

Auction Design for Selling CO₂ Emission Allowances Under the Regional Greenhouse Gas Initiative

Final Report · October 2007

Charles Holt
William Shobe
University of Virginia

Dallas Burtraw
Karen Palmer
Resources for the Future

Jacob Goeree
California Institute of Technology

Auction Design for Selling CO₂ Emission Allowances Under the Regional Greenhouse Gas Initiative

Final Report

October 26, 2007

Investigators:

Charles Holt, William Shobe

University of Virginia

Dallas Burtraw, Karen Palmer

Resources for the Future

Jacob Goeree

California Institute of Technology

Acknowledgements

This report was funded by the New York State Energy Research Development Authority (NYSERDA). The research benefited from outstanding assistance from Erica Myers, Danny Kahn, Anthony Paul and Susie Chung at Resources for the Future and Lindsay Osco, Susan Ivey, Courtney Mallow, A.J. Bostian and Angela Smith at the University of Virginia.

The authors want to express their appreciation to the many persons who provided comments and advice over the course of this investigation, especially to David Coup for project management along with many helpful insights, to NYSERDA's Technical Advisory Group and to the RGGI Staff Working Group.

Disclaimer

The statements and recommendations in this report are solely the responsibility of the authors and do not necessarily represent the views of NYSERDA or the RGGI Staff Working Group or others associated with the RGGI.

Obtaining Copies of This Report

Copies of this report can be obtained from:

www.rff.org or www.coopercenter.org/econ/rggi_final_report.pdf.

For more information about RGGI, see www.rggi.org.

Auction Design for Selling CO₂ Emission Allowances Under the Regional Greenhouse Gas Initiative

Executive Summary	5
-------------------------	---

Part 1: Motivation and Organization of the Project

1 Introduction	10
1.1 Methods of Analysis	12
1.2 Two Phases of Research.....	13
1.3 Organization of this Report.....	14
2 Background on Auction Types.....	15
2.1 Introduction	15
2.2 Auction Formats.....	16
3 Criteria for Selection of an Auction Type.....	21
4 Basic Experimental Approach for This Investigation	24
4.1 Production Technology and Market Structure	25
4.2 Measures of Performance.....	27
4.3 Important Aspects of Phase 1 Experiments	30
4.4 Important Aspects of Phase 2 Experiments	33

Part 2: Research Questions

5 Key Auction Design Features.....	35
5.1 Defining Allowance Vintages	35
5.2 Frequency, Timing, and Size of Auctions.....	35
5.3 Financial Assurance Mechanisms	41
5.4 Market Monitoring.....	42
5.4.1 Guidance on Monitoring in RGGI	44

6 Auction Performance in Collusion-Enhancing Environments and How to Limit Collusion	45
6.1 Effects of a Loose Cap in Uniform-Price and Discriminatory-Price Auctions.....	46
6.1.1 Motivation	46
6.1.2 Procedures	48
6.1.3 Aggregate Results	48
6.2 Collusive Environment with Standard Caps	50
6.2.1 Motivation	50
6.2.2 Procedures	50
6.2.3 Aggregate Results	51
6.3 Discussion and Extensions	53
7 Reserve Prices and Price Volatility	54
7.1 Price Volatility	54
7.2 Reserve Prices	55
8 Price Discovery	59
8.1 Price Discovery: Unanticipated Shift in the Demand for Permits	60
8.2 Relationship Between Auctions and Secondary (Spot) Markets	63
8.2.1 Experience in Existing Allowance Markets	64
8.2.2 Evidence From Experiments	66
9 Hoarding of Allowances.....	67
9.1 The Speculation and Market Manipulation Motives.....	71
9.2 The Market Disruption Motive	72
9.3 The Competitive Advantage Motive.....	72
9.4 The External Compliance Motive	73
9.5 Possible Approaches to Address Hoarding.....	74
10 Combining Auctions with Free Allocations	75

Part 3: Recommendations

11 Recommendations for Auction Design.....	77
Recommendation 1: Uniform-Price Auction	77
Recommendation 2: Single-Round, Sealed-Bid Format.....	77

Recommendation 3: Separate Auctions by Vintage.....	77
Recommendation 4: Quarterly Auctions.....	78
Recommendation 5: Auction Future Allowances in Advance.....	78
Recommendation 6: Reserve Price	78
Recommendation 7: Unsold Allowances.....	79
Recommendation 8: Lot Size.....	79
Recommendation 9: Open Auctions to All Qualified Bidders	79
Recommendation 10: Bids are Binding Contracts.....	79
Recommendation 11: Joint and Uniform Auction	80
Recommendation 12: Market Monitoring.....	80
Recommendation 13: Disclosure of Beneficial Ownership.....	80
Recommendation 14: Auction Information Disclosure	81
Recommendation 15: Statement of Intent.....	81
Recommendation 16: Ongoing Evaluation.....	81
12 References.....	82
13 Appendix A: Statement of Work Map.....	86
14 Appendix B: Annotated Bibliography.....	87

Auction Design for Selling CO₂ Emission Allowances Under the Regional Greenhouse Gas Initiative

Executive Summary

In 2009, the 10 northeastern states that comprise the Regional Greenhouse Gas Initiative (RGGI) will launch the first cap-and-trade program for greenhouse gas emissions within the United States. This program, which covers CO₂ emissions from electricity generators within the region, is the result of a multi-year cooperative effort among states from Maryland to Maine. CO₂ emissions will be capped at levels comparable to emissions levels at the beginning of this decade and then ramped down to 10% below initial cap levels by 2019. RGGI member states have developed an architecture that can serve as a model for a national program to limit greenhouse gas emissions.

The RGGI proposal represents a substantial break with the past. Rather than give the allowances away for free, as has been done in earlier cap-and-trade programs, the RGGI states agreed to allocate at least 25% of the emission allowances created by a cap-and-trade program to benefit consumers and to support strategic energy investments. An auction of allowances is the most straightforward way to implement this policy. More recently, several RGGI states have decided to auction 100% of their annual CO₂ allowance budgets. As the first greenhouse gas cap-and-trade program to start with a substantial auction of allowances, this major regional initiative will have a global impact.

A carefully designed allowance auction can help maximize the benefits of the RGGI program and can serve as a model for other states and, ultimately, for a federal program to control greenhouse gas emissions. The investigators on this project worked to develop an auction design that meets several key criteria:

- Low administrative costs, low transaction costs for bidders;
- Perceived as fair, transparent, and understandable to participants and the public;
- Economically efficient—that is, getting allowances to those who value them the most;
- Avoiding collusive behavior by bidders and providing good signals about market prices;
- Helping to minimize price volatility;

- Raising reasonable revenues from the sale of a valuable public asset; and
- Compatible with existing electricity and energy markets.

This research uses experimental economic methods, insights gleaned from the economics literature, and results from past experience with various types of auctions, including prior allowance auctions, to develop recommendations on the most appropriate design for auctions of RGGI CO₂ allowances.

The research was conducted in two phases. Phase 1 included the initial preparation of an annotated bibliography and a round of experiments to investigate the performance of a number of auction types considered to be likely candidates for use in a CO₂ allowance auction. The main auction types considered were the sealed-bid, increasing-price sequential (English clock), and decreasing-price sequential (Dutch) auction forms. We also examined whether sealed-bid auctions should use the pay-as-bid (discriminatory) or uniform-price rules. Our experimental findings in Phase 1 regarding the performance of these auction types did not reveal a clear winner; all of the formats performed well in these initial experiments.

In Phase 2, we continued experiments aimed at the basic question of identifying the auction type that performs best along an expanded set of performance measures and in a richer institutional setting. The auction formats were compared with respect to price discovery, that is, ensuring that the price of allowances at auction reflects their market value, and in limiting collusive behavior. We also examined the effect of reserve prices and allowance banking and did more analysis of how the auction combines with secondary (or spot) markets. We looked at the effects of allowing participation in the auction by brokers or other traders not needing allowances for compliance and of combining auctions with “grandfathering” of some allowances for free to generators. In addition, we performed some experiments to look at so-called “hoarding behavior” and the effects of different mechanisms that have been proposed to limit hoarding. In particular, our experiments examined whether holding auctions with participation limited to generators can reduce the effect of hoarding behavior.

Several recommendations on auction design follow from the findings of this study and they are summarized below.

1. The RGGI auction should use a uniform-price auction format, the clearing price for the auction being the value of the highest rejected bid. The uniform-price auction format has much to recommend it, including simplicity, relative transparency, and the observed tendency for bidders to

ensure purchases of needed allowances by bidding closer to use values. This auction design performed very well in our price discovery experiments. The uniform-price auction also is familiar to the electricity sector, as it is the auction form used in most Independent System Operator (ISO) electricity auctions.

