Cooperative and Market-Based Solutions to Pollution Abatement Problems

Thesis by

Leslie Rachel Fine

In Partial Fulfillment of the Requirements
for the Degree of
Doctor of Philosophy

California Institute of Technology
Pasadena, California

2001
(Submitted December 1, 2000)

Dedication

To my parents.

Who taught me I can do anything.

To Ed.

Who reminds me of that every day.

Acknowledgements

My most heartfelt thanks to John Ledyard. His valuable advice, numerous suggestions and his ability to contain his frustration at my trouble with concepts that come so naturally to him have made this possible. He never let me take the easy way out - of anything. Despite the many demands on his time, he devoted considerable effort and patience to encourage and help me.

I thank Simon Wilkie for allowing me to invade his office for hours at a time, and his willingness to slog through problems with me. Matt Jackson's intuition for, well, everything has been invaluable. Peter Bossaerts has taught me everything I now understand about financial economics, although I won't hold him responsible.

I have benefited from the help of other faculty, graduate students, and staff of the Division as well. Kim Border has been my LaTeX and math consultant throughout my work. Charlie Plott has been a continual fountain of support and advice. I am also indebted to Susan Davis, Christina Kitchen, Fred Boehmke, Chris Anderson and Valentina Bali for their help, and to the entire graduate student body for their support and for recognizing that education is not a zero-sum game. John Lin and Ian Swett made the experimental work possible by supporting and fixing an unsupportable and unfixable software package.

I am also indebted to Myrna Wooders, John Conley and Leo Hurwicz for their help with Chapter 2. Glen Cass and Anne Scholtz shared their data and intuition for Chapter 4, and for that I am grateful.

Finally, I would like to thank my friends, family, and especially my wonderful husband, Edward Fine, for their love and support. As John Ledyard once advised, "Befriend mathematicians." So, I married one. It was a good call.

Abstract

This work concerns itself with the possibility of solutions, both cooperative and market based, to pollution abatement problems. In particular, we are interested in pollutant emissions in Southern California and possible solutions to the abatement problems enumerated in the 1990 Clean Air Act. A tradable pollution permit program has been implemented to reduce emissions, creating property rights associated with various pollutants.

Before we discuss the performance of market-based solutions to LA's pollution woes, we consider the existence of cooperative solutions. In Chapter 2, we examine pollutant emissions as a transboundary public bad. We show that for a class of environments in which pollution moves in a bi-directional, acyclic manner, there exists a sustainable coalition structure and associated levels of emissions. We do so via a new core concept, one more appropriate to modeling cooperative emissions agreements (and potential defection from them) than the standard definitions.

However, this leaves the question of implementing pollution abatement programs unanswered. While the existence of a cost-effective permit market equilibrium has long been understood, the implementation of such programs has been difficult. The design of Los Angeles' REgional CLean Air Incentives Market (RECLAIM) alleviated some of the implementation problems, and in part exacerbated them. For example, it created two overlapping cycles of permits and two zones of permits for different geographic regions. While these design features create a market that allows some measure of regulatory control, they establish a very difficult trading environment with the potential for inefficiency arising from the transactions costs enumerated above and the illiquidity induced by the myriad assets and relatively few participants in this market.

It was with these concerns in mind that the ACE market (Automated Credit Exchange) was designed. The ACE market utilizes an iterated *combined-value call market* (CV Market). Before discussing the performance of the RECLAIM program in general and the ACE mechanism in particular, we test experimentally whether a portfolio trading mechanism can overcome market illiquidity. Chapter 3 experimentally demonstrates the ability of a portfolio trading mechanism to overcome portfolio rebalancing problems, thereby inducing sufficient liquidity for markets to fully equilibrate.

With experimental evidence in hand, we consider the CV Market's performance in the real world. We find that as the allocation of permits reduces to the level of historical emissions, prices are increasing. As of April of this year, prices are roughly equal to the cost of the Best Available Control Technology (BACT). This took longer than expected, due both to tendencies to mis-report emissions under the old regime, and abatement technology advances encouraged by the program. We also find that the ACE market provides liquidity where needed to encourage long-term planning on behalf of polluting facilities.

Contents

---

[1] **Co-authored with Peter Bossaerts and John Ledyard**

List of Figures

List of Tables

Chapter 1 Introduction

The focus of the following work is the existence of cooperative and market-based solutions to pollution abatement problems. In Chapter 2, we consider pollutant emissions as a transboundary public bad. That is, pollution is considered as a by-product of production, where the produced private good is enjoyed within the region that produces it but the pollution emitted crosses regional boundaries. While local public goods models are generally well-understood, few have considered the possibility of the public goods crossing from one locality to another. In the few models that have, distance between regions and the patterns in which air (and hence, emissions) moves has not been generally considered.

We show that for a class of environments in which pollution moves in a bi-directional, acyclic manner, there exists a sustainable coalition structure and associated levels of emissions. We do so via a new core concept, one more appropriate to modeling cooperative emissions agreements (and potential defection from them) than the standard definitions. The *Stackelberg core* assumes that a defecting coalition acts as a Stackelberg leader, assuming a best response on behalf of the non-defecting parties. This core concept is sequential in nature, reflecting that if a set of regions defects from a treaty it could be quite a while before the other treaty signatories observe his defection and can respond.

The work in Chapter 2 only considers the existence of a cooperative solution to emissions abatement problems. It leaves unanswered the questions of implementing reduced emissions levels once they have been agreed upon. In this thesis, we concern ourselves with pollutant emissions in Southern California and possible solutions to the emissions abatement problems enumerated in the 1990 Clean Air Act. In the case of Southern California, a tradable pollution permit program has been implemented to reduce emissions, creating property rights associated with various pollutants.

The property rights approach to environmental protection has its roots in a 1960 paper by Ronald Coase [8]. He argued that by making property rights explicit and transferable, the market could play a role in both valuing the rights and ensuring that they gravitate to their best use. This is achieved by allowing the market to value these rights.

In 1994, faced with the need to reduce pollutant concentrations in order to come into compliance with the ambient standards set by the Clean Air Act, the Southern California Air Quality Management District (SCAQMD) launched the REgional CLean Air Incentives Market (RECLAIM). This program was adopted to provide facilities with added flexibility in meeting emission reduction requirements and to lower the cost of compliance. The program shifted the burden of identifying appropriate control strategies from the control authority to the polluter. In part, this shift was a necessity, driven by the fact that traditional processes were incapable of identifying enough appropriate technologies to produce sufficiently stringent reductions. As a result of the flexibility inherent to the RECLAIM program, pollution prevention in LA has been given an economic underpinning. All strategies can compete on a level playing field.

Montgomery [27] formally solved the problem of proving the existence of a cost-effective permit market equilibrium in the case of location-specific emissions. Generally, those sources with higher marginal impacts on the environmental target would pay higher prices per unit of emissions. However, location-specific targets are extremely difficult, if not impossible, to implement, as the monitoring costs are prohibitively high. The administrative difficulties associated with the ambient permit system has sent economists searching for a more feasible approach. While these more feasible alternatives may not sustain Montgomery's least-cost allocation, they may represent an improvement over the traditional command-and-control approach.

One such alternative is a zonal permit system. This deals with the spatial dimension of the problem of ambient air quality control by dividing the control area into a grid containing a number of zones. In the most restrictive form of this approach, trades would be allowed within zones, but not between. Less restrictive forms of zonal permit systems allow trades between zones using predefined trading ratios. The zonal approach is appealing as it provides a middle ground between purely emissions-based systems (which have the dual problem of over-controlling distant sources and leading to hot-spots) and the location-specific rules suggested by Montgomery. In the case of RECLAIM, a two-zone system was implemented. Zone one is for coastal polluters, and zone two for inland. Because of the prevailing wind currents, coastal polluters are forbidden to use inland permits but inland polluters can use coastal permits.

Programs such as RECLAIM, however, are fraught with implementation problems. One such problem is that of transactions costs, including the costs of finding an appropriate trading partner, establishing terms of trade, and completing the arrangements. Theory tells us that in the absence of transactions costs, permit markets can reallocate control responsibility such that control is achieved at a minimum cost. At any point in time remaining lower cost control options create trading opportunities, until the minimum cost-effective allocation has been reached. However, when faced with significant control costs, permit markets may not be cost-effective (see Stavins [36] for an excellent discussion of the role of transactions costs).

The design of the RECLAIM program in part alleviated some of the potential for large transactions costs, and in part exacerbated them. In order to smooth trading behavior, two overlapping cycles of permits were issued. A cycle one permit for a given year is an effective credit between January 1st and December 31st of that year, whereas a cycle two permit is effective July 1st to June 30th of the following year. Additionally, as discussed above, separate permits were issued for each of two zones in the South Coast Air Basin. Any given trading credit is only valid for the year, cycle, zone, and pollutant for which it is designated. Initially, permits were issued for 1994 - 2010, which implies a total of 136 permits available in the initial market.

While these design features create a market that allows some measure of regulatory control, it establishes a very difficult trading environment with the potential for inefficiency arising from the transactions costs enumerated above and the illiquidity induced by the myriad assets and relatively few participants in this market. It was with these concerns in mind that the ACE market (Automated Credit Exchange) was designed. As we will discuss further in Chapter 4, the ACE market utilizes an iterated *combined-value call market* (CV Market). The market mechanism used by ACE is based on research conducted at Caltech in 1993 (Ishikida et al. [19]). ACE allows bidders to submit contingent orders (I want A if and only if I can also secure B), matches revealed surplus, and calculates prices that leaves bidders at least as well off for having participated as if they had not.

The ACE market is the first CV Market to be implemented in a tradable pollution permit market. While Ishikida et al. ran extensive testbedding of their mechanism in context of the RECLAIM market, a more general comparison of the performance of combined-value markets to traditional, parallel markets is in order.

Therefore, before discussing the performance of the RECLAIM program in general and the ACE mechanism in particular, in Chapter 3 we test experimentally a portfolio trading mechanism that overcomes market illiquidity. This work grew out of the key finding in work by Bossaserts, Kleiman and Plott (BKP) [4], that thin markets (markets with only a few participants) fail to completely equilibrate. That is, liquidity dries up before full equilibrium is reached. BKP conjectured that the problem with thin markets could have been caused by difficulties that traders faced when rebalancing their portfolios. Participants attempted to trade up to more desirable portfolios, yet were frustrated when unable to coordinate trade in the component permits. The risk that orders would only partially be filled made traders hesitant to engage further in the markets' activity.

In pollution emissions markets, the objects of interest for market participants are generally streams of emissions permits, not individual ones. Therefore, it is surprising that trading mechanisms are generally organized as a set of parallel markets, where it is impossible to submit an order that depends on events in other markets. The risk of seeing one's portfolio only partially re-balanced could be vastly or entirely diminished if one were given the opportunity to submit contingent cross-market orders, whose execution depends on what happens elsewhere in the marketplace.

The contribution of Chapter 3 is to explore the ability of a portfolio trading mechanism to overcome the aforementioned portfolio rebalancing problems, thereby inducing sufficient liquidity for markets to fully equilibrate. In particular, we study whether our portfolio trading mechanism leads markets to full equilibrium in situations where BKP discovered that equilibration generally becomes stuck along the way because volume dries up.

We measure the distance from equilibrium by means of two well-known asset pricing models, namely, the Capital Asset Pricing Model (CAPM) and the complete-markets Arrow-Debreu model. We complement graphical evidence of equilibration with formal tests. The tests give the probability that the actual evolution of our measures of distance from equilibrium could have emerged merely by chance. That is, we compute the probability that our results would have come about had prices been just random walks. Under a random walk, no economic forces (pressures of portfolio demand against a fixed supply) are at work. The only economic basis for a random walk is simple speculation, which would indeed make prices unpredictable.

We confirm our conjecture that a portfolio trading mechanism induces full equilibration. Equilibrium generally emerges well before volume decreases, indicating that any reduction in liquidity must be attributed to exhaustion of gains from trade. We record these results in an environment where BKP observed that markets did not fully equilibrate because liquidity dried up.

After exploring the impact of a combined-value mechanism on an illiquid market in an experimental setting, we turn to an empirical analysis of the RECLAIM market in general and the ACE mechanism in particular. Do our experimental results, indicating the ability of a combined-value market to overcome illiquidity, hold up in the real world? We find that as the supply of RECLAIM Trading Credits (RTCs) aligns with the historical level of emissions in the LA Basin, that prices of RTCs increase. However, these prices have historically fallen quite short of those predicted by economists before the market's inception. While this is in part due to mis-reporting of abatement costs on behalf of polluting firms under the previous command-and-control regime, it also reflects the increased abatement efforts that the flexibility of the RECLAIM program encourages. As the allocation of RTCs reduces to historical reported emissions, the prices of RTCs have been similarly priced to the cost of the Best Available Control Technology (BACT).

We also find that, while the long-term market does seem thinly traded, the combined-value market mechanism used in the ACE market overcomes this illiquidity, as predicted in Chapter 3. The combined-value design allows environmental engineers to plan for long-term production and emissions and to do so with more security as to the value of their investment. However, the CV capabilities of ACE are rarely used by short-term traders. This is because short-term purchases and sales are generally to cover anomalies in the production plan. However, since a single-asset bid is simply a degenerate form of combined-value bid, imposing the combined-value structure on the short-term market certainly does no harm. Indeed, the additional liquidity provided by those few traders who trade in bundles including both short-term and long-term RTCs provide an important bridge between the two markets, improving liquidity in both.

Chapter 2 Cooperative Solutions in the Presence of Inter-Regional Emissions

2.1 Introduction

In the traditional literature on local public goods, it is assumed that production of a public good (or bad) in a given region is enjoyed only by the inhabitants of that region. While an argument can be made for this assumption in the case of many traditional local public goods (swimming pools, new roads), this is certainly not always the case. This assumption is notably violated in the case of pollution. A region involved in production of a good may produce pollution as a by-product. While the item being produced benefits only the inhabitants of the region, the pollution knows no borders and will effect the region's neighbors as well. This type of *spillover* is explicitly ruled out by Tiebout's 1956, "A Pure Theory of Local Expenditures" [37], in which he demonstrates the existence of an equilibrium in a local public goods environment.

Two recent papers have explored this problem. In Chander and Tulkens [6], pollution is treated as a transboundary public bad. They show that the core of the game generated by these externalities is nonempty, doing so in a constructive manner by suggesting a mechanism which implements a solution in the core. However, the problem is considered in a fairly simple environment. They assume that pollution travels uniformly, and that the level of ambient air quality for all regions is the same, regardless of location. This fails to account for distance from one's neighbors, prevailing winds, water currents, and so on.

Conley and Dix [9] examine the problem of spillovers and its impact on the level of public goods provision and on region size. They consider the optimal and Nash solutions and explore the comparative statics as the level of spillovers changes. While they do allow more flexibility as to how pollution travels, they leave the question of sustainable cooperative solutions to this problem unanswered.

In this paper, we explore questions of the existence of a cooperative solution to production with public by-products that spill over into neighboring regions. Unlike Chander and Tulkens, we allow pollution to spread non-uniformly between regions. We do, however, restrict pollution to movement patterns to those representable by a tree on the set of regions in the economy. This restriction still includes a wide range of environments in which pollution spillover is a problem. We will discuss this restriction further in Section 2.2.1. Unlike other examinations of the nature of Nash and cooperative solutions to local public goods problems with spillovers, we explicitly treat the physical locations of the regions in relation to one another.

The choice of the appropriate core concept in this environment is a source of debate. After considering several alternative definitions (the alpha and gamma cores in particular), we define a new, more appropriate core concept. The *Stackelberg core* assumes that a defecting coalition acts as a Stackelberg leader would, as a first mover. The other players, upon observing the chosen actions of the defectors, best respond. This best response is factored into the defecting coalition's strategy. This core concept models behavior seen when environmental treaties are broken. The sequential nature of this strategy assumption reflects the non-immediacy of the effects of the defection. That is, if a group decides to pollute above a level indicated in a treaty, it could be quite a while before the fellow signatories can determine this for certain and take retaliatory action.

With this assumption we are able to show that there exists a stable coalition structure and therefore the core of the cooperative game is nonempty. That is, there exists no subset of the group that can defect and leave all the defectors just as well off with at least one better off. Therefore, a cooperative solution exists to the pollution problem. We also show that the Nash solution to this problem is consistently to overproduce and to overpollute. These properties rely crucially on assumptions about how pollution travels. The importance of the assumption of a tree structure is demonstrated in two counterexamples showing the emptiness of the core when this property is lost, and has implications for when a cooperative agreement may and may not be tenable.

The remainder of this chapter is organized as follows. Section 2.2 details the model. Section 2.3 solves the omniscient social planners problem, and Section 2.4 describes the non-cooperative game induced by the spillovers. Section 2.5 describes the Stackelberg core and presents the main result. Section 2.6 presents a counterexample demonstrating the importance of the tree structure. Finally, Section 2.7 presents possible extensions to the model to include non-tree structures and Section 2.8 concludes.

2.2 Model

Consider a traditional local public goods economy with one private good and one public good (pollutant emissions). We have N regions, denoted $i = 1, ..., n$. Each region produces a private good, denoted by x_i. Additionally, each region emits pollution e_i as a by-product of non-increasing returns-to-scale technology $x_i = g_i(e_i)$, strictly concave and differentiable. All other factors of production are subsumed into the $g_i(.)$ function. Moreover, there exists a $e_i^0 > 0$ such that

$$\frac{\partial x_i}{\partial e_i} = \begin{cases} > 0 \text{ if } e_i < e_i^0 \\ = 0 \text{ if } e_i \geq e_i^0 \\ = \infty \text{ if } e_i = 0 \end{cases} \tag{2.1}$$

The preferences of the residents of each region are characterized by a representative utility function of the form $u_i(x_i, a_i) = x_i + v_i(a_i)$. The a_i terms represent ambient air quality, a function of $\{e_i\}_{i=1}^{N}$ and the manner in which pollution flows between the N regions. All regions have identical utility functions, and vary only in their $a_i = -\sum_{j=1}^{N} f_{ij}(e_j)$. The $v_i(.)$ portion of the utility function is concave and differentiable, with $\frac{\partial v_i}{\partial a_i} > 0$ for all $a_i \leq 0$. The f_{ij} terms are discussed below.

2.2.1 Ambient Air Quality, Spillover Coefficients and Trees

The f_{ij} coefficients can be represented by a transformation matrix F (which describes all pollution flows between all members of the economy) with structure:

$$\begin{bmatrix} f_{11} & f_{12} & \cdots & f_{1n} \\ f_{21} & \cdot & \cdot & f_{2n} \\ \cdot & \cdot & \cdot & \cdot \\ f_{n1} & \cdot & f_{nn-1} & f_{nn} \end{bmatrix} \tag{2.2}$$

Before we can continue, we need a brief review of some graph-theoretic terms. A *graph* G on N is a set of unordered pairs of distinct elements of N. A graph is a *tree* if two distinct elements are linked by a unique path, as earlier discussed.

The structure flow of emissions from one region to another and its effect on ambient air quality is critical to understanding the possibility of an equilibrium in a spillovers environment. For example, if emissions travel in a cyclical manner and the polluter never has to be a victim of his own emissions, then this environment reduces to one similar to Laffont's Garbage Game [22] (which has no core). However, if the manner in which pollution spreads obeys that dictated by a tree structure, there are possibilities for the existence of a stable coalition structures. For the majority of this paper, we will consider only tree structures. In the extensions, we will consider a more general class of admissible acyclic graph configurations, particularly those where some regions are isolated.

What makes the f_{ij}s represent a tree? A tree is a connected graph having no cycles. If we have a graph on n vertices, it is a tree if and only if it is connected and has $n-1$ edges. A *neighborhood matrix* is an $n * n$ matrix where if nodes i and j share an edge of a graph or if $i = j$, the ijth entry is a one, zero otherwise. This describes proximity in a graph. To get from the above transformation matrix to a neighborhood matrix, place a one in the ijth entry of the adjacency matrix if and only if there is a non-zero entry in the ijth place of the transformation matrix. The neighborhood matrix of a tree must have exactly $2(n-1) + n = 3n - 2$ ones (that is, $n-1$ edges represented twice, plus n entries on the diagonal). This condition is necessary, but not sufficient for the associated graph to be a tree. In other words, just being symmetric (a nondirected graph), with ones on the diagonal and having $3n-2$ one's doesn't guarantee that it's a tree because it also has to be connected. We will discuss in the extensions the necessity (or lack thereof) of a connected structure.

So, environments in which we have a tree structure on pollution have the following properties. If f_{ij} is positive, so is f_{ji}. Furthermore, no more than $3n-2$ of the entries can be positive, including the f_{ii}s.

In simplest terms, the f_{ij}s could be a distance measure. However, our environment reflects a more general set of circumstances than regions whose spillovers are solely a function of distance. Imagine three regions in a line. Furthermore, imagine regions one and three are on hilltops, while region two is down in the valley between them. Regions one and three may very well have a greater impact on each other's pollution than does region two. The transformation matrix we use allows for this situation. As long as regions one and three are "pollution neighbors," then they are admitted in our model.

2.3 The Social Planner's Problem

For efficiency in the overall economy, the omniscient social planner maximizes the sum of the utility functions for each region.

$$\max_{\{x_i,e_i\}_{i=1}^N} \; x_i + v_i(a_i) \tag{2.3}$$

$$\text{subject to } g_i(e_i) \geq x_i \tag{2.4}$$

$$\text{and } a_i = -\sum_{j=1}^{N} f_{ij}(e_j) \tag{2.5}$$

Because of our concavity assumptions and the additively separable utility function, we can plug the constraints into the utility function and solve the following unconstrained maximization problem:

$$\max_{\{e_i\}_{i=1}^N} g_i(e_i) + v_i(-\sum_{j=1}^{N} f_{ij}(e_j)) \tag{2.6}$$

The solution to this maximization problem yields the Pareto Efficient vector of pollution levels, denoted $(e_1^*, ..., e_n^*)$. This vector satisfies the following equality, for all i:

$$g_i^{'}(e_i) = \sum_{j=1}^{N} v_j^{'}(-f_{ji}(e_i)) * f_{ji} \tag{2.7}$$

Each region sets their marginal product of pollution equal to the sum of the marginal utilities with respect to that region's pollution level *weighted by the size of the spillover effect.* This is a similar efficiency condition to the classical Samuelson Condition for provision of a public good [29]. The Samuelson condition, which keeps all regions isolated from one another, tells us that the ratio of the marginal rate of substitution between the public good and the private good should be equal to their marginal rate of transformation. What we have here is more general. This condition states that the marginal rate of substitution *multiplied by the impact that your emissions have on all regions' ambient air quality* should be equal to the marginal rate of transformation. As in Chander and Tulkens, for all Pareto Efficient states of this economy, the pollution vector is unique. This follows directly from the strict concavity of the utility function.

Note the comparative static of how the Pareto Efficient level of emissions changes with the magnitude of the spillover parameter. As the strength of the spillover effect increases, the Pareto Efficient level of emissions decreases. So, a region on the edge of the society would be allowed to pollute at a higher rate than one in a dense area.

2.4 The Non-Cooperative Game

We now turn to the issue of the noncooperative game induced by the emission spillovers. Rather than considering the level at which a social planner would set emissions for all localities, we now consider what a regional development entrepreneur would decide upon for his individual region. To define the noncooperative game, we must first define the strategy set available to players. The strategy set for player i is to choose a feasible level of output and emissions from:

$$T_i = \{(x_i, e_i) | 0 \leq e_i; 0 \leq x_i \leq g_i(e_i^0)\} \tag{2.8}$$

Each region takes the production of emissions in all other regions as given, and maximizes his own utility with respect to his production and hence emissions. He solves

$$\max_{e_i} g_i(e_i) + v_i(-\sum_{j=1}^{N} f_{ij}(e_j)) \tag{2.9}$$

$$\text{where } e_j = \overline{e}_j \text{ for all } j \in N, j \neq i \tag{2.10}$$

Each region solves this problem by setting $g_i'(e_i) = f_{ii}v_i'(a_i)$. Denote this solution by $\overline{e}_i$. Note that $\overline{e}_i \geq e_i^*$ for all $i \in N$. Only in the degenerate case, where $f_{ij} = 0$ for all $i \neq j$, is the Nash point Pareto efficient. That is, only if all regions are so isolated that no one region's emissions has an effect on another's ambient air quality ($f_{ij} = 0$ for all i, j in N), then the purely selfish action is socially optimal.

When a coalition S agrees to play together, their strategy set expands to the joint strategy space. This is represented by

$$T(S) = \{(x_i, e_i)_{i \in S} | 0 \leq e_i, \forall i \in S; 0 \leq \sum_{i \in S} x_i \leq \sum_{i \in S} g_i(e_i^0)\} \tag{2.11}$$

The joint strategy choice $[(x_1, e_1), ..., (x_n, e_n)] \in T(N)$ induces a feasible state (x, e, a) if $a_i = -\sum f_{ij}(e_j)$ for all i. Because the actions of the players outside of any given coalition S affects the feasibility of any strategy for that coalition, rather than simply defining a $[T, u]$ game, we define the noncooperative game $[N, T, u]$.

2.5 The Stackelberg Core

To study the properties of the core, we must understand the value of a coalition in the $[N, T, u]$ game. To do this in an environment with spillovers, we must make assumptions about the responsive action in which the members of N/S engage when coalition S forms. The α-core assumes that the players in N/S adopt a minimax strategy, responding to S's actions in the way that is worst for S. This is the solution concept employed in Laffont's version of the "Garbage Game" and in Maler's "Acid Rain Game" [24]. However, unlike these two examples, in our environment when a region pollutes it is a victim of its own pollution.

Therefore, in this model if region $j \in N/S$ plays a minimax strategy by increasing emissions, the action could hurt j more than it would $i \in S$, and therefore the α-core is inappropriate as a solution concept in our model.

