Journal of Environmental Economics and Management **38**, 176–194 (1999)
Article ID jeem.1999.1080, available online at http://www.idealibrary.com on IDEAL®

Heterogeneity, Irreversible Production Choices, and Efficiency in Emission Permit Markets*

Shaul Ben-David, David S. Brookshire, Stuart Burness, Michael McKee, and Christian Schmidt

Department of Economics, University of New Mexico, Albuquerque, New Mexico 87131

Received October 29, 1997; revised March 9, 1999

This paper investigates a class of market mechanisms for environmental regulation based on the Clean Air Act tradable discharge permit program. Laboratory market experiments capture some of the more salient institutional features and focus on issues of firm technological heterogeneity and irreversible investment regarding the operation of the permit market. Experimental results suggest that higher degrees of abatement cost heterogeneity may lead to reduced trade volume, while the implications of increased cost heterogeneity regarding price volatility are mixed. Finally, increased cost heterogeneity appears to result in decreased "laboratory efficiency." © 1999 Academic Press

1. INTRODUCTION

Economists have long espoused market-based environmental regulations on the grounds of superior cost-effectiveness [14, 15]. The potential cost savings of tradable discharge permits (TDPs or permits)[1] was a major reason for the development of a sulfur dioxide permit emissions market (SO_2 market) under Title IV of the 1990 Clean Air Act Amendments (CAAA). The potential gains from a system of tradable permits arises from the actions of traders with different marginal abatement costs reallocating abatement effort and permits until marginal abatement costs are equalized and total abatement costs are minimized.[2] The greater the initial difference in the marginal abatement costs, the more heterogeneous the market and, hence, the greater the potential gains from trade. It remains to be seen whether the potential gains from trade will provide an operational incentive to execute trades and actually achieve potential cost savings.

The performance of the EPA auction and the SO_2 emissions market in general, by far the most ambitious application of TDPs to date, has been and continues to

*This research was funded in part by the Electric Power Research Institute; thanks go to Vic Niemeyer and Bob Patrick. Shaul Ben-David is Professor Emeritus, University of New Mexico; David S. Brookshire, Stuart Burness, and Michael McKee are Professors in the Department of Economics, University of New Mexico; and Christian Schmidt is Assistant Director, Deutsche Post, Germany. Thanks also to Tim Cason, Janie Chermak, the Associate Editor, and reviewers for helpful comments.

[1]The terms "TDP" and "permits" are used interchangeably throughout the paper.

[2]See Dales [15] for a detailed discussion of the concept and Montgomery [30] for a rigorous proof of the cost-effectiveness properties of the tradable discharge permits.

0095-0696/99 $30.00

be closely scrutinized.[3] Substantial prospective savings were considered possible if the market operated efficiently. The General Accounting Office (GAO) estimated annual savings in abatement costs of up to $3 billion after the year 2000 if all potential gains from trades are realized [45]. At the time, Hausker [23] and Cason [8], among others, expressed concern that poor performance of the early SO_2 permit market might provide a basis for reluctance by legislators and administrators to utilize incentive-based regulations in the future.

A central concern and area of investigation has focused on the structure and incentive properties of the SO_2 market. For instance, Cason [10] argues the importance of the specifics of the SO_2 market design, and Cason and Plott [12] use experimental results to highlight inefficiencies in the current SO_2 market design. Recently, Cason and Gangadharan [11] reported on an experimental investigation of an alternative electronic bulletin board tradable emission permit auction format and find results comparable to the double auction trading institution utilized by the SO_2 market.[4] Thus, the SO_2 market structure continues to serve as a test case for future applications of incentive-based regulation.

Particular concerns regarding the efficacy of the SO_2 market arise over the market's ability to capture potential cost savings. For example, Hausker [23] and the U.S. GAO [45] suggested that if trading volume is low, then potential cost savings may not be realized. More recently, Burtraw [6], however, argues that in the context of the SO_2 market, low trade volumes per se do not necessarily imply market inefficiency; i.e., even if permits are not actively traded, the *option* of using permits already increases the utilities' flexibility in complying with the regulation and thereby may significantly lower compliance costs.

The literature suggests that there are a number of external factors that may artificially restrict the trading volume in the SO_2 market. These include the manner in which firms are designated as Phase I or Phase II by the EPA [45, p. 43].[5] In addition, several other reasons for low trading volumes have been cited. Hahn [21] and Stavins [37] suggest that high transaction costs lead to reduced trading volume. Cason [10] expresses concern that the incentive structures of the SO_2 market may generate misleading price signals. The structure of the electric power industry, where extensive regulation provides countervailing incentives for

[3] We distinguish between the EPA auction and the SO_2 trading market in general. The EPA auction is slightly less than 3% of annual permit allocations. As pointed out by a reviewer, sellers could contribute to this auction, but this has not been observed. Thus the volume is invariant in the EPA auction. Thus our discussion of the annual trading volume refers to the SO_2 market in general.

[4] Other elements of the SO_2 market also continue to receive attention. Lile et al. [28] report on an evaluation of the Allowance Trading System (ATS). This investigation was motivated by the view that much can be learned from SO_2 market operations.

[5] The GAO report [45, p. 43] argues that most of the firms with lower marginal abatement costs are among the Phase I utilities, while high-cost abaters are concentrated among Phase II firms. The report suggests that separating these two groups by including Phase II firms 5 years after the Phase I firms reduces the heterogeneity in the market and therefore the potential gains. A reviewer has correctly observed that with bankable permits Phase II firms are not precluded from participating in Phase I markets, thus questioning the strength of the GAO argument.

trading, is another potential cause [33, 3, 45].[6] Godby et al. [20] examined banking and tradable shares in an experimental setting with uncertainty over the control of discharges and found that banking tends to eliminate the price instability introduced by uncertainty and reduces market efficiency, while trading in rights to future entitlements (shares) tends to reduce trading volume, increase price stability, and improve market efficiency.[7] Laffont and Tirole [27] cite the initial low trading volume in new financial markets as an example where the existence of asymmetric information can lead to low initial trading volume as potential traders postpone transactions while waiting for additional information. They predict that trading volume will increase over time as additional information becomes available; moreover, given the long-term nature and the irreversibility of alternative compliance strategies that confront firms, they suggest that these initial low trading volumes may result in inefficiencies.

While it is not possible to define a benchmark for trade volumes, many of the parties involved expected a higher number of trades to occur.[8] During the first years of the auctions operation, trading volumes recorded[9] do not tell the whole story, as they ignore trading that occurs outside the EPA auction. The volume of trading in the SO_2 market has not been what was originally envisioned for the complete realization of cost savings [6, 13]. However, as time has passed, a literature is appearing that argues that significant adjustments are occurring in the industry and that the SO_2 market is alive and well. The SO_2 market is progressing in terms of achieving cost savings, and trading volume is increasing [4, 25, 35]. However, many factors beyond the SO_2 market per se have been identified and continue to play a role in explaining price behavior and trading volume. Bohi and Burtraw [4], Carlson et al. [7], and Solomon [36] note that substantial cost savings have been achieved. However, Bohi and Burtraw [4] and Solomon [36] also note that impediments to trading remain, and the realization of the potential cost savings from trading is far from complete. Furthermore, Carlson et al. [7] indicate that a significant portion of the cost savings comes from other industry factors, such as technological change and reduction in the cost of low-sulfur coal.

