 program, emission *rates* would have remained unchanged at the values observed in 1993 for Phase I units, and that total emissions would vary according to observed changes in heat input. As discussed more extensively in Ellerman et al. (2000) and especially in the appendix by Schennach, this simple counterfactual closely tracks aggregate emissions as estimated by econometric techniques that take trends in emission rates into account.

emissions without the program were thought likely to be (EPRI, 1993; ICF, 1989). Although growth in the demand for electricity from fossil-fuel-fired generating units was expected to be between 1.5% and 2.5% per annum, expectations for counterfactual SO_2 emissions varied greatly depending on assumptions about the retirement of coal fired units, the utilization of nuclear capacity, and the economic competitiveness of new gas-fired generating units. High emissions scenarios predicted SO_2 emissions growing at an annual rate of about 1.25% per annum through 2010, while the low emissions cases predicted either constant emissions after 2000 (EPRI, 1993) or emissions that are declining at rates between 0.5% and 1% per annum (ICF, 1989).

For the simulations in this paper, several values for expected growth in counterfactual emissions are used. For the initial expected banking program, we use a value of $g = 0.65\%$, drawn from an early-1990s EPA forecast (Pechan, 1995), to provide a single value representing expectations at the beginning of Phase I. This estimate is approximately half way between the high and low emission scenarios in the early EPRI and ICF analyses of expected emissions absent the SO_2 cap. The low and high emission scenarios are represented by $g = 0\%$ and $g = 1.25\%$, respectively.

3.2 The Cost Function

Two parameter values define the cost functions: the convexity parameter, γ, and the scaling parameters of the cost functions for the Phase I units, α_1, and for all units, α_2. In (7), only the ratio of the scaling parameters is needed and it can be easily shown using (6) that $(\alpha_2/\alpha_1)^{1/\gamma} = q_1/q_2$, or the ratio of abatement by Phase I units in Phase 2 to that by all units in the same year. This ratio can be estimated based on observed data for 2000, 2001, and 2002; and its value is 0.83.[13]

The convexity parameter indicates the rate at which marginal cost rises with the quantity of abatement, and values for this parameter can be inferred from several studies. The early EPRI study of abatement costs contains several charts of this relationship, and it is linear over the relevant range.[14] Analysis by the authors and colleagues at MIT concerning the cost of reducing SO_2 emissions from the 2000 levels by retrofitting scrubbers to presently unscrubbed coal-fired units indicates a similar relationship. Accordingly, we assume a linear relationship between the quantity and marginal cost of abatement, $\gamma = 1.0$.

With these assumptions, the marginal cost of the annual reduction required by the SO_2 cap without banking can be calculated as is shown by the solid

---

[13] This ratio varies between 0.81 and 0.84 over these years. The effect of this variation on the total amount of allowances banked at the end of Phase I is less than 2%.

[14] EPRI, 1993; Figures 5-4 and 6-15.

line in Figure 2.[15] This is the price dual of the quantity path given in Figure 1 and the opportunity to reduce cost or increase profit by a temporal redistribution of abatement effort is immediately evident.

Figure 2. Expected Optimal Allowance Price Paths (Constant 2001$)

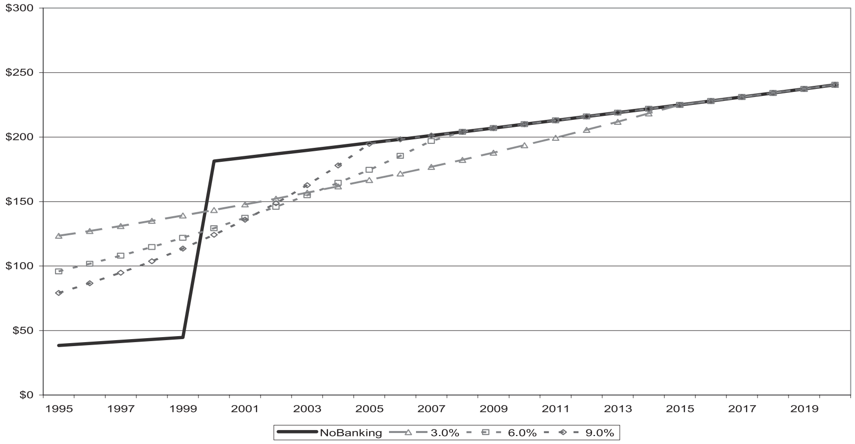

Since prices can be expected to be equal to marginal costs at all times during the banking period, equation (2) implies that the actual price path with banking will depend on the discount rate, which would be a real rate since the marginal abatement cost function is stated in today's dollars. The dashed lines in Figure 2 show the price paths for real discount rates of 3%, 6%, and 9%. The end of the banking period, τ, differs for each discount rate and higher discount rates are associated with shorter banking periods, lower initial prices, and greater increases in marginal abatement cost during the banking period. Once a banking program has ended, prices would cease rising at the discount rate and increase instead at a uniform lower rate characterizing the post-banking period.

3.3 The Risk-Adjusted Discount Rate

SO_2 allowances are financial assets that are readily tradable and can be turned into cash. As such, holding allowances implies foregoing the return that could be earned during the holding period by investing the cash in other financial assets having similar risk characteristics. The relevant criterion for determining

---

[15] In making the estimates of optimal price in Figure 2, the scaling parameter, α_2, is assigned a value of 26 that causes the simulated 2001 prices to approximate observed 2001 prices. Average monthly prices during 2001 ranged from a low of $159 in January to a high of $208 in August and the average monthly price for the year was $185. The price paths in Figure 2 will be shifted up or down to the extent that the true value of α_2 is higher or lower than 26, or that 2001 prices were below or above true equilibrium prices.

this return is the degree of undiversifiable risk associated with holding SO_2 allowances, indicated by the beta (β) coefficient in equation (1). With many years of SO_2 allowance price data now available (since mid-1994), the monthly returns from holding SO_2 allowances can be calculated and correlated with the returns from holding a broadly diversified portfolio of equities.