2. The RGGI auction should use a single-round, sealed-bid format. The literature suggests that multiple-round auctions can be more conducive to collusion, as they provide participants with opportunities for signaling and detecting when someone has reneged on a collusive agreement. In our preliminary recommendations, we had recommended that the first auction for each vintage be a clock auction (with a final, sealed-bid stage), but further examination suggests that clock auctions perform no better in terms of price discovery than single-round auctions. Ties in the auction should be resolved by a random process to help guard against collusive bidding.
3. Separate auctions should be held for different vintages. Since the yearly vintages within a compliance period are not identical assets, due to different first years of allowable use, they should be sold separately.
4. The auctions should be held quarterly. This schedule of auctions provides the benefits of periodic price discovery and enhanced liquidity without interfering with the performance of a secondary market.
5. Future allowances should be made available four years in advance of their vintage. Auctioning future vintages in advance should assist generators in their planning for future investments.
6. A reserve price should be used at each auction. In general, the reserve price should be announced publicly, with the possible exception of the first auction. How the reserve price is set in the auction interacts with other aspects of the program design. No bids for allowances should be accepted if the bid price falls below the reserve price.
7. Two options have been identified for what to do with unsold allowances because of insufficient demand or because the reserve price is triggered. One option is that unsold allowances should be rolled into a contingency reserve account. The allowances in the contingency account will not be released for sale until some RGGI auction closes above a specified value, such as the first offset trigger price. Once this condition is met, the

contingency reserve allowances would be available for auction on the next quarterly auction date. The contingency reserve account would help to minimize large fluctuations in allowance prices. Another option is that some or all of the unsold allowances could be rolled into the next auction.

8. Lot size at auction should be a minimum of 1,000 allowances. This will reduce administrative costs and bidding costs without placing significant burdens on bidders. The lot size should not be so large that it limits participation in the auction.
9. Auctions should be open to anyone willing and able to meet financial pre-qualification, but no single entity should be able to purchase (or take a beneficial interest in) more than 33% of the allowances for sale in any auction. Open auctions will enhance competition and limit opportunities for collusion. Limiting the share of allowances that a single entity can purchase in an auction raises the cost of using the auction to corner the market without placing too stringent a restriction on what generators can purchase.
10. Accepted bids should be treated as binding contracts, and bidders must provide strong financial assurance to cover the value of any bids. No bids above financial assurance levels should be allowed for any bidder.
11. There should be a joint and uniform auction for allowances of a given vintage sold from all RGGI states. Allowances should be completely identical, notwithstanding the state of origin. All contract and enforcement terms should be identical for all allowances, notwithstanding the state of origin. This structure helps keep transaction costs low and prevents some extraneous influences on prices.
12. RGGI market monitoring efforts should take advantage of existing monitoring activities by federal and state agencies and other interested parties. RGGI should coordinate with the Federal Energy Regulatory Commission, the U.S. Environmental Protection Agency (EPA), the Independent System Operators and the Commodity Futures Trading Commission (CFTC) in designing criteria for detecting market manipulation and for sharing of information regarding the performance of the allowance market and the detection of attempts to manipulate prices.
13. RGGI should require that the authorized account representatives be obliged to disclose the “beneficial ownership” of any allowance holdings. That is,

every participant would have to disclose the party sponsoring or benefiting from the agent's activities in the allowance market if it was other than themselves or their immediate employer. Currently this is not required in the EPA's Allowance Tracking System. This information is proprietary and should be kept confidential.

14. Information from the auction that should be publicly disclosed includes the auction clearing price, the identities of winning bidders and the quantity of allowances obtained by each winning bidder. The actual value bid by each auction participant should not be disclosed. Information about losing bidders should not be disclosed.
15. RGGI should articulate the auction goals in a "Statement of Intent" and ask all participants in the auction to acknowledge that statement and agree not to undermine these goals. The goals that might be articulated range from overall environmental integrity to specific behavior in the allowance market.
16. RGGI should evaluate the performance of the auction on an ongoing basis as part of their administrative oversight of the program.

Part 1: Motivation and Organization of the Project

1 Introduction

In 2009, the 10 northeastern states that comprise the Regional Greenhouse Gas Initiative (RGGI) will launch the first cap-and-trade program for greenhouse gas emissions within the United States. This innovative program, which covers carbon dioxide (CO₂) emissions from electricity generators within the region, is the result of a multi-year cooperative effort among states from Maryland to Maine. CO₂ emissions will be capped at levels comparable to emissions levels at the beginning of this decade and then ramped down to 10% below initial cap levels by 2019. Participants in the RGGI planning process have developed architecture for a successful cap-and-trade program that can serve as a model for a national program. A feature of this architecture, found in the original RGGI memorandum of understanding (MOU), specified that all states should allocate at least 25% of the emission allowances created by a cap-and-trade program to consumer benefit and strategic energy initiatives. An auction of allowances is the most likely way to implement this policy. Since signing the RGGI MOU, several RGGI states have made the decision to auction 100% of their annual CO₂ allowance budgets. By being the first greenhouse gas cap-and-trade program to start with a substantial auction of emission allowances, this program provides an opportunity to demonstrate how allowance auctioning can help the program to succeed.

The purpose of this research project is to assist with the design of auctions for the initial sale of CO₂ allowances in RGGI. The RGGI Model Rule specifies that each state must allocate at least 25% of its budgeted allowances to a consumer benefit or strategic energy purpose account. These “consumer benefit” allowances are to be sold or otherwise distributed to promote energy efficiency, to directly mitigate electricity ratepayer impacts, or to promote lower-carbon-emitting energy technologies. Some RGGI states have stated that they intend to auction 100% of their budgeted allowances. In July 2006, the authors of this report participated in workshop convened on behalf of stakeholders and state officials in RGGI to provide technical assistance on how to design an auction (Burtraw and Palmer 2006). That workshop included reports on past experiences with auctions for emission allowances and other commodities. The analysis provided in this study will illuminate how different auction design specifications relate to particular goals and criteria that RGGI has for the allowance auction. The report will conclude with recommendations on an auction

design for meeting RGGI goals and criteria. The analysis is applicable to a large, region-wide auction that involves the participation of all RGGI states. Guidance that is provided applies to the general case of an auction of emission allowances at any scale, recognizing that the size and other attributes of the market may affect the conclusions that are reached in the recommendations.

Considerable experience in the sale of assets by governments has led to the conclusion that careful attention to auction design can be critical to an auction's success in achieving the goals specified for the auction (e.g., McAfee 1996 and Whitford 2007). Sales have included a diverse array of rights, such as timber harvests, resource extraction, electromagnetic spectrum, securities, and, as in this case, emission allowances. Usually the goal specified in the sale is to maximize some combination of efficiency and revenues. Roughly speaking, efficiency is measured by whether the auction results in the rights being owned by those who value them the most. It is sometimes the case that choosing a particular element of auction design may require a tradeoff between revenues, efficiency, and other desirable attributes.

Factors particularly important to the success of an auction also include the auction's competitiveness, the ability of the auction to elicit bids that reflect actual valuations by bidders, and restricting bidder opportunities for acting strategically in a way that defeats the efficiency or revenue-raising function of the auction. In addition, there may be other, secondary characteristics that are of importance to policymakers or that may have an effect on the market into which the goods are sold, if not the efficiency of the auction itself. These factors may include price volatility, effects on related markets, transparency, low administrative and transaction costs, and perceived political risk in auction outcomes.

Auctions for RGGI CO₂ allowances will be taking place with the expectation that there will be an active secondary market for these assets and that a significant amount of allowances will be traded outside of the auctions themselves. Concerns about the role of an auction in CO₂ allowance trading can be roughly divided into three categories: those concerns that arise due to the existence of trading itself, concerns that arise due to the auction institution chosen, and some concerns that involve the interaction of the auction with an existing market. In this report, our primary focus will be on the latter two concerns. However, in the course of our research we also have looked in some detail at a few issues that arise not from the auction but from the tradability of allowances themselves. These particular issues would arise whether the allowances were distributed for free (grandfathered), auctioned, or allocated in some other way. Part of our report will look at potential opportunities for the auction design to address or mitigate concerns that would

exist even in the absence of an auction. Where possible, we also comment on adjusting the program design to address these larger market issues.

1.1 *Methods of Analysis*

The long experience with auctions and the substantial value of items sold at auction has led to the development of a large body of academic literature on the subject. This literature has three branches: theoretical, empirical, and experimental. By focusing on the abstract strategic form of auctions, the theoretical literature allows us to make predictions about how various types of bidders will respond to particular auction forms in particular situations. These theoretical predictions are put to the test in two ways.

First, there is substantial and growing experience with real-world auctions and an increasing value of goods that are sold by auction in both the public and private sectors. We can find empirical evidence of how well the theory fits the actual results by reviewing case studies and statistical examinations of the actual outcomes of real-world auctions (Hendricks and Paarsch 1995; Athey and Haile 2006).