The question remains open as to the robustness of the SO_2 market regarding efficient operation. To identify and specify these issues we consider a straightforward analysis of a model in which (otherwise identical) firms with differing and irreversible abatement technologies employ TDPs to achieve an optimal allocation of abatement/permits. The theoretical model yields three hypotheses, which are

[6] The concept of emissions trading is predicated on firms operating in a competitive market. Few industries are more regulated and sheltered from direct competition than electric power generation and therefore are less prone to respond to the incentives created by the TDP market to reduce abatement costs. Given the "test case" character of the current TDP market for the adoption of TDP markets for future environmental regulations, the choice of the electric power industry may be regarded as unfortunate. The deregulation process that has begun in the industry is expected to strengthen firms' incentives to engage in emissions trading [45, p. 47]. However, it remains to be seen if increased competition in the industry will occur early enough to alter the performance of the present market.

[7] Thus, interestingly, some market institutions may reduce trade volume and *increase* efficiency.

[8] See Burtraw [6] for the discussion of trade volumes.

[9] U.S. EPA Allowance Auction Result, 1993–98 [39–44].

examined in an experimental laboratory setting.[10] The design of the experimental market setting includes certain institutional features of the SO_2 market as dictated by the CAAA. Included are the mandatory EPA auction, the subsequent trading in the SO_2 market, and the reduction in the allocation of permits corresponding to the beginning of Phase II of the SO_2 market.[11]

2. MODEL AND HYPOTHESES

The relationship between heterogeneity and the degree to which potential gains from trade are realized has not been extensively studied in the literature.[12] In the finance literature, the question of market liquidity has received some scrutiny (see, for example, Pagano [32] and Allen and Gale [1]). This literature focuses primarily on the relationship between price volatility and market volume, where heterogeneity enters the analysis indirectly. Epps and Epps [17] provide one of the few cases where heterogeneity among traders is linked to market volume. They explain the positive correlation between price volatility and trade volume, which is frequently observed in speculative markets, as the result of an increase in the variance of the traders' reservation prices as the number of traders increases.

The models developed in the financial literature are not directly applicable to TDP markets. The commodities traded in financial markets are return-yielding assets that are mostly traded for speculative purposes, where one asset possesses a large number of substitutes. Speculative trading in the emission permit market can be expected to constitute a much smaller percentage of trades than in asset markets, as emission permits are inputs in the production process and have a limited number of substitution possibilities. The choice of whether to buy or sell permits in a TDP market also depends on long-term and often irreversible investments in different production technologies. Traders in TDP markets thus possess much less flexibility with regard to their trading strategies than in a typical financial market.

Furthermore, to be suitable for the analysis of a permit market, a model must explicitly address differences in marginal abatement costs between the various firms. The model developed in this section formalizes the notion of heterogeneity in terms of abatement costs and investigates the implications of varying degrees of heterogeneity in terms of market operation. Specifically, we develop testable

[10] Our analysis, which focuses on institutional features, thus differs from most previous work, which dealt more with the individual aspects of emissions trading, for example, Franciosi et al. [18, 19], Cason [8, 9], and Brown Kruse and Elliot [5]. Exceptions are provided by Mestelman et al. [29], Muller and Mestelman [31], and Godby et al. [20].

[11] By considering these institutional features of the SO_2 market, our focus is upon the testing of theoretical predictions, given a set market auction. As such, we acknowledge the extensive auction literature yet do not enter the debate over whether the auctions dictated by the CAAA were the appropriate choices. (For an overview of this extensive literature, see Davis and Holt [16] and Kagel and Roth [26].) Our research is focused on market behavior, given the dictated auctions within the overall institutional setting. As described by A. Roth [34], this constitutes an attempt to enter the dialogue between experimenters and policymakers. Roth terms this "whispering in the ears of princes," where the experimental environment is designed to closely resemble the natural environment that is the policy question at hand (p. 2). Here, the dialogue explores the efficiency of the TDP market.

[12] Discussion is often limited to heuristic references to this relation. See, for example, Huberman and Hogg [24].

hypotheses concerning price stability, levels of trading, and realized efficiency gains under different degrees of heterogeneity.

To focus on the theoretical implications of differential abatement costs, we consider a straightforward analysis of otherwise homogeneous firms that differ only in their abatement technologies. Let there be N such firms in the market for emission permits with firm i's abatement costs characterized by $A_i(q_i)$, where q_i is the quantity of the pollutant abated. Let $a_i(q_i)$ be the corresponding marginal abatement cost.

Consider first the case of no trading. For simplicity, suppose that uniform standards apply; then if q^0 is 100% abatement, then $q_i^0 = q_j^0 = q^0$, and if $\bar{q}$ is the actual level of abatement, then $\overline{q_i} = \overline{q_j} = \bar{q}$. Since there is no trading, each firm must hold $q^0 - \bar{q}$ permits. In the absence of trading, total abatement costs in the market will be $C^a = \Sigma_{i=1}^N A_i(\bar{q})$, and the total amount abated will be $N\bar{q}$. As shown by Montgomery [30], this is inefficient since generally $a_i(q) \neq a_j(q)$.

The theory then suggests that with efficient trading, firms will abate to the point q_i^* where marginal abatement costs equal the price of permits, p_A; that is, $a_i(q_i^*) = p_A \ \forall i$, where p_A is determined by the market clearing condition that $\Sigma_{i=1}^N q_i^* = N\bar{q}$. Since any change in abatement will result in firms either acquiring or releasing permits, the net quantity traded at an efficient outcome will be $\Sigma_{i=1}^N |\bar{q} - q_i^*|/2$, as each permit sold corresponds to a permit purchased.

To highlight the results and provide a parallel and practical framework for the experimental analysis, assume that each firm has a linear marginal abatement cost function $a_i(q_i) = a_i q_i$, with $a_i < a_j$ for $i < j$; this defines a technology A described as $\{a_i; i = 1, \dots, N\}$. With a well-defined and fully operational market for permits, at a competitive equilibrium, $\{p_A, q_i^*(p_A); i = 1, 2, \dots, N\}$, each firm gains $G_i = \int_{q_i^*}^{q} a_i q\, dq + p_A(q_i^* - \bar{q})$ from trading, since $q_i^* < \bar{q}$ for firms that buy permits and conversely for those that sell permits. The total gain from the permit market is then the sum of these gains, given by

$$G = \sum_{i=1}^{N} G_i = \sum_{i=1}^{N} \int_{q_i^*}^{q} a_i q\, dq,$$

where $\Sigma_{i=1}^N q_i^* = N\bar{q} = \overline{Q}$. Observe that G is the sum of avoided abatement costs.