The details of the calculation of the beta for SO_2 allowances and the appropriate discount rate are relegated to the Appendix. The results show that there is no evident *un*diversifiable risk associated with holding SO_2 allowances. Consequently, we consider the risk-free rate to be the most appropriate discount rate for SO_2 allowance banking. We estimate the real risk-free rate to have been around 3.5% during Phase I.

This result, which may strike readers as surprising, as it did us, is critical to the analysis that follows. Two considerations specific to SO_2 allowances help to explain this result. First, the asset beta associated with producing electricity, the joint product of the SO_2 emissions covered by the U.S. Acid Rain Program, is very low. During the late 1990s, equity betas, which take debt leverage into account, are typically around 0.5 for regulated electric utilities and 1.0 for the more highly leveraged, unregulated producers of electricity. When the observed equity betas for these two types of electricity producers are adjusted to account for the risk associated with their differing debt levels, the resulting asset betas are similar, 0.2, which is low, although not zero.

Second, and perhaps more importantly, the factors determining allowance prices are considerably different from those determining profits from generating electricity, not to mention the profits of the corporate sector as a whole. The profits of electricity producers will be influenced mostly by the price of electricity, the cost of fuel, regulatory treatment, and the growth in demand for electricity. The first three factors will have little direct influence on allowance prices, and the growth in the demand for electricity is only one of several factors determining counterfactual emissions. Factors that are at least as important in determining SO_2 emissions, and therefore the demand for allowances, are 1) the *relative* prices of fuels of differing sulfur content, in particular coal and natural gas prices, 2) the cost of scrubbing, and 3) non-Title-IV regulatory requirements affecting SO_2 emissions. While these factors may have some effect on the profits of the owners of electricity generating assets, the effect on equity returns for the market in equities as a whole is negligible.

Given these considerations, it is perhaps not so surprising that the returns from holding SO_2 allowances are uncorrelated, or at best weakly correlated, with market returns and that the beta for SO_2 allowances is zero or close to it.[16]

[16] As is conventional in evaluating risk premia, we started with a null hypothesis of $\beta = 1.0$, which was easily rejected. Then, the line of reasoning given in the text led us to view SO_2 allowances as assets without systemic risk and to a null hypothesis of $\beta = 0$, which could not be rejected. A null hypothesis of $\beta = 0.2$ might also have been chosen on the reasoning that allowances are more like electricity than the equity market as a whole. Although less appropriate in our view, this hypothesis also cannot not be rejected. The implied difference in discount rate between these two hypotheses is within the bounds of uncertainty presented in Figure 3 and this difference would not alter the conclusions of the paper.

Armed now with a discount rate, the basic specifications of an inferred abatement cost function, and reasonable estimates of expectations concerning counterfactual emissions, we can turn to an evaluation of the efficiency of banking in the U.S. Acid Rain Program.

4. IS THE EVOLUTION OF THE SO_2 BANK EFFICIENT?

One way to perform this evaluation is to simulate alternative banking programs as agents might have envisioned them in the early to mid-1990s when strategies to comply with the U.S. Acid Rain Program were being formulated. Figure 3 depicts alternative banking programs by the size of the accumulated bank at the end of each calendar year, as it is built up in Phase I and drawn down in Phase II.

Figure 3. Optimal Initial Banking Programs

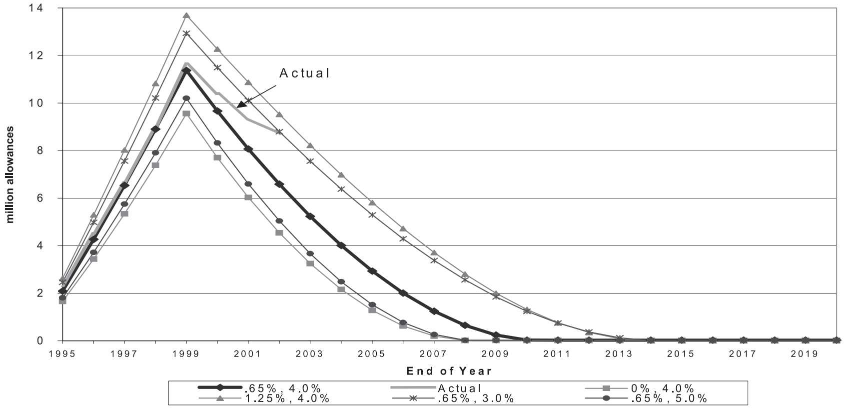

The shaded, fuzzy line tracks the actual evolution of the SO_2 allowance bank through 2002. All the other lines indicate efficient initial banking programs with plausible assumptions about discount rates and growth in counterfactual emissions in the early to mid-1990s. The bold line in the center, which closely tracks the observed banking path through 1999, represents a program with parameter values corresponding to a 4.0% real discount rate and growth in counterfactual emissions at 0.65% per annum. The lines above and below this banking path reflect combinations of higher or lower growth in counterfactual emissions (1.25% and 0% per annum, respectively) and higher or lower discount rates (5% and 3%), as indicated in the legend to Figure 3. These variations yield larger or smaller banking programs that reflect differences in expected costs over time because of differences in expected rates of growth in counterfactual emissions or in discount rates.

Figure 3 suggests that the actual banking path during Phase I resulted from the application of a discount rate of about 4.0% and expected growth of counterfactual emissions of 0.65%, or some approximately equivalent combination of discount rate and expected growth in counterfactual emissions. More importantly for evaluating the efficiency of banking, the accumulation of the bank during Phase I follows a path that would be optimal within a range of values for *g* and *r* that are reasonable for this period.

Beginning in 2000, the actual banking path departs from the constant-parameter-value path indicated through 1999 in a direction that would indicate a lower discount rate, higher counterfactual emissions, or some combination of the two. Thus, an important consideration in evaluating the efficiency of banking under the U.S. Acid Rain Program as a whole is whether such a change in parameter values can be reasonably assumed to have occurred.