Chander and Tulkens introduced the γ-core as a suitable solution concept for the spillovers environment. In this approach when a coalition deviates it plays fully cooperatively (maximizing the sum of the utility) and looks at the resulting equilibrium from its actions. This coalition S assumes that when it deviates, the members of N/S break up into single actors all playing a Nash strategy. In a spillovers environment this is more realistic than the minimax behavior assumed in the α-core. The joint strategy choice can be seen as a Nash equilibrium in which coalition S acts as one individual playing against the individual members of N/S. However, this constrains the group to play a fully cooperative solution (efficient with respect to the coalition) within the group, and assumes a simultaneous move on behalf of the members of N/S. These are not realistic constraints in the model we are exploring. Rather, we would like the defecting coalition to be able to adopt any strategy it wants, including one where some members of the coalition play differently than others, as will become important below.

So, we introduce a new core concept, the *Stackelberg core.* We assume that when a coalition deviates it plays like a Stackelberg leader. It can choose any strategy it wants, and when it calculates the value of this strategy it takes into account the best response that the members of N/S will have to this deviation *once they have observed it.* We can think of regions as playing a Cournot-like game against one another. The following example demonstrates the meaning of this core concept relative to those of the α and γ core when $N = \{1, 2\}$.[1]

Assume we have two identical regions with utilities represented by:

---

[1] This is quite different from the consistent set discussed in Chwe [7]. If an allocation a is in the consistent set, then it satisfies the indirect dominance property. That is, if a coalition S defects to point a, and some other coalition then defects from a to b (b directly dominating a), then this series of defections from defections will eventually land the group at allocation c, which is less-prefered than a. However, in the Stackelberg core, if a coalition S defects to a, the coalition members assume a best response on behalf of the individuals in N/S. These best responses will land the players all at allocation b. But, the coalition S predicted this move, and only moved to a because they WANTED to wind up at b.

$$u_1 = 3e_1 + (-\frac{1}{4})(-\frac{3}{4}e_1 - \frac{1}{4}e_2)^2 \tag{2.12}$$
$$u_2 = 3e_2 + (-\frac{1}{2})(-\frac{1}{4}e_1 - \frac{3}{4}e_2)^2 \tag{2.13}$$

If these two regions were playing cooperatively, they would choose $e_1^* = e_2^* = 6$, corresponding to a utility level of 18 for the two together. The α-core tells us that if region 1 defected from this, region 2 would do the worst thing possible for region 1 and set $e_2 = e_2^0$. As long as this threshold is sufficiently high, this will reduce utility to zero. In the γ-core, both players would employ a Nash strategy. This would yield $\overline{e}_1 = \frac{32-e_2}{3} = 8, \overline{e}_2 = \frac{32-e_1}{3} = 8$, yielding a utility level of 8 for each player. In the Stackelberg core, player 1 deviates assuming a Nash reaction on the part of player 2. So, he maximizes his utility function, assuming $e_2^S = \frac{32-e_1}{3}$. He maximizes

$$u_1 = 3e_1 + (-\frac{1}{4})(-\frac{3}{4}e_1 - \frac{1}{4}(\frac{32 - e_1}{3}))^2 \tag{2.14}$$

This yields $e_1^S = \frac{19}{2}$, $e_2^S = \frac{15}{2}$, with corresponding utility levels of $u_1 = \frac{33}{4}$, $u_2 = \frac{26}{4}$. Now that we know what each player can achieve under each core assumption, what does the core look like for each? As shown in Figure 2.1, if the α-core is used, the entire Pareto efficient frontier is in the core. For the γ-core, the highlighted section which makes each player at least as well off as they are at $u_i(e_i^N)$ is in the core. For the Stackelberg core, the set of points that make each player as well off as they are at $u_i(e_i^S)$ when i is the Stackelberg leader is in the core. Note that this core is smaller than either the α-core or the γ-core and is contained in both.

Once we have defined the appropriate responsive behavior assumptions, we can express the characteristic function of this game as:

$$V(S) = \max_{(x_i,e_i)_{i\in S}} \sum_{i\in S}[x_i + v_i(a_i(e))] \tag{2.15}$$
$$\text{subject to} \sum_{i\in S} x_i \leq \sum_{i\in S} g_i(e_i) \tag{2.16}$$
$$a_i = -\sum_{j\in N} f_{ij}(e_j) \text{ for all } i \in S \tag{2.17}$$

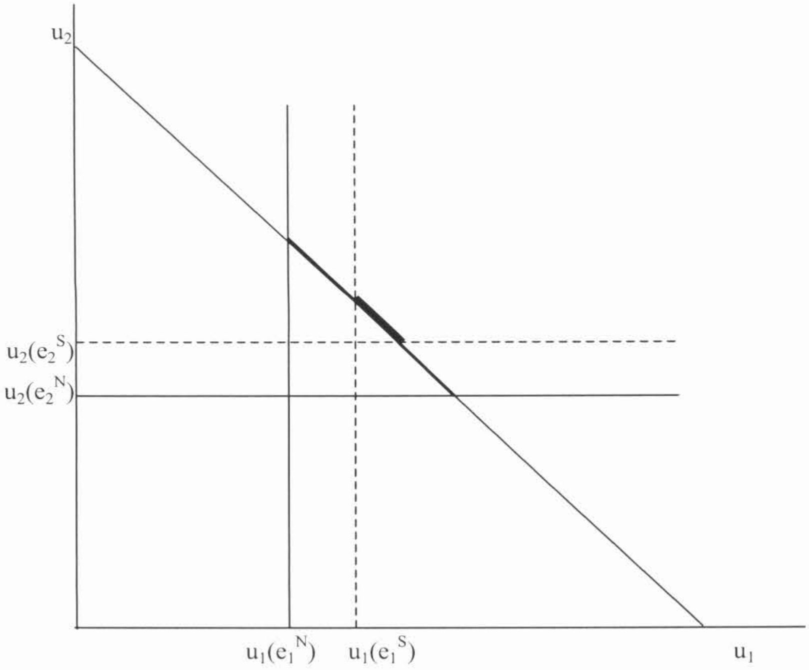

Figure 2.1: The α core, γ core and Stackelberg core

where all $j \in N/S$ maximize $x_j + v_j(a_j(e))$ subject to $x_j \leq g_j(e_j)$ and $a_j = -\sum_{i \in N} f_{ji}(e_i)$. Using this characteristic function, we can evaluate every strategy that a coalition may consider in terms of the aggregate payoff to that coalition.

A strategy of a coalition N is said to belong to the core of the cooperative $[N, V]$ game if the payoff it yields for each coalition $S \subset N$ is bigger than the payoff $V(S)$. We will now establish that there exist strategies that belong to the core of this game. The intuition of the nonemptiness of the core is as follows. We use a result by Demange [11] to show that the connected core is nonempty. That is, if we only allow coalitions that are representable as a G-connected subsets of a tree, there exists a connected coalition structure such that there is then there is no coalition that can improve on the value generated by it. We then show that if a connected coalition is not profitable, then neither is any non-connected coalition.

In order to appeal to Demange's theorem, we need a little more graph-theory. A subset

S of N is *G-connected* if the path between the two points of S is contained in S. So, this restricts us to subsets S that are connected, which in our environment means that we consider coalitions that consist of regions next to each other. For example, if we are considering North America, Mexico and Canada are not G-connected. However, if we add in the United States, these three regions make up a G-connected set. Lastly, a G-connected set is said to be *G-stable* if it is not blocked by any G-connected set. G-stability can be shown to be equivalent to the nonemptiness of the core when only G-connected subsets are allowed to consider defection.

Theorem 1 *If regions and their spillovers can be represented as a tree on N, the feasible set for any coalition S is compact and utility is continuous for all players then the Stackelberg core of the $[N, V]$ game is nonempty.*

Proof First we demonstrate that if only connected coalitions are considered, then the core of a general $[N, V]$ game is nonempty. Then we show that if a connected coalition cannot form profitably, then neither can any non-connected coalition.

1. By Demange [11], we know that if the feasible set for any coalition S is compact and utility is continuous for all players and G is a tree on N, then there exists a G-stable coalition structure. This implies that the connected core of the $[N, V]$ game is nonempty. The idea of the proof is that the existence of a G-stable coalition structure can be shown to be equivalent to the non-emptiness of the core a general cooperative game. The game is superadditive and its essential coalitions are all G-connected, and Demange then uses Scarf's theorem to show that these properties guarantee the nonemptiness of the connected core.

2. Next, we want to show that the nonemptiness of the connected core implies the nonemptiness of the core. We do this by showing that for any non-connected coalition S, the smallest connected coalition containing S can always do at least as well. But, that coalition is a connected coalition and can be expressed as a G-connected subset of a tree on N. Therefore, it cannot improve on what it receives under $V(N)$.

3. We know that a coalition S playing strategy e_S^C (that is, emissions levels that are

cooperative within the coalition) faces utility level $\sum_{i \in S} u_i(e_S^C, e_{N/S}^N)$. If this coalition is connected, there is no strategy profile e_S^C that makes the players in S better off than were they to play the fully cooperative solution. But, what if coalition S is not connected? Define T to be smallest set of players such that $\{S \cup T\}$ is a connected coalition. We know that at the very least $\{S \cup T\}$ could form a coalition in which all $j \in T$ continues to play the Nash strategy and all $i \in S$ play cooperatively. Therefore,

$$V(S \cup T) \geq V(S) + \sum_{j \in T} u_i(e_S^C, e_{N/S}^N) \tag{2.18}$$

So, connected coalition $S \cup T$ can always do at least as well as the non-connected coalition S could on its own. Note that this is the worst case scenario when $j \in T$ join. That is, since S could form profitably, the members of S could compensate the members of T by ϵ to change their output and make $V(S \cup T) > V(S) + \sum_{j \in T} u_i(e_S^C, e_{N/S}^N)$. But, we know that the connected core of this game is nonempty. Therefore, the core is nonempty.

□

2.6 Counterexamples Demonstrating the Importance of the Tree Structure

As emphasized earlier, the reciprocal and non-cyclical nature of the emissions spillovers is critical to our result. We require that regions and their spillovers are representable as a tree on the set N. This does not admit purely unidirectional pollution flows. The following two counterexamples demonstrate this fact, and show that there exists a stable coalition structure when we modify the neighborhood matrix to represent a tree.

2.6.1 Purely Unidirectional Pollution Flows

The following is a simple example of the emptiness of the core when we don't abide by the nondirected aspect of the tree structure. Imagine 3 regions arranged in a triangle, each with utility

$$u_i = x_i + (1 - \frac{1}{4}(-f_{i1} * e_1 - f_{i2} * e_2 - f_{i3} * e_3)^2) \tag{2.19}$$

and constant returns-to-scale production $x_i = 2e_i$. Furthermore, imagine a prevailing wind, such that pollution only travels counterclockwise around the circle (from 1 to 2, from 2 to 3, from 3 to 1) and that you only feel pollution effects from your closest neighbor. This example corresponds to a regional relationship with air currents that do not change seasonally. Specifically, consider the following pollution transformation matrix:

$$\begin{bmatrix} 1 & 0 & .6 \\ .1 & 1 & 0 \\ 0 & .4 & 1 \end{bmatrix} \tag{2.20}$$

The Pareto efficient level of pollution is $e_1 \approx 3.302, e_2 \approx 2.95$ and $e_3 \approx .616$. The corresponding utility levels are $u_1 \approx 4.23, e_2 \approx 4.21$ and $e_3 \approx 1.42$. However, in this situation of unidirectional flows, player 1 would like to be in a coalition with player 3, yielding utilities $u_1 \approx 4.81, e_2 \approx 4.21$ and $e_3 \approx .67$. Meanwhile, player 2 would like to be in a coalition with player 1, yielding $u_1 \approx 1.97, e_2 \approx 4.57$ and $e_3 \approx 1.97$. Finally, player 3 would like to be in a coalition with player 2, with utilities $u_1 \approx 1.25, e_2 \approx 3.94$ and $e_3 \approx 3.25$. There is no stable coalition structure.

Now, rather than having a unidirectional transmission pattern, imagine that pollution flows from player 1 towards both players 2 and 3 in one season, and then there is a shift, transmitting emissions from player 2 and 3 to player 1. This is the case in the below transformation matrix.

$$\begin{bmatrix} 1 & .3 & .6 \\ .1 & 1 & 0 \\ .4 & 0 & 1 \end{bmatrix} \tag{2.21}$$

In this case, both players 2 and 3 would like to collude with one another, yielding utility levels $u_1 \approx -1.9, e_2 \approx 4.89$ and $e_3 \approx 4.56$.

2.6.2 Acid Rain

We can also construct a counterexample corresponding to the situation of acid rain.[2]

Let pollution transmissions be represented by:

$$\begin{bmatrix} .9 & 0 & 0 \\ .4 & 1 & 0 \\ 0 & 0.04 & 1 \end{bmatrix} \tag{2.22}$$

One can think of player one as Minneapolis, two as Chicago, and three as New York. Once again, this non-tree structure leaves us without a stable coalition structure. Minnesota would like to be in a coalition with New York, Chicago with Minnesota, and New York with Chicago.

However, if we slightly perturb the transmission so that there is feedback from Chicago to Minneapolis and from New York to Chicago, the scenario changes.

$$\begin{bmatrix} .9 & .04 & 0 \\ .4 & 1 & .5 \\ 0 & 0.04 & 1 \end{bmatrix} \tag{2.23}$$

Now, New York and Minnesota would both like to collude with one another and leave Chicago out of the deal entirely.

---

[2] When the level of spillovers get very small between two regions, there is still potential for stable coalition formation. However, they are strange coalitions. For example, in this game, at small levels of flowbacks (from 3 to 2, and 2 to 1), both players 1 and 3 would like the coalition (Chicago, NY) to form. However, player 2 would prefer a coalition involving only players 1 and 3, leaving him out entirely. This makes intuitive sense. Since the spillovers are far stronger from player 1 to player 2 than they are from player 3 to player 2, and the spillovers are also quite small from player 2 to player 1, the 1,3 coalition makes player 2 very happy. He has to cut back very little to please player 1, and in return gets a major emissions reduction. however, both players 1 and 3 would prefer that the coalition 2,3 happens. Why? Player 1 doesn't have to cut back at all, and player 3 gets a major benefit for a small cutback. This anomaly will be explored further another paper.

2.7 Extensions to the Model

2.7.1 Non-Cyclic, Non-Tree Structures

As we have discussed extensively, the tree structure induced by the pollution transmission patterns is essential for our demonstration of the existence of a cooperative solution. However, there is an extension to another class of emissions configurations as well. Imagine regions in our economy that are so remote that their emissions effects the ambient air quality of no other region, nor are they affected by anyone else. In the extreme, as discussed in Section 2.4, when no one region is affected by another, the Nash level of pollution is the Pareto efficient level. However, in a more moderate case, when some regions are isolated, they will never be part of a coalition structure in any meaningful way (any way that changes their behavior).

Proposition 2 *If regions and their spillovers can be represented as a non-cyclical, nondirected graph on* N*, the feasible set for any coalition* S *is compact and utility is continuous for all players, then the Stackelberg core of the* $[N, V]$ *game is nonempty.*

2.7.2 Intermediate Preferences and Tiebout Equilibrium

Furthermore, there are many transformation matrices that, in our model, exhibit the property of *intermediate preferences.* That is, if an agent is between two others on the tree, whenever the latter agree on their ranking of alternatives a and b, so does the former. More precisely, intermediate preferences imply that there is a tree on N such that for any alternatives a and b, the sets $\{i \in N, u_i(a) > u_i(b)\}$ and $\{i \in N, u_i(a) \geq u_i(b)\}$ are G-connected.

Proposition 3 *If regions and their spillovers can be represented as a non-cyclical, nondirected graph on* N*, the feasible set for any coalition* S *is compact, utility is continuous for all players, the* $[N, V]$ *game exhibits increasing power of coalitions, and we have intermediate preferences, then the Stackelberg core of the* $[N, V]$ *game is nonempty and is a Tiebout equilibrium.*

2.8 Conclusion

We have shown that in a simple model of local public goods with spillovers, it is possible to find a coalition structure and an allocation for all coalitions in the economy such that no subset $S \subset N$ can do better for its members. While the type of public good we have in mind is pollution emissions, this model could extend to other public goods produced and primarily enjoyed in a given region, but with its effects felt in neighboring regions. For example, funding for the arts or building an airport. We have also shown that without the reciprocal feature of spillovers assumed here, it is possible that there is no sustainable cooperative solution. This observation has profound policy implications. For example, there may not be a sustainable cooperative solution to issues such as acid rain traveling unidirectionally. However, in regions where pollution spreads in a manner similar to a distance function, there is hope. For example, local ambient air quality problems may be solvable.

Applying the notion of the Stackelberg core has natural interpretations in the realm of cooperative environmental solutions. In particular, imagine a proposed treaty. When a region considers whether to ratify the treaty, it knows that if it doesn't it will wind up as a Stackelberg leader in pollution emissions. So, if it ratifies, it is because it cannot do any better for itself by defecting. Therefore, the ratification of the treaty as a signal of intent is the extent of its usefulness.

Chapter 3 Inducing Liquidity In Thin Financial Markets Through Combined-Value Trading Mechanisms[1]

3.1 Introduction

The key finding in a 1998 paper by Bossaerts, Klieman and Plott (BKP) [4] is that thin financial markets (markets with only a few participants) fail to completely equilibrate. Liquidity dries up before full equilibrium is reached. The observations in BKP contrast with those in an earlier paper by Bossaerts and Plott (BP) [5], where thick financial markets were found to fully equilibrate. It was conjectured in BKP that the problem with thin financial markets could have been caused by difficulties that traders faced when rebalancing their portfolios. Participants attempted to trade up to more desirable portfolios, yet were frustrated when unable to coordinate trade in the component securities. The risk that orders would only partially be filled made traders hesitant to engage further in the markets' activity. Whence the reduction in liquidity as the end of trading approached. Notice that this conjecture links nicely with asset pricing theory, which posits that investors are merely interested in (payoffs on) portfolios of securities, rather than the component securities. The value of an individual security is determined solely by its contribution to the risk and return of a portfolio. Beyond this, the risk and return of the security itself are irrelevant.

The objects of interest for market participants appear to be portfolios, and not individual securities. If so, it is surprising that financial markets are generally organized as a set of parallel markets, where it is impossible to submit an order that depends on events in other markets. The experimental financial markets in BKP were also set up as a system of parallel markets in individual securities. The risk of seeing one's portfolio only partially re-balanced could be vastly or entirely diminished if one were given the opportunity to submit contingent

---

[1] **Co-authored with Peter Bossaerts and John Ledyard**

cross-market orders, whose execution depends on what happens elsewhere in the marketplace. In the extreme, this calls for a drastic change in the trading mechanism of the marketplace towards a system where portfolios can be traded directly.

The contribution of this paper is to explore the ability of a portfolio trading mechanism to overcome the aforementioned portfolio rebalancing problems, thereby inducing sufficient liquidity for markets to fully equilibrate. In particular, we study whether our portfolio trading mechanism leads markets to full equilibrium in situations where BKP discovered that equilibration generally becomes stuck along the way because volume dries up.

We measure the distance from equilibrium by means of two well-known asset pricing models, namely, the Capital Asset Pricing Model (CAPM) and the complete-markets Arrow-Debreu model. The former makes a precise prediction about the relationship between prices of various securities for markets to be in equilibrium; the latter's prediction is less specific, being ordinal in nature. Both have been used successfully to measure equilibration in thick experimental financial markets (see [5]). We complement graphical evidence of equilibration with formal tests. The tests give the probability that the actual evolution of our measures of distance from equilibrium could have emerged merely by chance. That is, we compute the probability that our results would have come about had prices been just random walks. Under a random walk, no economic forces (pressures of portfolio demand against a fixed supply) are at work. The only economic basis for a random walk is simple speculation, which would indeed make prices unpredictable. The formal statistical tests are necessary because of recent concerns that experimental results can be generated by simple chance mechanisms, instead of the economic forces that they are usually attributed to.

Our trading mechanism is designed to cross heterogeneous portfolio orders. This is accomplished by a scale-back procedure that is reminiscent of the partial order filling in standard, one-security markets. The second ingredient of our trading mechanism is pricing. Markets need a clear, easily interpretable signal that reflects excess demand (price increases) or excess supply (price decreases). We use constrained, mixed linear-integer programming to determine prices and trades. The constraints are suggested by economic theory.

We confirm our conjecture that a portfolio trading mechanism induces full equilibration. Equilibrium generally emerges way before volume decreases, indicating that any reduction in liquidity must be attributed to exhaustion of gains from trade. We record these results in an environment where BKP observed that markets did not fully equilibrate because liquidity dried up.

The remainder of this paper is organized as follows. In the next section, we briefly describe how the results in BKP led to our conjecture that a portfolio trading mechanism may overcome an incomplete equilibration process. In Section 3.3, we discuss our measures of distance from equilibrium, which are used to determine the success of our portfolio trading mechanism. In Section 3.4, we describe the experimental setup. The portfolio trading mechanism itself is introduced in Section 3.5. Results and tests are reported in Section 3.6. We conclude in Section 3.7.

3.2 Market Structure And Illiquidity In Thin Financial Markets

The BKP paper reported on a set of experiments designed to test equilibration in a repeated, multiple-asset market of two risky securities and one risk-free security. While the experiments revealed slow, steady convergence towards equilibrium, volume invariably dried up, causing the equilibration process to stop short of the equilibrium. It was conjectured that equilibration halted because of subjects' hesitance to trade in the face of market thinness.

The basis for this conjecture is the following. When an agent wishes to improve the risk/reward characteristics of her portfolio, she generally has to simultaneously trade in several markets at once. If some or all markets are thin, it is possible that only a strict subset of her orders will be executed. The resulting portfolio may therefore be different from the desired one, and even be inferior to the initial portfolio before the trade.

With only a limited number of agents, the market mechanism in BKP (a number of parallel double auctions) forces traders to actively seek out transactions when they wish to rebalance their portfolios. This involves posting bids (if buying) or asks (if selling) that are aggressive enough to solicit trade. Portfolio reallocations, however, usually require one to simultaneously execute trades in several markets. Because of the thinness (size) of the markets in BKP, there was a fair chance that orders would only partially be executed, despite efforts to make the quotes as attractive as possible. Unfortunately, the resulting portfolio may easily have a worse risk/reward trade-off than the original one. The risk of ending up with an inferior portfolio may induce an optimizing agent not to try to improve her position at all, keeping her from participating further, and hence, generating illiquidity.

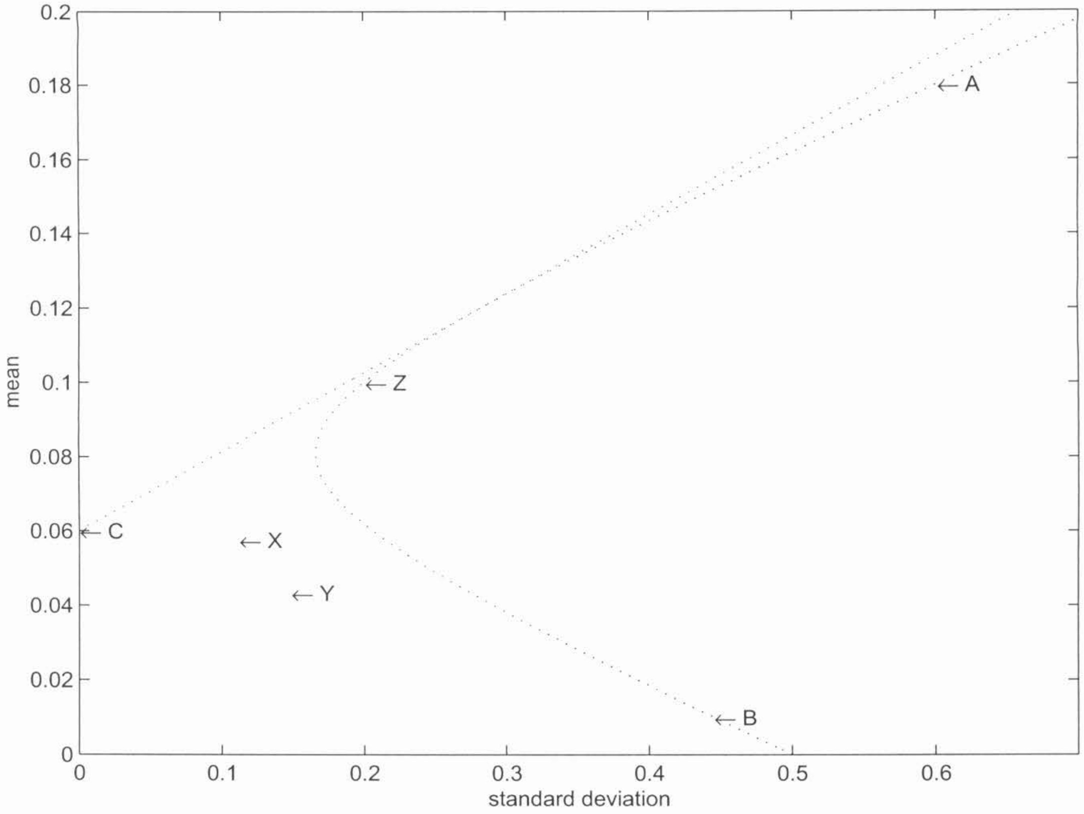

Figure 3.1: An investor who holds position X may want to move to point C, because it has a higher return and lower (no) risk. Since X consists of a combination of the positions A, B and C, both A and B have to be liquidated and additional units of C have to be bought. If the investor manages only to sell A, (s)he moves down to position Y, which is dominated by the original position (X). If (s)he can only sell B, (s)he moves to Z, which has higher expected return, yet is more volatile.