At a social optimum, one would choose $\{q_i; i = 1, \dots, N\}$ so as to maximize G subject to $\Sigma_{i=1}^N q_i = \overline{Q}$, where $\overline{Q}$ is the target level of abatement. At an interior maximum, necessary conditions include $q_i^S = \lambda/a_i$, where q_i^S is the socially optimal or cost-minimizing level of abatement for firm i, and λ is the Lagrange multiplier, interpreted here as the shadow value of an additional unit of abatement. Summing over i and using the constraint yields $\overline{Q} = \lambda \Sigma_{i=1}^N 1/a_i$ or $\lambda = \overline{Q}/\Sigma_{i=1}^N 1/a_i$. From the competitive solution we have $q_i^* = p_A/a_i$, which leads to the expected result that $\lambda = p_A$ and $q_i^S = q_i^*$. Thus the competitive equilibrium in the permits market is socially optimal.

Now consider a different abatement technology $B = \{b_i;\ i = 1, \dots, N\}$, where $b_i < b_j$ for $i < j$. Let $B_k = \Sigma_{i=1}^k b_i$ and $A_k = \Sigma_{i=1}^k a_i$ for $k = 1, \dots, N$. We say that two technologies A and B are *comparable* if $A_N = B_N$; comparability means that under uniform standards the cost of achieving any target level of abatement is the same under either technology. In addition, technologies A and B are *monotonic* in cost differences if $\{a_i - b_i;\ i = 1, \dots, N\}$ is monotonic; i.e., $a_i - b_i > a_j - b_j$ for all i, j such that $i < j$ (or vice versa). Monotonicity is a regularity condition that ensures that the two technologies have differential marginal abatement costs that change in a uniform fashion; alternatively, for comparable technologies, if the b_i are a linear transformation of the a_i, then technologies A and B are relatively monotonic.[13]

For two comparable technologies technology B is said to be *more heterogeneous* than technology A if $A_k > B_k$ for $k < N$.[14] It is now possible to state the following:

LEMMA A. *For two comparable monotonic technologies with linear marginal abatement costs, if technology B is more heterogeneous than technology A, then*

$$\sum_{i=1}^{N} \frac{1}{a_i} < \sum_{i=1}^{N} \frac{1}{b_i}.$$

Proof. Consider

$$S = \sum_{i=1}^{N} \left(\frac{1}{b_i} - \frac{1}{a_i} \right) = \sum_{i=1}^{N} \frac{a_i - b_i}{a_i b_i} = \sum_{i=1}^{N} S_i,$$

where $S_i = (a_i - b_i)/a_i b_i$. Since $a_1 = A_1 > B_1 = b_1$, monotonicity and comparability imply that there exists an integer $k < N$ such that $a_i - b_i \geq 0$ for $i \leq k$ and $a_i - b_i < 0$ for $i > k$. Equivalently, $S_i \geq 0$ for $i \leq k$ and $S_i < 0$ for $i > k$. Furthermore, since $a_i < a_j$ and $b_i < b_j$ for $i < k$, we have $a_i b_i < a_k b_k < a_j b_j$ for $i < k < j$. With the monotonicity axiom this implies that $S_i \geq (a_i - b_i)/a_k b_k$ for all i and strict inequality for $i \neq k$. Summing over all i and using $A_N = B_N$ yields

$$S = \sum_{i=1}^{N} \frac{a_i - b_i}{a_i b_i} > \sum_{i=1}^{N} \frac{a_i - b_i}{a_k b_k} = 0,$$

as desired. This leads to the following:

LEMMA B. *For any two comparable monotonic technologies with linear marginal abatement costs, if technology B is more heterogeneous than technology A, then the market clearing price is lower in the market with the more heterogeneous technology; i.e.,* $p_B < p_A$.

Proof. The proof is immediate from the equilibrium conditions stated above; i.e., $q_i^*(A) = p_A/a_i$ and $q_i^*(B) = p_B/b_i$. Summing over i and observing that total abatement is the same in either case leads to $p_A \Sigma_{i=1}^{N} 1/a_i = p_B \Sigma_{i=1}^{N} 1/b_i$. By Lemma A this yields $p_B < p_A$.

We also have

PROPOSITION 1. *For any two comparable monotonic technologies with linear marginal abatement costs, if technology B is more heterogeneous than technology A, then the potential gains from trade under technology B are greater than those under technology A.*

Proof. The potential gains from trade under technology A in this context are just the cost savings for each firm at the trading equilibrium aggregated over all

[13] Comparability imposes some restrictions on the parameters of the transformation; moreover, the converse is not generally true.

[14] If the a_i and b_i are normalized by dividing by $A_N = B_N$, respectively, then an increase in heterogeneity is analogous to a mean preserving increase in spread for probability distributions.

firms. This is given by

$$G_A = \sum_{i=1}^{N} \int_{p_A/a_i}^{q} a_i q\, dq,$$

since $q_i^*(A) = p_A/a_i$. Evaluating the integral using $\Sigma_{i=1}^{N} a_i = A_N$ yields

$$G_A = \frac{1}{2}\left[A_N \bar{q}^2 - p_A^2 \sum_{i=1}^{N} \frac{1}{a_i} \right].$$

Since $p_A \Sigma_{i=1}^{N} (1/a_i) = \overline{Q}$, this yields $G_A = (1/2)[A_N \bar{q}^2 - p_A \overline{Q}]$. A similar demonstration yields $G_B = (1/2)[B_N \bar{q}^2 - p_B \overline{Q}]$. Since $A_N = B_N$ and $p_A > p_B$, we have $G_B > G_A$, as desired.

PROPOSITION 2. *For any two comparable monotonic technologies with linear marginal abatement costs, if technology B is more heterogeneous than technology A, then the range of WTA/WTP is larger for technology B than for technology A.*

Proof. Each firm holds $q^0 - \bar{q}$ permits and abates in the amount $q_i^*(p)$, so that it needs to hold $q^0 - q_i^*(p)$. If the firm needs to divest itself of excess permits, it will sell $q_i^*(p) - \bar{q}$. Since $q_i^*(p) = p/a_i$, firm i's minimum WTA is $a_i \bar{q}$; i.e., the first increment of abatement is the least expensive, so the minimum of WTA occurs when $q_i^*(p) = \bar{q}$. As the a_i are monotonic increasing, a_1 is the lower bound for WTA for technology A. Similarly, if the firm needs additional permits, it will purchase $\bar{q} - q_i^*(p)$. As above, the maximum of WTP occurs for the permits purchased to cover the first unit of decreased abatement, when $q_i^*(p) = \bar{q}$. Again, since $q_i^*(p) = p/a_i$, the maximum WTP for each firm buying permits is $a_i \bar{q}$. Monotonicity of the technology implies that $a_n \bar{q}$ is the upper bound of WTP for firms in technology A. Since technologies A and B are comparable, $a_1 > b_1$ and $a_n < b_n$, which proves the desired result.