There was a significant decline in the real discount rate over this period: from approximately 4.0% in 2000 to 3.30% in 2001 and 2.75% in 2002. However, the decline in the draw-down rate was greater than what can be justified by lower discount rates as shown in Table 2, where the three data columns show the optimal draw-down rates for three cases: 1) if no change in expectation concerning the growth in counterfactual emissions or the discount rate had occurred, 2) if the observed changes in discount rate are allowed but growth in the counterfactual remained unchanged, and 3) the observed draw-down rate.

Table 2. Comparison of Drawdown Rates, 2000-02 (million allowances or tons)

Year	No change in *g* or *r*	Actual *r*, no change in *g*	Observed Drawdown
2000	1.71	1.71	1.25
2001	1.60	1.43	1.06
2002	1.48	1.21	0.90

The most likely explanation by our model is the value of *g*, the unobserved growth in counterfactual emissions, which is a matter of expectation. Participants in the market will tend to formulate such a change in condition as an increase in the expected demand for allowances over the remainder of the banking period due to a change in expectation concerning the growth of demand for electricity or any other condition that would make allowances more valuable than previously thought. In the late 1990s, electricity demand was growing at a higher rate than expected and the implied demand for allowances would have been higher as well.[17] Changes in the expected growth rate of emissions can have a powerful effect on banking behavior. For instance, the slower draw-down rate observed in 2000-02 would be efficient, given the lower discount rate, if expectations concerning the value of *g* increased from the constant value of 0.65% used in Figure 3

---

[17] The actual average annual increase in heat input (2.4% between 1993 and 2002), a reasonable proxy for electricity output, has been within but near the upper end of the expected range (1.5% to 2.5%) for growth in output. More importantly, the growth rate in counterfactual emissions, given the observed increase in output, over the same period, 1.63%, has been greater than the 1.25% rate that was seen as the high end of the likely range for emissions growth.

to 1.25% around 2000. Other factors that may also have contributed to a slower draw-down rate were the high level of natural gas prices first observed in 2000-01 and the announcement of legislative proposals in Congress to lower the SO_2 cap by about two-thirds. The importance of these unobserved factors remains speculative, but all would have worked to extend the banking period and to slow down the rate of draw-down in what had become by 2000 a fixed stock. A more satisfactory explanation of the post-2000 behavior is needed, but it will depend importantly on how changes in expectation develop and affect markets. Our purpose here is only to point out that the slower post-2000 draw-down is consistent with reasonable assumptions about changes in expectation that would have influenced banking behavior.

5. WHAT WAS WRONG WITH THE EARLIER VIEW OF INEFFICIENT BANKING?

Our finding of reasonably efficient SO_2 allowance banking is surprising in view of the earlier and widespread perception that, if anything, there had been too much banking during Phase I. An explanation of the error in assumption reveals two features affecting banking behavior whose importance was not properly evaluated: 1) the effect on banking behavior of an error in expectation concerning the level of counterfactual emissions, and 2) the extent of irreversibility in the early abatement decisions. In this section of the paper, we first describe the mistaken, earlier view of inefficient banking to make clear the assumptions behind it. Then, we proceed to an explain why a change in expectations concerning the level of counterfactual emissions, such as occurred at the beginning of Phase I, would have had little effect on banking behavior assuming sufficient reversibility in abatement decisions. We refer to this condition as robustness in the remainder of this paper. In the third subsection, we present the evidence that abatement decisions have been sufficiently reversible to permit efficient banking.

5.1 The Earlier View of Inefficient Banking

The behavior of SO_2 allowance prices in the early years of Phase I strongly suggested the existence of an error in expectation concerning the level of counterfactual emissions and therefore the amount of abatement required to comply with the new SO_2 caps. The allowance price revealed by the first EPA auction in March 1993, \$131, was about half the level of the few trades then reported at prices ranging between \$250 and \$300 and of an informed estimate of the expected Phase I price at \$250 (EPRI, 1993). Although these unexpectedly low prices were attributed initially to the design of EPA's annual auction (Cason, 1993; Cason, 1995; and Cason and Plott, 1996), the clearing price in the first EPA auction in 1993 proved to be a remarkably good predictor of Phase I prices. Subsequent research indicated that the scope for the perverse seller incentives identified by Cason and Cason and Plott were limited (Joskow, Schmalensee, and

Bailey, 1998) and that SO_2 emissions in the high-sulfur-consuming Midwest had fallen for reasons unrelated to the U.S. Acid Rain Program, namely, the reduction in the delivered cost of western low-sulfur coals to the Midwest due to the deregulation of railroads during the 1980s (Ellerman and Montero, 1998). As evidence accumulated with the quarterly emission reports in April and July of 1995 that SO_2 emissions at affected sources were much lower than expected, allowance prices fell even more to an all-time low, slightly under \$70, in early 1996. Subsequently and at least through early 2004, the SO_2 allowance price has risen steadily but erratically along what can be described as a stochastic Hotelling path as shown in Figure 4 below.[18]

Figure 4. Current Vintage SO_2 Allowance Prices, 1993-2004

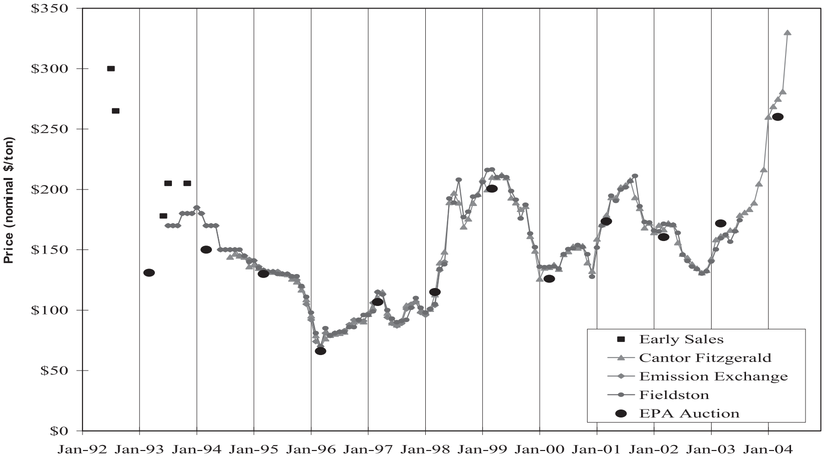

It is also evident that the abatement decisions made by the owners of the units subject to the initial Phase I cap were to a considerable extent irreversible. Almost half of the abatement in the first years of Phase I resulted from investments in scrubbers, a capital-intensive abatement technique which requires significant lead-time in construction and is irreversible for long periods of time. Moreover, an undetermined part of the alternative technique for abatement, switching to low-sulfur coal, involved multi-year contracts some of which were signed prior to 1995 as a hedge against an expected increase in low-sulfur coal prices once the U.S. Acid Rain Program became effective in 1995. These contracts would also have contributed to irreversibility by limiting the extent to which affected units could be switched back to high-sulfur coal when the expected higher prices for low sulfur coal did not materialize.