Second, increasingly auction theory is being tested and refined in the economics laboratory (Kagel 1987; Rassenti, Smith et al. 2002; Staropoli and Jullien 2006). An economics laboratory generally comprises a group of human participants at a set of computers that are linked together with specialized software that allows the participants to be presented with a set of carefully designed decision tasks where the incentives, choices, information, and other characteristics are carefully controlled. Economics experiments are increasingly used for analyzing public policy, economic theories, and institutions. By allowing one factor to vary while holding all other factors constant in the laboratory, experimentalists can test theoretical predictions about how that “treatment variable” affects outcomes. Due to their availability and suitability, college students frequently are recruited to participate in the experiments in economics laboratories. Generally, the experiment is structured so that these student participants earn a payment based on the outcome of their choices.

Auctions, in particular, are well-suited to experimental investigation because of their compact institutional form. Experiments have proven valuable not only in testing strategic relationships and the likely market outcomes that might be obtained, but also because the formalized experimental setting requires a detailed specification of the institution that will embody the auction format. Thinking through all of the details

associated with that institution helps planners to anticipate potential problems (Holt, Shobe and Smith 2006).

During the past few years, the FCC has used laboratory experiments to guide the work on spectrum auctions (Goeree and Holt, 2005; Goeree, Holt and Ledyard, 2006, 2007). In particular, the high visibility 700 MHz auction scheduled for early 2008 will implement a Hierarchical Package Bidding (HPB) format that was proposed by two members of the RGGI auction design team (Goeree and Holt). When this format was first proposed to the FCC last February, the agency reacted enthusiastically and asked for laboratory tests. The FCC Public Notice (released 8/31/07) that invited comments about our proposed HPB format mentions the importance of the experiments and the relative success of HPB.¹ The more recent Procedures Public Notice (released 10/05/07) states “...we will use *HPB* in part because the mechanism for calculating prices is significantly simpler than other package bidding formats...”²

Experiments with human subjects are resource intensive and original experiments may not be necessary or appropriate for addressing all the questions that RGGI has about how to design the allowance auction. Thus, in this research we draw on the results of a rich literature on past auction experiments to help to inform our judgments about certain auction design features. We also draw on our own experience and familiarity with the design of past emissions cap-and-trade programs and allowance auctions in other settings, such as the Virginia NO_x auction, and the limited experience of others in Europe with auctioning of CO₂ emission allowances under the European Union Emissions Trading Scheme (ETS). In addition, we draw upon the empirical literature analyzing the performance of past auctions mentioned above to help inform our conclusions and recommendations regarding the design of a RGGI auction.

1.2 Two Phases of Research

This research was conducted in two phases. Phase 1 included the initial preparation of an annotated bibliography of key papers on auctions generally, emissions auctions in particular, and other related auction topics. Phase 1 also included a round of experiments to investigate the performance of a number of auction types considered to be likely candidates for use in a CO₂ allowance auction. The Phase 1 experiments measured the

¹ See http://fjallfoss.fcc.gov/edocs_public/attachmatch/DA-07-3415A1.pdf.

² See http://fjallfoss.fcc.gov/edocs_public/attachmatch/DA-07-4171A1.pdf.

efficiency and revenues from several different auction types in the abstract. For the most part, they did not take into account asymmetries in the cost of compliance among participants or the likely presence of a secondary (spot) market for allowances. Based on results from the academic literature, experience with previous auctions of allowances, and the laboratory experiments, the Phase 1 interim report, completed in May 2007, provided preliminary recommendations concerning which auction forms are likely to provide the best fit for the auction of RGGI allowances and related recommendations on the frequency of auctions, the role of reserve prices, and facilitating forward markets, among other things. The results and recommendations of the phase 1 analysis largely are folded into this report, although they have been modified and extended substantially as the result of additional experimental evidence and research.

In Phase 2 of the research, we expand considerably upon the literature review initially presented as an appendix to the interim report. Because our experimental findings in Phase 1 regarding the performance of the different auction types did not reveal a clear winner, we continued experiments aimed at the basic question of identifying the auction type that performs best along an expanded set of performance measures, including price discovery and limiting collusive behavior. We also examine the effect of reserve prices and banking, do more analysis of combining auction with secondary markets, and look at the effects of participation by “non-compliance entities.” In addition, to respond to questions that have arisen as the states begin to conduct rulemaking to implement RGGI, we perform some experiments to look at so-called “hoarding behavior” and the effects of different mechanisms that have been proposed to limit hoarding.

1.3 Organization of this Report

The remainder of this report is organized in two parts. The first part introduces the motivation and experimental methodology to be used. In section 2, we describe the general types of auctions that were considered in this investigation. Section 3 describes the criteria that are used to evaluate these auction types. In section 4, we describe the basic approach to the use of experiments, including how the cost structure and incentives in the experiment mirror in a stylized way the situation in RGGI. Section 4 also describes the ways that we evaluate the experiments in quantitative terms and describes the types of experiments that were completed.

The second part of the report focuses on results from the literature and from the experiments that were conducted as a part of this study. Section 5 addresses an assortment

of issues specific to RGGI, some of which do not hinge on the auction directly but all of which affect the implementation of the auction. One of the topics addressed here is how the auctions and the allowance markets could be monitored. Section 6 addresses collusion in an auction and how to limit it, one of the important issues in any auction design and one that informs our recommendations also. Section 7 addresses price volatility in general and the role for a reserve price. Section 8 addresses the ability of the auction to move to an equilibrium price, especially when new information about underlying costs of market participants is revealed, and describes the relationship between the auction and the spot market. Section 9 addresses the potential hoarding of allowances as it has been discussed in the context of potential behavior by parties external to RGGI. We investigate whether combining auctions with free allocation changes the performance of the auction in section 10. We return to a more comprehensive description of our recommendations in section 11. We have two appendices that provide a mapping of the questions from the original Statement of Work to places in the text where these questions are addressed and an annotated bibliography of much of the relevant literature.

2 Background on Auction Types

2.1 Introduction

Sources covered by the RGGI program will be required to surrender one emission allowance for each ton of CO₂ they emit into the atmosphere. Allowances are identical except for their vintage, which determines the first year in which they may be used. Once an allowance vintage year has been reached, the allowance may be used in that year or banked for use in a future year. Since many allowances will be sold in a single auction, this study is limited to auction forms appropriate to the sale of multiple units of an identical commodity. Only one allowance vintage should be sold in any auction. This is because there likely will be differences in the market value (and price) of allowances of different vintages.

Multi-unit auctions usefully can be categorized in two dimensions. The first dimension is the **number of rounds of bidding**, one or more than one, before the final determination of the sale price is achieved. Single-round auctions sometimes are known as sealed-bid auctions, meaning that after the bidder submits a bid there is no further interaction and the bidder simply awaits an announced outcome. In contrast, a multiple-round auction involves interaction because the bidder has a chance to change the bid in

response to information that is learned after each round. The second dimension is **how the price is set** for the buyers. A uniform-price or single-price auction identifies a single price for all transactions. A discriminatory-price (or “pay as bid”) auction yields final prices that differ among buyers and depends on the amount of each buyer’s bid. These different auction forms have different properties and may be used in combination.

These two characteristics allow us to make a useful, if necessarily incomplete, categorization of four basic auction types: uniform price sealed-bid, discriminatory price sealed-bid, uniform price multi-round, and discriminatory price multi-round. A very large academic literature has explored various aspects of auction performance; however, relatively few papers have examined the relative merits of each of these auction forms in multi-unit auctions. Moreover, most of the prior experimental research on multi-unit auctions pertains to the special case of only two units per bidder, which highlights the strategic incentives but is of limited relevance for the RGGI setting. How the auction types rank in economic efficiency and in raising revenue varies depending on numerous factors, including competitiveness, risk aversion of bidders, reservation prices, the presence of resale markets, and disclosure of bid information.

In experiments we considered eight alternative auction formats, including three mentioned in the statement of work that have been used previously in the field to auction emission allowances. Two of the types that have been used previously are single round, including the discriminatory price sealed-bid (used for SO₂ allowances under Title IV of the 1990 Clean Air Act Amendments) and the uniform price sealed-bid (used in Ireland for auctioning EU ETS CO₂ allowances). The third type that has been used for emission allowances is an English clock (ascending bid) multi-round uniform-price auction (used for the Virginia NO_x auction of allowances to comply with the NO_x SIP Call). These auction formats are described in section 2.2, along with the other auction types that were tested, including: a Dutch (declining price) multi-round auction with discriminatory pricing (“buy now”) and an English clock followed by a final sealed-bid discriminatory-price “shootout,” which we called the “shot clock.” In addition, we conducted other trials with a multi-round discriminatory-price auction, a continuous discriminatory-price, and continuous uniform-price auctions. Each of these is described below.