This can easily be illustrated graphically if risk can be measured by return volatility and reward by mean return. Imagine a risk/reward trade-off as in Figure 3.1. The straight line provides the best risk (return volatility) – reward (mean return) tradeoff that anybody could get in the marketplace. The curved line provides the best risk/reward tradeoff of portfolios consisting only of risky securities. Imagine that these portfolios are constructed on the basis of three securities: A, B and C. The risk and reward of these securities are indicated in Figure 3.1 as well. C has zero return volatility, and, hence, is riskfree. Now consider position X in the plot. It depicts a portfolio consisting of a combination of the riskfree security (C) and the two risky securities (A and B). This portfolio is clearly dominated by the riskfree security, which earns more despite lower (no) risk. So, one could easily improve the risk/reward trade-off of X by moving towards an all riskfree portfolio, by selling the risky securities. If, however, one succeeds in selling only holdings of A, the risk/reward trade-off worsens, because we would move towards, e.g., position Y.[2] The new position is dominated by the old one. On the other hand, if only B is sold, the new position (e.g., Z in Figure 3.1) is not dominated in mean-variance space by the old one, but it incurs more risk, which the agent may not be willing to bear. Consequently, a rational agent may decide not to rebalance the portfolio. Not only is an opportunity to gain from trade missed, but also, markets will be less liquid.

Figure 3.1 illustrates that the risk of obtaining inferior portfolio positions is less if (i) we submit orders in a particular sequence (as opposed to simultaneously), (ii) the agent is less risk averse. If risk aversion is sufficiently low, position Z will be superior to the initial position (X). By first submitting orders to sell A, a risk-seeking agent guarantees a better outcome. After A is traded, she can turn to selling B. In large markets, chances are greater that there are enough agents with low risk aversion,[3] who will not refrain from trading. This generates the necessary order flow that will attract others to the market. Liquidity ensues. In small markets these chances are slim, so that illiquidity may follow.

Notice that the market mechanism is conjectured to be at the heart of liquidity problems in small markets. The mechanism in BKP was one of parallel double auctions, where agents could not submit orders in one market contingent on events in another. To avoid moving from X to Y in Figure 3.1, one would like to submit an order to sell all of B conditional on a sale of all of A. If such orders are possible, then even a risk averse agent would not refrain

---

[2] The exact location would depend on the amount of cash generated by selling A.

[3] Provided, of course, that agents have heterogeneous attitudes towards risk.

from participating in the trading. This calls for a portfolio trading mechanism: a system whereby agents can submit orders to trade packages of securities, instead of just one.

Much of the difficulty in implementing a test of our conjecture stems from the absence of flexible portfolio trading mechanisms.[4] So, the challenge was to design such a mechanism, which we will refer to as a *combined-value trading* system (CVT). This mechanism eliminates the natural tendency of parallel markets to decompose a portfolio and its valuation into its constituent parts. Rather, we allow traders to submit bids which reflect their desired multi-security portfolio transactions. Standard trading mechanisms involve orders by a player i which could be described with the pair (b_i, q_i). These are orders of the form "i is willing to pay up to b_i (respectively accept no less than b_i if the order is for a sale) for q_i units of a security." In contrast, CVT allows for orders which can be represented by the $N+1$-tuple $(b_i, \overrightarrow{q}_i)$, meaning "$i$ is willing to pay up to b_i for the vector of units of N securities $\overrightarrow{q}_i$." Additionally, agents can submit a scaling parameter, F_i. This scale indicates the minimal acceptable level at which a bid can be filled. So, now a bid is to be represented by the $N+2$-tuple $(b_i, \overrightarrow{q}_i, F_i)$, to be understood as "$i$ is willing to pay up to fb_i for the vector of units of N securities $f\overrightarrow{q}_i$, for any f between F_i and 1." We will discuss CVT in more detail in Section 3.5.

3.3 Measuring Equilibration

Given the limited size of the stakes in a typical experiment, we can safely assume that subjects' preferences towards risk can be approximated by quadratic utility functions (provided of course that expected utility theory describes their attitudes towards risk in the first place). Hence, the Capital Asset Pricing Model (CAPM) would obtain in equilibrium, at least approximately. The CAPM model predicts that in equilibrium prices will be such that the market portfolio (i.e., the aggregate supply of risky securities) is optimal for quadratic preferences. That is, the market portfolio generates maximum mean return for its volatility. See

---

[4] See [33, 34] for a mechanism that is related to ours. [33] reports experimental results from multi-security experiments, but mean-variance preferences were induced, effectively eliminating uncertainty. In our experiments, risk was explicit. Hence, subjects' natural inclinations in the face of uncertainty were the basis of trading and pricing. See also [19, 25].

Sharpe's 1964 work [32] for a detailed discussion. This prediction is independent of agents' levels of risk aversion, which is significant, because these cannot readily be measured, and moreover, may change during the course of the experiment.

Consequently, to determine whether experimental markets have equilibrated, we compare the reward-to-risk trade-off of the market portfolio against the maximum possible trade-off available at market prices. This trade-off is usually referred to as *Sharpe ratio*, to be defined as follows. Let R_{Ft} denote the return on a risk-free security at time t; let R_{mt} be the return on the market portfolio and let σ_{mt} denote its volatility. The Sharpe ratio of the market portfolio is then

$$\frac{E[R_{mt} - R_{Ft}]}{\sigma_{mt}}$$

The maximal Sharpe ratio is the ratio of mean return in excess of the riskfree rate over volatility (when a riskfree security exists, the maximal Sharpe ratio is constant for all levels of volatility).[5]

At any moment in our experiments, we measure how far markets are from equilibrium by computing the difference between the market Sharpe ratio and the maximal Sharpe ratio. Markets reach equilibrium when the difference becomes zero. This measure was successfully employed in an experimental setting in BKP, as well as the large-scale follow-up study (BP). Notice that our distance measure can be computed without observation or estimation error. In field studies, lack of observability of the market portfolio makes it difficult to assess whether markets have equilibrated. Likewise, the payoff distribution and its parameters (means, variances, covariances) have to be estimated, and, hence, sampling error must be dealt with. In the laboratory, both the market portfolio and the payoff distribution are under control of the experimenter, and, therefore, known.[6]

The significance of our experimental results is further enhanced because (i) we made every possible effort to teach subjects beforehand about the nature of the payoff distribution,

---

[5] In our computation of the maximal Sharpe ratio, we did take into account constraints on shortselling of risky securities. When shortsale constraints are binding, the maximal Sharpe ratio is not independent of volatility anymore. In that case, we use the maximal Sharpe ratio corresponding to the volatility of the market portfolio.

[6] There are additional problems with field studies, such as the necessity to assume that the payoff distribution can be estimated from observed payoff frequencies, and that agents knew this payoff distribution when setting prices, etc. None of these affect our interpretation of experimental data.

by means of pre-experiment learning sessions – see details below; (ii) we did not make endowments common knowledge, and avoided their becoming common knowledge by changing them across experiments (like in BP) as well as across periods within an experiment (unlike in BP). The latter means that subjects could not have deliberately used the CAPM to set prices and determine optimal portfolios.[7] Consequently, if we observe equilibration, it is only because of the economic forces that are at work, and not because of subjects' deliberate usage of the CAPM.

Because our experimental financial markets are complete (there are an equal number of states and securities), we can also study a more general prediction about equilibrium. Unlike the CAPM prediction, however, it is ordinal. Implicit in the prices of traded securities are the prices of primitive state securities (also called Arrow-Debreu securities) that pay one unit of currency in one state and zero in others. General equilibrium theory predicts that the ranking of these so-called *state prices* will be inverse to the ranking of the aggregate wealth (payout), provided that the states are equally likely. We will investigate this prediction as well. We will normalize the state prices to add up to 1. This way, we work with the state-price probabilities that have become popular in mathematical finance (where they are also often referred to as equivalent martingale probabilities). The prediction about ranking obviously holds for state-price probabilities as well.

3.4 Experimental Design

We conducted a total of seven experiments, indexed in Table 3.1 by the date of the session. Subjects were recruited from the Caltech community, primarily undergraduates and a few graduate students from the natural sciences and engineering. Because of the complexity of the trading interface and of the intricacies of fixed-income trading, recruiting was limited to those who had taken or were in the midst of taking courses related to finance. Some subjects, particularly in later sessions, participated in more than one session. The number of subjects varied from a low of six to a high of fourteen.

---

[7] Once in CAPM equilibrium, optimal portfolios can readily be constructed as a combination of the market portfolio and the riskfree security.

Table 3.1: List of Experiments and Parameters

Date	Period	Number		Endowment			Exchange
		Subjects	Type	A	B	C	Rate
11/1/99	1 - 5	6	3	9	1	0	0.02333
			3	1	9	0	0.02333
11/4/99	1 - 6	14	7	9	1	0	0.02333
			7	1	9	0	0.02333
11/11/99	1 - 4	11	4	9	1	0	0.02333
			7	1	9	0	0.02333
	5 - 6	10	4	9	1	0	0.02333
			6	1	9	0	0.02333
11/16/99	1	14	9	9	1	0	0.02333
			5	1	9	0	0.02333
	2	14	10	9	1	0	0.02333
			4	1	9	0	0.02333
	3	14	9	9	1	0	0.02333
			5	1	9	0	0.02333
	4	13	10	9	1	0	0.02333
			3	1	9	0	0.02333
	5	13	9	9	1	0	0.02333
			4	1	9	0	0.02333
	6	13	7	10	10	0	0.02333
			1	18	2	0	0.02333
			1	1	9	0	0.02333
			2	9	1	0	0.02333
			2	2	18	0	0.02333
11/30/99	1 - 4	15	10	9	1	0	0.02333
			5	1	9	0	0.02333
	5	14	10	9	1	0	0.02333
			4	1	9	0	0.02333
	6	14	9	9	1	0	0.02333
			5	1	9	0	0.02333
12/2/99	1 - 5	12	6	6	3	0	0.03
			3	3	6	0	0.03
			3	4	5	0	0.03
	6	8	5	6	0	3	0.03
			3	4	5	0	0.03
	7	7	3	6	3	0	0.03
			3	3	6	0	0.03
			1	4	5	0	0.03
	8	7	3	6	3	0	0.03
			1	3	6	0	0.03
			3	4	5	0	0.03
12/7/99	1 - 2	14	9	9	1	0	0.03
			5	1	9	0	0.03
	3	13	10	9	1	0	0.03
			3	1	9	0	0.03
	4	13	8	9	1	0	0.03
			5	1	9	0	0.03
	5	13	8	9	1	0	0.03
			5	1	9	0	0.03
	6	13	10	9	1	0	0.03
			3	1	9	0	0.03
	7	13	8	9	1	0	0.03
			5	1	9	0	0.03
	8	6	5	9	1	0	0.03
			1	1	9	0	0.03

We created three securities, denoted A, B, and C, each with a life of one period. At the end of the period, each security paid a single dividend and was then retired. The magnitude of the dividend depended on a random draw of one of three equally likely states, X, Y and Z. The state was drawn after the period was closed, so there was no insider or asymmetric information. The payoff table was as follows:

	State		
Security	X	Y	Z
A	170	370	150
B	160	190	250
C	100	100	100

Notice that the dividend of A varies dramatically from state to state (with an expected value of 230), the dividend of B varies less and has an expected value of 200, and the dividend of C is constant at 100.

Each experiment consisted of multiple periods of similar trading conditions, varying only by the initial allocations given to subjects. At the beginning of a period, each player was supplied with some securities A and B, and some francs cash (an experimental currency whose conversion rate into dollars is indicated in Table 3.1). The period proceeded in a number of rounds, the first round three minutes in length and subsequent rounds ninety seconds long. At the end of each round, all trading was stopped while the allocation algorithm (described in Section 3.5) solved for all trades. These trades were executed, and net trade and market prices were announced. Then another round began. The number of rounds in a period varied between experiments. The November experiments had 10 rounds per period. As we will report in Section 3.6, it appeared that markets equilibrated and subjects slowed their activity after 5 to 8 rounds. Therefore, the December experiments had only seven rounds per period. At the end of a period, one of the three states was chosen and announced, players were paid their dividends according to the payoff table, and the period ended. Then, we began another period, with new allocations. In the case of the November experiments, the session consisted of six periods, and the December sessions were eight periods long. Subjects knew about the length of the session they were in.

Note that while the aggregate supply of securities A and B varied from period to period, security C was always in zero net supply. Additionally, while no short sales were permitted in A and B, players were allowed to short up to 8 units of C. Thus, the ability to buy A and B is not limited by the number of francs on hand. If a subject wishes to expand holdings beyond the bounds indicated by the endowment of francs, she could do so by selling units of C and paying the dividend. A sale of C is equivalent to borrowing an amount equal to the sale price, in exchange for a repayment of 100 francs at the end of the period. To the buyer, the difference between the price paid and the dividend is a riskfree return since the payment is guaranteed. This sale and purchase of asset C determines the risk-free rate simultaneously with the rates of return on the risky securities. Since it is possible for players to lose money due to unfortunate draws of the state and to then declare bankruptcy, we needed to ensure the integrity of the incentive system. Any subject who lost money in two consecutive periods in an experimental session was asked to leave the experiment.

The initial allocations varied widely from experiment to experiment, as detailed in Table 3.1. While players were aware of their own initial allocations, they were not told the allocations of others. Therefore, the size of the market portfolio was unknown to any player. This is an important design consideration. As discussed in the previous section, it ensured that participants could not use the CAPM and Arrow-Debreu pricing model to deliberately set prices, thereby artificially generating the outcomes that we were looking for.

Subjects were informed that the experiments would last about three hours. Before assembling in the Caltech Social Science Experimental Laboratory, all subjects were given a URL with the instructions. At the end of each page of instructions is a short quiz, which had to be correctly completed before moving on to the next page (The instructions and quizzes can be found in Appendix A). Once all pages were completed, subjects were given a URL for a practice experiment. Subjects were not allowed to participate in the three-hour in lab experiment if they did not enter at least five practice bids in the practice experiment. Each subject who completed these tasks and arrived at the lab on-time received a $10 bonus to their final payoff.

3.5 The Combined-Value Trading (CVT) Mechanism

We used a Combined-Value Trading (CVT) system in our experiments. In this environment, subject i submits an order of the form $(b_i, \overrightarrow{q}_i, F_i)$, read "$i$ is willing to pay up to fb_i for the vector of goods $f\overrightarrow{q}_i$, for any f between F_i and 1." With bids of this form, determining the payments (prices) and allocations that maximize gains from trade (surplus) and provide the incentives for traders to reveal their true willingness to pay is not as straightforward as in single-asset markets.

In our implementation, we allowed for a variety of order types. A multi-market order for agent i is a vector $< b_i, q_i^A, q_i^B, q_i^C, F_i >$ where $b_i > 0$ means agent i is willing to pay at most b_i to buy the order and $b_i < 0$ means agent i is willing to accept at least b_i for the order. Similarly, $q_i^j > 0$ means agent i wants to purchase up to q_i^j units of j in the order, and $q_i^j < 0$ means agent i wants to sell up to q_i^j units of j in the order. F_i is a scale factor $(0 \leq F_i \leq 1)$ which indicates that agent i is willing to accept an order of the form $< f_i \cdot b_i, f_i q_i^A, f_i q_i^B, f_i q_i^C >$ where $f_i \in [F_i, 1]$. For an "all or none" bid, $F_i = 1$.

Orders for single securities can thereby easily be combined in one order as a package (portfolio). For example, a package order can specify a willingness to pay up to 1000 francs for 4 units of asset A, 2 units of asset B and 3 units of asset C, with full flexibility. This bid would be written $< 1000, (4, 2, 3), 0 >$. One form of packaged order is a *swap*, buying units of some assets and selling units of other assets. For example, a swap can specify a willingness to pay up to 100 francs to supply 3 units of asset A if and only if 1 unit of asset B is received, as well as a willingness to accept any scaled version of this bid. This bid would be $< -100, (3, -1, 0), 0 >$. Note that although the market system will not allow traders to bid more francs then they currently possess, a packaged bid with asset C can overcome this problem. Suppose a player has zero francs and is willing to pay 90 francs to supply of 1 unit of A for 2 units of B, $< 90, (1, -2, 0), 0 >$, which is a bid he cannot afford. In order to cover this transaction, he can offer to sell a unit of asset C, effectively borrowing the 90 francs for a promise of 100 francs at the end of the period. Now, his bid is $< 0, (1, -2, 1), 0 >$, which is within the player's budget constraint.

After the bids are called, we solve for the allocations and market prices. We will now use

i to denote orders as opposed to agent identification to reduce the notational burdens. The *allocation* at each iteration is determined by solving the integer program:

$$\max_{f_i} \sum_i |b_i| \, f_i \tag{3.1}$$

$$\text{subject to} \quad \sum_i q_i^k f_i \leq 0, \quad k = \text{A,B,C} \tag{3.2}$$

$$f_i \in [F_i, 1] \cup 0, \quad \forall i \tag{3.3}$$

Note that the first constraint implies that the market will admit a net surplus. If there is a surplus of an asset in the solution, the market offers it for sale at the market price in the next round of the period.

While solving the allocation problem is fairly straightforward, calculating appropriate *market prices* is not as simple. After the allocation is computed in the above maximization problem, we know which bids will be matched and completed. But we also need to compute what each matched bid will pay or receive. That is, we need to compute the transaction prices. The principles we used in designing the pricing rule were that, (i) payments equal receipts among the bidders, (ii) no one pays more (receives less) than she bid (offered), (iii) there are incentives to reveal one's true willingness to pay, and (iv) everyone pays the same price per unit unless there are significant reasons for deviating. As we will see below, deviations from these principles will occur only in cases with important inflexibilities.

After the allocation problem is solved, there are three categories of orders: (i) orders that were *accepted* by the allocation, (ii) orders that were *rejected* by the allocation, and (iii) orders that were *partially accepted* (that is, $0 < f_i^* < 1$, where f_i^* is the fraction of order i actually allocated). We treat the accepted part as an accepted order and the rejected part as a rejected order. A little economics will now take us a long way. Think about the entire collection of submitted orders as a quasi-linear economy and ask what a market equilibrium (a competitive equilibrium) would look like. It is easy to show that, if a competitive equilibrium allocation and price vector p exist, then the allocation would solve the maximization problem above and the prices p (one price for each item) would solve the following problem:

$$b_i - p \cdot q_i \geq 0 \quad \text{for all accepted orders}$$

$$b_i - p \cdot q_i \leq 0 \quad \text{for all rejected orders}$$

$$p \cdot \textstyle\sum_{i \in A} q_i = 0 \quad \text{(Walras' law)}$$

If such an equilibrium price exists, it would satisfy our principles and would be the natural price to set. Unfortunately, neither the uniqueness nor the existence of such a price vector p is guaranteed. If the market equilibrium prices exist but are not unique, there are many ways to pick one. We use a reference pricing rule, which minimizes the difference between the equilibrium price and the reference price subject to satisfaction of our principles.

Non-existence is a deeper problem requiring somewhat more finesse. These problems can occur when subjects submit bids with flexibility levels other than zero. This causes a nonconvexity of the player's preferences, which may in turn lead to non-existence of an equilibrium. To see why, consider the example in Table 3.2, as well as the corresponding Figure 3.2 , which detail a bid schedule for a single asset. Suppose that bid 7 is to sell 3 units for at least 3 francs per unit and is an inflexible order. Further suppose bid 2 is to pay up to 4 francs per unit for 3 units and is fully flexible. Lastly, suppose bid 3 is to pay up to 2 francs per unit for 2 units, also flexibly. Surplus is maximized, given the flexibility constraints, if bids 1,2,3,5,6, and 7 are filled (3 is partially filled). There is, however, no competitive equilibrium. To see this, notice that at any price above 2 francs per unit bidder 3 is unwilling to buy units, and at any price below 3 per unit bidder 7 is unwilling to sell any units. There is no price such that demand equals supply, because the supply function effectively jumps where the demand function would cross it.

To price the allocation when a market equilibrium price does not exist, we construct a "pseudo-competitive equilibrium price." First, we ignore rejected orders and consider only the accepted orders, i.e., orders i such that $f_i^* > 0$. We then calculate a fully flexible allocation by maximizing the surplus, subject to no excess demand, with $f_i^* \in [0, 1]$. This is the allocation that would occur if all accepted orders were fully flexible. Next, we find

Table 3.2: Example of a set of orders generating the supply-demand schedule in Figure 3.2.

Bid Number	Units Offered	Price
1	2	5.0
2	3	4.0
3	2	2.0
4	2	1.0
5	-1	0.5
6	-2	1.5
7	-3	3.0
8	-3	3.5

prices for this allocation exactly as we did before; splitting the difference if the competitive equilibrium price is not unique. (In this case it is easy to show an equilibrium price will exist.) In the case of our example in Table 3.2, this means that the price is 3.50 francs; the price that would obtain if all orders were fully flexible ($F_i = 0, \forall i$).

But if we were to charge and pay every bid according to the price of 3.50, buyer 3 would be paying more than the maximum her bid indicated she was willing to pay per unit (2 francs). Further, even though seller 7 created the non-existence problem by requiring his bid be all-or-none, he would receive a surplus of 0.50 francs on the extra unit sold that way. To provide the right incentives, to minimize all-or-none bids where they are unnecessary, and to not over-charge or under-pay, we charge or pay each part of an originally accepted bid that was rejected in the fully flexible allocation exactly what they bid. So in Table 3.2, seller 7 will receive 3 francs for the last unit and buyer 3 will pay 2 francs for her unit.

We may still have a problem. If seller 7 receives 3 francs for the last unit sold and buyer 3 pays 2 francs for the unit bought, we will have to pay out more than we receive. In fact, the area indicated as negative surplus in Figure 3.2 is exactly the amount we will be short. Let V^f denote the surplus from the fully flexible allocation, V^* denote the surplus from the original matching procedure, and let $dV = V^f - V^*$, the added surplus from flexibility. This dV is exactly the negative surplus in Figure 3.2. We need to collect this amount from those bids accepted in both the original and in the fully flexible allocation. There are many ways to carry out this accommodation of the inflexible bidders. We choose to charge the player

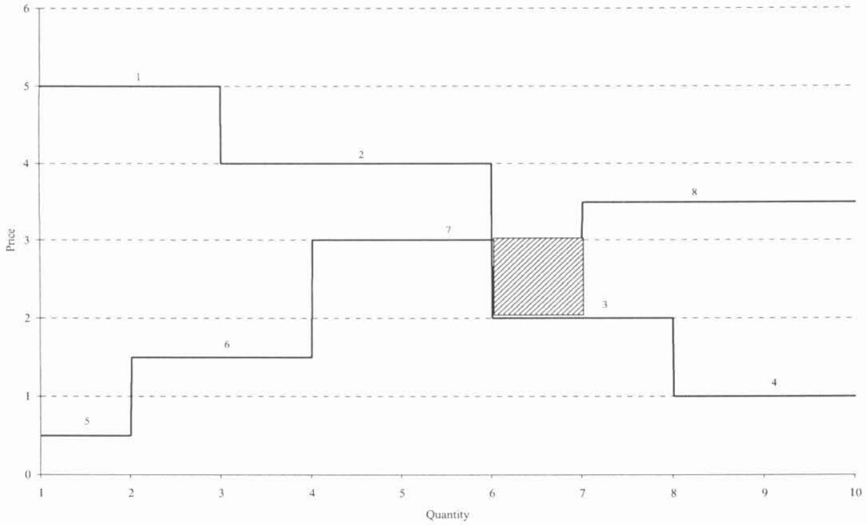

Figure 3.2: Supply-Demand graph generated by bidbook in Table 3.2. The shaded area represents the deficit caused by player 7's inflexibility.

that caused the accommodation to be necessary. In this case, $dV = 1$ franc. Charging this to seller 7 still leaves him better off than if he hadn't traded at all, and brings the market into balance. It is possible that even after charging the inflexible players for their accommodation, the market could still be out of balance. In this case, it becomes necessary to split prices. That is, we charge different prices to the buyers and sellers, with the aggregate difference designed to make up the market deficit.[8]

It is important to note that there may be trading rounds in which no transactions take place. In these cases, the mechanism calculates prices at which no bid would want to trade because it would make negative profits. These prices take the lack of trade as a signal that the current allocation is a competitive equilibrium, and generates prices which support it in the same manner as detailed above. This is to ensure that even when no trade occurs, subjects receive information about the markets that enable trade to occur in the future.

---

[8] Of the 412 rounds in our experiments, there were only eight instances where accommodation was necessary because of inflexibilities.

The above may seem an extremely complex way to match and price combined value trades. It is made particularly difficult when one allows all-or-none trades. But the complexity is invisible to the agents, and is merely meant to provide them with clear signals: execution of trades as well as prices that reflect the economics of the situation. The success of the CVT mechanism is ultimately an empirical question. We measure its success in terms of the liquidity it generates in thin markets. We now turn to the results.

3.6 Results

We will discuss four aspects of the experimental outcomes: comparison of equilibration between CVT experiments and experiments with parallel markets, patterns in volume (liquidity), formal tests of equilibration, and final holdings.

3.6.1 Equilibration

As described earlier, equilibration of financial markets will be measured using the predictions of the CAPM and the Arrow-Debreu pricing model. As far as the CAPM is concerned, we examine the dynamic evolution of the difference between the maximal Sharpe ratio and that of the market portfolio. To measure the Sharpe ratio, we wait until at least one transaction has occurred in each market. Using the most recent transaction prices, we compute the maximum Sharpe ratio as well as the market's Sharpe ratio, and take the difference. This is then repeated after any call in which a transaction occurs. This produces a plot of the evolution of the difference between the maximum Sharpe ratio and that of the market. Figure 3.3 shows plots of the Sharpe ratio differences for our seven experiments.