---

[18] As noted earlier, in the months subsequent to those shown on this chart, the price rose to a high of almost \$1600 in late 2005 and fallen to \$550 by late 2006.

The combination of an unexpectedly lower level of counterfactual emissions and of what seemed to be significant irreversibility in abatement decisions led many observers, including ourselves (Montero and Ellerman, 1998), to believe that more allowances were being banked than would have occurred in the absence of the error in expectation. In sum, an efficient level of banking would not be achieved: a conclusion that is directly contradicted by the finding of this paper.

5.2 The Robustness of Banking to Errors in the Counterfactual

The first error in the earlier conclusion lies in the implicit assumption that changes in expectation concerning the initial level of counterfactual emissions, u_0, would have a significant effect on the amount banked during the accumulation phase. In any year, during the accumulation phase, a careful distinction must be made between abatement to meet the cap and the additional abatement motivated by banking. How a change in the counterfactual, or more generally, an error in expectation concerning counterfactual emissions, might affect each of these components was not analyzed earlier and has not received the attention it deserves in the literature on permit banking.

This question can be answered using the model developed in section 2 and comparing outcomes for different levels of counterfactual emissions, u_0, assuming complete reversibility in abatement and holding constant the values for the discount rate, the expected growth in counterfactual emissions, and the cost function. Table 3 reports the results of such an experiment for changes in the initial counterfactual that range from +20% to -20% of the value used in the simulations reported earlier.

Table 3. Effect of Variations in the Level of Initial Counterfactual Emissions on 1995 Price and Quantities

	1995 Price (1995$/ton)	Abatement in 1995 (million tons)	Allowances Banked in 1995 (millions)	1995 Emissions (million tons)
+20%	$162	5.41	2.11	5.51
+10%	$134	4.48	2.09	5.53
Base Case	$107	3.57	2.09	5.53
- 10%	$80	2.67	2.10	5.52
-20%	$54	1.81	2.15	5.47

There is a remarkable difference between the effects expressed in the first two data columns and the last two. The initial price and quantity of abatement fall significantly with lower initial counterfactual emissions; yet, the amount of banking and the level of predicted initial emissions with the cap hardly change. It would appear that, when abatement decisions are completely reversible, all of the adjustment to this type of error in expectation will be made in current abatement, not abatement motivated by banking.

A more complete picture of the effect of a change in the counterfactual on banking behavior is given by Figure 5, which shows the optimal emission paths for affected sources for the same variations in the initial counterfactual.

Figure 5. Optimal Emission Paths with Variation of the Initial Counterfactual

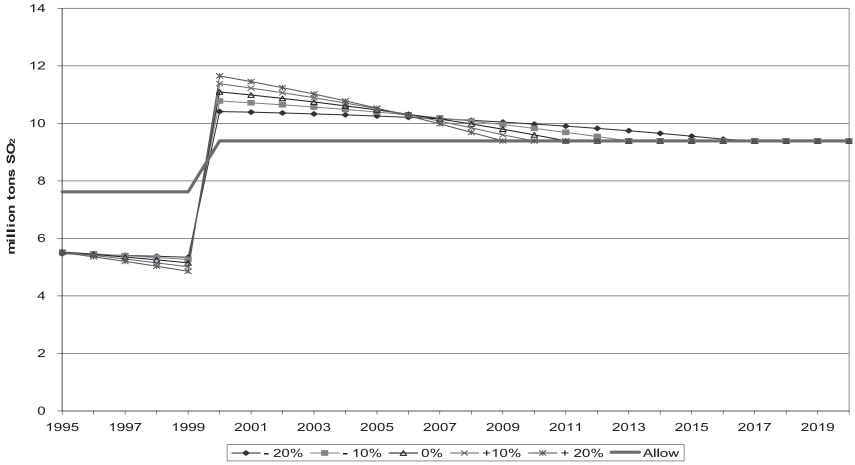

In 1995, at the start of the program, the level of emissions and the amount of banking is virtually unchanged, as in Table 3. In the later years of Phase I, slightly lower emissions and more banking are associated with a higher counterfactual; however, the accumulated bank at the end of Phase I, which ranges from 11.06 to 12.14 allowances or tons of SO_2 over these cases, is affected little by what are large changes in the level of the counterfactual. The main effect of a change in the level of the counterfactual (not in the growth rate, *g*) is seen in the draw-down phase where higher and lower than expected levels of counterfactual emissions are associated with faster and slower draw-downs, respectively, of the accumulated bank.