2.2 Auction Formats

All of the auction formats tested in this study are multi-unit auctions for a fixed number (Q) of allowances. Each bidder is assigned a production capacity and each unit of production requires some number of emission allowances, which varies among participants

reflecting a distribution of technologies. Bidders' values for allowances are determined by the profit margins on their production capacity and by the number of allowances needed to cover the production activity.

Each bidder is given an "activity constraint" that restricts the number of allowances on which they can bid. In practice, this activity maximum could be infinity (i.e., no limit) or it could be determined by financial pre-qualifications. In experiments when banking is not allowed, bidders only bid on allowances that they need to support their production activity. In this case, the activity limits are not binding, except in the English-clock (ascending price) auctions where, as a feature of that format, a bidder is not allowed to increase the number of units requested as the price rises in subsequent rounds. Also, in the shot-clock format, the activity constraint that the bidder ends up with in the penultimate round constrains the number of allowances that they can bid for in the final round shootout.

Discriminatory Sealed-bid: This is a single-round (sealed-bid) auction in which the bidders can submit multiple offers to purchase allowances with bids at different prices. The highest bids for the Q allowances to be sold obtain allowances at their own bid prices. This is the type of design used for the annual auction of SO₂ emission allowances by the U.S. Environmental Protection Agency (EPA). The auction is "discriminatory" because the price paid varies among bidders in relation to their bid price. In the implementation we used for experiments, ties at the cutoff price were decided at random.

The discriminatory-price auction is a simple auction to conduct and understand. Discriminatory-price auctions have been used frequently by governments to sell assets such as timber, securities, oil leases, and real estate. They also are used in procurement where participants in the auction bid a price at which they are willing to supply goods to the government. Under fairly stringent assumptions about bidder characteristics, the discriminatory- and uniform-price single-unit auctions should, in theory, raise the same revenue and should be equally efficient (Vickrey 1961). However, this revenue equivalence fails to hold up in more realistic environments, and, in particular, it fails to hold in the case of multi-unit auctions (Ausubel and Crampton 1998).

According to theory, in auctions for single-prize "units" where participants bid to obtain a single unit and many units are sold, such as hunting licenses, the presence of bidder risk aversion will cause revenues to be higher in a discriminatory-price auction than

in a uniform-price auction.³ However, in multi-unit auctions, where participants bid to obtain multiple units, the amount of revenue collected can be greater than or less than revenues from a uniform-price format auction. In experiments, revenue comparisons between the discriminatory- and uniform-price formats depend on the nature of the distributions of the bidders' values (willingness to pay) for the items being auctioned, the experience of the bidders, the information available to bidders, and possibly other factors. (Miller and Plot 1985; Cox et al. 1985)

Uniform-Price Sealed-bid: This is also a single-round, sealed-bid auction in which bidders can submit multiple bids at different prices, but the price paid by all bidders with the highest bids for the Q available units is equal to the highest rejected bid. This is the type of auction that was used for CO₂ allowances in Ireland in 2006. It is also transparent and easy to conduct. In auctions for a single unit, Vickrey (1961) showed that in theory the bids will reveal bidders' true values and will produce an efficient outcome, with allocation of the allowances going to the bidder with the highest value. However, Ausubel and Crampton (1998) showed this result does not necessarily carry over to a multi-unit case due to a phenomenon known as demand reduction. This implies that bidders may attempt to manipulate the clearing price by bidding low on "marginal" units in the hopes of bringing down the market-clearing price. Their theoretical results indicate that there is no clear ranking as between discriminatory- and uniform-price auctions when multiple units are sold. In experiments, the uniform-price auction yields a variety of results depending on the setting.

Uniform-price auctions also may involve some embarrassment for the seller if some bidders with very high bids obtain units at low prices. This may leave the seller open to criticism that buyers obtained goods at prices substantially below what they were willing to pay. On the other hand, discriminatory-price auctions put the buyer in a similar position. They may find that they have paid much more than others for the goods purchased.

English Clock: This is a multi-round auction in which the auctioneer posts a sequence of increasing (ascending) prices, usually at regular time intervals, and in response the bidders state the quantity they are willing to buy at the specified price. The "provisional" price starts at a price low enough so that the quantity demanded at that price

³ Cox et al. (1985) model an auction where multiple units are sold but where each bidder submits a bid for a single unit. Under simple theoretical models of bidder behavior in this auction setup, the presence of bidder risk aversion may be expected to cause revenues to be higher in a discriminatory-price auction than in a uniform-price auction. Experimental tests, however, reject the hypothesis of higher revenues from discriminatory auctions. Revenues from the two auction types were statistically indistinguishable.

is expected to be greater than the amount the auctioneer has to sell. The price is raised, as if by the hand of a clock, in response to the excess demand (e.g., the amount that the total quantity bid (Q^*) is greater than the available number of allowances (Q)). At each stage, the provisional price is announced and bidders state how many units they desire. The auction stops when the demand falls below the amount offered for sale (i.e. $Q^* \leq Q$).

There is an issue of how to deal with the possibility of unsold units if $Q^* < Q$ in the final round. In the experiments, we follow the procedure used in the Virginia NO_x auction of rolling back the price by one bid increment and selling all Q units at the lower price if to do so would increase revenue. In case of a rollback, all who expressed a willingness to buy at the higher price are included, and the remaining $Q - Q^*$ units are allocated on the basis of the chronological order in which the bids were submitted in the penultimate round, which provides bidders with an incentive to bid early in each round. If rolling back the price by one bid increment would not increase revenue, then the permits are sold at the final price determined by the clock and there remain unsold permits.

In order to force bidders to bid actively, each bidder's activity limit (limit on the quantity that can be bid and purchased) falls to the number of units requested in a round and cannot be raised in subsequent rounds, so activity has a "lose it or use it" feature that prevents bidders from hiding their interest in early rounds. Another issue that arises in a multi-unit auction is what information to provide bidders after each round. The experience with the Virginia NO_x auction and in other settings that we have reviewed suggests that it is best not to reveal the total number of allowances requested in each round so that bidders will not be able to determine whether unilateral demand reductions on their part will stop the clock. Providing less information also will tend to discourage collusion among bidders (Klemperer 2002). Multi-round auctions have the advantage of giving bidders an opportunity to think carefully and possibly acquire more information (see Jehiel and Compte 2007) as the prices develop. Also, at each stage, each remaining bidder knows that there are other bidders who value the items at least as much as they do. Delays can be minimized by providing an incentive to bid early.

Shot Clock: This is also a multi-round English-clock auction with the same activity constraints described above. The clock price rises in successive rounds, and it stops when the total number of units requested falls to a cutoff level that is a specified fraction higher than the number of units being auctioned (i.e., $(1+x)Q$, where $x > 0$). For example, if a 10% cutoff trigger were used in an auction of 1,000 items, the cutoff could be triggered when the quantity bid drops below 1,100 items. When the clock stops, all bidders may submit a final set of sealed bids in the form of quantities and prices into a

discriminatory-price, sealed-bid auction. The final set of bids is subject to two constraints: 1) all purchase orders at the final clock price are submitted as bids unless they are raised, so any bid must be greater or equal to the final clock price, and 2) the number of allowances bid for may not exceed a bidder's activity level in the previous round of the clock phase of the auction.

Allowances are awarded to those making the Q highest bids and bidders pay their own bid prices, so this is a hybrid between an English-clock and a discriminatory, sealed-bid auction. If there are any units leftover in the shootout round, the final clock price applies. It is sometimes called an "Anglo-Dutch auction," since the ascending-price phase is like an English auction and the final shootout has a discriminatory flavor, as does the multi-round Dutch auction to be described next (Klemperer 1999). As in the case of the English-clock auction, incentives for collusion and strategic manipulation may be reduced by not revealing the numbers of allowances requested after each round. In addition, incentives for collusion may be reduced by not revealing the exact level of the cutoff. The presence of the final shootout stage is intended to reduce the effectiveness of collusion and strategic manipulation (Klemperer 1999; Goeree and Offerman 2004).

Dutch: This multi-round discriminatory-price auction starts with a high provisional price, which falls by predetermined increments. The auction is discriminatory in price because in each round the bidder can "lock in" some purchases at the current provisional price (analogous to a "buy now" provision in an online auction at eBay) and/or the bidder can wait for the price to fall. The auction stops when the number of allowances locked in is greater than or equal to Q , with ties in the final round decided by the time at which a bid was entered, again providing an incentive for bidders to act early in each round.