PROPOSITION 3. *Consider two comparable monotonic technologies with linear marginal abatement costs where technology B is more heterogeneous than A. For the case of two firms in each technology, the ex ante optimal volume of trading is greater for technology B than for technology A.*

Proof. The quantity of permits bought by firm 1 is $q_1^*(p) - \bar{q}$. Likewise, the quantity of permits sold by firm 2 is $\bar{q} - q_i^*(p)$. Since $q_i^*(p) = p^A/a_i$, setting supply equal to demand yields $p^A = 2\bar{q}a_1 a_2$, where the technologies are normalized so that $\Sigma\, a_i = 1$. Similarly, $p^B = 2\bar{q}b_1 b_2$. It is straightforward to demonstrate that $a_1 a_2 > b_1 b_2$ for comparable monotonic technologies when $n = 2$. Substituting the price values into either the demand function or the supply function yields the quantity sold by firm 1, which is $\bar{q}(2a_2 - 1) > 0$ in the A technology and $\bar{q}(2b_2 - 1)$ in the B technology. Since $b_2 > a_2$, this yields the desired result.

Proposition 1 states that potential gains from trade are greater in the more heterogeneous market. While this is not a laboratory-testable hypothesis, it does raise the question of whether realized gains from trade follow the theoretical prediction. We formalize this as

HYPOTHESIS 1. *The realized gains from trade are greater in more heterogeneous markets than in more homogeneous markets.*

The theory suggests that the incentives for achieving cost savings are proportional to the potential cost savings. Hypothesis 1 purports to test whether these incentives are carried over into the experimental market. We refer to this property as the "laboratory" efficiency of a market: the degree to which an actual experimental market structure realizes the potential gains from trade in that setting.[15] This may have possible implications vis-à-vis the concerns raised in the GAO report cited in footnote 5 concerning the separation of the market into Phase I and Phase II, as well as possible implications for market operation resulting from the creation of a national market versus several regional markets.[16]

Proposition 2 indicates that there are a wider variety of possible outcomes under more heterogeneous technologies. Consequently we propose

HYPOTHESIS 2. *The observed price variance is larger in a more heterogeneous market than in a more homogeneous market.*

However, in this context it is not quite obvious what the implications of greater price volatility are. On the one hand a failure to accept Hypothesis 2 might seem to add support to the possibility that the greater potential gains from trade are translated into incentives to achieve potential cost savings in the experimental market. However, as will be seen below, price variability is viewed by some to be related to the volume of market trading.

It appears to be the common consensus that the volume of trading would be greater in more heterogeneous markets. Proposition 3 provides support for this conjecture for the case where there are two firms (and the result generalizes to the case of three firms). Unfortunately, Proposition 3 does not appear to generalize in any readily apparent fashion, even under more restrictive assumption sets. Nonetheless, on the basis of the intuition suggested by Proposition 3, we propose the following[17]:

HYPOTHESIS 3. *The ex post observed volume of permits traded is higher in a more heterogeneous market than in a more homogeneous market.*

Within the context of our theoretical analysis, Hypotheses 1 and 3 are clearly consistent with each other. However, there are conflicting views regarding the relationship between price volatility and the volume of trading relative to the heterogeneity of the market.[18]

[15] We thank a reviewer for suggestions regarding the separation of these issues.

[16] In the literature, a trade-off is often perceived between creating a market as large as possible to realize as many potential gains from trade versus keeping the market small to account for regional differences in geographical or meteorological conditions [2, 38].

[17] Proposition 3 considers only trades in equilibrium, while in the experimental setting trades out of equilibrium are a central part of the institutional framework. For this reason we distinguish between these two notions of trade volume as ex ante optimal and ex post observed, respectively.

[18] For example, Epps and Epps [17] explain a positive correlation between price variability and market volume by an increase in traders' heterogeneity as the market volume increases. Allen and Gale [1] argue, however, that price volatility is inversely correlated with the level of market participation: if participation is higher as market heterogeneity increases (because of greater gains from trade), then price volatility will be lower in the more heterogeneous market. Finally, Laffont and Tirole [27] argue that the lack of sufficient information can lead firms to postpone trades and lead to low volumes, implying that trade volume and price variability should be inversely related.

3. EXPERIMENTAL DESIGN

To date, experimental work on the SO_2 market has largely focused on individual features of emissions trading: the auction mechanism [18], its incentive structure [8, 9, 12], or differing degrees of market power [5]. Exceptions include the work of Mestelman et al. [29], Muller and Mestelman [31], and Godby et al. [20], who examine the behavior of the traders in a fully specified institutional setting. Recently, Cason and Gangadharan [11] have examined the TDP market in Los Angeles, focusing on an electronic bulletin board trading institution. Our experimental laboratory design also focuses on the performance of the SO_2 market as a whole and, thus, incorporates two features central to our inquiry: the mandatory EPA auction, with additional voluntary trades among firms, and the planned reduction in the allocation of permits corresponding to Phase II of the SO_2 market.[19,20]

The subjects are told that each represents a firm producing a good requiring two inputs, capital and permits.[21] Firms receive an allocation of permits at the beginning of each period and revenues from the sale of the good at the end of each of the 10 periods in the session.[22] Each firm has the choice of three technologies (A, B, and C) to produce the good, each representing distinct combinations of capital and permits.[23] Permits and capital costs are inversely related in the technologies. Technology C represents the cleanest but most expensive mode of production. At the beginning of the session all firms employ technology A, which requires the largest number of permits. The implied marginal cost for switching between

[19] The experimental instructions are available from the authors at the University of New Mexico.

[20] To focus on the market features of interest to us, the experimental setting necessarily abstracts from many additional market characteristics that would exist in the naturally occurring setting. Failure to abstract would possibly make the experimental setting too complex. Thus, the firms are not required to specify an output level or price, nor is the quality of the output permitted to vary. In addition, entry and exit are not permitted, nor do any traders possess market power. Furthermore, the number of firms is held constant for Phases I and II of the SO_2 market, while in the actual market, in addition to the reduction of the number of permits allocated to each firm, the number of firms covered under the program is expanded greatly.

[21] In the experiments, value-neutral terms "pesos" and "coupons" were used.

[22] Prior to the 10-period lab experiment, two practice periods gave the subjects the opportunity to familiarize themselves with the market setting. The experimental design allowed for up to 10 traders to participate in the market. In the experiments reported here, three sessions had 10 traders, eight sessions had nine traders, six sessions had eight traders, and one session had seven traders. In the cases where the number of traders was lower than 10, the abatement cost schedule in the market was adjusted to ensure that the equilibrium prices were the same regardless of the number of traders.

[23] The experimental market attempts to mimic the actual SO_2 market, where the good being produced is electricity. In this context technology A represents the production without any investment in abatement technology, such as flue gas scrubbers, or the use of a fossil fuel source with a high sulfur content, such as Appalachian coal. Technology B represents the investment into some form of pollution control, e.g., a low-capacity scrubber, or the washing of high-sulfur coal prior to burning. Technology C represents the investment in a high-capacity scrubber or switching to a low-sulfur fuel such as natural gas or Western coal, which may be more expensive because of higher transportation costs or extensive refitting of the boiler. In either case, the investment in new abatement technology is irreversible because of the difficulty of moving it to another power plant, which is due to the plant specific setup. The same goes for switching to new fuel supplies, as fuel contracts are often made for several decades at a time, including termination clauses that make the cancellation of a contract prohibitively costly.

technologies[24] is different among the traders. The design allows for the creation of otherwise identical markets differing only in their degree of heterogeneity by varying the differences in marginal abatement costs across treatments.