With few exceptions, the (absolute) Sharpe ratios differences are less than 0.15. There is a tendency for the difference to be wider in earlier rounds of a period, but subsequently it narrows, which indicates that the markets equilibrate. However, in later rounds of a period the difference sometimes widens, which seems to imply that the market moves off its equilibrium. We will document below that less trade occurs in later rounds of a period, making prices more sensitive to the few orders that are executed, or in the absence of trade, November 1st Experiment

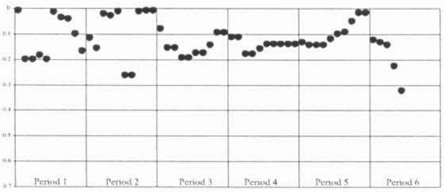

November 4th Experiment

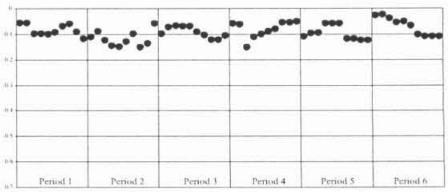

November 11th Experiment

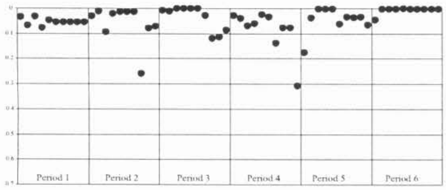

November 16th Experiment

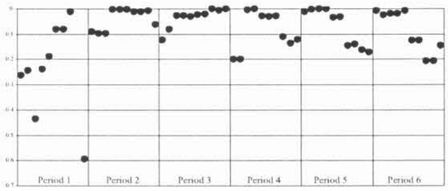

November 30th Experiment

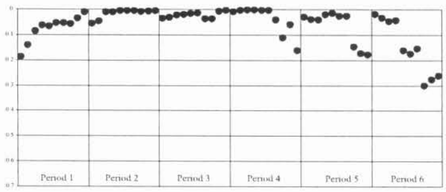

December 2nd Experiment

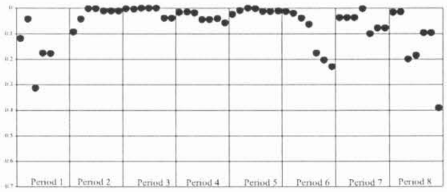

December 7th Experiment

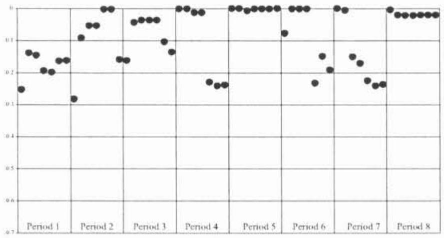

Figure 3.3: Evolution of the difference between the maximum Sharpe ratio and the Sharpe ratio of the market portfolio. The theoretical maximum difference is zero. Time is measured in periods.

to the bid and ask prices. This causes the market to move away from CAPM equilibrium in later rounds. We will also report that reducing the number of rounds forces trading to occur earlier, implying that the prices in later rounds are no less arbitrary. This is confirmed in the last two frames of Figure 3.3, where the number of rounds is seven instead of ten (December experiments).

Subjects did have the tendency to submit orders that attempted to exploit mistakes that others may have made. This was especially true for the riskfree asset (security C): while the theoretical no-arbitrage price is 100, some subjects invariably bid a low price for security C, and sometimes got it when others inadvertently submitted a low ask. When no orders for a security are filled (which occurred often in later rounds), the quoted price is affected by these bids (see the description of the mechanism in the previous section). Because such speculative bids did impact in particular the price of security C, we decided to ignore the quoted prices for this security, and set its price equal to its no-arbitrage value (100). The Sharpe ratios were computed on the basis of this theoretical price.

It should be emphasized that equilibration is far from a foregone conclusion in our markets. Subjects did not know the composition of the market portfolio, and, hence, could not use the CAPM to price securities, or to determine optimal investment strategies. As far as the latter is concerned, we will later present evidence on subject's actual end-of-period holdings.

We set out to study how well our CVT mechanism does *relative to* the usual system of parallel, continuous double auctions with the same number of subjects. So far, the figures document that our experimental markets did generally equilibrate all the way towards CAPM, unlike in the experiments reported in BKP. Still, we would like to know the extent to which there is improvement. To gain perspective, Figure 3.4 shows the evidence from the thin-market experiments in BKP. The experiments are concatenated to fit in one plot, so the time scale is much bigger. The difference with the results from the CVT mechanism (Figure 3.3) is pronounced. Equilibration in the CVT mechanism generally occurs in the early rounds of many periods, whereas equilibration in the thin-market continuous double auctions effectively occurred only in the later experiments, where only subjects who had been in earlier sessions participated.

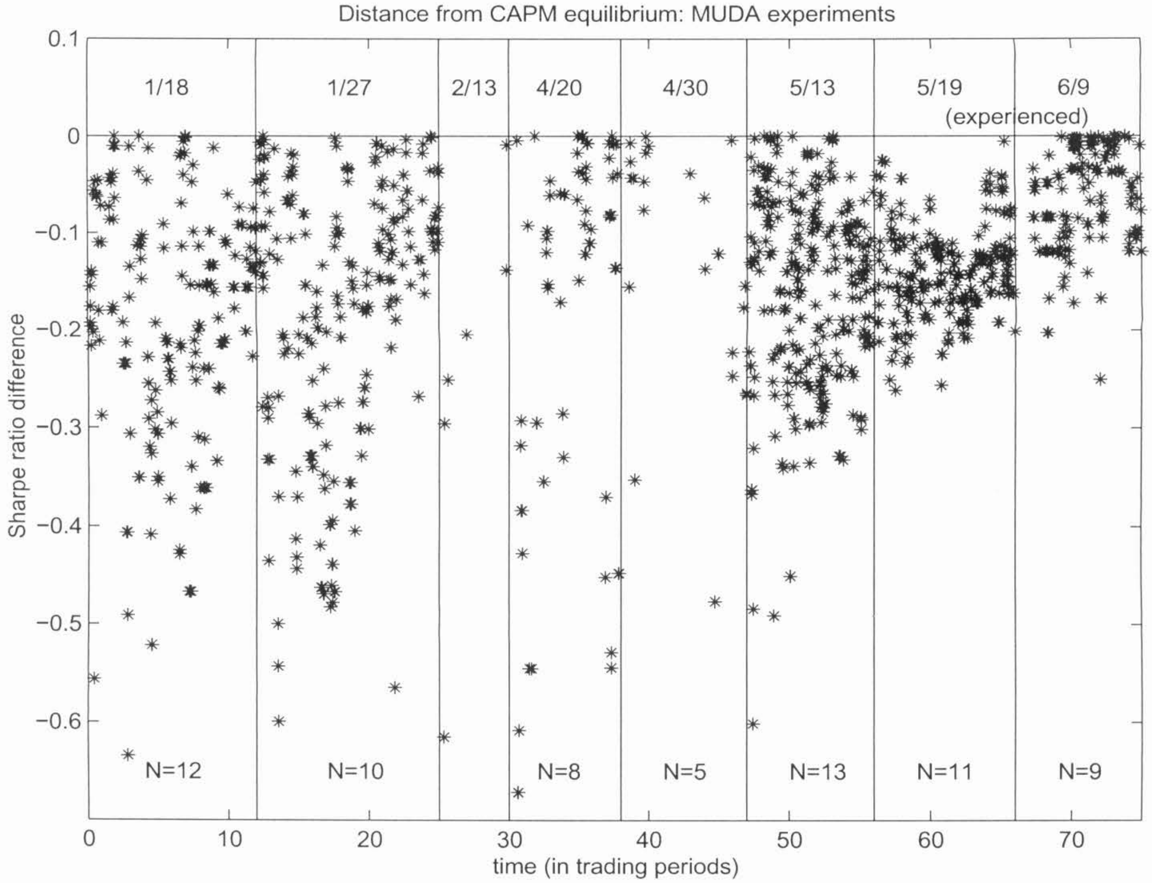

Figure 3.4: Evolution of the difference between the maximum Sharpe ratio and the Sharpe ratio of the market portfolio for in thin market experiments with parallel double auctions. The experiments are concatenated to fit in one plot.

As a further benchmark, Figure 3.5 plots the evolution for the thick-market experiments in BP. Again, the experiments are concatenated. The thick-market experiments generally produce Sharpe ratio differences of the same magnitude as the CVT markets (in the range 0 to -0.2), but require far more trade (between 18 and 65 subjects participated). So, both thickening and a change in the market structure facilitate equilibration.

Figure 3.6 provides evidence on the second prediction that asset pricing theory makes, namely, that state-price probabilities should be ranked inversely to the aggregate wealth in the corresponding state. From the payoff matrix and the aggregate allocation of securities

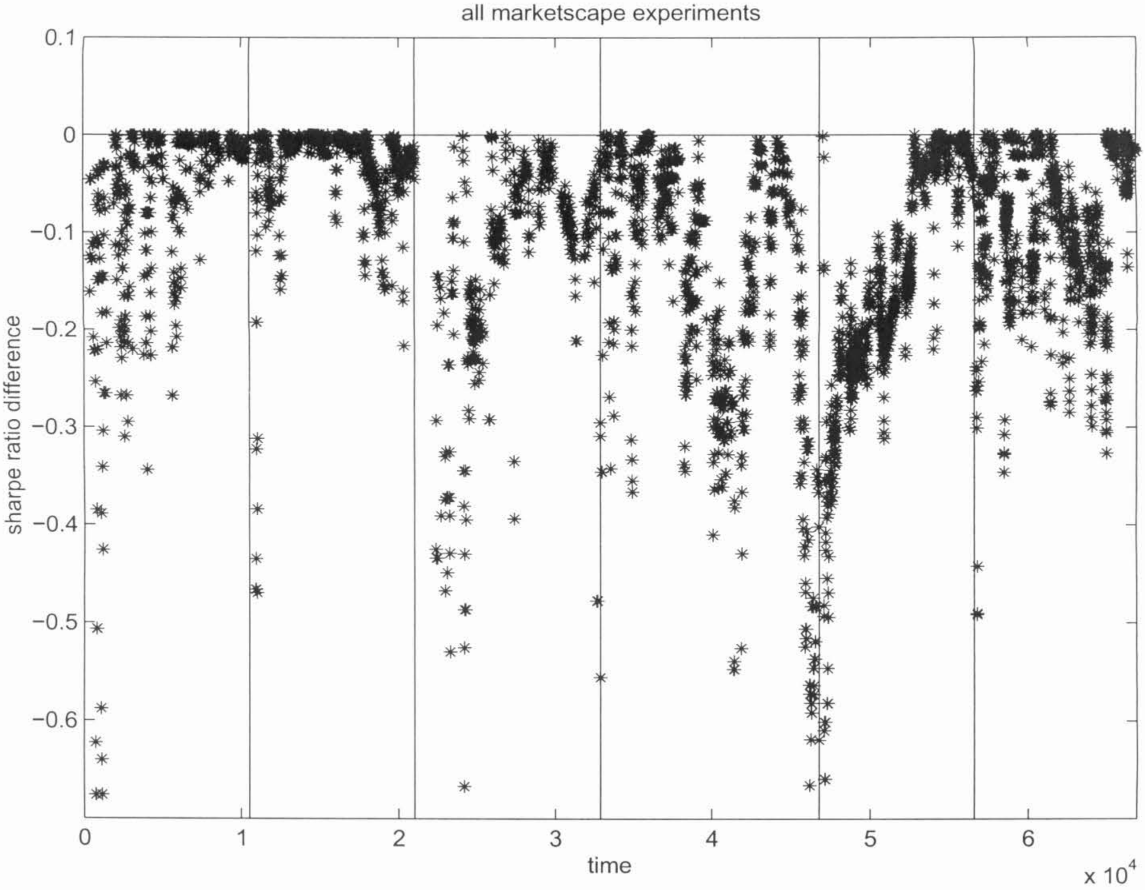

Figure 3.5: Evolution of the difference between the maximum Sharpe ratio and the Sharpe ratio of the market portfolio for in thick market experiments with parallel double auctions. The experiments are concatenated to fit in one plot.

(Table 3.1), one can infer that aggregate wealth was highest in state Y and lowest in state X, implying that the state-price probability of state X be highest, and that of Y be lowest. The figure generally confirms this prediction. There are aberrations, but it is not clear how significant these are. Any formal test would run into the difficulty that the prediction is only ordinal, and that the source of the randomness is not obvious. Later on, we introduce one approach to formally determine the significance level of the visual evidence.

Again, the finding that state price probabilities generally rank as predicted by the theory is not a foregone conclusion, because subjects did not know aggregate wealth, and hence,

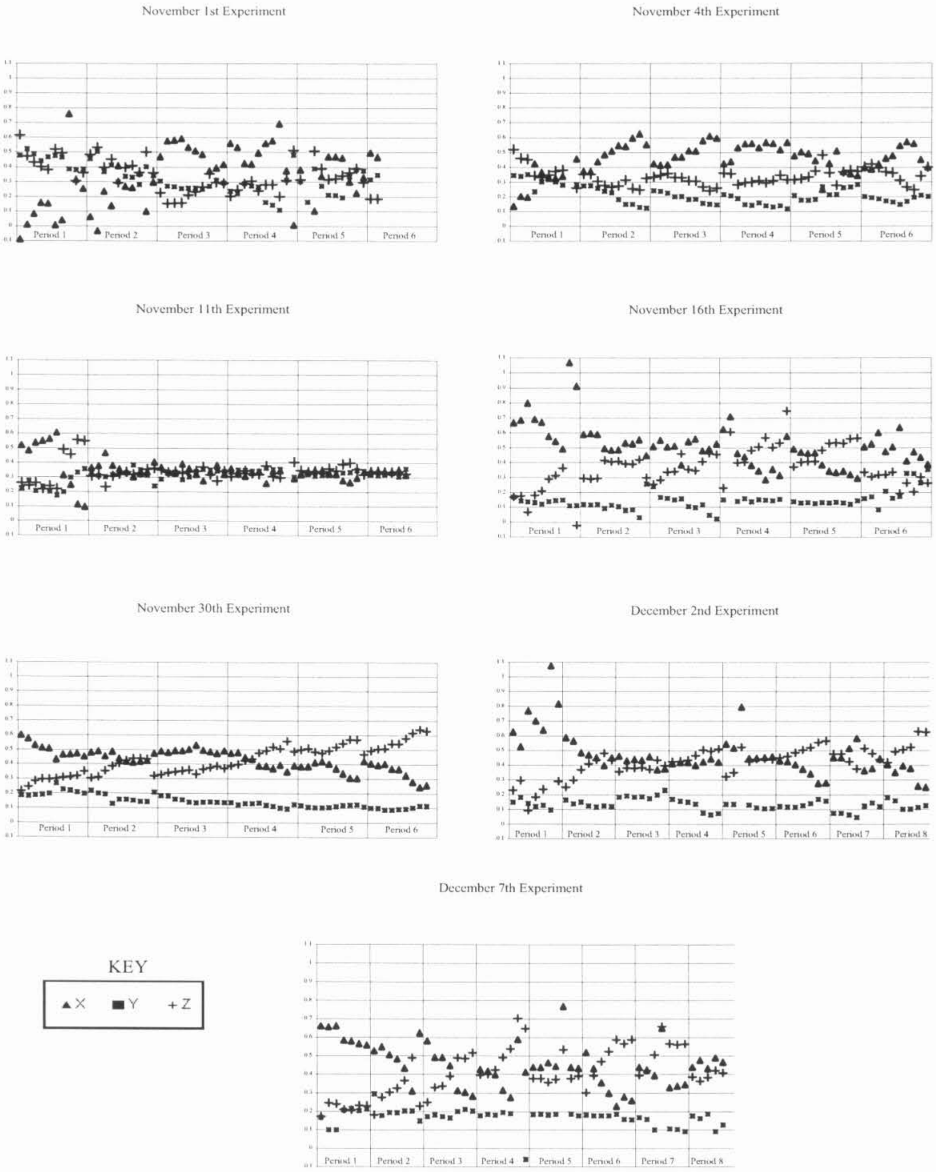

Figure 3.6: State-price probabilities for the seven experiments. The payoffs and allocations were such that the state Y yielded the highest aggregate wealth and state X the lowest. Therefore, state X should have the highest state-price probability and state Y the lowest. Values outside of the [0,1] interval indicate a clear arbitrage opportunity.

could not determine the equilibrium ranking of state-price probabilities (if they cared at all).

3.6.2 Volume

Figure 3.7 depicts the evolution of volume (in francs traded per subject) over time. In the ten-round experiments (November experiments), volume remains even over the first six rounds, and declines subsequently. Cross-inspection with Figure 3.3 reveals that equilibration (if it occurs) generally completes before round six, and hence, before volume declines. This means that the reduction in liquidity (volume) in later rounds could only reflect exhaustion of gains from trade. In the seven-round experiments (December experiments), volume declines after two rounds, but does not dry up subsequently. Again, cross-inspection with Figure 3.3 reveals that equilibration is usually completed before volume declines.

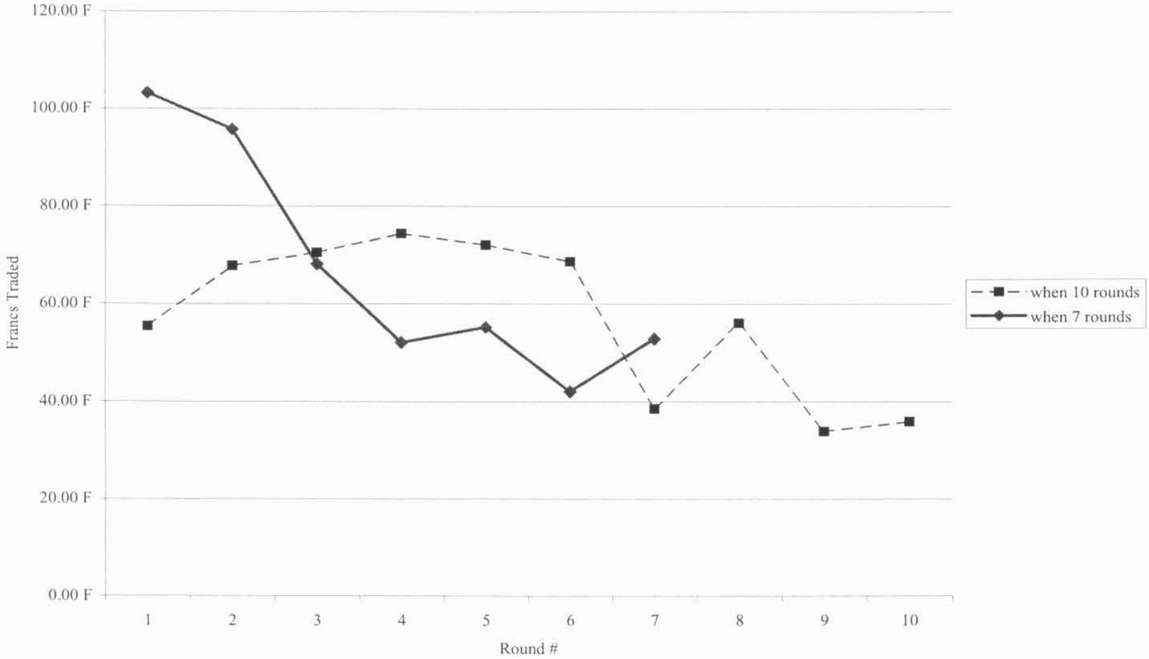

Figure 3.7: Per Capita Transaction Volume (in Francs), by round, averaged across periods and experiments.

We postulated that the CVT mechanism induces liquidity in thin experimental financial markets because subjects can rebalance portfolios more easily. To verify that subjects do indeed avail themselves of the added portfolio trading flexibility provided by CVT, Figures 3.8 and 3.9 report the percentage of orders submitted that are combined-value (meaning that they involve at least two securities, as opposed to orders for one security against cash only). Both by round and by period, between 20% and 30% of the orders are combined-value. The vast majority of these combined-value orders happen to be swaps (exchange of one or more securities for another).

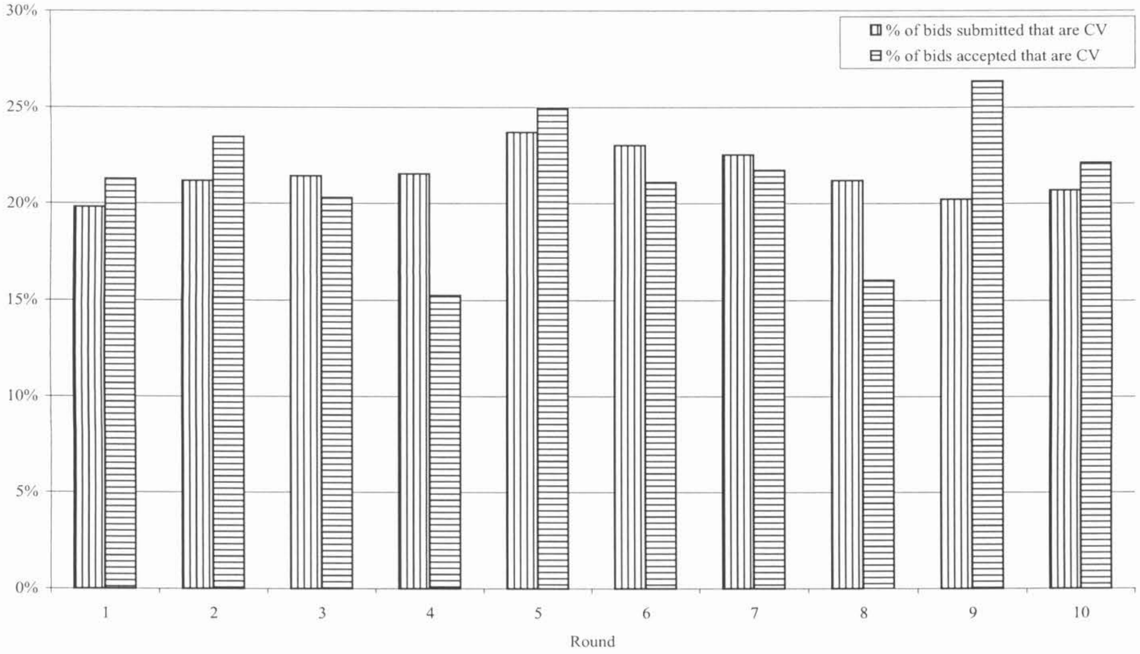

Figure 3.8: Percent of bids submitted and accepted by subjects that involve at least two securities, by round. Note that submitting a combined-value bid does not decrease the likelihood of the bid's acceptance.

The success of our portfolio trading mechanism is further gauged in the finding that combined-value bids are no less likely to transact than are single-asset bids. Figures 3.8 and 3.9 report that the percentage of orders accepted that are combined-value is about equal to the percentage of orders submitted as combined-value.

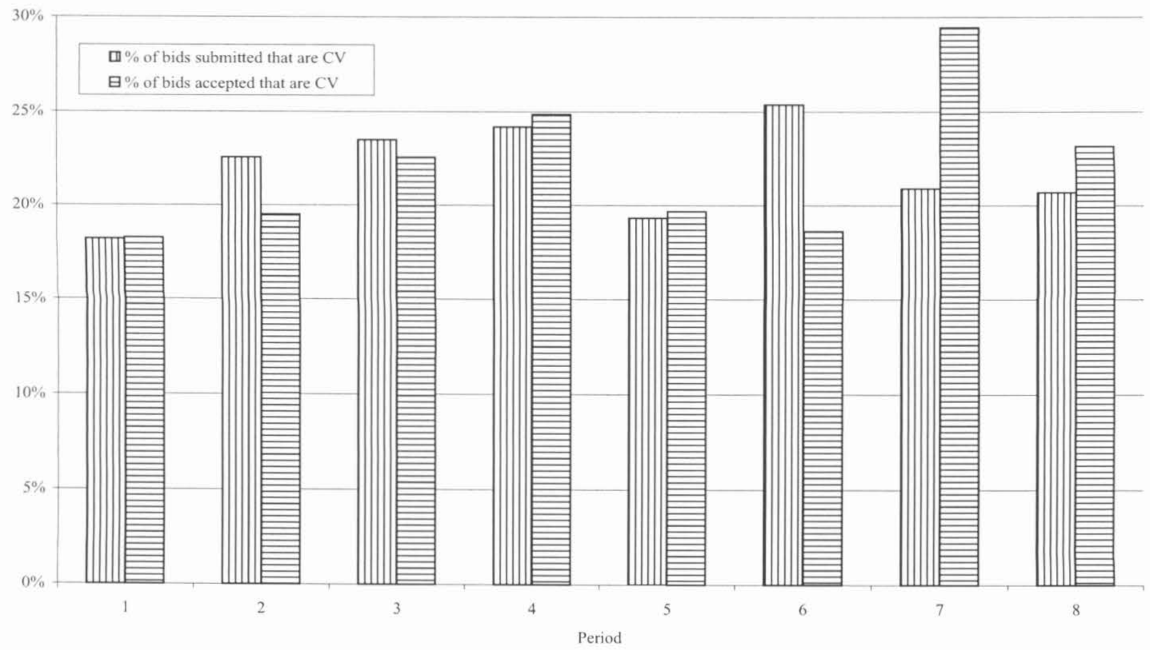

Figure 3.9: Percent of bids submitted and accepted by subjects that involve at least two securities, by period.

3.6.3 Formal Tests of Equilibration

So far, we have produced graphical evidence of equilibration. Such evidence does not reveal its significance. To gauge the significance, we are going to introduce a statistical test that formally distinguishes two hypotheses: (i) prices are unpredictable, i.e., they form a random walk (the null hypothesis), and (ii) prices are driven by the economic forces predicted by asset pricing theory, i.e., they are attracted by equilibrium (the alternative hypothesis). The test effectively asks: what are the chances that we read too much asset pricing theory in the data when in fact there is nothing going on besides speculation. Speculation will eliminate arbitrage opportunities, and, if subjects are risk neutral, will cause prices to behave like random walks. Even if prices are a random walk, with only three securities it is likely that the market portfolio accidentally becomes mean-variance efficient as predicted by the CAPM, or that state-price probabilities happen to rank in accordance with Arrow-Debreu equilibrium. We want to rule out that our observations are cases of mere luck. Let us first consider formal tests of the CAPM.