Other Auction Types: In addition to these five auction types, we conducted less comprehensive tests on various other auction formats. One was a **multi-round, discriminatory-price auction** with increasing prices, which was loosely patterned after the Federal Communications Commission (FCC) simultaneous, multi-round auction. In each round, bidders submit bids and the Q highest bids are announced as provisional winners. These winning bids need not be raised in the following round, but provisionally rejected bids must either be raised or withdrawn (thereby reducing a bidder's activity). In tests, this format required more than five times as many rounds of bidding to reach convergence as a simple clock auction, since bid increases for a small number of rejected bids tended to rotate across bidders, thereby slowing the overall degree of price increases. An alternative would be to have a fixed number of stages (say 2–3) and to have bids in the initial stage(s) determine eligibility and lower limits for bids submitted in a final stage.

This format is similar to the shot clock in that early-round bidding determines eligibility and may provide some price discovery information, but we decided to use the shot clock instead for the purpose of investigating performance of this type of hybrid approach.

Another auction format that was explored was a **continuous-price auction**. Bidders can submit multiple bids with different prices for different quantities, and they have the opportunity to raise (but not lower) their bids continuously during a specified timeframe in which bidding is open. At any point in time, bidders can observe which of their own bids are “provisionally winning,” but they cannot observe others’ bids. One variation is a continuous discriminatory-price auction, in which bidders pay what they bid. In contrast, winning bidders in a continuous uniform-price auction pay the highest rejected bid at the end of the auction.

3 Criteria for Selection of an Auction Type

The specification of an auction design is guided by a number of performance criteria and principles. These criteria come from various sources, including the statement of work for this project, the report on the July 2006 RGGI Auction Workshop in New York City, conversations with the RGGI Staff Working Group, and input from RGGI stakeholders. Several of these criteria relate both to the auction and to the performance of the larger allowance market. Each of the criteria is listed below with a brief description.

- a. *Efficiency*: The allowance auction should be designed in a way that results in those bidders who have the highest value for a RGGI CO₂ emission allowance obtaining that allowance. On the producer side, emissions are, for the most part, reduced by reducing production differentially for different types of fossil-fired facilities, e.g. coal and gas. An efficient allocation of allowances means that CO₂ emission reductions are being made at lowest cost to society.
- b. *Price discovery*: A market for CO₂ emission allowances should result in the allowance price being equal to the marginal cost of reducing CO₂ emissions (either through fuel switching or by reducing electricity generation), and that cost will be approximately the same for all firms. This outcome results in the most cost-effective distribution of CO₂ emissions across firms. Accurate price discovery in an auction can help identify a market price close to the marginal cost of reducing emissions. Once the market has reached this equilibrium, then the spot market will provide a continuous summary of current opinions about the current value of allowances and, hence, the current marginal cost of reducing emissions. This price

will adjust daily as expectations change concerning fuel prices, electricity demand, and other factors.

- c. *Liquid allowance market/no interference with secondary market*: The auction should not impede the liquidity of the larger allowance market. Liquidity refers to the ability to convert emission allowances into cash through sale or to purchase additional allowances. Liquidity is not the same thing as the volume of trade in the allowance market. Liquidity is ensured by having many buyers and, in the secondary market, many sellers. The auction should not inhibit the smooth functioning of the secondary market by limiting options for trading or by yielding a price that differs systematically from the secondary market.
- d. *Minimize price volatility*: The allowance auction should be structured in a way that seeks to limit the variation in price over time. Large variations in allowance prices can be an impediment to firm-level investment planning. Allowance prices will vary in response to new information about interest rates or fuel costs, but these variations should not be augmented or amplified by market responses to uncertain features of the regulation or auction design.
- e. *Guard against collusion and/or market manipulation*: The allowance auction should be designed in a way that limits opportunities for bidders to actively or tacitly collude to keep prices low. To the extent possible, the auction also should limit opportunities for bidders to bid up the price of allowances above the competitive price, which we refer to as hoarding. Because collusion and hoarding are potential issues in the allowance market, and not just the auction, there may be a limit to the ability of an auction design to limit incentives for hoarding.
- f. *Fairness and transparency*: The auction rules should be transparent and available to everyone who might want to participate. The rules should not discriminate against any potential qualified participants.
- g. *Revenue*: In most cases when auctioning a publicly owned asset, such as treasury bonds or surplus property, the government prefers an auction design that maximizes government revenues. This is not always the case. In the electromagnetic spectrum auctions, the FCC also was responsible for ensuring competitive communications markets in the United States and, in some cases, was required to give preference to particular classes of bidders (McMillan 1994; McAfee and McMillan 1996). Although some stakeholders have indicated that maximizing revenue is not a priority for them, evaluating the effectiveness of the

auction in raising revenue is an important piece of information for comparing potential auction designs. A revenue reduction that results from success bidder collusion is something to be avoided.

- h. *Minimize administrative and transaction costs*: The auction should not impose large administrative costs on the RGGI states. The auction should also not impose large transaction costs on auction participants.
- i. *Familiarity to industry*: If two auction designs perform equally well on a host of the criteria, it might be reasonable to favor a design that is more familiar to the electricity industry in the region. This would reduce the cost of learning the institution and could raise participants' comfort level with the auction.
- j. *Align well with wholesale energy and capacity markets*: Emission allowances are an input to the production of electricity and thus it is important that generators have an opportunity to purchase or contract for future purchase of emission allowances before they place offers to supply electric energy or generating capacity in the relevant markets. Since RGGI compliance periods are three, or possibly four, years long, the need to secure allowances before actually producing electricity is less critical. However, for some purposes, such as obtaining financing for new investment, it may be important to obtain allowances that have a future vintage or that can be banked for the future to provide assurance to investors. Also, the auction may provide useful cost information to agencies that monitor cost and performance in the wholesale energy and capacity markets.

Electricity producers, electricity regulators, independent-system operators, and other RGGI stakeholders have experience with supply-side auctions that are used to set prices of electricity and generating capacity, and some of that experience is relevant for the context of allowance auctions as well. However, there are important differences between energy markets and allowance markets that make this context quite different. For one thing, allowances are storable and bankable for future use but electricity is not. Thus, strategies to manipulate prices and supply in the allowance market ultimately will be less effective than such strategies would be in electricity markets. Electric-energy auctions are held very proximate to the time of use, so the post-auction secondary market is likely to be less important in terms of revealing information about price or reallocating the resource efficiently if necessary.

- k. *Be open to all participating states, with participation optional*: The RGGI auction should be open to all participating states, providing an opportunity to reduce transaction costs and make the auction as large and liquid as possible.

4 Basic Experimental Approach for This Investigation

Auction experiments were conducted in a stylized setting that was intended to capture key aspects of the market for allowances, while keeping the setup simple enough to be relatively transparent for participants. Each experimental session typically involved 12 participants recruited from the undergraduate population at the University of Virginia. Each participant was given the role of a firm with multiple “units” of capacity that could be used to produce a product that sold at a known price. The use of each capacity unit required that the person obtain permits. Participants were given a financial reward for showing up in addition to earnings from purchasing the auctioned “permits”⁴ at prices below their value.

Subjects earned money from participating in the experiments, and they had a financial incentive to improve their earnings by improving their payoffs in the experiments. Subjects were paid \$6 for showing up. Alternates who showed up but were not needed were paid \$10 and dismissed. On average, participating subjects earned \$27 per session. We conducted more than 100 experiment “sessions,” each of which lasted from 1-2 hours and involved 6-12 subjects. In total, over 1000 subjects participated in these experiments.

Experiments were pursued in two phases. In the first phase, which was completed in May 2007, the primary focus was on a basic setup without spot markets or structural conditions that would facilitate collusion or impede price discovery. All of the five auction formats used performed comparably well, yielding efficiency measures of near 100% and revenues that were close to competitive equilibrium Walrasian benchmarks. In the second phase, conducted between May and October 2007, these auction designs were investigated with a richer informational and strategic setting, which better served to “stress-test” and differentiate the performance of the auction types. In the next section, we report on the basic production technology and market structure that were common throughout the experiments.

⁴ The word “permits” was used in the experimental sessions to abstract somewhat from the specific context of pollution trading.

4.1 Production Technology and Market Structure

To represent the technological characteristics of the market in the laboratory experiments we draw on information that the emission rate (tons CO₂ per MWh) for gas-fired generation is about 0.428 times that for coal-fired generation. Currently, electricity generation from coal- and gas-fired generation is roughly equal in the 10-state RGGI region.