Each experimental session consisted of 10 periods, and each period involved a three-stage market setting.[25] In stage I a percentage of permits are withheld by the authority. These permits are offered to all firms in a first-price, sealed-bid, discriminative revenue-neutral auction. Mimicking the EPA design, this (discriminative) auction required that successful bidders pay the price they bid (see, e.g., Kagel and Roth [26, p. 362]), but that revenues from the auction would be prorated back to firms in proportion to the percentage of total permits they surrendered. After the auction closes firms are informed of their revised holdings as well as the high, low, and average permit prices. In stage II the firms choose to *either* buy *or* sell permits in a double-oral auction. In this auction subjects can accept any bid/offer posted, and market clearing occurs on a bid-by-bid and offer-by-offer approach. If a firm wishes neither to use nor to sell a permit, it may be banked for future use/sale. As the market progresses firms see a "Market Watch Box" on their screen, which continuously displays the highest bid and lowest offer and associated quantities. Once tendered, bids/offers cannot be withdrawn (although they can be revised upward in quantity), and firms cannot offer more permits than they hold; i.e., firms must honor their bids/offers. Once the double-oral auction is closed, the firms are informed of their current balance of capital and permits. In stage III each firm chooses its production technology for the next period. A movement from technology A to B or from technology B to C is irreversible. If a firm chooses a technology for which it does not hold a sufficient number of permits, it is automatically moved to the cleaner technology to maintain compliance. At the end of each period, the firms are informed of their individual trading history, listing the balance of capital and permits as well as the trading profits for each round. At the end of the (10-period) session the balance of permits is converted into pesos, and the resulting peso balance is converted into dollars at a preannounced exchange rate.[26]

This is a relatively complicated experimental setting. A natural question concerns the manner in which optimal training is effected in such a complex setting. Previous research indicates that experience obtained over multiple sessions is more reliable than that obtained by lengthening the laboratory session [22].[27] In addi-

[24] Each trader has three abatement technologies, A, B, and C, in order of increasing costliness. In addition, the cost of changing from A to B is less than the cost of changing from B to C. In effect, the marginal cost schedule is given by a step function that approximates, in a discrete sense, the linear increasing marginal cost function in the analysis.

[25] The subjects remained in the same group during an entire session. The production parameters remained constant during the entire session and across replications, with the possible exception of the number of permits allocated. While subject mixture remained constant during a single session, subjects were mixed between different sessions. This was done to avoid possible collaboration and strategic gaming between subjects belonging to the same group over a set of sessions.

[26] The exchange rate was $0.02 for each peso.

[27] The subject pool consisted of a group of 80 University of New Mexico undergraduates, recruited primarily from among economics and political science majors. The subjects were divided into two groups of 40 each. One group pledged to be available on Monday/Wednesday afternoons, the other on Tuesdays/Thursdays. Each subject had one session a week over a periods of 7 weeks. The subjects had responded to an invitation by mail and were screened for their ability to make economic decisions in the laboratory setting. The screening sessions utilized a monopoly experiment. The subjects chosen ranked in the upper half of the screening process.

tion, the subjects participated in two training sessions.[28] To ensure a common level of experience and consistency of sequence effects, subjects were not allowed to skip a session.[29] A subject failing to show up for a session was eliminated from the subject pool.

Four experimental treatments were used to test the impact of the change in two variables on the experimental market. The first variable is the number of permits allocated to each firm during a session. The second variable is the degree of heterogeneity in the underlying abatement cost schedules. For the number of permits allocated, we introduced a treatment in which the number of permits allocated remained constant in one case (constant treatment), and one in which the number of permits allocated was reduced beginning in period 6 (reduction treatment). The former was done to isolate the effects of irreversibility of moving to cleaner production choices by keeping constant the number of permits that were allocated. The latter mirrors the actual SO_2 market, where a reduction in the total volume of permits is planned for the year 2000.

In the constant treatments, each firm receives an allocation of permits that is sufficient to cover production at technology A as long as the firm buys back the quantity of permits that is automatically deducted from its allocation for the stage I auction.[30] The initial allocation remains constant for periods 1 through 5. In period 6, a preannounced reduction of 40% of the initial level of permits is implemented, and the allocation remains reduced for the remainder of the session. To accommodate this reduction, firms, depending on their situation, have the option of adopting a cleaner technology, buying or selling the permits, utilizing permits, banking permits, or any feasible combination thereof.

Heterogeneity among traders is introduced via the marginal abatement cost of switching from one technology to another, with differences between the individual firms' marginal abatement costs being much larger in the heterogeneous market. One set of abatement cost parameters representing a less heterogeneous market (LH) was used. In the other, a set of abatement cost parameters representing a highly heterogeneous market (HH) was used. Except for these marginal abatement costs, the market settings are identical across the two markets. This allows for a direct comparison of market performance under different degrees of heterogeneity. For the constant treatment, five replications were conducted for both the LH and

[28] In the first session, the subjects were familiarized with the individual elements of the market, such as the discriminatory auction and the double-oral auction. In the second session, several short sessions of the entire market were run to train the subjects in the interaction of the different stages of the market. For the training sessions, subjects received a flat payment, rather than making pay-offs conditional on the subjects' performance. This encouraged the subjects to test different strategies without being penalized for mistakes.

[29] To keep attrition to a minimum, the subjects were paid 75% of their earnings at the end of each session. The remaining 25% of the subjects' pay-offs were held in an account and were released once a subject completed all sessions. This was forfeited by those who failed to participate in all experimental sessions. In addition to the 25% withheld, subjects were also paid a $50 completion bonus. This payment structure seemed to work well; fewer than 10% of the subject pool dropped out over a period of 7 weeks.

[30] This structure mimics the current EPA auction. As the constant treatment is also the first treatment in which the subjects participated, the subjects had ample opportunity to become thoroughly familiarized with the trading institutions before moving on to more complex market institutions.

the HH treatments. For the reduction treatments, four replications were conducted for both the LH and HH treatments.[31]

4. EXPERIMENTAL RESULTS

For expositional purposes we consider the hypotheses in reverse order. Hypothesis 3 predicts that increased market heterogeneity will lead to increased trading volume. In the experimental markets the theoretical total volume of trades is identical in the HH and LH treatments, but the potential gains from trade are lower in the LH setting. Trade volume is defined as the sum of permits traded during the mandatory auction (stage I) and during the double-oral auction (stage II).[32] Table I presents the results in terms of average trade volume. In the constant treatment, the average trading volume is 152 permits for the HH and 208 for the LH market. In the reduction treatment, the average trade volume is 161 permits for the HH and 154 for the LH market.[33]