Testing Random Walk Pricing Against The CAPM

We take the random walk hypothesis as the null, and test it against the hypothesis that the market is pulled towards the CAPM. Our test works as follows. Let $\Delta_{M,t}$ denote the distance between the Sharpe ratio of the market and the maximum Sharpe ratio, the subscript t denoting time (period and round). Consider the projection of the change in $\Delta_{M,t}$ onto $\Delta_{M,t-1}$:

$$\Delta_{M,t} - \Delta_{M,t-1} = \kappa \, \Delta_{M,t-1} + \epsilon_t. \tag{3.4}$$

where κ is such that ϵ_t is uncorrelated with $\Delta_{M,t-1}$. CAPM implies $\Delta_{M,t} = 0$; convergence to CAPM pricing implies $\kappa < 0$. We then determine the distribution of the least squares estimates of κ under the null hypothesis of a random walk, by randomly drawing from (bootstrapping) the empirical joint distribution of changes in transaction prices. The null hypothesis of a random walk is rejected in favor of stochastic convergence to CAPM if the least squares estimate of κ is beyond a critical value in the left tail of the ensuing distribution. This testing procedure is a variation of *indirect inference* (see [15]): we summarize the data in terms of a simple statistical model (in our case, a least squares projection) and determine the distribution of the estimates by simulating the variables entering the statistical model. Instead of simulating off a theoretical distribution, we bootstrap the empirical distribution, however.[9]

For each experiment, we estimated κ using OLS. We determined 5% and 10% critical values under the random walk null hypothesis by bootstrapping from the empirical joint distribution of price changes (we generated 200 price series of the same length as the sample used to estimate κ).[10]

Table 3.3 reports the results. The null of a random walk is rejected in five of our experiments. It is not surprising that we fail to reject the null in two of our experiments. To see

---

[9] [14] also uses indirect inference, but, instead of matching an arbitrary statistical model, they match the scores of the likelihood function.

[10] We bootstrapped the mean-corrected empirical distribution, in order to stay with the null hypothesis of a random walk.

why, look at the plots in Figure 3.3. The November 4 Sharpe ratio distance behaves itself as a random walk. In the November 30 session the volatility is so low that given the opening prices one could easily have stuck on the frontier accidentally. In the five experiments in which we can reject the null of the random walk, the statistical analysis indicates that there is only a tiny probability to accidentally obtain, for instance, the plots in Figure 3.3 if prices were indeed a random walk.

Table 3.3: Test of Random Walk Pricing Against CAPM Equilibrium

Experiment	Attraction Coefficient κ		
	Estimate[a]	Critical Value[b]	
		5%	10%
11/01/99	-0.0893^{*}	-0.1086	-0.0789
11/04/99	-0.0400	-0.1335	-0.1028
11/11/99	-0.3900^{**}	-0.1153	-0.0662
11/16/99	-0.3700^{**}	-0.0978	-0.0702
11/30/99	-0.0300	-0.0942	-0.0786
12/02/99	-0.2900^{**}	-0.1325	-0.0818
12/07/99	-0.2800^{**}	-0.1001	-0.0812

---

[a] Meaning of superscripts: ** = significant at the 5% level, * = significant at the 10% level.
[b] Based on 200 bootstrapped samples of the same size as used to estimate κ.

The rejections of the random walk hypothesis reported in Table 3.3 do not imply that the subjects ignored profit opportunities from speculating on price changes. This is because our rejection of the random walk is based on information on which subjects could not condition, namely the Sharpe ratio of the market portfolio. Recall that this ratio is not readily determined from the history of price movements, as computation requires knowledge of the composition of the market portfolio. As emphasized before, subjects did not know this, and hence could not determine in which way prices would move when the market was still out of equilibrium, even if they believed in the CAPM.

Testing Random Walk Pricing Against The Arrow-Debreu Model

To a large extent, the graphical evidence in Figure 3.6 suggested that state-price probabilities moved in the direction predicted by Arrow-Debreu equilibrium, even if their ranking contradicted the theory for long periods of time. One would like a formal test to confirm the visual evidence of "movement." In particular, one would like to determine whether indeed state-price probabilities adjust towards Arrow-Debreu equilibrium even when their ranking is not as predicted by the theory. We again take random walk pricing as our null hypothesis. That is, we determine the probability of observing the dynamics in state-price probabilities in our experiments if prices were merely random walks. Under the alternative that markets are attracted by Arrow-Debreu equilibrium, we expect specific changes in the state-price probabilities when they are not aligned appropriately. In particular, we expect the following. Let $P_{X,t}$, $P_{Y,t}$ and $P_{Z,t}$ denote the time-t state-price probabilities for states X, Y and Z, respectively. Recall that these should be ranked inversely to the aggregate wealth in each state, namely $P_{X,t} > P_{Z,t} > P_{Y,t}$. Time is measured in number of transactions.

Ranking of State-Price Probabilities at t	Expected Effect
$P_{X,t} > P_{Y,t} > P_{Z,t}$	$(P_{Z,t+1} - P_{Y,t+1}) - (P_{Z,t} - P_{Y,t}) > 0$
$P_{X,t} > P_{Z,t} > P_{Y,t}$	Anything is Possible
$P_{Y,t} > P_{X,t} > P_{Z,t}$	$(P_{X,t+1} - P_{Y,t+1}) - (P_{X,t} - P_{Y,t}) > 0$
$P_{Y,t} > P_{Z,t} > P_{X,t}$	$(P_{X,t+1} - P_{Y,t+1}) - (P_{X,t} - P_{Y,t}) > 0$ or $(P_{Z,t+1} - P_{Y,t+1}) - (P_{Z,t} - P_{Y,t}) > 0$
$P_{Z,t} > P_{Y,t} > P_{X,t}$	$(P_{X,t+1} - P_{Y,t+1}) - (P_{X,t} - P_{Y,t}) > 0$
$P_{Z,t} > P_{X,t} > P_{Y,t}$	$(P_{X,t+1} - P_{Z,t+1}) - (P_{X,t} - P_{Z,t}) > 0$

These predictions are weak, because they only concern the sign of the change in the difference between two state price probabilities. The question is: are economic forces strong enough that the expected effects can be detected sharply?

As test statistic, we compute the frequency of observing (transitioning to) the expected outcome for each state (ranking of state-price probabilities). We subsequently average across states. The averaging is mandated by the fact that, in finite samples, not all states need occur, in which case some transition frequencies are undefined. Let π denote the mean transition frequency. Notice that the second frequency will always be 1 (100%). We include this frequency, so that outcomes where the Arrow-Debreu prediction holds ($P_{X,t} > P_{Z,t} > P_{Y,t}$) receive more weight. In analogy with our formal test of the CAPM, we compute the distribution of π under the null hypothesis by bootstrapping the empirical joint distribution of price changes in each experiment (we generated 200 price series of the same length as the sample used to estimate π).[11] Table 3.4 reports the results. In addition to the estimated mean transition frequencies (second column), we report 5%, 90% and 95% critical values under the null of random walk pricing.

Table 3.4: Test of Random Walk Pricing Against Arrow-Debreu Equilibrium

Experiment	Mean Transition Probability π			
	Estimate[a]	Critical Value[b]		
		5%	90%	95%
11/01/99	0.9056**	0.237	0.7896	0.8266
11/04/99	0.7460	0.3692	0.8077	0.8732
11/11/99	0.9000**	0.5259	0.8576	0.8896
11/16/99	0.8605**	0.2542	0.7874	0.8082
11/30/99	0.7667	0.4109	0.8748	0.8909
12/02/99	0.7568	0.4614	0.8572	0.8798
12/07/99	0.5220	0.2182	0.6153	0.6662

[a] Meaning of superscripts: ** = significant at the 95% level

[b] Based on 200 bootstrapped samples of the same size as used to estimate π.

The magnitude of the estimated mean transition frequencies allow us to reject the null in only three of our experiments. Upon examination of Figure 3.6, the results are not surprising. The state-price probabilities on November 4 are generally properly ranked. On November 30, we witnessed an interesting phenomenon. Although all three states were equally likely, state Z never occurred in this session. By the fourth period, subjects seem to have believed that this state was "due," and therefore gave it a higher probability weight than warranted. This phenomenon is also observed in thick-market experiments. See BP for further discussion.

With the exception of the December 7 session, the critical values are quite high. There are pronounced differences across experiments in terms of beginning prices and empirical

---

[11] We bootstrapped the mean-corrected empirical distribution, in order to stay with the null hypothesis of a random walk.

distribution of price changes, which translate into marked differences in the distribution of π under the null of a random walk. Note that if the initial price configuration satisfies or is close to satisfying the Arrow-Debreu equilibrium restriction, the simulated πs will be high (the predicted outcome if $P_{X,t} > P_{Z,t} > P_{Y,t}$ obtains with unit frequency). This explains the high level of the 5% critical value most of our experiments. Our procedure penalizes experiments that happen to start out with prices that satisfy Arrow-Debreu equilibrium restrictions. Overall then, Table 3.4 does provide formal evidence that Arrow-Debreu equilibrium predicts price movements in experimental financial markets better than the random walk hypothesis.

Again, the rejections of the random walk hypothesis reported in Table 3.4 do not imply that the subjects ignored profit opportunities from speculating on price changes. This is because our rejection of the random walk is based on information that subjects could not condition on, namely, the distribution of aggregate wealth across states.

3.6.4 Final Holdings

We already pointed to one commonality with the thick-markets experiments reported in BP, namely, markets fully equilibrate. There is another commonality. In BP, the end-of-period securities holdings of subjects were investigated. CAPM theory predicts that subjects should all hold risky securities in the same proportion, namely, the proportion given by the market portfolio. The reason is simple: the market portfolio is the optimal portfolio of risky securities in the CAPM equilibrium. BP documents that subjects' holdings are not as predicted by the theory, and even more puzzling, that there is no convergence towards the theoretical prediction in later periods. This is particularly paradoxical, because the pricing result (CAPM pricing) is generally understood to depend critically on the allocational prediction: the market portfolio becomes mean-variance optimal only because every agent demands mean-variance optimal portfolios.

We discover the same pricing-allocation paradox in the CVT experiments. Figure 3.10 plots the holdings of risky securities in one of the experiments (December 7). The composition of the market portfolio is indicated with circles. Subjects' positions generally differ markedly from the market portfolio, and the differences do not diminish in later periods.

Furthermore, there is no correlation between those subject's who hold well-balanced portfolios and those who take advantage of the bids admitted by the CVT. This obviously begs the question: why is it that the market portfolio is priced to be mean-variance optimal, while subjects obviously are not holding mean-variance optimal portfolios?

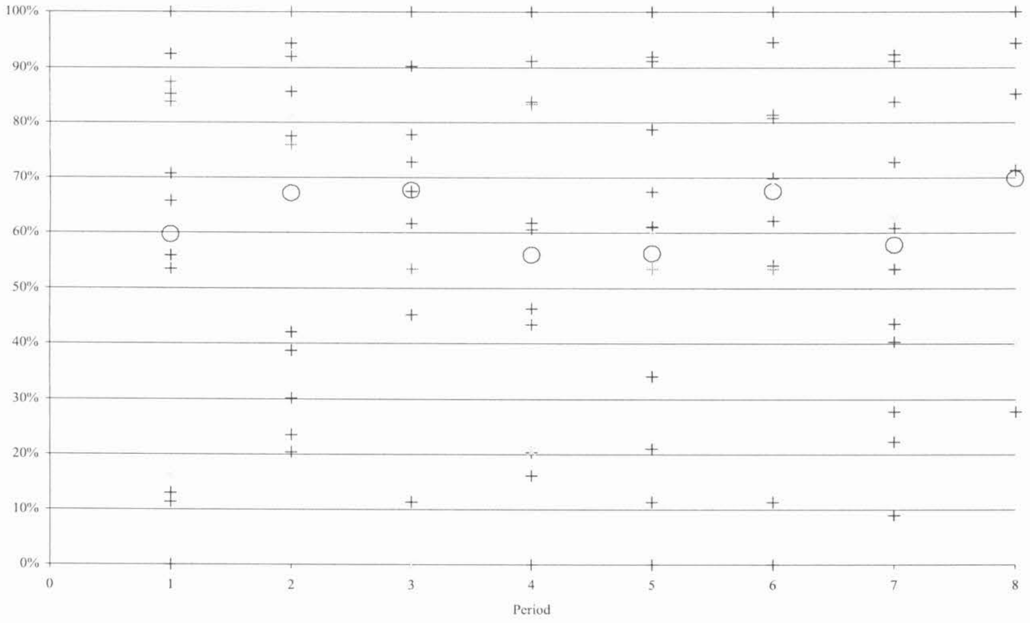

Figure 3.10: Percent of wealth allocated to asset A in both the market and individual portfolios for the December 7 experiment. Circles represent the market portfolio. While the market portfolio is near optimality, individual subjects' holdings are quite extreme.

3.7 Conclusion

In this paper, we report results from seven small-scale experimental financial markets where order submission and trading took place through a portfolio trading mechanism – the combined-value trading (CVT) system. The results were compared to those from earlier experiments when markets were organized as a set of parallel double auctions. The new mechanism eliminated the tendency of the equilibration process to stop short of the equilibrium. Volume did not disappear before the market reached full equilibrium. Distance from equilibrium was measured relative to the Capital Asset Pricing Model (CAPM).

Our results suggest that, to avoid illiquidity in thin (small-scale) financial markets, a portfolio trading mechanism should be used. The link between volume and portfolio trading is predicted by asset pricing theory, which posits that agents are not interested in securities individually, but in portfolios (packages of securities). Unconnected, parallel double auctions do not allow agents to readily trade up to desired portfolio compositions, unless markets are sufficiently thick.

This paper confirms a conjectured link between liquidity and equilibration. The finding has implications for empirical studies of asset pricing, where illiquid assets are often thought of as generating an "equilibrium liquidity premium" over and above the usual risk premium. In view of the results of this paper, it is odd to think about an equilibrium liquidity premium, because illiquid markets appear to be associated with markets that do not equilibrate, and hence, an equilibrium liquidity premium cannot be envisaged (nor can one think of an equilibrium risk premium).

In many respects, the results reported in this paper resemble those from large-scale experimental markets with parallel double auctions. This includes the allocation-pricing paradox: while prices are found to converge to the predictions of general equilibrium asset pricing theory, final holdings are markedly at odds with it. That is, the CVT mechanism is capable of generating the same qualitative results that one obtains in markets with a significantly larger number of subjects (19 to 63, instead of 6 to 15), at a significantly lower cost.

Chapter 4 An Empirical Analysis of Combined-Value Markets in the RECLAIM Permit Trading Program

4.1 Introduction

For nearly thirty years, the efficient management of pollution abatement through market solutions has been a favorite topic of public economists. In the seminal paper in the field, Montgomery [27] establishes the possibility that a tradable pollution permit system could effectively internalize the public bad of ambient pollution. Emissions trading programs allow many facilities the flexibility to choose the most cost-effective means of achieving their emissions target. But, it wasn't until the 1990 Clean Air Act that the general public was aware of market-based approaches to environmental protection. Since then, several trading systems for pollution permits have been launched around the country. While this paper will not describe the details and merits of market-based environmental regulations in general, it will explore the implementation and performance of one of these markets in particular.

Los Angeles is notorious for its poor air quality. In the Clean Air Act it was the only region in the country classified as an extreme non-attainment area for exceeding the National Ambient Air Quality Standards for ozone. Since then, the South Coast Air Quality Management District (SCAQMD) has launched a program for trading permits in Nitrogen and Sulfur Oxides (NOx and SOx) in the Los Angeles basin. This program, the REgional CLean Air Incentives Market (RECLAIM), was initialized in October of 1993 and has been operating since early 1994. The program placed minimal restrictions on how permits could be traded. Most transactions are either conducted by a broker or through one of two permit markets. In this study, we will examine the trading patterns in one of these two markets, that run by the Automated Credit Exchange (ACE) in Pasadena, California.[1]

The remainder of the paper will proceed as follows. After discussing the design of RE-

---

[1] We thank ACE for allowing us access to their data. More information on ACE can be found at http://www.acemarket.com.

CLAIM and the ACE market mechanism, we will discuss some predictions about and early results from the market. Then, we will explore the overall market evolution from April 1996 to the present. We will look at pricing and participation data over the first four years of ACE trading. We will consider patterns within the individual assets, as compared relative to the proximity of the period in which they are valid. As will become clear, there are really two markets operating in one mechanism; one for spot and short-term permit procurement and sales, and one for future planning. It is in this latter market that the sophistication of the ACE mechanism is critical.

4.2 RECLAIM and ACE

When RECLAIM was initialized, it defined a market in NOx and SOx for all of the Los Angeles Basin. These two pollutants are particularly noxious as they contribute heavily to the ozone problem. By setting emissions caps for facilities and a yearly reduction mandate for that facility-wide cap, RECLAIM was adopted to provide facilities with added flexibility in meeting emission reduction requirements and to lower the cost of compliance. The designers wanted to ensure that these caps reflected typical (recession neutral) production activity (at the time of the initial allocation, California's recent recession had caused many facilities to operate at below average production levels). Therefore, they allowed firms to set their facilities' initial baselines on the basis of actual (reported) emissions in one of four years between 1989 and 1992. This cap then declines from 1994 to 2003, after which it is constant. The 2003 target goal is approximately 58% of 1994 reported SOx emissions levels (40% of the initial RTC allocation) and 48% of the reported NOx 1994 emissions levels (30% of the initial RTC allocation). A facility is allowed to emit pollutants up to the number of RECLAIM Trading Credits (RTCs) it has been issued. Additionally, they are permitted to sell or purchase any surplus or deficit. A facility may do so either through bilateral negotiation or through the marketplace. Note that in contrast to previous emission reductions programs, there is no banking of RTCs. They are good for one year only.

In order to smooth trading behavior, two *cycles* of permits were issued. A cycle one permit for a given year is an effective credit between January 1 and December 31 of that year, whereas a cycle two permit is effective July 1 to June 30 of the following year. Half of the participating firms were assigned to each, but they are free to move between the cycles as they see fit to cover their pollution. After the expiration of an RTC, firms have an additional 60-day reconciliation period in order to review compliance and purchase any permits they still require. Additionally, separate permits were issued for each of two zones in the South Coast Air Basin. Zone one is for coastal polluters, and zone two for inland. Because of the prevailing wind currents, coastal polluters are forbidden to use inland permits but inland polluters can use coastal permits. Any given RTC is only valid for the year, cycle, zone, and pollutant for which it is designated. Initially, permits were issued for 1994 - 2010, which implies a total of 136 permits available in the initial market.

The ACE market utilizes an iterated *combined-value call market* (CV Market) to trade in quarterly trading sessions.[2] This mechanism is designed to alleviate some of the illiquidity caused by the myriad assets and relatively few participants in this market. The market mechanism used by ACE is based on research conducted at Caltech in 1993.[3] At the time of design, there had never previously been an operating combined-value call market. The mechanism allows bidders to submit contingent orders (I want A if and only if I can also secure B), matches revealed surplus, and calculates prices that leaves bidders at least as well off for having participated as if they had not. It is an iterative call market that picks a standing allocation at the end of each round, lasting a total of three to five rounds. To determine the exact number of rounds, ACE uses an improvement rule mandating that if surplus and volume do not increase by at least 5% after each iteration, the market will end and transactions will be made based on the last iteration.

Recently, the ACE market has accounted for approximately 75% of non-zero price trades in the RECLAIM market, and this percent is increasing (as we will discuss in Section 4.4.) In 1999, RECLAIM recorded 239 non-zero price trades (219 in the NOx market and 20 in the SOx market). ACE conducted 213 of these trades, 208 of which were for NOx permits. We are only interested in trade for a price, since zero price trades are either inter-facility

---

[2] Although there are sometimes as many as five or six sessions per year, and recently monthly markets have begun

[3] For a detailed description of the design and experimental testbedding of the ACE mechanism, see Ishikida et al. [19].

or transfers to or from a permit broker, and are therefore not truly part of the competitive market. We will use ACE data in this paper for two reasons. First, since it conducts the vast majority of priced trades in RECLAIM, it offers a rich environment in which to study the program. Second, since ACE is the first combined-value market of its kind, a careful examination of these data will tell us about the potential of combined-value market mechanisms for tradable pollution permit markets.

We will examine both bidding and pricing behavior exhibited by the participants in ACE and the ACE market mechanism, respectively. We conduct this analysis both over the life of the ACE Markets and across the permit vintages. It is important to look at the two time trends separately. First, prices of a given future are increasing over time. This reflects the decreasing number of RTCs available relative to reported emissions, as planned by the RECLAIM project. The second, is at any given session, price increases in the length of time until the RTC vintage is effective for covering emissions. For example, in 1995 the cheapest RTCs available were to cover emissions in 1995. The price for a 1996 RTC in the 1995 market was more expensive, and so on. Permits sold in 1995 for emissions beyond 1998 were fairly stable. We will discuss some reasons for this in Section 4.8.

The ACE data explored in this paper cover the markets from April 1996 to January of 2000. Markets were held at-least quarterly, with some years having 5 or 6 markets. Specifically, the data are from the following auctions: April 1996, July 1996, August 1996, October 1996, February 1997, April 1997, July 1997, October 1997, January 1998, April 1998, July 1998, October 1998, January 1999, April 1999, July 1999, August 1999, October 1999, and January 2000. Very recently, some pricing data have become available for the April 2000 market, and we will use these data when applicable.

4.3 Predictions

Karl Hausker, Chief Economist for the U.S. Senate Committee on Energy and Natural Resources, conducted extensive pre-market research on what we should expect from emissions trading programs such as RECLAIM. He suggested that demand and supply in spot and short-term markets arises from fluctuations in load growth, plant availability, and the price and availability of various fuels. In the long-term market, demand and supply are driven by load-growth and the need for new plants, and would therefore likely be for streams of allowances and demand in the long-term market. The builder of a new plant would seek a 10 or 20-year supply of permits in order to guarantee his planned production stream. In a similar fashion, a selling plant could generate a stream of surplus allowances at a later date.

In the former scenario, the short-term planner relies on the short-term market and accepts the risk of price changes. In the latter, he locks in prices by turning to the long-term market. Hausker commented that, "In reality, the market is unlikely to function in this idealized manner." His concern was that many sources of market inefficiency (cost-of-service regulation, regulatory uncertainty, etc.) would tend to reduce transactions. This reduction would be felt most acutely in the long-term market. Additionally, markets would be too thin, particularly in early years. If firms see allowances as a key to load growth, the number of sellers in the long-term market would shrink more quickly than in the short-term markets. A risk-averse utility will be far more willing to sell allowances to be issued in the near future than to sell a 20-year stream (much less a perpetual stream).

Furthermore, risk aversion plays a major role in environmental planning. In face of regulatory uncertainty, a player may want to avoid major changes in allowance holdings. This further encourages short-term buying and selling and discourages transactions in the long term market. Additionally, the Public Utilities Commission must approve many long-term changes at emitting facilities, and the approval process at the time of the transaction discourages long-term transactions relative to short-term ones. A transaction in the long-term market will typically involve a new plant or major emission-control decision and involve far more money than a short-term transaction. Ironically, risk aversion would also make long-term transactions a necessity, as a firm is unlikely to build a new plant without possession of the necessary stream of allowances. Similarly, a potential seller would be unlikely to rely solely on short-term markets to help recoup its investment in emission control. Hausker was concerned, therefore, that long-term market is most susceptible to inefficiencies.

4.4 Buying Patterns and the Evolution of the ACE Market

The RECLAIM market is a very new one, and as such it is still evolving. As firms learn about the transaction costs and pricing inherent in the various trading options available to them, we expect to see movement of trade towards those outlets where trade is most easily facilitated. Therefore, an important measure of the appeal of the various market (and non-market) mechanisms is the percentage of firms that chose one outlet for trading RTCs over another. The number of participants in the RECLAIM universe has not changed dramatically since its inception, and has in fact declined.[4] As time progresses, more and more of the members of the RECLAIM universe are using the ACE markets than any other trading option, as the below description of Figure 4.1 will demonstrate.

When examined on a session-by-session basis, it appears that the number of participants in the markets is not changing over time. In the upper frame of Figure 4.1, we show the number of bidders and winners participating in the each session of the ACE markets from April 1996 to January 2000. Even though the number of RECLAIM facilities has not dramatically changed over this time period, the number of bidders and winners in the ACE markets is increasing when examined on a yearly basis. In fact, while only 10% of 1996 RECLAIM universe members conducted trades-for-a-price through ACE, 20% did so in 1999. Based on January 2000 data, we can expect ACE to conduct trades on behalf of 31% of RECLAIM universe members in 2000. This is even more dramatic when coupled with the fact that 213 of 219 (over 97%) of trades for a price in 1999 were conducted via ACE. As the RECLAIM program progresses, environmental engineers learn that the non-market mechanisms for trading involve extreme transactions costs relative to the ease and efficiency of the ACE market. As shown in the bottom frame of Figure 4.1, over time they are more likely to trade with ACE. Why are the number of participants, when considered on a session-by-session

---

[4] When RECLAIM was adopted, 394 facilities were identified as the initial "universe." Between inception and June 30, 1998, 18 facilities have been included, 61 excluded, and 25 facilities have ceased operation. Thus the RECLAIM universe consisted of 326 facilities on July 1, 1998. By the end of the 1998 compliance year (June 30, 1999), the RECLAIM universe contained 331 facilities. See the SCAQMD Annual Report for further details [31]).