We introduced an asymmetric cost of compliance by requiring some subjects to obtain more permits to operate capacity than others. For the laboratory experiments, it is important to have a correct stylized representation of the underlying technology, but it is not important to achieve precision. We assume that coal-fired generation requires two emission allowances for every one allowance required by gas-fired generation and that capacity for generation exists in equal proportions. Specifically, half of the subjects were “low users,” who needed one permit for each capacity unit, and half were “high users,” who needed two permits to operate each capacity unit. One can think of low users as using natural gas and the high users as using coal. Experiments typically had equal numbers of low and high users, which was intended to roughly mimic the relative proportion of coal and gas generators in the region. In most of the experiments, except those to explore the possibility of collusion, production costs for each unit of capacity were randomly generated for each new auction in order to ensure that comparisons among auctions were not driven by particular configurations of costs. To keep the experiment from becoming too complicated, we used a relatively small number of permits. Typically, 60 permits were sold in each auction. Thus, each permit in the experiment corresponds to a block of “allowances” in the market; the size of the block being determined by the “lot size” chosen by RGGI.

In all of the experiments, the product price—the price of electricity—is certain, and it is not affected by any decisions made by subjects in the experiment. This characteristic removes any strategic incentive to affect the price of emission allowances in order to affect the market price of electricity. The possibility of manipulating allowance price or allowance holdings in order to manipulate product price is a concern that has been addressed in the economics literature (Misiulek and Elder 1989). Although it is a potential consequence in the allowance market, it is not a consequence of the auction. Nonetheless, there are some features of auction design that might help mitigate the manipulation of electricity prices, which we discuss below. With a fixed and certain output price there is no production-related motive for non-emitters (nuclear, hydro) to acquire permits. However,

in some experiments we include subjects with non-emitter roles to see how it may affect the dynamics of the auction and spot market.

When banking of permits was not allowed in the experiments, each auction was a separate strategic situation. In some experiments, banking over auction rounds was allowed to provide an opportunity for inter-period planning. With all experiments with banking there also was the opportunity to trade in a spot market.

The profit margin (payoff) to subjects is determined by the difference between the known price of the product and its cost of production. The cost of production involves the randomly generated cost and the cost of acquiring permits. The value of permits to a subject is determined by taking the profit margin and dividing by the required number of permits to operate a unit of capacity. For example, with a production cost of 6 and a price of 12, the profit margin is 6. The permit value would be 6 for a low user who requires one permit to operate the capacity unit, whereas the value of each permit would be $3 = 6/2$ for a high user who is required to have two permits to operate.

The costs of operating capacity for low users were set to be roughly twice as high as the costs for high users to reflect the higher costs associated with natural gas-fired generation. This cost difference also served to approximately equalize earnings across subjects with different roles; that is, subjects with relatively higher emission costs had relatively lower production costs. The costs for low users were randomly drawn from the interval [5, 10], with all values in this interval being equally likely, and the costs for high users were drawn from the interval [2, 6].

With a fixed-output price, a “wide” distribution of costs determines a wide range of permit values. For low users, costs are drawn from the range [5, 10], with all draws in this range being equally likely; then a product price of 12 will result in a range of permit values between 2 ($= 12-10$) and 7 ($= 12-5$). The values for high users are obtained by dividing profit margins by the required number of permits (2) per capacity unit, so a cost distribution from the range [2, 6] results in values between $3 = (12 - 6)/2$ and $5 = (12 - 2)/2$.

Note that a narrow range of costs would determine a narrow range of values and a relatively flat (“elastic”) demand for permits, whereas a wide range of costs would determine a wide range of values and a more inelastic demand. We used narrow ranges of values in some sessions to induce a more elastic demand for permits, an attribute that is identified as important in the experimental literature (Miller and Plott 1985). The Porter et al. (2007) analysis of the Virginia NOx Allowance Auctions suggested that an English

clock auction, which charges a uniform price, might perform better with a narrow cost range, while discriminatory auction formats might perform better with a wider cost range.

In the narrow-range treatment, the cost distributions of [2, 6] and [5, 10] for high and low users were reduced to [3, 4] and [7, 8], respectively. The narrow range of values creates a more competitive situation, with low earnings, so participant earnings were doubled for the narrow-range treatment by doubling the conversion rate between lab earnings and cash earnings paid at the end of the experiment. Table 4.1 reports the parameters that are used.

Table 4.1. Experiment Parameters

	Wide Cost Range	Narrow Cost Range
Low User Cost Distribution	[5, 10]	[7, 8]
High User Cost Distribution	[2, 6]	[3, 4]
Product Price	12	12

4.2 Measures of Performance

Section 3 outlines important criteria for evaluating the alternative designs that we have considered. Some of these criteria can be informed by a review of the literature and previous experience, and some can be informed by the laboratory experiments. To analyze the laboratory results, we rely on two primary numerical measures of performance: **efficiency** and the ability to raise **revenue**.

To understand how efficiency is measured in the experiments, we refer to Figure 4.1. The vertical axis in the figure indicates the value of permits to a firm, represented by a participant in the auction. The value of a permit relates directly to the value of production. To simplify this discussion, let us assume one permit is required for each unit of production. (In the experiments, sometimes two permits are required for one unit of production.) In this case, the value of a permit is determined by the difference between the price received for production and the participant's cost of production, not including allowance costs. This value is indicated on the vertical axis, and the quantity of permits (emissions) is on the horizontal axis. Each step on the demand schedule shows the permits that are worth as least as much as the associated value on the vertical axis. The area under the schedule is a measure of "economic surplus," or in this case, simply economic profit.

The schedule indicates the aggregate willingness to pay for permits across the industry. In the absence of a constraint on the number of permits (emissions), firms would expand production until the marginal value of production equaled zero. However, with the addition of an emission constraint, profit maximizers would be willing to buy a permit as long as the value of producing a unit is greater than or equal to the cost of a permit. One of the virtues of using experiments to test auction performance is that the exact values of each production unit are known to the experimenter.




Figure 4.1. Measures of Efficiency

With a cap on emissions, as indicated by the vertical line in the figure, the quantity of production is effectively limited, and there will be production units that would be profitable in the absence of the environmental constraint that will not be produced. An efficient allocation of emission permits will maximize value so that the production units that are used are the most profitable units. Efficiency measures the extent to which this maximum production value is realized. Panel A illustrates a fully efficient allocation of permits. However, if there were any skips in the allocation of permits and if any permits were allocated to lower portions of the demand schedule, then permits would not be going to their highest valued use, and the allocation would be inefficient. Panel B illustrates the loss of efficiency by having some amount of permits allocated in a way that is not their highest valued use.

For the purposes of the experiments, the measure of revenue indicates how much of the maximum production value, or profit, is captured as revenue.⁵ It is unrealistic to expect 100% revenue, as there would be no profits for producers.

In a competitive market, there will be a price that causes the quantity demanded to be just equal to the quantity supplied; this price–quantity pair is known as Walrasian equilibrium. In Walrasian equilibrium for pollution permits, the production value would be maximized because producers would exchange permits for money until the permits were owned by those who value them the most. In other words, the Walrasian equilibrium would be an efficient allocation of permits.




Figure 4.2. Walrasian Equilibrium and Revenue from the Auction

The Walrasian equilibrium is pictured in Figure 4.2. The Walrasian revenue is the amount of money that would be raised if all of the permits were purchased at the equilibrium price. In our analysis, we use a Walrasian revenue prediction to measure the ability of the auction to raise revenue. The Walrasian revenue prediction is the percent of

⁵ It is noteworthy that the profit associated with each unit of production in the presence of an emission cap is not the same as it would be in the absence of an emission cap because the emission cap effectively imposes a barrier to entry that raises the returns to those who hold permits. When the auction captures a portion of the profits of producers, it is capturing in large part the change in profits that would accrue to producers if permits were given away for free.

maximum production value that would be captured as revenue if permits were sold at a competitive equilibrium price.

Several other qualitative measures of the performance of the auction types come into play. One is **price discovery**, which indicates how well the market prices in each auction track the equilibrium price. In particular, we focused on price discovery in several different auction types when there was an abrupt change in equilibrium price due to an unannounced shift in subjects' costs. If price discovery is high, then the auction prices move relatively quickly to the new equilibrium in repeated experiments. Another is the proximity of the auction outcome to the **spot-market price**. In theory, the auction price should be very close to the spot-market price when subjects can both participate in the auction and trade freely in the spot market. Finally, **debriefing** of subjects provides important insights into the transparency of auction design. The auction is transparent if subjects inform us that they easily understood the rules of the auction and how the auction equilibrium was achieved. We also debrief subjects about their strategies, especially with respect to collusive behavior. We also refer to the recorded correspondence among subjects when they were able to communicate using instant messaging to provide "cheap talk" making suggestions to the group about how to bid.