[31] To provide the reader with a better understanding of the sequence of treatments and the opportunity for subjects to learn in this complicated experimental setting, it may be helpful to observe how a given subject (referred to as subject X for expositional purposes) experienced the entire set of treatments. Belonging to the Monday/Wednesday group, X would show up at the experimental laboratory on Wednesday afternoon together with 19 other subjects. X would be assigned a seat in the lab. She would be told that to make the set of experiments as efficient as possible, two markets would be run at the same time, each consisting of 10 subjects (left half and right half of the lab). The same instructions were read out for the two groups. The sessions in week 1 and week 2 were training sessions: subjects were introduced to individual parts of the experimental market. Training session 1 was on Wednesday of week 1, including the interaction of the different stages of the experimental market; training session 2 was on Monday of the following week. In week 3, X would show up again on Monday, again with 19 others, was assigned a seat in one of the two groups (one of them with less heterogeneous parameters and the other with highly heterogeneous parameters), and would read the instructions for the first treatments (constant LH and constant HH simultaneously). The subjects were aware that the cost structures in the two groups were different, but not that the difference in the cost structure was one treatment variable and resulted in different potential gains from trade. In week 4, X would show up on a Wednesday and was assigned to a group (but not knowing whether it was a HH or LH treatment. The assignment to the LH or HH group was random. Again, the subjects were read the instructions of the treatment (reduction LH and HH simultaneously). The treatments in weeks 5 and 6 are not within the scope of this paper.

In week 7, to increase the number of replications, the subjects were randomly assigned to participate in an institution that they had already participated in. While the subjects were facing a familiar institution, care was taken to assign subjects to a treatment in which they had not previously participated, e.g., subject X may have participated in the constant LH treatment in week 3 but would participate in the constant HH treatment in week 7. Learning in terms of increasing familiarity with the general market institutions did occur from week to week. This type of learning cannot be avoided in repeated general market settings. To the contrary, because of the complicated experimental setting, this type of learning was actually desired and the sequence of the treatments—which was identical for all subjects—was chosen to move the subjects from simpler (constant) to more complex (reduction) settings. Nevertheless, facing a new institution at each session (bar the session in the last week) and not knowing whether she was in a LH or HH treatment, X had little chance of using what she had learned in a previous session to strategically influence the market in the current session.

[32] To make possible a direct comparison of trade volumes across the treatments and sessions, observed trade volumes in the various sessions are normalized and weighted. The reported trade volume is averaged across the 10 market periods.

[33] The market sessions are not independent over market periods, since the design maintains the same subjects in each market throughout the session and feedback is provided at the end of each round.

TABLE I
Average Trade Volume

	High heterogeneity (HH)	Low heterogeneity (LH)
Constant treatment		
Session 1	131	221
Session 2	156	151
Session 3	129	205
Session 4	208	189
Session 5	137	276
Average	152	208
Reduction treatment		
Session 1	146	144
Session 2	155	167
Session 3	187	173
Session 4	156	131
Average	161	154

Note. Trade volume is adjusted for differences in endowments of TDP by dividing the total endowment into the per capita trades.

The observed trading volume is higher for the LH market in the constant treatment, as is the standard deviation across replications: 41 in the LH market versus 30 in the HH market. In the reduction treatment, the standard deviation is slightly higher again in the LH markets, but the differential is much smaller than in the constant treatment: 7 for the LH versus 16 for the HH market. While trade volume appears to be higher in the constant treatment for LH markets, no clear difference is observed in the reduction treatment. An ANOVA analysis comparing the variation within treatments with the variation across treatments rejects the null hypothesis of the means being equal across the constant treatments (at the 10% level) but fails to reject for the reduction treatments. The first result leads to a rejection of Hypothesis 3, while the second is inconclusive regarding Hypothesis 3. LH markets do not appear to be detrimental to trade volume. However, there is a great deal more variability in the behavior of the markets in the LH setting, and this may compromise market efficiency and convergence.

Hypothesis 2 states that price variability will be higher in HH markets than in LH markets. Table II presents the variance in prices for both treatments. To determine the variance, the prices from all market transactions for the 10 periods from both the stage I and the stage II auctions are included. For the constant treatment, the price variance across sessions is 1.400 in the HH market and 0.094 in the LH market, and this is consistent with Hypothesis 2. In the reduction case, the relative magnitudes are reversed: the price variance in the HH market is 0.379 versus 1.268 in the LH market. The ANOVA analysis rejects the null hypothesis that the means are constant for both the constant and reduction treatments at the 5% level. Thus, the experimental evidence supports Hypothesis 2 for the constant treatment, while the opposite occurs for the reduction treatment. Overall the experimental evidence is inconclusive.

Markets with LH traders seem to be less able to cope with the reduction in the number of permits, and this is manifested by the higher variability in the price of

TABLE II
Price Variability within Sessions

	High heterogeneity (HH)	Low heterogeneity (LH)
Constant treatment		
Session 1	1.080	0.125
Session 2	1.015	0.023
Session 3	3.366	0.135
Session 4	1.082	0.129
Session 5	0.458	0.058
Average	1.400	0.094
Reduction treatment		
Session 1	0.313	1.262
Session 2	0.067	1.384
Session 3	0.927	1.034
Session 4	0.209	1.392
Average	0.379	1.268

Note. The price variances are calculated using prices from both Stage I and Stage II.

the permits. This result is consistent with the participation argument advanced by Allen and Gale [1], who argue that when the gains from trade are smaller, fewer agents can be expected to participate in the market. The smaller participation rate in the LH market makes it more vulnerable to exogenous shocks. The results from the tests of Hypotheses 1 and 2 suggest the possibility of an inverse relationship between trade volume and price variability: the larger the price variance, the lower the trade volume.

Hypothesis 1 predicts that the realized gains from trade will be higher in the HH markets. The computation of gains from trade is straightforward. One need only calculate the reduction in abatement costs associated with post permit trading. With reversible technology (investment) choices testing, Hypothesis 1 would be equally straightforward. One would merely compare actual compliance (abatement) costs with the compliance costs associated with the cost-minimizing solution developed in Section 2. However, this metric is flawed when technology choices are irreversible. In particular, if a firm chooses a cleaner abatement technology than required for the cost-minimizing solution, the firm *cannot* return to the original technology. Therefore, a new cost-minimizing solution must be defined relative to this new technology profile,[34] and then subsequent actual compliance costs must be compared with those of the revised minimum cost solution. Likewise, if a firm later chooses a technology that precludes the revised minimum cost solution, then the process must be repeated again. Such a procedure is necessitated, as overall (laboratory) market efficiency is not determined by the technology profile of the last period but by technology choice in all 10 periods of the session. Therefore, when one or more firms make an inefficient technology choice, the underlying

[34] The set composed of the specific technology of each firm for all firms in a given period will be referred to as the *technology profile* for the industry.

cost-minimizing solution must be redefined for the remainder of the session.[35] Thus, a number of sequential metrics may have to be developed—an ex ante optimum metric based on the original cost-minimizing solution and any number of possible contingent optimum metrics necessitated whenever a firm makes a suboptimal but irreversible technology choice that precludes the prior cost-minimizing technology profile.