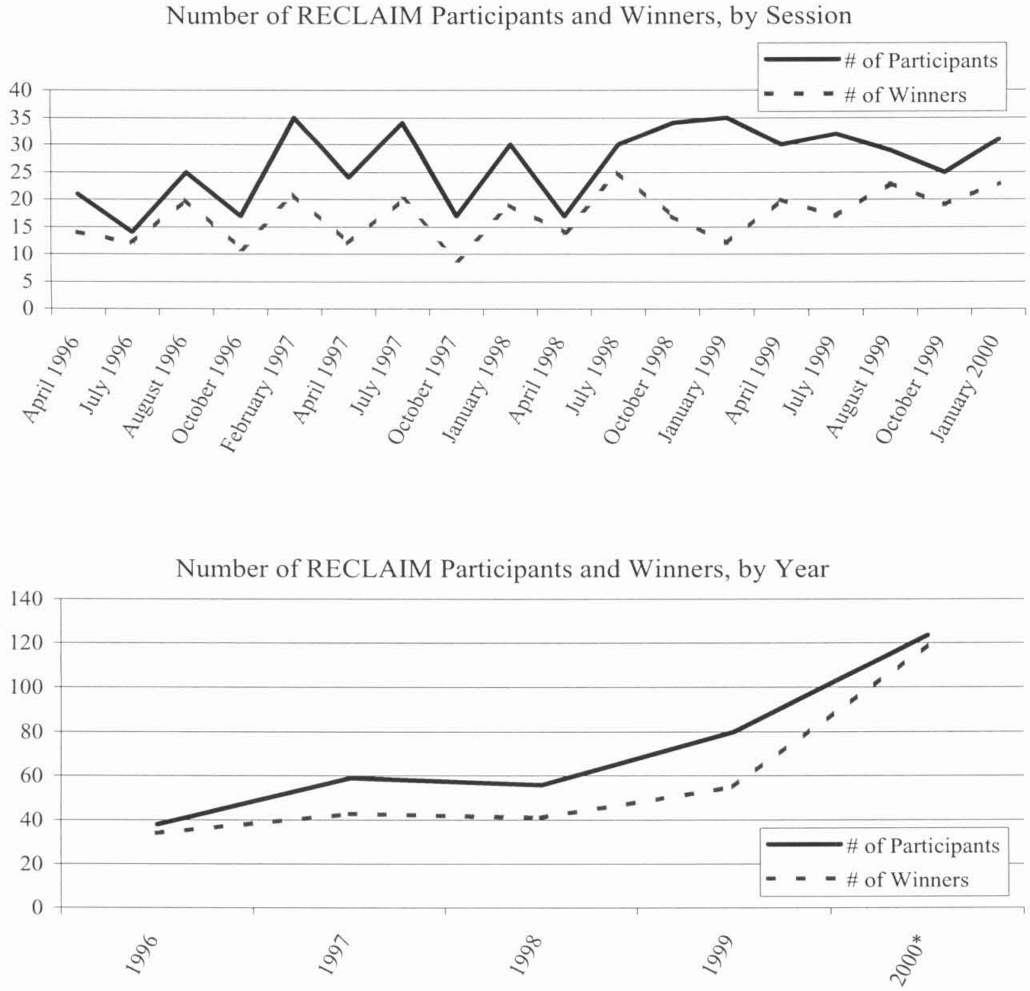

Number of RECLAIM Participants and Winners, by Session

Number of RECLAIM Participants and Winners, by Year

Note: 2000 values are based on January 2000 data, divided by the proportion of participants and winners represented relative to the entire year in 1997, 1998, and 1999.

Figure 4.1: Participation in the ACE market, by session and by year

basis, so *constant*, and the number of yearly participants increasing so dramatically? The two annual cycles established by RECLAIM were designed exactly with this goal in mind. That is, having two cycles has encouraged smoother trading behavior over the years, as was the intention.

As discussed in Section 4.2, the combined-value mechanism employed in ACE allows bidders to express very complex preferences in their bids. Figure 4.2 describes the percent of bids submitted and of those fulfilled that are for more than one asset (combined value or *package* bids). Notice that the number of package bids submitted over time is relatively steady, and that package bids are roughly as likely to transact as are single-asset bids. On average, 18% of bids submitted are for packages, as are 14% of fulfilled orders. As we will see later, we are far more likely to observe package bidding in the long-term planning market than in the short-term market.

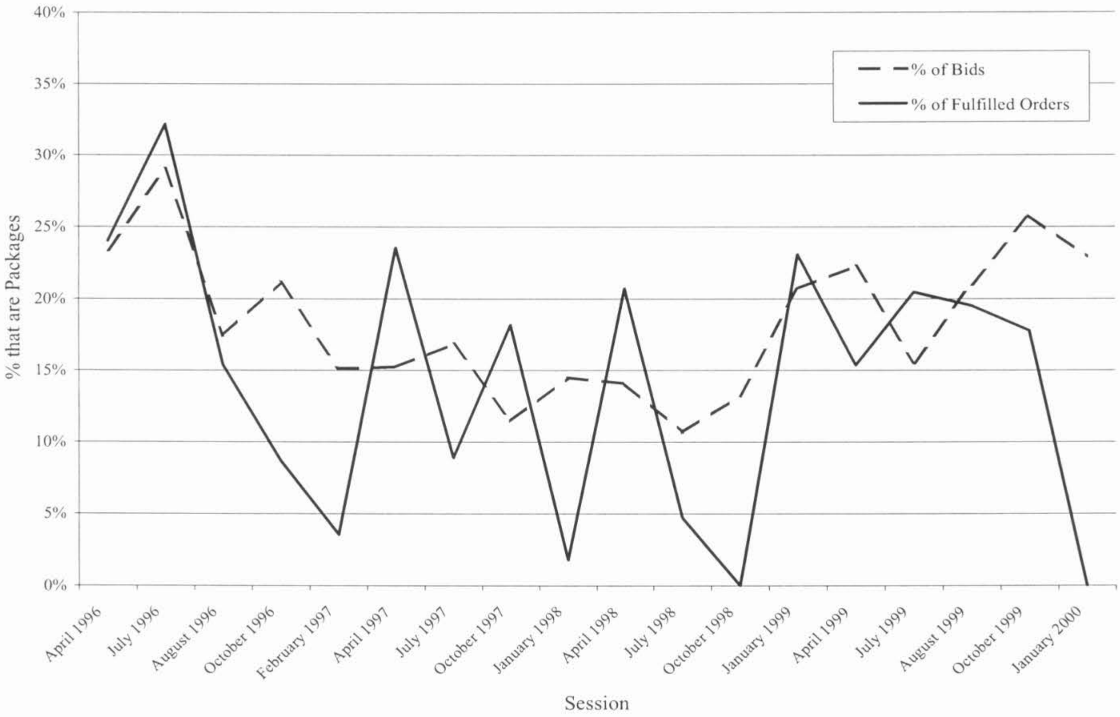

Figure 4.2: Percentage of bids and fulfilled orders that are for packages

Figure 4.3 shows the reported emissions and total RTCs issued in both the NOx and SOx markets. The initial supply of RTCs to the market was based on facility-chosen emission levels between 1989 and 1992. As you can see, the initial supply of RTCs was in excess of the emissions in 1994. The supply was reduced each year in order to achieve compliance with California and U.S. ozone standards. In 1997-98, the supply of NOx RTCs decreased below the levels declared in 1994 and this cross for SOx occurred in 1999. As Figure 4.4 shows, the volume of buy, sell, and swap bids submitted to the ACE Market since its inception reflect the underlying supply and demand. With the exception of the first few trading sessions, the volume of sells is decreasing as the excess of RTCs over emissions decreases.

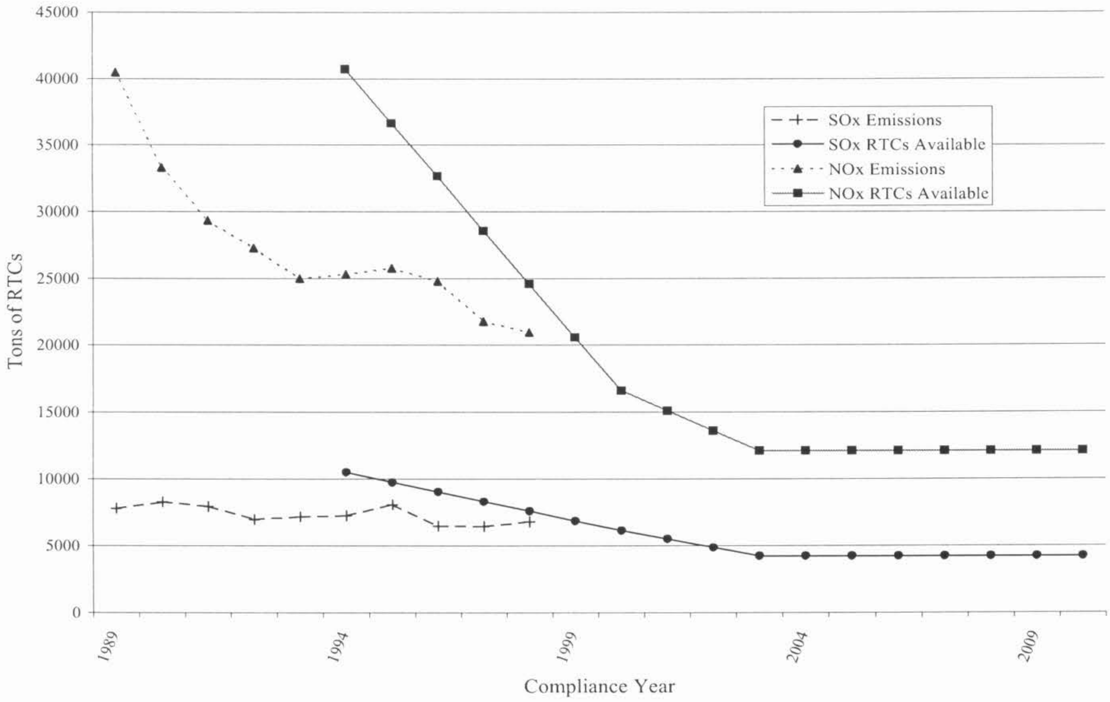

Figure 4.3: Reported emissions and initial allocations in the RECLAIM NOx and SOx markets

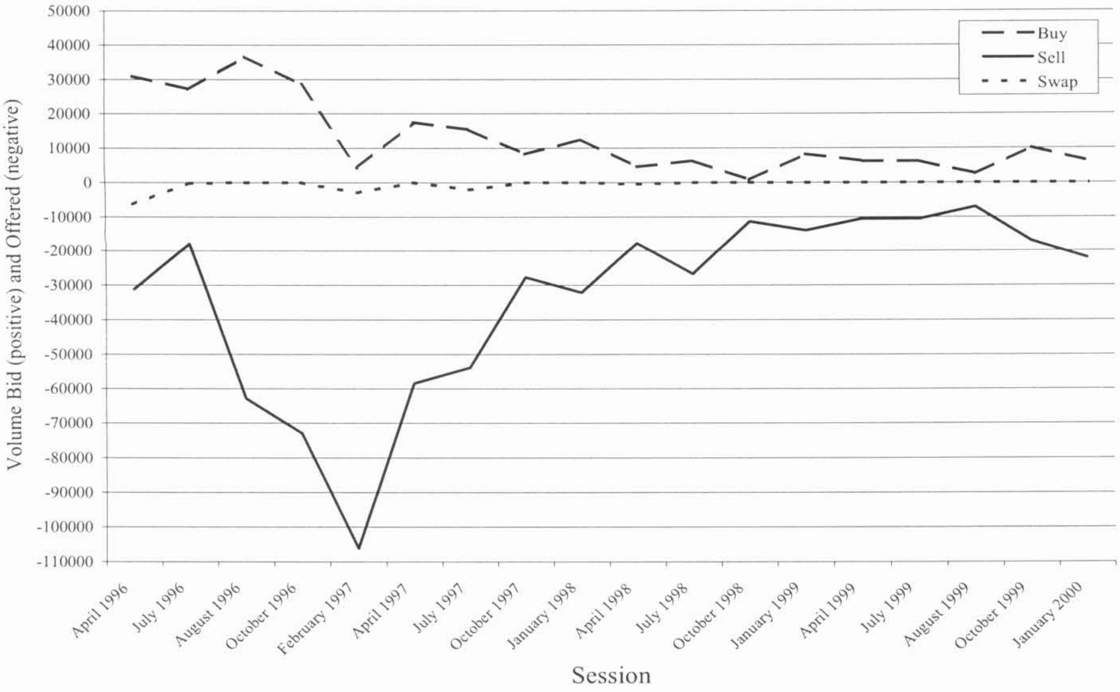

Figure 4.4: Number of buy, sell, and swap bids submitted, by session

4.5 Pricing Patterns During Market Evolution

One way to get a better understanding of market participants' behavior is by focusing on a subset of the overall market activity. For the moment, we will focus on the market in the most current RTCs available for purchase, known as the *spot market*. As previously mentioned and shown in Figure 4.3, the initial supply of RTCs was in excess of the reported emissions at the market's inception in 1994. The supply of RTCs was reduced each year (an average of 8.3% per year for NOx and 6.8% for SOx), and in 1997-98 this supply dipped below the total 1994 emissions level. At this point, prices began an uphill climb. Figure 4.5 shows the evolution of the NOx and SOx spot market price over the time period April 1996 to January 2000. As you will notice, there is almost no trade in any of the markets except for NOx Zone 1 (we will discuss some reasons for this below), so henceforth we will use the data from this asset to demonstrate the patterns that emerge. As is evident, the price of RTCs dramatically increases shortly after the expected cross of RTC availability and reported emissions. However, even this highest price falls quite short of the AQMD's anticipated price of $11,257 per ton. Furthermore, the price acceleration we observe is far steeper than that predicted by economists prior to market inception. We will discuss these predictions and the discrepancy between them and the realized prices now.

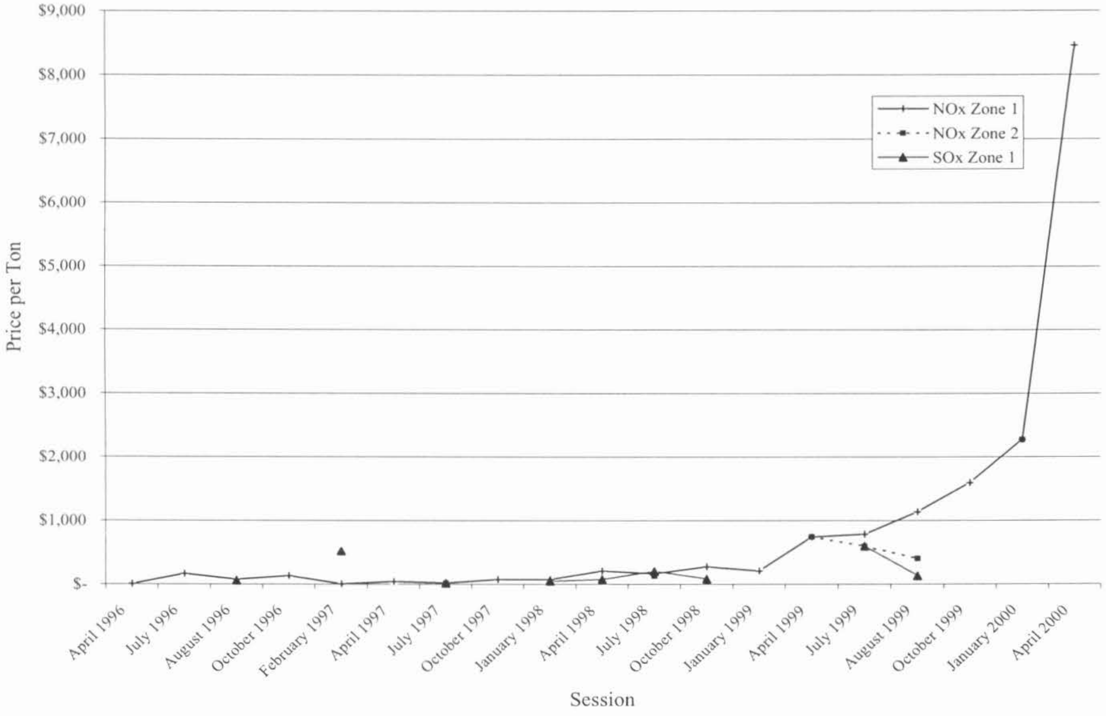

Figure 4.5: Spot market pricing in ACE markets, April 1996 to April 2000

After the first two years of trading, Klier et al. [21] authored a study predicting patterns we might observe and examining the early data on NOx trading in the RECLAIM program. They expected slow activity in the market at first, as the initial allocation of permits was well above the reported emissions levels (see Figure 4.3). Additionally, they suggest a correlation between the price of RTCs and the marginal cost of pollution abatement, stemming from the dual effect of the presence of permits in the market. Permits both decrease the marginal cost of pollution abatement and increase the marginal cost of production (the latter effect is because holding an RTC represents an opportunity cost to the firm). Therefore, the price of RTCs will always be less than or equal to the marginal cost of abatement at any production level. They find that in excess of 54% of trading is for the permit just expiring (so, the majority of market activity is to balance the books). Lastly, they predict that market activity will increase as the emissions cap decreases over time.

Prior to RECLAIM's implementation, Johnson and Pekelney built the Emissions Trading Model (ETM) to assess the potential economics and environmental impacts of RECLAIM's emissions trading program. It estimates trades that are likely to occur under the program and links to a general equilibrium model of the regional economy. Figure 4.6 shows the predicted price patterns they expected to observe in the first few years of RECLAIM permit trading, translated to 1999 dollars.

In reality, permit prices have not yet achieved these predicted levels. Figure 4.7 shows prices for various vintages of permits, both SOx and NOx, for zones one and two.[5] As we can see (by comparing Figures 4.6 and 4.7), prices have stayed quite low. Assuming these prices accurately reflect the marginal control costs as argued in Section 4.3, it appears that the planning flexibility offered by the RECLAIM program resulted in lower than expected marginal control costs, perhaps from the shift to a facility-wide performance standard (see Bohi & Burtraw [3] for an excellent discussion of the impacts on control costs). Additionally, prior to the start of RECLAIM, facilities had incentive to misrepresent their emissions and true costs of abatement, as we will discuss below. However, April 2000 data indicates that prices are now approaching the predicted levels. The marginal cost of "Best Available Control Technology" (or BACT) for NOx is believed to be in the range of \$3.50 to \$4.50 per pound. The spot market price for permits in the most recent (April 2000) market was \$4.23. We are now just reaching the market transition point that the ETM believed would occur in 1996.

Why is it that the SOx market has been so thinly traded and the prices are so much

---

[5] Because there are two different RTCs that cover a given year (cycle one and cycle two), a *compliance year* is defined to be both cycle one and cycle two of that year. Therefore, there is a 6-month overlap between each annual data point. That is, 1998 comprises January 1, 1998, to June 30, 1999. The RECLAIM data is taken from the SCAQMD Annual Report for 1999, and includes all reported non-zero trades.

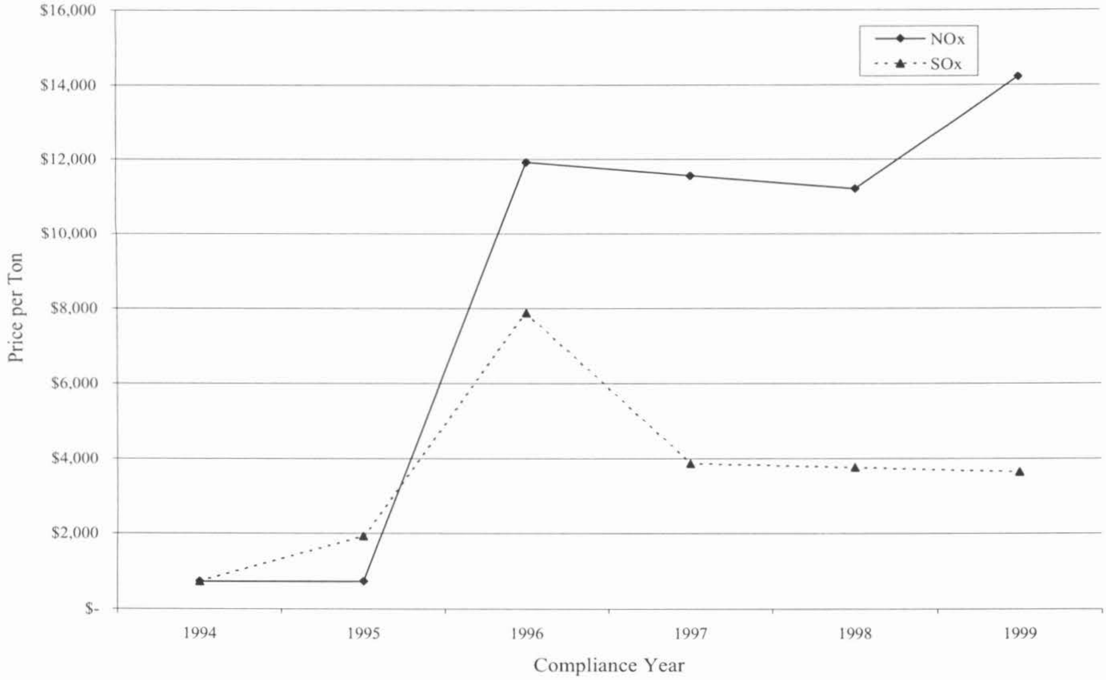

Figure 4.6: Prices predicted by the Emissions Trading Model (1999 dollars)

closer to predicted levels (although still significantly lower)? In the last 30 years, the Los Angeles area has engaged in an extremely aggressive SOx abatement program, and the levels of SOx emissions are low and stable, particularly when compared to the NOx emissions levels. Furthermore, partially due to this aggressive program, sulfur dioxide abatement technology is well-established. Since all SOx emitters are well aware of the available abatement technology and have made their decision about abating versus procuring permits, there are no lower-cost abaters in the permit market to drive trade with higher-cost abaters.

On the other hand, the NOx problem in Los Angeles and the lack of available effective abatement technology was still a major issue at the inception of the RECLAIM market. Technology is still advancing, and as it does so, the technology gap between firms drives trade. Additionally, prior to the existence of the permit trading program, source owners

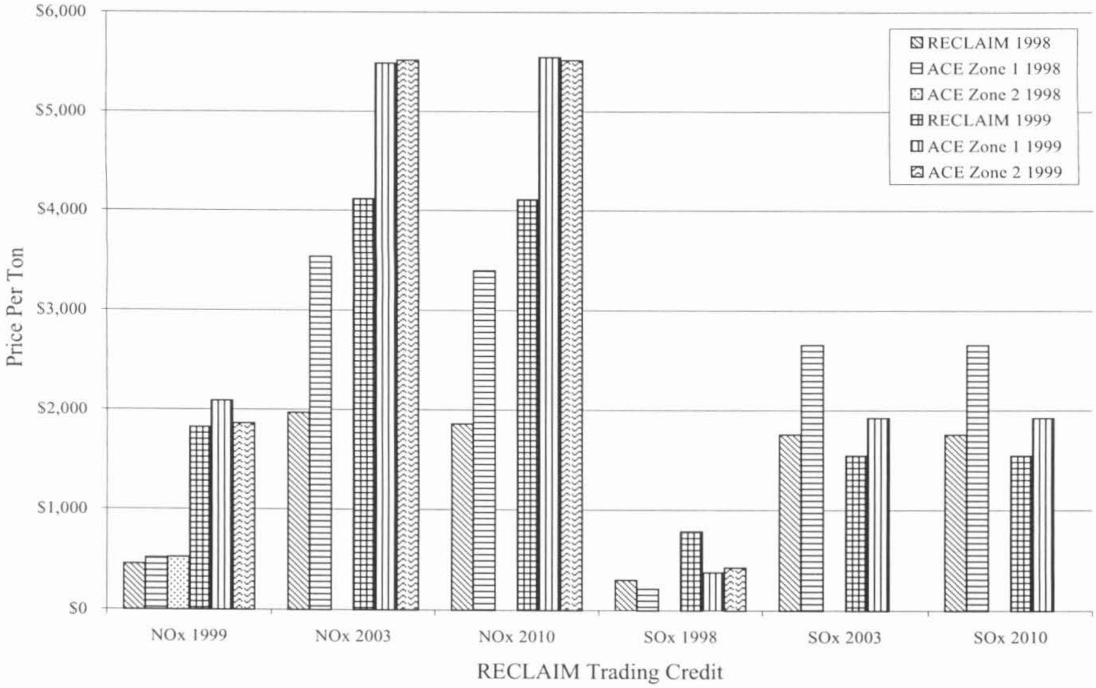

Figure 4.7: Realized prices of various vintages of RTCs in 1998 and 1999

had little incentive to reveal the true cost of abatement; indeed the incentive historically has been to overstate the control costs during public hearings in order to deflect proposed command-and-control type regulations. This strategic reporting further accounts for the gap between predicted and actual RTC prices.[6]

4.6 Bidding and Trading Time Horizon

In each session, there are over 100 assets available for purchase. Because at any given time, a trading session allows trading for permits spanning at least the next decade, these myriad available assets are constantly changing. If we consider the RTCs as different assets because

---

[6] Existing source demands for RTCs is projected over time and are based on historical reported emissions under the command-and-control system.

the date of effectiveness is different, we have a very thin data set from which it is difficult to perceive a pricing pattern. Rather, we consider these not as dated objects (in terms of their year of validity) but as a stream of future contracts dated *relative* to the current market.