The very different effect of a change in the level of counterfactual emissions on current abatement and on extra abatement for banking during Phase I reflects the differing motivation of each. A change in the level of counterfactual emissions operates first and foremost on the quantity of current abatement required to achieve the cap. If there were no banking, there would be nothing else to consider: more or less abatement would be required in the current and all future periods and expected prices would be correspondingly higher or lower. The effect on the extra abatement for banking is more complicated. It depends on the extent to which the incentive to bank, that is the *difference* in marginal cost at T, the transition from Phase I to Phase II, is affected by the change in the counterfactual. Using the notation from the earlier part of the paper, this incentive is given by:

$$\Delta mc_T = \alpha_2 \, (u_0 e^{gT} - a_2) - \alpha_1 \, (\varepsilon u_0 e^{gT} - a_1) \tag{8}$$

In (8), the two terms in parentheses on the right-hand-side express the quantity of required current abatement at that moment in continuous time, *T*, when the transition from Phase I to Phase II occurs. The alpha coefficients provide the scalars translating abatement into marginal cost for the two marginal abatement cost curves. If the composition of units in Phases I and II were the same, then $\varepsilon = \alpha_2/\alpha_1 = 1.0$ and the two marginal cost functions would collapse into one. The difference in marginal cost would then be $\alpha\ (a_1 - a_2)$, or the change in marginal cost of abatement caused by the tightened cap.[19] In the case of the Acid Rain Program, the relation is more complex because the units included in the two phases are not the same. The change in marginal cost depends not only on the relation between the caps and counterfactual emissions in the two phases, but also on the abatement and emissions characteristics of the units included in the first phase.

The effect of a change in the counterfactual on the banking incentive when the units included in the two phases differ is given by:

$$\frac{\partial \Delta mc}{\partial u_0} = \alpha_1 e^{gT} \left(\frac{\alpha_2}{\alpha_1} - \varepsilon\right) \tag{9}$$

The level of the counterfactual drops out of the derivative to create a constant that depends on design features (*T* and ε), cost parameters (α_1 and α_2), and the expected rate of growth in the level of counterfactual emissions, *g*. The critical variable, determining the sign of the effect, is the term in parentheses. The ratio, α_2 / α_1, is that of the slope of the Phase II cost function to the slope of the Phase I cost function and it expresses the ratio of Phase II marginal cost to Phase I marginal cost for the same level of abatement. Since Phase II abatement includes Phase I abatement, this ratio is always less than unity and, as noted in section 3.2, its value is about 0.83. The term, ε, is the ratio of counterfactual emissions from Phase I units relative to the counterfactual for all units, and it is equal to 0.57. The derivative at (9) is then positive because the sources included in Phase I include a greater proportion of the abatement potential for all sources that would be covered in Phase II than they do of the emissions. It is easy to see that this difference, which reflects the characteristics of the generating units included in Phase I, could be very small and even negative for programs with a different set of units in the first phase.[20]

Changing the incentive to bank is, however, not the end of the story. A banking program must meet other conditions, notably, the quantity accumulated in Phase I must equal that drawn down in Phase II and marginal cost must rise at the discount rate throughout the banking period. The complete interaction can be

[19] Note that this result, in which u_0 drops out of the equation entirely, is true only for a linear marginal abatement cost function, such as we assume.

[20] Simple algebraic manipulation of (8) and the corresponding geometric representation shows that a positive value for the term within parentheses in (9) is a precondition for banking.

illustrated by the following differential, which expresses the change in cumulative Phase I emissions at t ($\leq T$) that is associated with a one-unit change in the counterfactual at t = 0.

$$dX_t = \left[\frac{\partial U_t}{\partial u_0} - \left(\frac{\partial Q_t}{\partial u_0} + \frac{\partial Q_t}{\partial \tau} \frac{\partial \tau}{\partial u_0} \right) \right] du_0 \tag{10}$$

In this formula, the capital letters represent the time integrals for actual emissions (*X*), counterfactual emissions (*U*), and abatement (*Q*). The derivation of this cumulative differential for any *t* within Phase I is complicated, but most of the terms can be understood without reference to the more complete derivations.[21] The first term within brackets is obviously positive. With no change in the emission cap, the first term within parentheses must also be positive and, if there is no banking or borrowing, equal to the change in counterfactual emissions. Consequently, a change in the counterfactual will have no effect on emissions, which are determined by the unchanged cap. With banking, cumulative emissions will be lower than they would be without banking and the change in emissions resulting from a change in the counterfactual will not necessarily be equal to the change in counterfactual emissions for any t ≤ T. The complete effect on emissions depends also on the second term within parentheses, which states the effect on abatement motivated by banking. This effect is the product of the partial derivatives expressing the chained effect of a change in the counterfactual on τ and of the change in τ on abatement. For the parameters characterizing the Acid Rain Program, this product is positive and the magnitude of the two terms within parentheses are such that the sign of dX_t is negative and the magnitude small. A simulated one million ton increase in the initial counterfactual leads to a cumulative reduction in emissions, and therefore an increase in the end-of-Phase I bank, is about 230,000 tons, or approximately 5% of the five million ton cumulative increase in the counterfactual.

In summary, the requirements of an optimal banking program imply that most of any change in counterfactual emissions will be absorbed by the change in current abatement required to reduce emissions to the Phase I cap. The complicated effect on abatement undertaken for banking depends on the shares of counterfactual emissions and abatement possibilities of the units included in Phase I of the program. If the units in both phases are the same and the marginal abatement cost function is linear, banking will be completely robust to a change in the counterfactual. Where the composition of units differs in the first phase, as in the U.S. Acid Rain Program, the robustness of banking to changes in the counterfactual depends on the shares of abatement potential and emissions for the Phase I units. In the Acid Rain Program, those shares are such that banking behavior is relatively robust to changes in the level of counterfactual emissions.

---

[21] A memo providing the complete derivation is available from the authors upon request.

5.3 The Condition of Sufficient Reversibility

Banking behavior would remain robust when a mistaken view of the counterfactual is recognized only if abatement decisions are sufficiently reversible to allow the required adjustment in *current* abatement to occur. For instance, if all abatement decisions were irreversible, it would not be possible to change the current abatement component of total abatement and all of the excess abatement would be banked. The relevant question then is whether the undeniable elements of irreversibility in the U.S. Acid Rain Program were sufficiently important to prevent the required adjustment in current abatement during Phase I. If a sufficient proportion of the initial intentions and commitments to abatement were reversible in the course of Phase I, and especially in 1995 as prices were falling, agents would be able to adjust and no excess banking would occur.

The clear implication of the finding of reasonably efficient banking in this paper is that sufficient reversibility existed. Evidence in support of this conclusion is provided by the choices of abatement technique since 1995 by the units continuously subject to the Acid Rain Program since 1995, as shown in Table 5.