In applying these measures, the ratio of actual abatement costs to the ideal abatement costs for the current technology profile is determined. The market is (laboratory) efficient relative to a given technology profile if the ratio of actual abatement costs to ideal abatement costs is 1. If the ratio of actual abatement costs to ideal abatement costs is less than 1, then insufficient abatement occurs in that period and will have to be made up during later periods using more expensive abatement technologies. Alternatively, if the ratio is larger than 1, then from that period on the industry is committed to a technology profile that entails unnecessarily high abatement costs. In either case, laboratory efficiency will be lowered.[36] To combine the efficiency losses resulting from deviation either above or below the benchmark, the absolute values of deviations of the calculated ratios from 1 are summed in each period. Consequently, the higher the value of the resulting metric, the lower the laboratory efficiency.

Laboratory efficiencies are presented in Table III. In the constant treatment the average cumulative deviations of the contingent optimum metric are 0.572 for the HH market versus 0.138 for the LH market. For the reduction treatment, the value of the metric is 0.802 in the HH market versus 0.175 in the LH market. In both treatments, the efficiencies are significantly lower in the HH markets. The ANOVA analysis rejects the null hypothesis that the mean efficiencies are identical for both the constant and reduction treatments at the 5% level. Consequently, the experimental evidence supports the converse of Hypothesis 1. Looking at the standard deviation of the efficiency metric across the replications, the standard deviation is consistently higher in the HH market for both treatments. For the constant (reduction) treatment, the standard deviation is 0.28 (0.54) for the HH markets and 0.03 (0.08) for the corresponding LH markets. One possible explanation for this is that under the HH scenario more (or, perhaps, more costly) errors are made by traders.

[35] Several readers have commented that this process of redefining the optimum when firms make irreversible "mistakes" biases this measure of efficiency relative to one wherein no such adjustment is made. This is a valid argument, as it permits the candidate efficiency measure to escape its past mistakes in periods after the one in which the mistake was made. However, there is a valid counterargument. Once a firm makes such an irreversible mistake, the attainable optimum is revised, generally involving a different technology profile for the remaining firms. If the efficiency measure continues to be based on the original technology profile, then this efficiency measure could signal efficiency gains for firm choices that represent mistakes vis-à-vis the revised technology profile. So the choice is between a measure of laboratory efficiency that includes the cumulative effect(s) of past mistakes and rewards future mistakes versus one that ignores the effects of past mistakes (in periods after the one in which the mistake is made) but rewards the effects of future correct decisions. (In fact we encountered such anomalies, which led to the construction of the revised metric.) While the choice between these metrics may not be clear to all, we prefer the more forward-looking properties of the latter.

[36] In anticipation of potential conflicting messages from these potential phenomena, efficiency measures parameters were chosen so that it was very unlikely that insufficient abatement could occur when firms were making inefficient (and irreversible) technology choices. In addition, inspection of the experimental results indicated that there was a general tendency for subjects in aggregate to engage in excessive abatement via inefficient technology choice.

TABLE III
Production Efficiency: Contingent Optimum Metric

	High heterogeneity (HH)	Low heterogeneity (LH)
Constant treatment		
Session 1	0.602	0.148
Session 2	0.602	0.162
Session 3	0.315	0.102
Session 4	1.058	0.189
Session 5	0.284	0.089
Average	0.572	0.138
Reduction treatment		
Session 1	1.550	0.0441
Session 2	0.430	0.2337
Session 3	0.177	0.2231
Session 4	1.052	0.1659
Variance	0.288	0.007

Note. To capture efficiency losses resulting from either under investment or overinvestment in a given period, the contingent optimum metric sums the absolute values of deviations from the benchmark in each period. Consequently, the higher value of the metric, the lower the market efficiency.

5. DISCUSSION AND CONCLUSION

The analysis of the laboratory experiments leads to three findings. First, no experimentally determined relationship has been confirmed between heterogeneity and trade volume. Second, some evidence for an inverse relation between price variability and trade volume exists. Third, efficiency is inversely related to the degree of heterogeneity in the experimental market. The experimental results have implications for both the current SO_2 market and for the design of future permit markets.

The investigation of the effects of heterogeneity on trade volume was motivated by the concern that the similarity of Phase I firms might yield low trade volumes leading to efficiency losses [23, 45] and that the separation of the market into two less heterogeneous groups might exacerbate this tendency [45, p. 28]. While the experimental design does not address the latter issue, the data do suggest that a lack of heterogeneity may not have an adverse effect on trade volume, so that the GAO's concerns may have been unfounded.

The experimental results suggest the possibility that price variability inversely affects trade volume, so attempts to increase market activity might focus on reducing price variability. One factor that increases price variability in the current market design is the type of auction used. A discriminatory rather than a uniform price auction produces a range of prices at which permits are traded in the auction. Thus, the signals to the firms may be noisy. The experimental results are suggestive of the GAO position [45, p. 53] that a uniform price auction would increase market activity. Additionally, Joskow et al. [25] detail the evolution of price variability in the SO_2 market. Their work suggests that as the SO_2 market has operated over time and as more information has become available, price variability has decreased [25, p. 673].

The results concerning Hypothesis 3 indicate that market heterogeneity adversely affects efficiency in the experimental market. An inverse relation between heterogeneity and efficiency might be considered counterintuitive, as greater heterogeneity suggests larger potential gains from trade. However, heterogeneity not only increases potential gains from trade, but also potential noise in the market. With the exception of trade volume, the variance across replications is considerably higher in the HH markets, and this supports the existence of noise in the HH markets. The HH markets, on average, perform not much differently from LH markets, but potential deviations from the average are much larger. With larger surpluses suboptimal decisions apparently become more costly, reducing the efficiency of the market. This would be consistent with Laffont and Tirole's [27] concern that postponed trading may lead to resource misallocation. The greater the potential gains from trade, the greater the potential misallocation in more heterogeneous markets when markets are more prone to deviate from efficient performance. The higher variance in market efficiency across replications in the HH markets supports this notion.

In the emissions trading literature, a trade-off is often perceived between the desire to differentiate a market regionally to accommodate differing geographical, meteorological, or administrative conditions and the desire to create permits markets as large as possible to maximize potential gains from trade [2, 38]. In the presence of irreversible production choices, the results presented here would suggest that this trade-off may not actually exist, but that smaller markets may be more desirable on grounds of greater market efficiency as well as the ability to better adjust for local conditions.