Once we make this adjustment and consider the permits in terms of their relative vintage, patterns emerge that indicate there are really two markets in one here. The bidding patterns and prices in the spot and short-term market are in stark contrast to those exhibited by the long-term planning market. We will explore these differences now. We will only consider the spot and futures markets up to 17 permits in the future (for a total of 9 years of forward contracts in any given trading session).[7]

There is a significant stream of futures available to a RECLAIM trader. It is not clear, however, that all traders take advantage of this futures market. As mentioned earlier, the vast majority of bids and trades are in the spot market. In fact, the number of bids entered is inversely proportional to the time horizon. In Figure 4.8, we show the number of bids and fulfilled orders in the market as a function of the vintage of the earliest item in the bid. As we will see, the short-term, heavily traded market exhibits quite different attributes than the long-term. They operate as two simultaneous, but distinct, markets in one.

4.7 Bidding Behavior

The ACE market is a powerful institution, allowing for bids that describe very complex preferences. Are the bidders using these options, or are the complex bidding features of the ACE mechanism simply window-trimming? Ishikida et al. [19] provides a detailed explanation as to the motivation a firm may have to submit contingent bids. The myriad options a firm's environmental engineer faces present a very complex financial problem. Given his chosen level of emissions, he has a large set of choices as to how he complies to the RECLAIM rules. He can purchase abatement equipment, buy and sell permits, or some combination of these. For example, he could delay abatement and buy short-term permits until installation of abatement equipment and sell futures after the abatement installation date, or any other

---

[7] Beyond 17 permits in the future, the data set is quite thin. For later sessions we do not have data for as long a time horizon as in the earlier sessions.

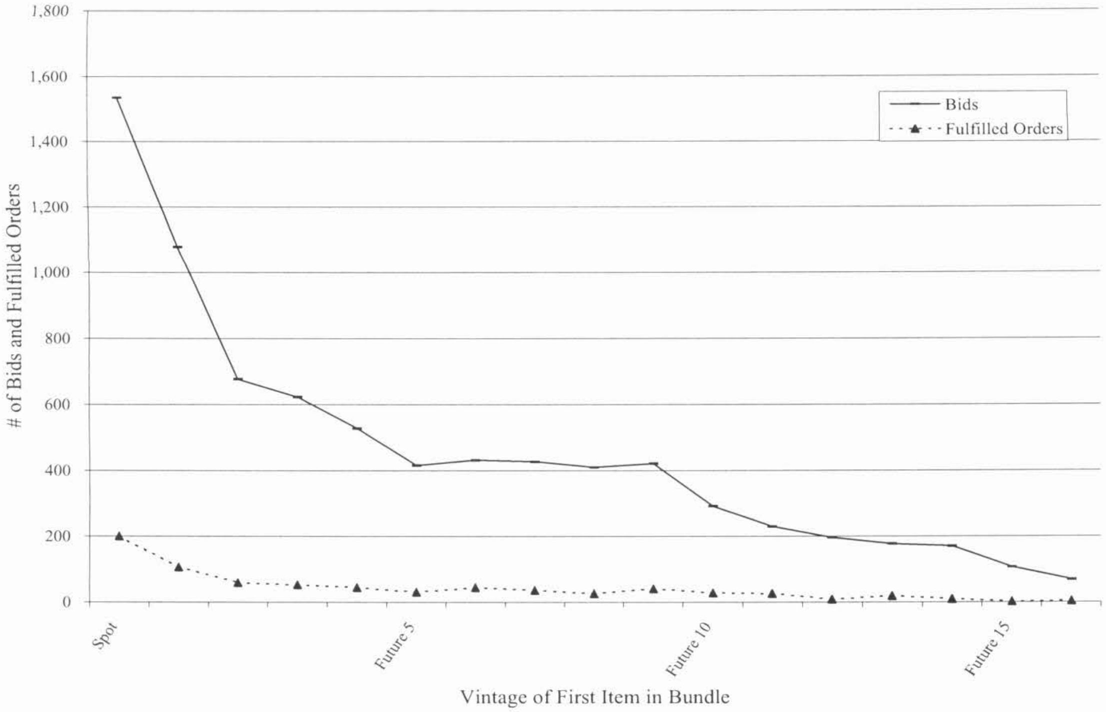

Figure 4.8: Relative vintage of assets in bids and fulfilled orders

combination that leaves him compliant. The optimal decision depends on the prices of the permits, as he could choose many intertemporal combinations. Further, to avoid market risk, he would like to make these decisions as a package so that he does not expose himself to the risk overpaying for the entire plan as the prices change. For example, if the ability to cover the pollution from a new plant for the next 20 years is worth X dollars, he can bid up to X for the package of permits and know that he will not be charged more than this if he receives the desired stream. Whereas, in a traditional set of parallel markets, one for each asset, he would have to bid in each of 20 separate markets and hope to pay an average price no greater than $X/20$ for each. Since in parallel markets it is impossible to submit an order that depends on events in other markets, long-term planning is a far riskier venture. A planner might be hesitant, for example, to build that new plant under a parallel market structure.

Given that agents have this power to express their preferences well via the ACE mechanism, it is surprising to see the modest use its power receives. But, upon studying bidding behavior we see two distinct patterns of bidding emerge. Figure 4.2 shows that, on average, 18% of submitted bids and 14% of transactions are for packages. However, as Figure 4.9 shows, for bids in the spot market (the earliest vintage currently available), approximately 91% are for a single asset. Additionally, there are no swappers at all in the spot market. That is, of those 9% of bids that are for more than one asset, the bids are either pure buys or pure sells. The spot market is never a time for rebalancing.

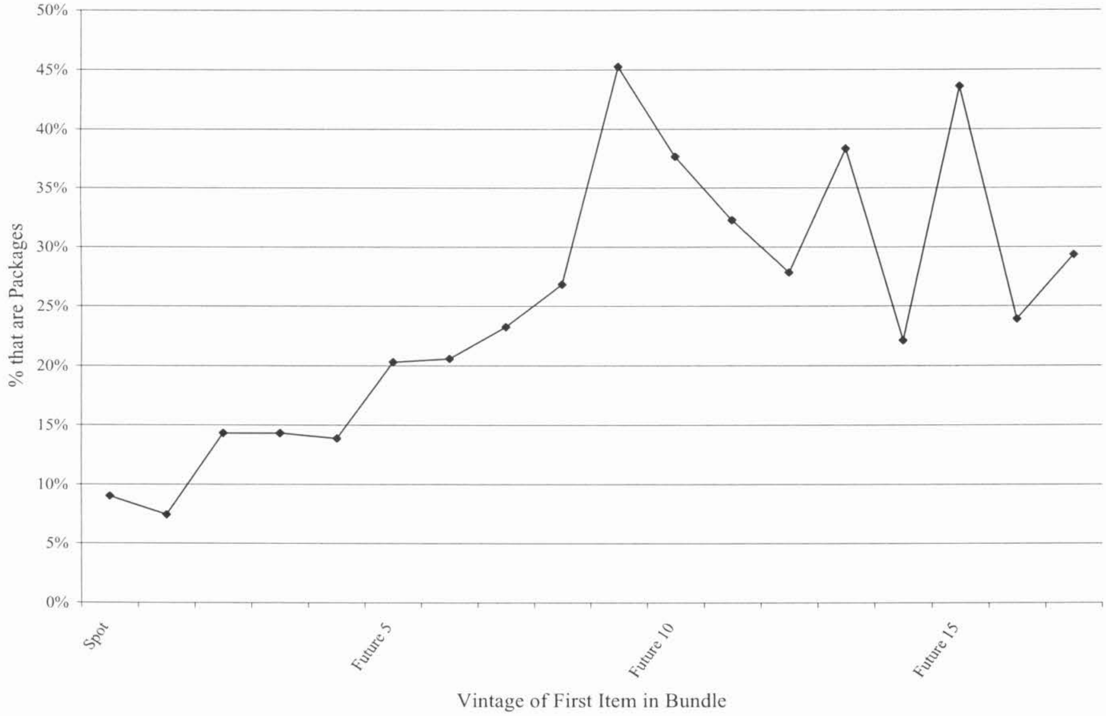

Figure 4.9: Percent of bids that are packages, by relative vintage

It is important to note the discrepancy in bidding styles between the short- and long-term markets. If we only look at the long-term market, 32% of bids are combined value (for permits more than 3 years in the future). So, we have a lot of simple trades in the spot and short-term markets not using combined-value bidding to do so. In the longer-term market the bidding is far more frequently for packages. In fact, the frequency of package bidding in the ACE market is quite similar to that observed in the experiments discussed in Chapter 3.

The pattern of bidding behavior in the short and long-term ACE markets confirms Hausker's predictions of appropriate market behavior, and refutes some of his concerns about risk aversion and market thinness in the long term market (see Section 4.3). Indeed, we see two markets emerge, with very different bidding behavior. By Hausker's argument, it is reasonable to expect bids in the short term market to be simple, whereas the long-term bids are complex as they represent production and facility planning. As we shall now demonstrate, the pricing patterns in these two markets are also quite different from one another. Furthermore, the long-term market exhibits the stable pricing Hausker feared impossible.

4.8 Pricing Patterns in Spot and Future Markets

As Figure 4.10 shows, asset pricing is dramatically different in the short and long-term markets. From the spot market to about 7 permits ($3\frac{1}{2}$ years) into the future, the price of permits increases. Once we look to 4 years in the future, and for the entire foreseeable future thereon, the pricing is remarkably stable. This is because the type of bids that are submitted in the long-term market are for packages, and often a host of prices will satisfy the ACE mechanism. If there are multiple prices that will satisfy the mechanism's constraints, it picks the vector of prices that distributes surplus evenly. That is, ACE finds the competitive equilibrium price that maximizes net surplus to the buyers, find the price that does so for the sellers, and take the midpoint asset-by-asset. The combined-value mechanism provides liquidity necessary for stable, meaningful prices in the long-term market.

Among other patterns, the data reveal a consistent phenomenon – the nearer the permit is to expiration (defined as the end of the 60-day reconciliation period), the lower the purchase price. Why might this be?

Although permits are in abundant supply, a firm's pollution is not always entirely predictable. In addition to the production and input uncertainty mentioned above, firms' pollu-

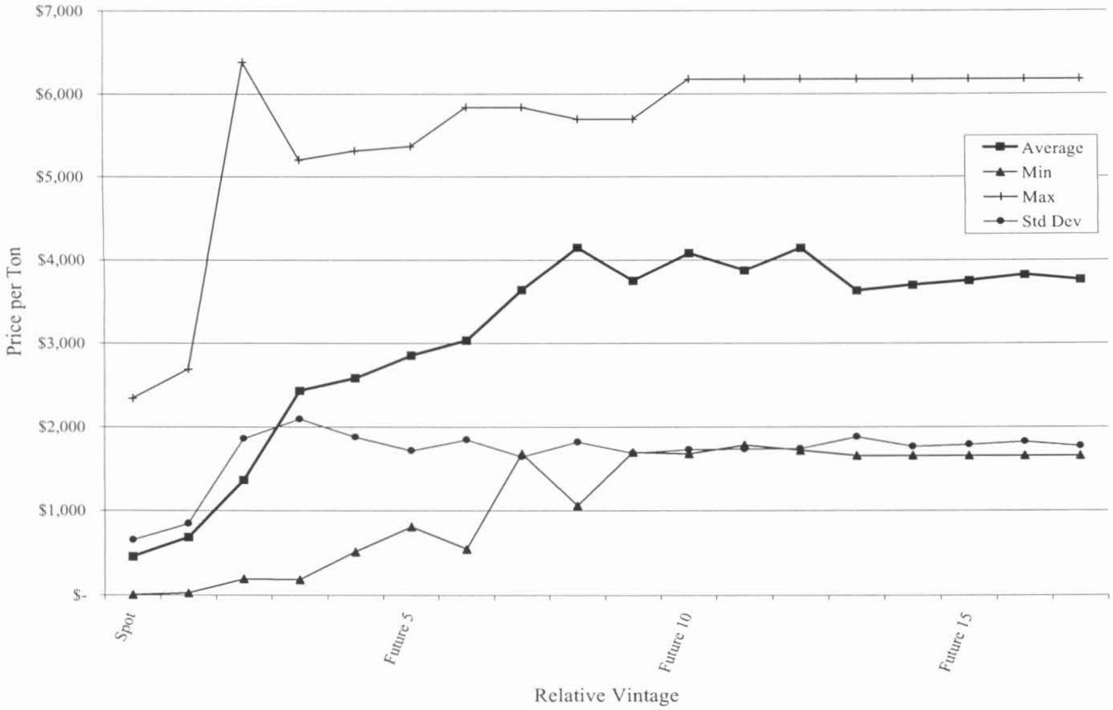

Figure 4.10: Price of NOx Zone 1 RTCs, as a function of relative vintage

tion levels are susceptible to variables such as weather and prevailing air currents. Therefore, a permit buyer may want to purchase extra permits early in the process, but late enough that he can estimate needs well. Furthermore, a short-term buyer is more likely to be covering an anomaly (production spike, weather patterns, etc.), since he has already completed his horizon planning, and has accounted for all routine needs. This is borne out by the infrequency of short-term CV bids, as explained in the previous section.

If, during the reconciliation period, all RTC needs for the cycle expiring are known, those facilities that are selling have no need for their excess RTCs. Furthermore, as Figure 4.11 shows, there is generally a higher volume of RTCs for sale in the spot market than there are bids. Therefore, we would expect the price of RTCs in the reconciliation period to be driven to zero. While prices are still quite low, they are still definitely different from this theoretical price.

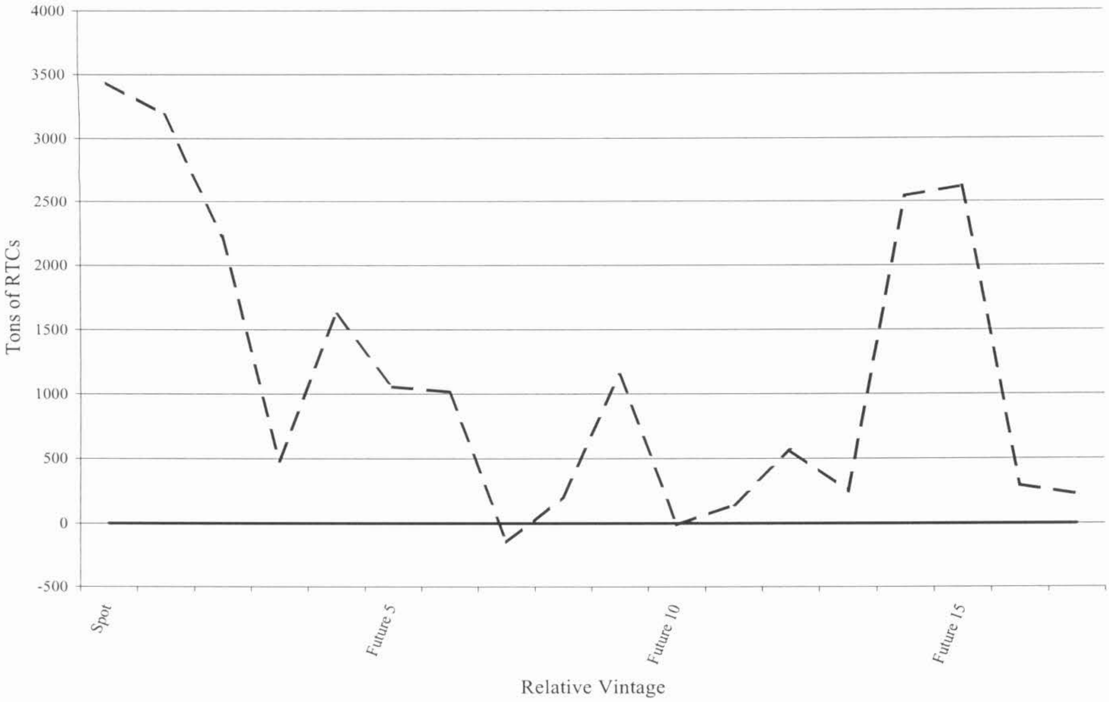

Figure 4.11: Surplus of volume offered, as a function of relative vintage

Why do we observe non-zero prices in the spot market? While the number of RTCs available *across both cycles* has historically been in excess of reported emissions, this does not mean that in any given market there is an excess supply. Firms were either assigned to Cycle 1 or Cycle 2, and then allocated RTCs. Therefore, it is quite possible that there may not be enough permits of a given cycle to cover that entire year's emissions without benefit of permits from another cycle.

We, and all facilities in the RECLAIM universe, know that there will eventually come a time when there are less RTCs allocated than the historical emissions level. At this point, there will be a shortage of RTCs, and the price will be driven up. Call the first cycle in which this occurs time t. RTCs for the previous cycle, call them $t-1$, are abundant relative to time t permits, and can be used to cover pollution in that period. Therefore, owners of t RTCs will purchase $t-1$ RTCs in that spot market, supporting a positive price for $t-1$ permits even though they are in abundant supply.

Now, induct backwards. The facility holding $t-1$ permits recognizes the value of his permits to the t facilities. He will buy spot permits in the market prior to his, loading up on $t-2$ RTCs. Now, he can in part cover his $t-1$ emissions with $t-2$ permits, and is free to sell his $t-1$ permits to the t players. This supports a positive price for $t-2$ permits in the spot market.

This begs the question, why doesn't the owner of the $t-1$ permits submit a package bid in the $t-2$ spot market, swapping the purchase of $t-2$ permits for the sale of $t-1$ permits? There are two reasons, both due to the increasing prices we reported earlier. The first is that, as prices are increasing (for all vintages), he may prefer to hold some extra RTCs rather than being exposed to the prices in the $t-1$ spot market should he find himself needing more permits to cover unanticipated emissions. The second is that, if the $t-1$ facility believes that the price of $t-1$ permits in the spot market will be greater than the price they would fetch in the $t-2$ spot market (when the $t-1$ permits are a future permit), then it may pay to wait and sell in the latest market possible.

Since July 1997, the spot market price has been strictly increasing over time. This is because facilities look down the road and recognize the impending shortage of RTCs, and begin rolling permits forward in the manner described above. Since this time, half of the sessions have exhibited the second above property (that the price of a permit in its spot market is greater than that same permit fetched in the previous market, when it was for the future). Furthermore, 4 of the last 6 markets have exhibited this property. Since abatement technology has been more rapidly deployed than expected (as discussed in Section 4.5), the shortage has been pushed further in to the future than facilities may have anticipated. Therefore, it would have been ex ante reasonable for facilities to believe that the pricing property would have held earlier, more often, and with greater force. In fact, we have now just entered the time when the shortage is binding and the prices are skyrocketing, as discussed in Section refsec:pricing and exemplified in Figure 4.5.

4.9 Conclusion

The RECLAIM market was implemented in the early 1990s with the goal of providing a market solution to the problem of ozone-depleting toxins in the Los Angeles Basin. As the emissions cap has reduced since then, the prices of RTCs has increased. This reflects the increasing marginal costs of abatement at these lower levels of emissions. While prices are increasing, they are still well below the prices predicted by the Emissions Trading Model. While this is in part due to mis-reporting of abatement costs on behalf of polluting firms under the previous command-and-control regime, it also reflects the increased abatement efforts that the flexibility of the RECLAIM program encourages. The stable pricing and low transaction volumes in the SOx market (with its well-established abatement technology) in contrast to the volatility and high volume in the NOx market further emphasizes the impact of abatement technology penetration on RTC pricing. However, this abatement technology can only take the market so far. As of the April 2000 market, the spot market permit price of \$4.23 per pound is aligned with the current beliefs of Best Available Control Technology (BACT) prices of \$3.50 to \$4.50 per pound. The purpose of the RECLAIM program is to allow facilities to minimize their cost of compliance by choosing to abate, purchase RTCs, reduce or shift production, or some combination of these. Any of these options other than the purchase of RTCs requires advance planning and can take years to implement. Those who have not begun their planning for RTC prices that are now in line with BACT will pay exorbitant prices for RTCs in the near term, and we can expect the price to increase further, past the price of emissions control. However, once capital equipment adjustments have been made, prices should equilibrate to the marginal cost of the BACT.

Karl Hausker eloquently voiced many concerns relevant to a permit market such as RECLAIM. In particular, he was concerned that the long-term market would suffer from extreme thinness due to uncertainty, transactions costs, and other sources of market inefficiency as enumerated in Section 4.3. Although, as Figure 4.9 shows, the long-term market does seem thinly traded, the combined-value market mechanism used in the ACE market overcomes this illiquidity, as predicted in Chapter 3. The combined-value design allows environmental engineers to plan for long-term production and emissions and to do so with more security as to the value of their investment.

The ACE market mechanism has, over the first four years of the RECLAIM emissions credit trading program, become the market venue of choice for polluting facilities. The short-term market and long-term markets exhibit dramatically different bidding and pricing patterns. The short-term planner places single-asset bids, making one-time adjustments to his predicted emissions levels. For the long-term planner, the combined value market provides clear, stable pricing and the ability to plan a pollution stream without the risk of only assuring part of the stream. Since a single-asset bid is simply a degenerate form of combined-value bid, imposing the combined-value structure on the short-term market certainly does no harm. Indeed, the additional liquidity provided by those few traders who trade in bundles including both short-term and long-term RTCs provide an important bridge between the two markets, improving liquidity in both.

Chapter 5 Conclusion

The previous three chapters have concerned themselves with the possibility of solutions to pollution abatement programs and the manner in which they should be implemented. In Chapter 2, we established the existence of a cooperative solution to transboundary pollution problem in certain pollution configurations, those that can be represented as a tree structure. This type of pollution flow is observed in the LA basin, as well as other regions where ambient air quality is an issue of growing concern.

For years, economists have favored the idea of tradable permit markets to implement pollution abatement. However, a well-designed pollution permit program has to overcome many sources of market illiquidity if it is to perform well. By performing well, we mean a market that distributes pollution emissions evenly over an area, brings the prices of permits in line with the marginal cost of the Best Available Control Technology (BACT) at any given time, provides clear pricing signals and encourages sufficient liquidity to encourage long-term planning and trading. It was with these principles in mind that the ACE market mechanism, a combined-value trading program, was designed.

We tested the ability of a combined-value mechanism to overcome market thinness in Chapter 3. As compared to other experimental evaluations of market designs, the combined-value mechanism performs remarkably well. Liquidity is reached in markets with only 6 to 14 people in environments where a parallel market structure requires 50 or more participants. Our results suggest that, to avoid illiquidity in thin (small-scale) financial markets, a portfolio trading mechanism should be used.

After exploring the ability of a combined-value market mechanism to overcome illiquidity, we turned to the implementation of such a mechanism in the Southern California RECLAIM market. Does the combined-value mechanism perform as well in the field as it did in the lab? In reality, the RECLAIM market has performed very well, alleviating many predicted problems of market thinness and risk aversion in long-term planning market. As time progresses and the number of permits available reduces below historical emissions, the prices for permits are stabilizing. The spot market prices is now almost exactly equal to the Best Available Control Technology prices. While we might expect these prices to exceed the BACT in the short-term (during installation and other adjustments), it seems clear that in the future the prices will stabilize at a level equal to the marginal cost of control of emissions, as desired. Of course, as the number of available RTCs reduces, this marginal control cost will increase as will the price of RTCs.

Moreover, the combined-value market mechanism has created the software foundation for creating systems of multiple-item barter. The items in this case are tradable pollution permits, but we need not limit ourselves to this application. Emerging market debt, cargo space and freight, virtually any marketplaces where the assets in question have synergistic values can benefit from a combined-value market mechanism. We now have the ability to use package bidding to create trades that, in the extreme, need not involve any money whatsoever.

Bibliography

[1] Arrow, K. (1953): "Le rôle des valuers boursières pour la repartition la meilleure des risques," translation published in *Review of Economic Studies* 31 (1964), 91-96.

[2] Banks, J., J. Ledyard, and D. Porter (1989), Allocating uncertain and unresponsive resources: an experimental approach, *RAND Journal of Economics*, Vol. 20, No. 1, Spring.

[3] Bohi, D., and D. Burtraw (1997): "SO_2 Allowance Trading: How do Expectations and Experience Measure Up," *The Electricity Journal* 10 (7).

[4] Bossaerts, P., D. Kleiman and C. Plott (1998): "Price Discovery In Financial Markets: The Case Of The CAPM," Caltech working paper.

[5] Bossaerts, P. and C. Plott (2000): "Basic Principles Of Asset Pricing Theory: Evidence From Large-Scale Experimental Financial Markets," Caltech working paper.

[6] Chander, P. and H. Tulkens (1997): "The Core of an Economy with Multilateral Environmental Externalities," *International Journal of Game Theory*, 3.

[7] Chwe, M. (1994): "Farsighted Coalitional Stability," *Journal of Economic Theory* 63, 299-325.

[8] Coase, R. (1960): "The Problem of Social Cost," *The Journal of Law and Economics*, 3.

[9] Conley, J. and M. Dix (1999): "Optimal and Equilibrium Membership in Clubs in the Presence of Spillovers," Work in Progress.

[10] Culyer, A.J. and Todd Sandler (1982): "Joint Products and Multijurisdictional Spillovers," *The Quarterly Journal of Economics* November, 707-16.

[11] Demange, Gabrielle (1994): "Intermediate Preferences and Stable Coalition Structures," *Journal of Mathematical Economics* 23, 45-58.

[12] Dwyer, John P. (August 1992): "California's Tradeable Emissions Policy and Greenhouse Gas Control," *Journal of Energy Engineering,* 118, 59-76.

[13] Foster, Vivian and Robert W. Hahn (April 1995): "Designing More Efficient Markets: Lessons from Los Angeles Smog Control," *Jounral of Law and Economics,* XXXVIII, 19-48.

[14] Gallance, R.A. and G. Tauchen (1996): "Which Moments to Match?" *Econometric Theory* 12, 657-681.

[15] Gouriéroux, A. Monfort and E. Renault (1993): "Indirect Inference," *Journal of Applied Econometrics* 8, S85-S118.