4.3 Important Aspects of Phase 1 Experiments

The first phase of experiments was completed in May 2007, with the aim of identifying a preferred basic auction type. The Phase 1 experiments are described in table 4.2. We ran 15 sessions with wide cost distributions: 3 sessions for each of the five auction types (uniform price sealed-bid, discriminatory price sealed-bid, English clock, Dutch clock, and shot clock) using a total of 180 subjects (= 15x12). In the second set of sessions with the narrow cost ranges, we focused primarily on the three auction types listed in the statement of work. We ran three sessions using each of the three main auction types (uniform price sealed-bid, discriminatory price sealed-bid, and English clock) and one session each for the other two auction types (Dutch clock, shot clock), for a total of 11 sessions and 132 subjects. The participants for the second set of sessions in Phase 1 were recruited from those who had participated in the first set of sessions.

Table 4.2. Phase 1 Experiments

<i>Experiment</i>	<i># of Sessions</i>	<i># of Subjects</i>
Auction Performance: Wide Cost Distribution	15	180
Auction Performance: Narrow Cost Distribution	11	132
Baseline for Phase 2 with Spot Market and Banking	4	48

Finally, as part of Phase 1 we also conducted four sessions with a spot market following each auction (wide cost range).⁶ The spot market was structured so that participants could submit limit orders that specify a maximum quantity of permits and a maximum purchase price or a minimum sales price (e.g., sell up to six permits for at least \$4). Buy orders were arrayed from high to low, sell orders were arrayed from low to high, and the price determined by the intersection of these arrays was the price at which transactions were executed. Then after the spot market cleared, subjects decided how many permits to use in production and whether to bank permits or incur a deficit. It was announced that any deficit in permits was penalized at a rate of \$9 (about three times the predicted price) after the spot market that followed every third round of auctions (i.e., after rounds 3, 6, 9, etc.). That is, the experiment characterized a compliance or true-up period that occurred after every three auctions. In fact, there were only eight auction/spot-market pairs, but this was not announced in advance, so subjects behaved as though the experiment would continue beyond number eight.

The main result from the Phase 1 experiments was that all auction formats are reasonably efficient, and the revenues for the two single-round, sealed-bid formats (discriminatory and uniform price) were at least as high as those for the multi-round formats (Figure 4.3). In addition, we found no clear support for the conjecture that the uniform and English-clock auctions would perform better with narrow cost ranges. In particular, we did not see the dramatic revenue increase of more than 15% reported by Porter et al. (2007) for the treatment with the narrow range of bidder values (Figure 4.4). One procedural difference is that, in their experiment, subjects were put into a situation in which the nature of demand (elastic or inelastic) switched randomly from one auction to

⁶ In addition, we ran another 12 pilot sessions for testing purposes (some done with fewer than eight auctions or 12 participants) to refine the instructions and procedures.

the next. In contrast, our comparisons were between sessions with a series of auctions using the same demand structure (elastic or inelastic).




Figure 4.3. Revenues and Efficiencies by Session with Wide Cost Range; Revenue Near Walrasian Levels of 79%




Figure 4.4. Revenues and Efficiencies with Narrow Cost Range; Revenues Near Walrasian Levels of 94%

4.4 Important Aspects of Phase 2 Experiments

The second phase of experiments was completed between May and October 2007. Because the first phase did not lead to a definitive recommendation over the basic auction type, the second phase of experiments added institutional or strategic complexity to test the basic auction types in a multi-unit setting. In the second phase, we ran 68 sessions involving 600 subjects in total, and these experiments are summarized in Table 4.3.

Table 4.3. Phase 2 Experiments

<i>Experiment</i>	<i># of Sessions</i>	<i># of Subjects</i>
Loose Cap	18	108
Collusion	16	96
Price Discovery	18	216
Partial Grandfathering and Outside Sales	7	84

The most important variations included the following. One was modeling of a spot market, extending the approach initiated in Phase 1, to better understand how the auction interacts with the spot market. One version involved the existence of a standing-outside option for sale of allowances. Another version involved the participation of subjects who did not participate in the auction and had no capacity for production and who only participated through the secondary spot market. These subjects, referred to as “brokers,” could buy permits in the spot market and sell them in the “world market” for a known price. The world-market price exceeded the equilibrium price of the permit market if it were closed to outside sales.

Two additional rounds of experiments explored the opportunity for collusion. In addition to comparing auction formats, we explored the difference in settings when subjects only had the opportunity for tacit collusion with no communication among themselves and when they had the opportunity to communicate, so called “cheap talk” because they could make promises to collude but were not necessary committed to doing so. These experiments involved 96 subjects. Another round of experiments tested auction performance in the presence of a relatively loose cap on emissions. Having a loose cap creates a less competitive environment, since most bids would be accepted. These experiments involved 18 sessions and 108 subjects.

A third type of extension was to explore the ability of auction prices to track unanticipated changes in market conditions. This “price discovery” experiment involved a sharp increase in demand in the middle of a sequence of auctions. These experiments involved 216 subjects.

A fourth extension was to explore the effects of partial grandfathering of allowances on the performance of the auction. We conducted only a few sessions under this set of assumptions.

Part 2: Research Questions

5 Key Auction Design Features

The experiments performed for this research focus on questions of auction type and variations of auction performance depending on stringency of the cap, unanticipated shifts in allowance demand, and other factors. Several important features of an auction program are less amenable to experimental testing in the laboratory, but decisions about these features can be informed by past experience with auctions and by sound economic reasoning. In this section, we address several of these key auction design features. Each feature is related to a particular question or set of questions from the Statement of Work, which is highlighted at the beginning of each section.

5.1 Defining Allowance Vintages

The vintage of an allowance determines when it can first be used for compliance with the RGGI CO₂ cap. In the interim report, we made a preliminary proposal that allowance vintages be defined by three-year compliance periods and not on an annual basis. This way of defining vintages would eliminate differences between allowances that aren't really meaningful, since all allowances from within a three-year compliance period should have equal value at the time of compliance. We have concluded that this multi-year approach to defining allowance vintages will not work because of changes in the length of the compliance period that could be triggered in the event of sustained high allowance prices. Specifically, the length of the compliance period will be extended by 12 months if the average allowance price for the prior 12 months hits the second-stage price trigger of \$10 per ton (expressed in real 2005 dollars) at any time after 26 months into a particular compliance period. Once this provision takes effect, the starting year for future compliance periods also will shift ahead by a year, and thus it becomes necessary for allowances to have annual vintages in order to be able to accommodate these potential changes.

5.2 Frequency, Timing, and Size of Auctions

How often and when should RGGI allowance auctions be held?

The timing, size, and frequency of auctions all are important considerations that were raised in the RGGI auction workshop in the summer of 2006 and that have been

raised in our conversations with RGGI generators, RGGI Staff Working Group members, emissions brokers and traders, and others. In many cases, the concerns of different groups are similar, and in all cases, finding the best way to address those concerns involves trading off the costs and benefits across possible approaches. In our interim report, we presented some preliminary recommendations with respect to the timing and frequency of auctions, and we elaborate on those here.

Timing and frequency of auctions are a key concern for generators. Generators have expressed a desire for some degree of certainty regarding future allowance prices and allowance availability to assist in their planning for future investments. They want auctions of allowances of current and future vintages to occur before regional Independent System Operators (ISO) capacity auctions to allow generators to be able to secure the allowances they would need to perform those contract obligations in the future. Generators also want auctions to be held frequently enough that they will be able to get the allowances they need in order to cover the emissions from their operations in the near term. They want allowances to be auctioned in small enough lots so that inadequate cash flow or potential high costs of borrowing do not impede their ability to purchase the allowances they need. On the other hand, companies don't want to have auctions be so frequent and lots so small that the transaction costs of participating in the auctions becomes burdensome relative to what they can reasonably expect to purchase in a particular auction.

Frequent, small auctions have the virtue of limiting the number of allowances auctioned at one time and thereby limiting the likelihood that a buyer could use the auction to manipulate the market. This is particularly true given that RGGI compliance periods are slated to be three years in length. Frequent auctions also will contribute to the liquidity of the allowance market by making allowances available for purchase on a regular basis. Frequent, small auctions also will limit the potential for the allowance auction to disrupt the spot market by dumping large quantities of allowances on the market at a particular time.

When deciding how frequently auctions should be held, the desirable features of frequent auctions need to be weighed against the administrative costs of conducting multiple auctions and the transaction costs to the firms of participating on a frequent basis. Past experience suggests that a significant proportion of the administrative cost of holding auctions is related to the initial set-up of the auction, including the development of auction rules, deploying auction software, and establishing the mechanisms for prequalifying bidders (discussed below) and that the incremental costs of repeating a particular auction

type will be low in comparison to these initial costs. The costs that may increase with the number of auctions include:

- Costs to bidders and auction administrators of establishing prequalification
- Costs to bidders and of preparing bids or bid strategies
- Administrative costs of setting up, advertising, and running an auction
- Costs to bidders of posting bids.

Some costs will fall over some range as the frequency of auctions increases:

- Costs that may arise from noncompetitive behavior
- Costs to smaller firms of providing financial prequalification
- Costs of purchasing larger blocks of allowances.