Table 5. Emission Reduction due to Scrubbing and Switching (Phase I units only)

	Scrubbing (million tons)	Switching (million tons)	Scrubbing Share
1995	1.71	2.19	44%
1996	1.83	2.20	45%
1997	1.92	2.29	46%
1998	1.89	2.47	43%
1999	1.80	2.74	40%
2000	1.94	3.38	36%
2001	1.99	3.27	38%
2002	1.97	3.56	36%

Switching, the form of abatement requiring the least lead time and having the least irreversibility, increased by about 63%, or by almost a million and a half tons, while the amount of abatement from scrubbing increased relatively little from 1997, when all the Phase I scrubbers were first operating for the full year. The magnitude of the initial error in expectations can only be estimated, but the magnitudes are such that earlier intentions to abate one or even two million tons more by switching in 1995 could be presumed either to have been cancelled or to have been quickly reversed as prices fell from around $150 in late 1994 to the all time low of $70 in early 1996. Then, as allowance prices increased in the ensuing years to highs of as much as $200, much of what may have been initially cancelled abatement by switching was restored.

The pattern of incremental abatement since 1995 also identifies the cost and incidence of the inefficiency created by the elements of irreversibility in the

U.S. Acid Rain Program. Instead of the approximately 45:55 split between abatement by scrubbing and by switching observed in 1995, the efficient split—given perfect foresight or a complete absence of irreversibility—might have been closer to the 35:65 split observed in 2002. The error in expectation led to an earlier commitment to irreversible investments in abatement than would otherwise have occurred; however, this inefficiency dissipated as allowance prices rose over time to the level of the erroneous expectation that justified the investments in the first place. The cost of the inefficiency associated with the irreversible commitments showed up not in the level of emissions or in the amount of banking but in the supply of the reversible technique, in this case, switching to low sulfur coal. The suppliers of low sulfur coal were faced with greater reduction in demand than would have been the case had the initial expectations been correct or had it been possible to undo the irreversible abatement decisions. In sum, while sufficient reversibility will allow efficient banking, it does not imply that the shares of abatement by competing techniques are optimal.

6. CONCLUSION

The finding that the temporal pattern of abatement under the U.S. Acid Rain Program has been reasonably efficient is both reassuring and surprising. It is reassuring in affirming once again that properly constructed markets produce good results. The surprise results from the discrepancy with the commonly accepted view that banking in this program was excessive. Reconciling the results of this paper with that view reveals the importance of understanding both how changes in expectations concerning counterfactual emissions affect banking and the extent of irreversibility in abatement decisions. In the case of the U.S. Acid Rain Program, efficient banking is relatively robust with respect to changes in expectation concerning counterfactual emissions. Also, the elements of irreversibility in abatement decisions were not so large as to prevent a reduction in abatement approximately equal to the error in expectation concerning the level of initial counterfactual emissions. The undeniable elements of irreversibility were, however, not without a cost in efficiency. That cost showed up not in the amount of banking but in the choice of abatement technique during the early years of Phase I when less switching to lower sulfur coal occurred than would have been optimal given perfect foresight or complete reversibility.

Several of the findings in this paper suggest more general implications for banking behavior. First, the correlation of the returns from holding SO_2 allowances with those from holding a diversified portfolio of equities is very low if not zero. This means that the discount rate will be relatively low and the role of banking relatively more important than it would be with a higher discount rate reflecting greater correlation with market returns. Caution is appropriate in generalizing this result to other tradable permit systems, such as for NO_x or CO_2 emissions; however, neither should it be assumed that the allowance beta and the discount rate are high. Second, changes in the level of counterfactual emissions,

u_0, (as distinct from the expected rate of growth in counterfactual emissions, g) will always have a much larger effect on the quantity of abatement and on allowance prices than on banking behavior which is determined by a number of other factors. The effect of changes in the level of counterfactual emissions on banking will tend to be small and the magnitude and sign will depend on design features specific to each program. Third, when allowed in phased-in cap-and-trade programs, banking can be expected to produce more abatement and higher allowance prices in the early phases of the program in what can be seen as voluntary "early action." This implication will be particularly important in designing CO_2 cap-and-trade programs. Since the level of the counterfactual is inherently uncertain and the initial CO_2 caps are likely to be relatively undemanding, the expectation (or the reality) of later, more stringent caps will tend to produce a positive price and early abatement even if the initial cap is non-binding.

The results of this paper should not be read as asserting that SO_2 allowance banking has been efficient in any exact sense. The uncertainties about discount rates, growth in counterfactual emissions, and abatement cost functions are too great to allow such a statement; and few real-world examples of economic behavior meet this test. Nevertheless, these uncertainties can be bounded within relatively narrow ranges and when the likely values are used, reasonably efficient banking is indicated. Some agents may have hoarded or even dumped banked allowances in a manner that could not be judged to be economically efficient, but these exceptions have not been important enough to affect aggregate behavior noticeably. More importantly, the U.S. Acid Rain Program has shown that firms will use banking provisions in a rational and predictable way. In particular, this program provides no support for the often expressed fear that firms will create emission spikes by using banked permits all at once. To the contrary, abatement is moved forward in time and the temporal pattern of emissions is remarkably smooth and stable. Allowance prices are considerably more variable than emissions, reflecting the many vagaries that affect any market, but this property should be of less importance to environmental regulators.

APPENDIX. DETERMINATION OF THE REAL RISK-FREE DISCOUNT RATE

Table A1 provides estimates of β when the monthly returns from holding SO_2 allowances are regressed by ordinary least squares on the monthly returns from holding the S&P 500 or NASDAQ indices over the same period.

Table A1. Estimation of Beta for SO2 Allowances (August 1994-December 2003)

Market Index	Beta coefficient	Standard error
S&P 500	0.1725	0.1398
NASDAQ	0.0890	0.0756

In both cases β is not significantly different from zero and the same result occurs when the same regression is made over shorter periods, for instance, leaving out the earlier observations when it could be argued that markets were not as well formed, or even for periods as short as two years (24 observations). The use of robust variance estimators and corrections for serial correlation do not change the result, nor does including later data for 2004-05. Since the results shown in Table 1 indicate no evident *un*diversifiable risk associated with holding SO_2 allowances, we use the risk-free rate as a reasonable approximation of the appropriate discount rate for SO_2 allowance banking.