[16] Hahn, R.W. and R.G. Noll (1982): "Designing a Market for Tradable Emission Permits," *Reform of Environmental Regulation,* (W. A. Magat, ed.) Ballinger Press, Cambridge, Mass.

[17] Hall, Jane V. and Amy L. Walton (1996): "A Case Study in Pollution Markets: Dismal Science vs. Dismal Reality," *Contemporary Economics Policy,* XIV, 67-78.

[18] Hausker, Karl (1992): "The Politics and Economics of Auction Design in the Market for Sulfure Dioxide Pollution," *Journal of Public Analysis and Management,* 11, 553-72.

[19] Ishikida, T., J. Ledyard, M. Olsen and D. Porter (forthcoming): "Experimental Testbedding of a Pollution Trading System: Southern California's RECLAIM Emissions Market," to appear in *Research in Experimental Economics*, (R. Mark Isaac ed.) JAI Press, Greenwich CT.

[20] Johnson, Scott L. and David M. Pekelney (1996): "Economics Assessment of the Regional Clean Air Incentives Market: A New Emissions Trading Program for Los Angeles," *Land Economics,* 72 (3), 277-97.

[21] Klier, Thomas H., Richard H. Mattoon, and Michael A. Prager (1997): "A Mixed Bag: Assessment of Market Performance and Firm Trading Behavior in the NOx RECLAIM Programme," *Journal of Environmental Planning and Management* 40 (6), 741-74.

[22] Laffont, J.J. (1977): "Effects Externes et Theorie Economique," in *Monographies du Seminaire d'Econometrie Editions du Centre*, Elsevier Amsterdam, 231-52.

[23] Ledyard, John O. and Kristin Szakaly-Moore (1994): "Designing Organizations for Trading Pollution Rights," *Journal of Economic Behavior and Organization,* 25, 167-96.

[24] Maler, K.G. (1989): "The Acid Rain Game," *Valuation Methods and Policy Making in Environmental Economics,* (H. Folmer and E. v. Ierland, eds.) Amsterdam, Elsevier Science Publishers.

[25] Marron, Donald B., and Carlton W. Bartels, "The Use of Computer-Assisted Auctions for Allocating Tradable Pollution Permits," in *Market-Based Control: A Paradigm for Distributed Resource Allocation*, (S. Clearwater, ed.), World Scientific Pub. Co., New York. 1996.

[26] McCann, Richard J. (July 1996): "Environmental Commidities Markets: 'Messy' versus 'Ideal' Worlds," textitContemporary Economic Inquiry, XIV, 85-97.

[27] Montgomery, W. David (1972)): "Markets in Licenses and Efficient Pollution Control Regimes," *Journal of Economic Theory,* 5, 395-418.

[28] Neelin, Janet, Hugo Sonnenschein and Matthew Speigel (1988): "A Further Test of Noncooperative Bargaining Theory: Comment," *American Economic Review* 824-36.

[29] Samuelson, P. (1954): "The Theory of Public Expenditure," *Review of Economics and Statistics* 37, 387-9.

[30] Scarf, H.E. (1967): "The Core of an N-Person Game," *Econometrica* 35, 50-69.

[31] SCAQMD (1999): *Annual Report.*

[32] Sharpe, W. (1964): "Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risk," *Journal of Finance* 19, 425-442.

[33] Srinivasan, S. (1999): "Trading Portfolio Electronically: An Experimental Approach," OptiMark working paper.

[34] Srinivasan, S. (1999): "Portfolio Trading And Electronic Networks," OptiMark working paper.

[35] Stahl L. (1972): "Bargaining Theory," *Economic Research Institute.*

[36] Stavins, Robert N. (1995): "Transaction Costs and Tradeable Permits," *Journal of Environmental Economics and Management,* 29, 133-48.

[37] Tiebout, C. (1956): "A Pure Theory of Local Expenditures," *Journal of Political Economy* 64, 416-24.

[38] Tietenberg, Thomas (1995): "Tradeable Permits for Pollution Control when Emission Location Matters: What have we Learned?" *Environmental and Resource Economics,* 5, 95-113.

[39] Williams, A. (1966): "The Optimal Provision of Public Goods in a System of Local Governments," *Journal of Political Economy,* 74.

Appendix A Experiment Instructions

A.1 Overview

This session consists of simultaneous market trading periods. There will be three commodities traded. Two of these commodities are securities. The third commodity is NOTES. Notes allow you to borrow and loan money. The securities, SECURITY A and SECURITY B, and the NOTES have a one period life. That is, they pay a single dividend and are removed from the system at the end of the period. The currency used in all markets is called francs. Each franc is worth AMOUNT HERE.

A.2 Timing of Events

At opening of each period you will be given, as "working capital," a portfolio of units of SECURITY A, units of SECURITY B and some francs (cash). You are free to place offers to purchase and sell as many securities as you want within the limits imposed by your working capital. You are not allowed to purchase securities if you do not have the francs to pay for them. Furthermore, you are not allowed to sell securities unless you actually have them in inventory. That is, you cannot go short in securities and you cannot purchase securities beyond the limits of your francs.

During the first round of a period, you will have 3 minutes to submit bids. Then, the bidding will be stopped and all orders will be called. That is, the offers to buy and sell will be collected. Orders that can be matched will be completed and payments will be determined. After these orders are completed, another call round begins. All remaining rounds in the period will be 90-seconds long. To see which of your orders were completed in the last call round, press either the "Results" button, or the "Previous Round Results" button. There will be 10 rounds in each period.

At the end of 10 rounds (15 minutes), the dividend payments will be determined and distributed according to the portfolio you hold at the close of the period. Your income for the period consists of the dividends on the securities you hold at the close, as well as the cash you hold at that point, minus a predetermined payment (a "loan" repayment) for the working capital you were given at the beginning of the period. This income will be recorded as your earnings for the period. Your working capital will be automatically refreshed and a new period will open. There is no carry-over of Securities or francs from period to period.

QUESTIONS ON SECTIONS 1 AND 2:

1. How long is the first round in each period?
   - (a) 90 seconds
   - (b) 3 minutes
   - (c) 10 minutes
2. How long are subsequent rounds?
   - (a) 60 seconds
   - (b) 90 seconds
   - (c) 15 minutes
3. How many rounds are there in a period?
   - (a) 10 rounds
   - (b) 15 rounds
   - (c) determined randomly

A.3 Dividend Determination

The actual dividends paid by Security A will depend on a randomly drawn STATE. There are *three* possible states, referred to as X, Y, and Z. States are equally likely, and the state drawn in a given round has no effect on its likelihood of it being drawn in another. The dividend that a security pays in each state is uniquely determined from the table that follows.

	State		
Security	X	Y	Z
A	170	370	150
B	160	190	250
C	100	100	100

For example, if the state is X for some period, then the holder of Security A will receive 170 francs for each unit held at the end of that period and the holder of Security B will receive 160 francs for each unit held at the end of that period. If the state drawn for the period is Y then the numbers are 370 and 190 respectively, and if the state is Z they are 150 and 250.

Notes are like bonds or IOUs that allow you to borrow and loan francs. Selling a Note is like borrowing the amount of the sale price. Buying a Note is like loaning the amount of the sale price. For each Note you sell you must pay the holder 100 francs at the end of the period. This payment will automatically be deducted from your francs and dividend payments at the end of the period. For example, if you sell a Note for 75 francs you have borrowed 75 francs and for this loan you repay a total of 100 at the end of the period. Effectively "you repay the loan" of 75 francs plus a 25-franc "interest" payment. If you buy a note for 80 francs then you have loaned the seller 80 francs until the end of the period and for this loan you will be repaid a total of 100 francs; in essence the 80 francs plus a 20 franc "interest" payment. Clearly no one should buy a Note for more than 100 francs because it means that the buyer loaned more than would be repaid. At the beginning of the period you will be given no Notes. If you sell Notes your inventory will be listed as negative, indicating that you must pay 100 francs on each of these, and if you buy Notes your inventory will become positive indicating that you will be paid 100 francs for each Note you hold. You are not allowed to sell over eight (8) Notes (net) in any given period.

QUESTIONS ON SECTION 3:

1. In state X, which commodity is most valuable?
   - (a) Security A
   - (b) Security B
   - (c) Notes
2. If you purchase a Note for 60 francs, how much is the effective interest payment?
   - (a) 200 francs
   - (b) 40 francs
   - (c) 60 francs
3. What is the maximal number of Notes you can sell (net) in any given period?
   - (a) 8 Notes
   - (b) 0 Notes
   - (c) unlimited

A.4 Initial Portfolio and Working Capital

At the beginning of each period you will be given working capital, which consists of shares of Security A, shares of Security B, and francs. You will not be given any Notes. The working capital you receive can be different from the working capital received by others. For this working capital you must repay a predetermined number of francs (a "working capital loan" repayment) at the end of the period.

A quick calculation will allow you to calculate the expected payoff of the working capital. On average each Security A will pay 230 francs per period [(1/3)(170+370+150)] reflecting the fact that each of the dividends will be paid 1/3 of the time. A similar calculation for Security B shows that on average each will pay 200 francs per period. The expected end-of-period payoff of your working capital is thus 230 times the number of units of Security A plus 200 times the number of units of Security B plus the number of francs. Subtract from it the required "working capital loan" repayment and you get your expected income from the working capital. For instance, if you are given 4 Securities A, 5 Securities B and 400 francs cash, and are required to repay 1000 francs for it at the end, your expected income would be 1320 francs (=4*230+5*200+400-1000). Note that losses are permitted. However, if you suffer a loss for two consecutive periods, you are in default and can no longer participate.

QUESTIONS ON SECTION 4:

1. Which has the highest expected value?
   (a) they're the same
   (b) Security B
   (c) Notes
   (d) Security A
2. If you have 3 securities A, 2 securities B, 500 francs cash, and a 1000 franc loan repayment, your expected income is:
   (a) 590 francs

(b) 1500 francs

(c) 1000 francs

	State		
Security	X	Y	Z
A	170	370	150
B	160	190	250
C	100	100	100

A.5 How the Market Trading System Works

The market is divided into 90-second rounds (with the exception of round 1, which is 3 minutes) during which you can submit orders to the market. An order is a listing of the units of Security A, Security B and notes, along with a franc amount. A buy order implies that you want to buy the specified number of units at the specified price or less, and a sell order implies that you want to sell the specified number of units at the specified price or more. From the main window, you can choose the BIDBOOK button to view your bidbook. An example of a bidbook window is below. This window shows your budget, your current assets, the value of your current portfolio by state (net of your working capital loan repayment), and the status of any created bids. In order to create a new bid, press the CREATE button.

When you press the CREATE button, the create bid window appears. There is a column for each commodity under the heading ASSETS in the left half of the window. Enter the number of units that you would like to buy or sell in the box labeled "Request." Note that orders for fractions are allowed, and orders will not be rounded. Below the unit request, indicate whether this is a buy or sell order. Enter the price per unit at which you are willing to buy or sell each commodity in the box labeled "Price/unit." DO NOT enter a price under the box labeled "Portfolio Price," as the portfolio price will be automatically generated from the individual buy and sell prices. Finally, enter a scale in the box labeled "Scale" (explained below).

Below is an example of an order to sell 10 units of A and buy 2.5 units of B and 4 units of Notes with a scale of 0.

There are several features of the above that are important:

1. The order is for a package of items. The scale of 1.0 (explained below) means that it is an all or nothing order, in that the entire amount (10A, 2.5B, and 4 NOTES) requested must be filled if the order is to be a valid transaction.

2. If the order is a buy order, Portfolio Price = (number of units) * (price per unit). If the order is a sell order, Portfolio Price = -(number of units) * (price per unit). The Portfolio Price of a buy order will always be positive, and the Portfolio Price for a sell

order will always be negative.

3. Regardless of whether you enter a price under "Portfolio Price," the computer will use the "Price/unit" price to calculate the portfolio price. If you enter only a portfolio price, then the computer will take the Price/unit, and therefore the portfolio price as \$0.

4. A buy offer is an upper bound on what you are willing to pay. A sell offer is a lower bound on the amount of francs you are willing to receive for the order.

FOR A DETAILED EXAMPLE EXPLAINING THESE PROPERTIES, CLICK HERE. (This information is available in Section A.6.)

A.5.1 Tailoring your Order

Your order can be customized in two ways that may assist you in obtaining a desired transaction.

1. Scalable Orders: Recall that the order placed in the above example was all or nothing. You can request your order to be less restrictive. You do this by submitting a scale number between 0 and 1 by which you would be willing to scale your order to have it accepted. The best way to explain this feature is with an example. Suppose a scale of 0.5 is submitted with the order. This means that you are willing to sell 5A, buy 1.25B and 2 Notes for a package price of no more than 3 francs. That is, one-half scales all units. The choice of a 0.5 scale also means that you are willing to scale your order to any factor between 0.5 and 1.0. A choice of a 1 scale means that you are willing to trade only the entire amount of your order. A choice of a 0 scale means that you are willing to trade any fraction of your order.

2. Open Orders: To submit your bid, click on the "Open bid," "Partially open bid," or "Closed bid" button. If you submit a closed order, no other participant can view what you send to the market. If you would like others to see what you sent to the market, you can send the order to the open book. You can send your order as either partially

or fully open. A partially open order sends the quantities requested only. A fully open order sends both the quantities and the prices requested. All participants can view orders in the open book and will be part of the orders submitted to the market. If you would like to view the open or partially open bids that others have submitted, click on the "Public" button in your main window.

A.5.2 Order Restrictions

Your order can be submitted to the market if it does not violate your credit line and you do not try to sell more units than you have in your account. The constraint on your orders is cumulative. That is, each time an order is submitted your account is reduced by the amount in that order. Sell orders count against the item account and buy offers are subtracted from your franc account. However, sell orders do not increase your franc account and buy orders do not increase your unit accounts.

Below is an example of a player's bidbook after two bids have been submitted. The first order is a buy order for 5 units of A at 50 francs/unit, with a portfolio price of 250 francs. The order is all or nothing (scale=1), and the order is a partially open order, as denoted by the half sun. The second order is a sell order for 2 units of A at 75 francs/unit, with a portfolio price of -168.75 francs. This player is willing to sell any fraction of the order (scale=0), and the order is a closed order (no sun). Note that positive units imply a buy order and negative units imply a sell order. To change an order, click the EDIT button next to that order. To delete or duplicate an order, check the box labeled "Select," and then click either the DELETE or DUPLICATE button, depending on your choice of action.

Call Results:

The amount you will pay (receive if the market price is negative) for an accepted order is:

Your payment = (market buy prices * amounts you buy) - (market sell prices * amounts you sell) + a possible inflexibility penalty

Your payment will have the following simple properties:

1. Your payment will always be less than or equal to your offer amount.

2. In general, you are charged a penalty only because your unwillingness to scale to zero means we have to charge others less than the market price in order to accept your offer.

3. If your buy (sell) order is accepted and your payment is equal to your offer, but less (more) than it would be at market, you have been included because you helped complete an inflexible trade.

4. If your order is accepted and you pay (receive) less (more) than your offer and you pay no penalty, then raising your offer price would not have changed whether the order was accepted or your payment for the order.

5. In general, your order will stand a better chance of being accepted the higher your offer is and the more flexible it is. The most flexible order is a simple order in only one asset and a scale of 0.0.

Even if no trade takes place in a round, a market price may be calculated. This price reflects a bid price that may have resulted in trade. Note, however, that even if you submit a buy (sell) order with a price higher (lower) than the market price listed, there are many scenarios in which trade still may not occur. For example, you may have submitted a fully flexible buy order for one unit of security A at a price of 175. The market price was 170, but your trade did not execute. This may be because the only sell order was an inflexible bid for 2 units of asset A. Or, the sell order may have been a package bid for units of A and B, and there was no order with which the B side could be matched. So, the market price should be used only as an indicator of the market, not a market clearing price.

After the round is completed, your Bidbook will update automatically, indicating which of your bids were accepted and which are "out." The solver will not consider any bids labeled "out" in the following round. However, if you would like to resubmit all of your bids, you can do so by clicking the "Renew Bids" button. In order to edit a bid, click edit button for that bid. If you would like to delete or duplicate a bid, click the select button for that bid and then select delete or duplicate.

QUESTIONS ON SECTION 5:

1. A package bid to buy 2 units of Security A at 100 francs and sell 3 Notes at 150 francs will list a package price of:
   (a) 250 francs
   (b) -250 francs
   (c) 650 francs

2. If you submit an order to buy 4 Securities A and 3 Securities B with a scale of .25, which of these purchases are you willing to accept?
   (a) 1A, .75B
   (b) 3A, 2 B
   (c) 4A, 1 B

3. If you choose "Partially Open Bid" when creating a bid, what information is sent to the open book?
   (a) quantities
   (b) quantities and prices
   (c) quantities, prices and bidder ID

	State		
Security	X	Y	Z
A	170	370	150
B	160	190	250
C	100	100	100

A.6 Illustration of the Five Pricing Properties

Suppose the following orders have been submitted:

Bid Number	Quantity	Unit Offer
1	2	5
2	3	4
3	2	2
4	2	1
5	-1	.5
6	-2	1.5
7	-3	3
8	-3	3.5

If all were scaled at 0, they would yield the following demand-supply picture of the book:

The first step of the market trading system is to accept the collection of orders that maximize the gains from trade. Here, orders 1,2,5, and 6 are fully accepted and 2 units of order 7 are accepted. The gains from trade are 12.50 francs.

The next step of the market trading system is to set prices. Here, the price will be 3.50. This satisfies **PROPERTY 1**: your payment is less than or equal to your accepted offer. Here any price between 3 and 4 would work, so 3.50 is chosen in the middle.

Accepted Orders	Quantity Accepted	Pay (receive)	Revealed Gain
1	2	2(3.5)=7	3
2	3	3(3.5)=10.5	1.5
5	-1	-1(3.5)=-3.5	3
6	-2	-2(3.5)=-7	4
7	-2	-2(3.5)=-7	1
TOTAL	0	0	12.5

PROPERTY 4 is illustrated by bid 1 (as well as 5 and 6). If 1 had bid for 2 units at 10f/unit, the order would still have been accepted and would still have paid 7f. If order two

had set a buy price of 4.5, it would still have been accepted but would have paid 2(3.75)=7.5, because it would have affected the price.

Now let us consider the same example, but assume that order 5 is all-or-none. That is, its scale is set at 1. Then the story is the same as before. The accepted orders and price are identical.

Now, suppose order 7 is scaled at 1. If it is not accepted, then 2 units of order 2 will not trade and gains from trade would be 10.5, but if order 3 is used to help complete order 7, gains would be 11.5. So, maximizing the gains from trade means accepting orders 1,2,5,6,7 and 1 unit of 3.

So far so good. But now there is no single price satisfying property 1. Such a price would have to be greater than 3 (for order 7) and less than 2 (for order 3). The market trading system deals with this in 3 steps. In step 1, it acts as if all orders are scaled at 0 and then computes trades and prices as previously described. In this case, as before, 1,2,5,6 and part of 7 trade at a price of 3.5.

In step 2, all orders that would have faced a loss at this price are identified. Here only order 3 has this problem, since (2)-(3.5).

If we were to stop now the market would have to pay out more than it receives. In step 3, this shortfall is collected from the inflexible orders that cause it. In this case, order 7 must pay an inflexibility penalty of 1.5. This is the amount order 3 is short. This illustrates **PROPERTY 2**. The final result of the market call is:

Accepted Orders	Quantity Accepted	Pay (receive)	Revealed Gain
1	2	2(3.5)=7	3
2	3	3(3.5)=10.5	1.5
5	-1	-1(3.5)=-3.5	3
6	-2	-2(3.5)=-7	4
7	-3	-3(3.5)-1.5=-9	0
3	1	1(2)=2	0
TOTAL	0	0	11.5

Notice that orders 1,2,5 and 6 are unaffected by the inflexibility. Order 7 compensates order 3 the 1.5 necessary to complete a trade. Order 7 received 1 in revealed surplus when scaled at 0 and receives 0 when scaled at 1. But 7 may really be better of if, for example, he gets a true benefit of 12 with 3 units and 6 with 2 units. Then under a 0 scale 7 gets 6-7=-1 and under a 1 scale 7 gets 12-9=3.

A.7 The Mechanics of the Interface

The interface has many powerful properties for you to use in making decisions and placing your bids:

1. Once you have entered your bids, you may edit them in the bidbook window. Select a bid (or many) by checking the "Select" button on the end of that bid's row. You may then edit, delete, undelete, or duplicate the bid.

2. After the round ends and bids are submitted, your bidbook will tell you which bids were accepted (labeled "In") and which were not (labeled "Out"). You can re-submit all of your "out" bids by pressing the "Renew Bids" button in the bidbook.

3. The time remaining in the current round and the time remaining until the next round begins are displayed on your main window, your bidbook, and you create bid page. As discussed earlier, all rounds will be 90 seconds long, with the exception of the first round of each period. There will be a pause (30-50 seconds) between rounds for you to review the results of the previous round. If the timer ceases to function (refreshing every few seconds), PLEASE ALERT THE MONITOR.

4. You can check how much you have earned in the experiment thus far in the experiment by selecting the "Earnings" button on the main window.

5. You can check the results from previous rounds in the current period by selecting the "Previous Rounds" button on the main window.

6. If you would like to view the open or partially open bids that others have submitted, click on the "Public" button in your main window.

7. Along with the Internet interface, you will also have a portfolio calculator available to you, in an Excel spreadsheet. Using this spreadsheet, you can enter prices and a potential portfolio. The sheet will tell you the value of the portfolio as well as the return on your investment.

A.8 How You Make Money

The following examples will help you understand the various possibilities.

Example 4 Do nothing. △

Suppose you start and end with 4 units of Security A plus 400 francs. Your expected dividends from the securities are 920 francs (230 each times 4 securities) from the A Securities. Adding the 400 francs yields a total of 1320 francs you expect to hold at the end of the period. From this subtract the required working capital loan repayment for the capital advance, say, 1100. Your earnings would be 220 francs (on average).

Example 5 Play it safe. △

This means that you sell all of your securities and loan out the money (buy Notes) you receive in payment. Of course, in this case your earnings depend upon the prices at which you transact. Assume that you are given a working capital of 4 Securities A, 5 Securities B and 400 francs cash. Assume your working capital loan is 2100 francs. Suppose you are able to sell all securities at their expected value, that is, at 230 for Securities A, and at 200 for Securities B. This would produce 1920 francs revenue. Adding the 400 francs cash with which you started gives you a total of 2320 francs. This guarantees you an income of 220 francs (2320 cash, minus the 2100 francs repayment for working capital). Now, these 2320 francs can be invested in Notes. You could buy 23 Notes if you paid the full repayment value of 100 each (of course you would make no profit in this case), but if you purchase at less than 100 you make the difference on each Note you buy. Suppose you could buy 25 Notes at 90 each (a total investment of 2250 francs of your 2320). Each repays 100 so you make 10 on each for a profit of 250. Your period earnings would be 220 + 250 = 470 francs for the period. Of course this calculation makes important assumptions about the prices you received for the securities and the prices that you paid for the Notes.

Example 6 Speculate on price changes. △

Anything you buy can be resold. If you sell at a price higher than you paid, you make the difference. This difference is in the form of increased francs which you can then invest in

Securities and receive the returns.

Example 7 Leveraging. △

Sell all of your Security B and use the francs to purchase Security A. In addition, use your initial francs to purchase Security A. Finally, sell all the Notes you are able and use the francs to purchase Security A. Suppose that prices are at the expected value of 230 for A and 200 for B, and that the price of Notes is 90. Suppose you have 2 Securities A and 7 Securities B initially, and that you were given 400 francs cash. Assume your working capital loan is 1700 francs. Selling all of B generates 1400 francs; selling 3 Notes generates 270 francs. Together with the initial cash allocation, you now hold 2370 francs in cash. With it, you can purchase 10 Securities A, at a cost of 2300 francs, leaving you with 70 francs cash. At the end of the period, you will be holding 13 Securities A and 70 francs in cash. You will be required to pay 300 francs for the (short) sale of the 3 Notes, in addition to the 1700 francs that you owe for the working capital. If the state is X, Y or Z, your dividends would be 2210, 4810, 1950 respectively; the corresponding period earnings would be 280, 2880, 20 respectively.

Example 8 Fully invested hedge. △

You might use your francs cash to purchase Securities B and also sell Notes and use the proceeds to purchase even more Securities B. Since Security B dividends are low in those states in which Security A dividends are high, your overall variability of returns is reduced. Security A pays a large dividend of 370 if state Y occurs, but it only pays 150 if state Z occurs. By purchasing some of the Security B you will be able to "insure" your holdings of Security A somewhat because Security B pays 250 francs in state Z.

A.9 Why Submit a Package Bid

The following example will illustrate the benefit of submitting a bid for a bundle of goods rather than bids for single assets. Suppose you are attempting to execute a fully invested hedge, as explained above. Suppose further that you wish to buy 2 units of A for 200 francs only if you can hedge against the risk by purchasing 3 units of B at 200 francs. If you submit two separate orders (one to purchase A and one to purchase B), you are not guaranteed that both will bids will be accepted. For example, if the bid for A transact but the bid for B does not and state Z occurs, then you will suffer a loss of 100 francs (2*200 - 2*150 = -100). However, if this bid had been submitted as a package and transacted, the state Z value would be a profit of 50 francs (2*150 + 3*250 - 2*200 - 3*200 = 50).

QUESTIONS ON SECTION 7:

1. What method of making money is the riskiest, but yields the highest reward?
   (a) fully invested hedge
   (b) leveraging
   (c) speculate on price changes
2. Which commodity is the riskiest?
   (a) Security A
   (b) Security B
   (c) Notes

	State		
Security	X	Y	Z
A	170	370	150
B	160	190	250
C	100	100	100