REFERENCES

1. F. Allen and D. Gale, Limited market participation and volatility of asset prices, *Amer. Econom. Rev.* **84**, No. 4, 933–955 (1994).
2. S. Atkinson and T. Tietenberg, The empirical properties of two classes of designs for transferable discharge permit markets, *J. Environ. Econom. Manag.* **9**, 101–121 (1982).
3. D. R. Bohi and D. Burtraw, Utility investment behavior and the emission trading market, *Resources and Energy* **14**, Nos. 1 and 2, 129–153 (1992).
4. D. R. Bohi and D. Burtraw, SO_2 allowance trading: How experience and expectations measure up, *The Electricity Journal* **10**, No. 7, 67–75 (1997).
5. J. Brown Kruse and S. Elliot, Strategic Manipulation of Pollution Permit Markets: An Experimental Approach, working paper, Department of Economics, University of Colorado, Boulder, CO (1990).
6. D. Burtraw, The SO_2 emissions trading program: Cost savings without allowance trading, *Contemp. Econom. Policy* **14**, 79–94 (1996).
7. D. Carlson, D. Burtraw, M. Cropper, and K. L. Palmer, "Sulfur Dioxide Control by Electric Utilities: What Are the Gains from Trade?" Discussions paper 98-44, July 1998, Resources for the Future.
8. T. N. Cason, Seller incentive properties of EPA's emission trading auction, *J. Environ. Econom. Manag.* **25**, No. 2, 177–195 (1993).
9. T. N. Cason, An experimental investigation of the seller incentives in the EPA's emission trading auction, *Amer. Econom. Rev.* **85**, No. 4, 905–922 (1995).
10. T. N. Cason, Market masked regulation, *Regulation* **1997**, Summer 14–16 (1997).
11. T. N. Cason and L. Gangadharan, An experimental study of electronic bulletin board trading for emission permits, *J. Regulatory Econom.* **14**, 55–73 (1998).

12. T. N. Cason and C. Plott, EPA's new emission trading mechanism: A laboratory evaluation, *J. Environ. Econom. Manag.* **30**, No. 2, 133–160 (1996).
13. K. Conrad and R. E. Kohn, The U.S. market for SO_2 permits, *Energy Policy* **24**, No. 12, 1051–1055 (1996).
14. T. Crocker, The structuring of atmospheric pollution control systems, *in* "The Economics of Air Pollution," W. W. Norton, New York (1966).
15. J. Dales, "Pollution, Property and Prices," Univ. of Toronto Press, Toronto, Canada (1968).
16. D. D. Davis and C. A. Holt, "Experimental Economics," Princeton Univ. Press, Princeton, NJ.
17. T. W. Epps and M. L. Epps, The stochastic dependence of security price changes and transaction volumes: Implications of the mixture-of-distribution hypothesis, *Econometrica* **44**, 305–321 (1976).
18. R. Franciosi, M. Isaac, D. Pingry, and S. Reynolds, An experimental investigation of the Hahn-Noll revenue neutral auction for emissions licenses, *J. Environ. Econom. Manag.* **24**, 1–24 (1993).
19. R. Franciosi, M. Isaac, D. Pingry, and S. Reynolds, Extension of research into marketable emission permits, working paper, Department of Economics, University of Arizona, Tucson, AZ (1993).
20. R. W. Godby, S. Mestelman, R. A. Muller, and J. D. Welland, Emissions trading with shares and coupons when control over discharges is uncertain, *J. Environ. Econom. Manag.* **32**, 359–381 (1997).
21. R. W. Hahn, Trade-offs in designing markets with multiple objectives, *J. Environ. Econom. Manag.* **3**, 1–12 (1986).
22. G. W. Harrison, M. McKee, and E. E. Rutstrom, Experimental evaluation of institutions of monopoly restraint, *in* "Advances in Behavioral Economics," Vol. 2 (J. Kagel and L. Green, Eds.), Ablex, Norwood, NJ (1990).
23. K. Hausker, The politics and economics of auction design in the market for sulfur dioxide pollution, *J. Policy Analysis Manag.* **11**, No. 4, 553–572 (1992).
24. B. A. Huberman and T. Hogg, Distributed computation as an economic system, *J. Econom. Perspect.* **9**, No. 1, 141–152 (1995).
25. P. L. Joskow, R. Schmalensee, and E. Bailey, The market for sulfur dioxide emissions, *Amer. Econom. Rev.* **88**, No. 4, 669–685 (1998).
26. J. H. Kagel and A. E. Roth (Eds.), "Handbook of Experimental Economics," Princeton Univ. Press, Princeton, NJ (1995).
27. J. Laffont and J. Tirole, Environmental policy, compliance and innovation, *Eur. Econom. Rev.* **38**, Nos. 3 and 4, 555–562 (1994).
28. R. D. Lile, D. R. Bohi, and D. Burtraw, An assessment of the EPA's SO_2 emission allowance trading system, Discussion Paper no. 97-21, Resources for the Future, Washington, DC (1996).
29. S. Mestelman, R. Moir, and R. A. Muller, Emissions trading with shares and coupons: A laboratory test of Canadian proposals, *in* "Economic Experiments in Marketable Emissions Trading" (C. Holt and R. M. Isaac, Eds.), Westview Press, Boulder, CO (1993).
30. D. Montgomery, Markets in licenses and efficient pollution control programs, *J. Econom. Theory* **5**, 395–418 (1972).
31. R. A. Muller and S. Mestelman, Emission trading with shares and coupons: A laboratory experiment, *Energy J.* **15**, No. 2, 185–211 (1994).
32. M. Pagano, Endogenous market thinness and stock price volatility, *Rev. Econom. Stud.* **56**, 269–288 (1989).
33. R. C. Raufer and S. L. Feldman, "Acid Rain and Emissions Trading: Implementing a Market Approach to Pollution Control," Rowman and Littlefield, Totowa, NJ (1987).
34. A. E. Roth (ed.), "Laboratory Experimentation in Economics: Six Points of View," Cambridge Univ. Press, Cambridge, UK (1987).
35. R. Schmalensee, P. L. Joskow, A. D. Ellerman, J. P. Montero, and E. M. Bailey, An interim evaluation of sulfur dioxide emissions trading, *J. Econom. Perspect.* **12**, No. 3, 53–68 (1998).
36. B. Solomon, Five tears of interstate SO_2 trading: Geographic patterns and potential cost savings *The Electricity J.* May, 58–70 (1998).
37. R. N. Stavins, Transaction costs and tradeable permits, *J. Environ. Econom. Manag.* **29**, No. 2, 133–148 (1995).
38. T. H. Tietenberg, "Emissions Trading, An Exercise in Reforming Pollution Policy," Resources for the Future, Washington, DC (1985).
39. U.S. Environmental Protection Agency Acid Rain Division, "1993 EPA Allowance Auction Results," March, Washington, DC (1993).

40. U.S. Environmental Protection Agency Acid Rain Division, "1994 EPA Allowance Auction Results," March, Washington, DC (1994).
41. U.S. Environmental Protection Agency Acid Rain Division, "1995 EPA Allowance Auction Results," Washington, DC (1995).
42. U.S. Environmental Protection Agency Acid Rain Division, "1996 EPA Allowance Auction Results," Washington, DC (1996).
43. U.S. Environmental Protection Agency Acid Rain Division, "1997 EPA Allowance Auction Results," Washington, DC (1997).
44. U.S. Environmental Protection Agency Acid Rain Division, "1998 EPA Allowance Auction Results," Washington, DC (1998).
45. U.S. General Accounting Office, "Air Pollution: Allowance Trading Offers an Opportunity to Reduce Emissions at Less Costs," GAO/RCED-95-30, December, Washington, DC (1994).