Treasury notes provide the standard for determining risk-free rates of return in the U.S. economy. Inflation-indexed Treasury notes are an obvious source of real rates; however, these notes have been offered only since the beginning of 1997 and in limited maturities, typically ten and thirty years. For the years prior to 1997, and especially for the years 1993-1995 when the owners of affected facilities were developing and deciding compliance strategies, real risk-free rates must be inferred from nominal rates less assumed inflationary expectations.

The appropriate horizon for the risk-free rate applicable to allowance banking is not obvious. Although initially the end of the banking period would have been reasonably expected to be between 2010 and 2015, most of the banked allowances would be held for a shorter time. The holding period for a banked 1995 vintage allowance would be at least five years and perhaps more than fifteen, but the expected holding period would diminish with each succeeding vintage of Phase I allowances. Moreover, the holding period could be expected to be lengthened or shortened as the discount rate or expectations change. We use a ten-year maturity, partly because of the availability of a ten-year, real, risk-free rate, but also because it appears to be a reasonable average holding period for allowances banked during Phase I.

Figure A1 compares an observed 10-year real risk-free rate calculated from 10-year Treasury Inflation Protected Securities (TIPS) from 1997 through 2003 and an inferred real rate from 1992 through 2003 based on a nominal constant 10-year maturity rate less an adjustment for inflationary expectations. The TIPS rate is constructed by taking the monthly average return for each successive ten year note so that the rate in 1997 reflects the 2007 note, that for 1998 reflects the 2008 note, etc. The Federal Reserve provides average monthly constant maturity 10-year nominal rates. These rates are adjusted for inflation by subtracting the average of the trailing 1-year and 10-year inflation rates in the all-item consumer price index for urban consumers. The use of this average assumes that inflationary expectations reflect equally the experience of the recent and the more distant past.

It is evident from the period when the two overlap that the inferred real rate fluctuates more than the actual real rate. Also, the generally higher level of the actual real rate suggests that the inflationary adjustment used in this paper has generally been greater than warranted. For the period prior to 1997, when initial banking decisions would have made, the inferred rate fluctuates between 2.0% and 5.0%. Based on this data, we assume a rate of about 3.5% as an appropriate

Figure A1. Actual and Inferred Real, Risk-Free Interest Rates, 1991-2003

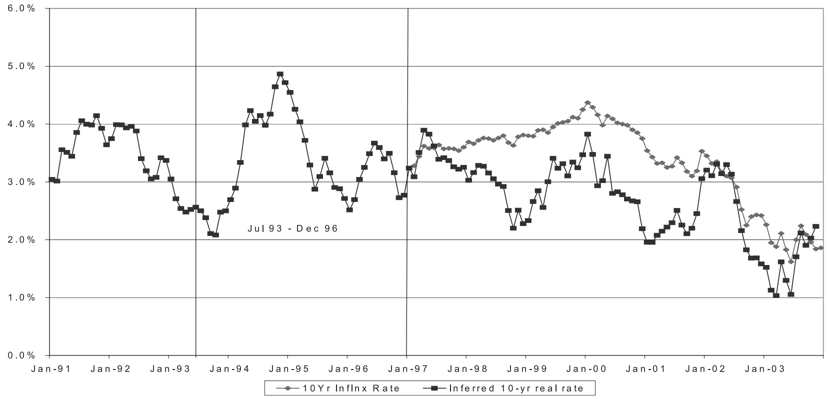

Source : Federal Reserve Bank of St. Louis, Historical series for Inflation-indexed Treasury Securities, http://www.research.stlouisfed.org/fred2/categories/82.

real risk-free rate for the years 1993-96 when initial banking decisions were decided and implemented.

REFERENCES

Cason, Timothy N. (1993). “Seller Incentive Properties of EPA’s Emissions Trading Auction.” *Journal of Environmental Economics and Management* 24(2): 177-95.

Cason, Timothy N. (1995). “An Experimental Investigation of the Seller Incentives in EPA’s Emission Trading Auction.” *American Economic Review* 85(4): 905-22.

Cason, Timothy N. and Charles R. Plott (1996). “EPA’s New Emissions Trading Mechanism: A Laboratory Evaluation.” *Journal of Environmental Economics and Management* 30(2): 133-60.

Cronshaw, Mark B. and Jamie Brown Kruse (1996). “Regulated Firms in Pollution Permit Markets with Banking.” *Journal of Regulatory Economics* 9:179-189.

Electric Power Research Institute (EPRI 1993). *Integrated Analysis of Fuel, Technology and Emission Allowance Markets: Electric Utility Responses to the Clean Air Act Amendments of 1990*. EPRI TR-102510.

Electric Power Research Institute (EPRI 1997). *SO_2 Compliance and Allowance Trading: Developments and Outlook*. EPRI TR-107897.

Ellerman, A. Denny (2002). *Analysis of the Bush Proposal to Reduce the SO_2 Cap*. MIT/CEEPR Working Paper 2002-002. Accessed at: http://web.mit.edu/ceepr/www/2002-002.pdf.

Ellerman, A. Denny (2004). “The US SO_2 Cap-and-Trade Programme.” In *Tradeable Permits: Policy Evaluation, Design, and Reform.* Paris: Organization for Economic Development and Cooperation (97-2004-07-1-P). pp. 71-97.

Ellerman, A. Denny, Paul L. Joskow, Richard Schmalensee, Juan-Pablo Montero, and Elizabeth M. Bailey (Ellerman et al., 2000). *Markets for Clean Air: The U.S. Acid Rain Program*. Cambridge University Press.

Ellerman, A. Denny and Juan-Pablo Montero (1998). “The Declining Trend in Sulfur Dioxide Emissions: Implications for Allowance Prices.” *Journal of Environmental Economics and Management* 36: 26-45.

Farrow, Scott (1985). "Testing the Efficiency of Extraction from a Stock Resource." *Journal of Political Economy* 93: 452-487.

Hotelling, H. (1931). "The Economics of Exhaustible Resources." *Journal of Political Economy* 39: 137-175.

ICF Resources Inc. (ICF 1989). "Economic Analysis of Title V (Acid Rain Provisions) of the Administration's Proposed Clean Air Act Amendments (H.R. 3030/S. 1490)." prepared for the U.S. Environmental Protection Agency.

Joskow, Paul L., Richard Schmalensee and Elizabeth M. Bailey (1998). "The Market for Sulfur Dioxide Emissions." *American Economic Review* 88(4): 669-85.

Montero, Juan-Pablo and A. Denny Ellerman (1998). "Explaining low sulfur dioxide allowance prices: The effect of expectation errors and irreversibility." MIT/CEEPR Working Paper 98-011. Accessed at: http://web.mit.edu/ceepr/www/98011.pdf.

Pechan, E.H., and Associates, Inc. (1995). *The Acid Rain Data Base Version 1 (ARDBV1)*. Contract #68-D3-0005, prepared for the U.S. Environmental Protection Agency's Office of Atmospheric Programs, Acid Rain Division, Washington.

Resources Data International, Inc. (RDI, 1995). *RDI's Phase I Databook—Performance under the Clean Air Act Amendments of 1990*. Boulder, CO.

Rubin, Jonathan D. (1996). "A Model of Intertemporal Emission Trading, Banking, and Borrowing." *Journal of Environmental Economics and Management* 31: 269-286.

Rubin, Jonathan and Catherine Kling (1993). "An Emission Saved is an Emission Earned: An Empirical Study of Emission Banking for Light-Duty Vehicle Manufacturers." *Journal of Environmental Economics and Management* 25: 257-274.

Schennach, Susanne (2000). "The Economics of Pollution Permit Banking in the Context of Title IV of the Clean Air Act Amendments." *Journal of Environmental Economics and Management* 40: 189-210.

Slade, Margaret and Henry Thille (1997). "Hotelling Confronts CAPM: A Test of the Theory of Exhaustible Resources." *Canadian Journal of Economics* 30: 687-708.

Smith, Anne E., Jeremy Platt and A. Denny Ellerman (1998). "The Costs of Reducing Electric Utility SO_2 Emissions may not be as Low as You Might Think." MIT/CEEPR Working Paper 98-010. Accessed at: http://web.mit.edu/ceepr/www/98010.pdf.

United States General Accounting Office (USGAO 1994). *Allowance Trading Offers an Opportunity to Reduce Emissions at Less Cost*. GAO/RCED-95-30.

MIT CENTER FOR ENERGY AND ENVIRONMENTAL POLICY RESEARCH REPRINT SERIES

CEEPR Reprints are available free of charge (limited quantities.) Order from MIT Center for Energy and Environmental Policy Research, E40-428, Cambridge, MA 02139-4307.

177. Inscrutable OPEC? Behavioral Tests of the Cartel Hypothesis, James L. Smith, The Energy Journal, Vol. 26, No. 1, (2005)
178. U.S. oil and natural gas reserve prices, 1982-2003, M.A. Adelman and G.C. Watkins, Energy Economics, Vol. 27, pp. 553-571, (2005)
179. Panel data analysis of U.S. coal productivity, Thomas M. Stoker, Ernst R. Berndt, A. Denny Ellerman, and Susanne M. Schennach, Journal of Econometrics, Vol. 127, pp. 131-164, (2005)
180. Characteristics of North Sea oil reserve appreciation, G.C. Watkins, The Quarterly Review of Economics and Finance, Vol. 42, pp. 335-372, (2002)
181. Petroleum Prospect Valuation: The Option to Drill Again, James L. Smith, The Energy Journal, Vol. 26, No. 4, pp. 53-68, (2005)
182. Markets for Power in the United States: An Interim Assessment, Paul L. Joskow, The Energy Journal, Vol. 27, No. 1, (2006)
183. Pollution markets with imperfectly observed emissions, Juan-Pablo Montero, RAND Journal of Economics, Vol. 36, No. 3, pp. 645-660, (2005)
184. Forward Trading and Collusion in Oligopoly, Matti Liski and Juan-Pablo Montero, Journal of Economic Theory, Vol. 131, No. 1, pp. 212-230, (2006)
185. Do Markets Reduce Costs? Assessing the Impact of Regulatory Restructuring on US Electric Generation Efficiency, Kira R. Fabrizio, Nancy L. Rose and Catherine D. Wolfram, AER, Vol. 97, No. 4, pp. 1250-1277, (2007)
186. Retail Electricity Competition, Paul Joskow and Jean Tirole, RAND Journal of Economics, Vol. 37, No. 4, pp. 799 - 815, (2006)
187. Reliability and Competitive Electricity Markets, Paul Joskow and Jean Tirole, RAND Journal of Economics, Vol. 38, No. 1, pp. 60-84, (2007)
188. The Effect of Variability in Industrial Emissions on Ozone Formation in Houston, Texas, Mort Webster, Junsang Nam, Yosuke Kimura, Harvey Jeffries, William Vizuete and David T. Allen, Atmospheric Environment, Vol 41, No. 40, pp. 9580-9593, (2007)
189. A Simple Auction Mechanism for the Optimal Allocation of the Commons, Juan-Pablo Montero, AER, Vol. 98, No. 1, pp. 496-518, (2008)
190. Capacity Payments in Imperfect Electricity Markets: Need and Design, Paul L. Joskow, Utilities Policy, Vol. 16, No. 3, pp. 159-170, (2008)
191. The Efficiency and Robustness of Allowance Banking in the U.S. Acid Rain Program, A. Denny Ellerman and Juan-Pablo Montero, The Energy Journal, Vol. 28, No. 4, pp. 67-92, (2007)

Massachusetts Institute of Technology
Center for Energy and Environmental Policy Research
One Amherst Street (E40-428)
Cambridge, Massachusetts 02139-4307