Generators have lengthy planning horizons and can get significant benefits from an active market for future allowance vintages. An active futures market provides valuable signals about the relative scarcity of allowances in the future. Firms wishing to insure against the price risk or against the potential unavailability of allowances in the future can do so by purchasing allowances early. If current allowance vintages are relatively scarce, then purchasing future vintages may lower the costs of long-term compliance planning. The ability to purchase future allowances also could help with capacity planning and with demonstrating to the ISOs that the generator has the ability to perform if called upon to supply capacity, a precondition for bidding to supply in capacity auctions that may go out several years into the future. Auctioning some allowances prior to their vintage year will create a market in allowance futures. Once the first auction has taken place for a particular vintage, spot market trades will start to occur, and they will provide a more continuous signal of how the market price is evolving.⁷

One risk to bidders of purchasing allowances in advance is the potential that RGGI could be superseded by a federal program. So, while there are benefits to providing current access to future allowances, this must be balanced against the regulatory risk in an environment where large changes in the regulatory framework appear increasingly likely.

⁷ It is possible that spot market or contractual trades in future allowances could take place before the first auction occurs, but it is likely that this type of pre-auction trading activity, should it arise, will be thin because none of the parties to the trade could own the asset in trade. The considerable additional risk of contracting over an asset that only will become available some years in the future lowers the net value of the trade because of the much higher risk of nonperformance.

The greater the perceived risk, the greater will be the price discount applied to future vintages and the lower their value in managing risk for generators.

Balancing the costs, risks, and benefits leads us to conclude that a regular program of quarterly allowance auctions would be the most appropriate frequency. This provides the benefits of periodic price discovery and enhanced liquidity without interfering with the performance of a secondary market. As we discuss in section 8.2, experimental evidence and evidence from other allowance auctions is persuasive that auction and spot market prices will track each other closely.⁸ A regular sequence of auctions for allowances will be built into spot-market participant expectations and is unlikely to cause disruption.

Again, a balancing of risks and benefits leads us to conclude that selling allowances four years in advance of their vintage is likely to generate significant gains without imposing large regulatory risks on the value of the future vintage allowances. Current vintage allowances would be sold at each quarterly auction during their vintage year. Future vintages would be sold only at one of the quarterly auctions. For example, in the first quarterly auction of a given year, two separate auction sessions would be held: one auction for the current vintage and one auction for the vintage one year ahead. In the second quarter, one auction would sell the current vintage and one would sell the vintage two years ahead, and so on for the four quarterly auctions. Table 5.1 gives the details of proposed auction timing. The table reflects the possibility that auctions would be held before the program becomes binding, labeled as year 0.

⁸ This conclusion is consistent with the observed outcome in the Virginia NO_x auction of June 2004. NO_x allowance prices had risen in the weeks leading up to the Virginia auction. Prices at auction were on the order of 5% above the morning spot price, and prices continued on an upward trend for the following weeks.

Table 5.1. Schedule of Sales by Vintage

Yr	Q	Regular Auction		Forward Auction	
		Vintage sold	% of vintage sold	Vintage sold	% of vintage sold
0	1	Y1	12.5%	Y2	16.67%
0	2	Y1	12.5%	Y3	14.3%
0	3	Y1	12.5%	Y4	12.5%
0	4	Y1	12.5%	Y5	12.5%
1	1	Y1	12.5%	Y2	16.67%
1	2	Y1	12.5%	Y3	14.3%
1	3	Y1	12.5%	Y4	12.5%
1	4	Y1	12.5%	Y5	12.5%
2	1	Y2	16.67%	Y3	14.3%
2	2	Y2	16.67%	Y4	12.5%
2	3	Y2	16.67%	Y5	12.5%
2	4	Y2	16.67%	Y6	12.5%
3	1	Y3	14.3%	Y4	12.5%
3	2	Y3	14.3%	Y5	12.5%
3	3	Y3	14.3%	Y6	12.5%
3	4	Y3	14.3%	Y7	12.5%
4	1	Y4	12.5%	Y5	12.5%
4	2	Y4	12.5%	Y6	12.5%
4	3	Y4	12.5%	Y7	12.5%
4	4	Y4	12.5%	Y8	12.5%
5	1	Y5	12.5%	Y6	12.5%
5	2	Y5	12.5%	Y7	12.5%
5	3	Y5	12.5%	Y8	12.5%
5	4	Y5	12.5%	Y9	12.5%

How large should each auction be?

The size of each auction is determined completely by the number of auctions for allowances with vintages from each compliance period. Once fully implemented, the auction plan outlined in the previous section would result in eight auctions for each vintage: four in its vintage year and one each in the previous four years. At each auction, a minimum of 12.5% of the allowances to be auctioned for that vintage would be for sale.⁹

How does staggered implementation by RGGI states of auctions affect the efficiency and fairness criteria laid out below?

Other things equal, market mechanisms generate their greatest gains when risks and transaction costs are kept to a minimum and when opportunities for strategic behavior unrelated to true asset values are minimized. Careful attention to the details of market implementation can aid in minimizing these costly impediments to gains from trade. First, every reasonable effort should be made to ensure that all RGGI allowances are identical from the buyer and user points of view. There should be no differences in contract language, enforcement terms, reporting, or fungibility. Insofar as possible, there should be a single point of contracting and one auctioneer. If this were not true, then allowances from different states would have different values and there would be a multiplication of assets and prices, resulting in a large increase in transaction costs and risk for market participants.

A joint regional auction is far preferable to separate state auctions for several reasons. Differences in auction design and implementation may lead to confusing and irrelevant differences in price signals. States would be tempted to choose the timing of auctions, reserve prices, or other parameters in ways that favor them. In addition, multiple auctions almost certainly will raise the administrative costs of making allowances available to the market and the transaction costs for firms seeking to acquire them.

The same cautions do not necessarily apply to different choices among the states concerning the proportion of allowances sold rather than given out for free. As long as state policies on allocation are announced in advance and are applied in a predictable way, differences among the states are not likely to disrupt the performance of auctions or the subsequent secondary markets for allowances. It is well understood from economic theory

⁹ It is possible that the percentage could be more if unsold allowances from earlier auctions were added to the one-eighth standard share. Also, if some allowances are being allocated for free to generators, the rule for that allocation should be announced as far in advance as possible. Actual free allocations should be announced as early as possible so that generators know what they will get from the free allocation and can plan their purchases accordingly.

and evidence from market activity that the market price of allowances is not determined by how the allowances are handed out but rather by the cost of reducing CO₂ emissions at the margin.¹⁰

5.3 Financial Assurance Mechanisms

The prequalification of auction participants is essential to the integrity of the auction, and the auction literature provides strong evidence that careful thought must be given to the bonding mechanism that certifies eligibility in the auction (McMillan 1994; Binmore and Klemperer 2002; Klemperer 2002; Borgers and Dustmann 2005). Financial assurance mechanisms provide a way for RGGI to ensure that all auction participants can and will perform on their bids. Use of such mechanisms is standard procedure in all types of auctions, including energy auctions, to ensure the ultimate performance of those bidding to supply energy, such as in the New Jersey Basic Generation Service auction and in the FCC spectrum auctions. In the FCC case, bidders who were exempted from the standard financial assurance requirements were the ones who ended up defaulting on their bids (Burtraw and Palmer 2006). Financial assurance mechanisms typically include the posting of a bond, deposit, or letter of credit up front that would cover a substantial fraction of the bidders ultimate payment should she win in the auction. This mechanism would limit bidders to bids that don't exceed the level of financial assurance provided prior to the auction. An additional, penalties for default or non-payment on the part of the winning bidders would provide further incentive for winning bidders to perform on their bids.

Past allowance auctions, including the Virginia NO_x auction and the Irish auction of CO₂ allowances, have required financial assurance. In the case of Ireland, they discovered *ex post* that they should have set their deposit requirement for bidders at a much higher level to be sufficient to cover the upside potential for the price of EU CO₂ allowances. In Virginia, bidders with high credit ratings could use their rating to provide financial assurance. Most other bidders were required to deposit all of the money they were willing to spend on allowances in escrow with a designated third-party bank prior to the auction. The auction software automatically prevented bidders from bidding more than the amount of assurance posted. The share of the money placed in escrow actually owed for allowances by winning bidders was transferred by the broker directly to the state at the

¹⁰ The exception to this would be if allowances were awarded for free on the basis of output or emissions and allowance allocations were updated over time. Such an approach to free allocation would provide generators with an incentive to increase their generation in order to obtain a larger share of the total allowance "pie," and this would tend to increase the price of emission allowances (Burtraw et al. 2001).

conclusion of the auction upon presentation of evidence that a particular bidder had won in the auction (Burtraw and Palmer 2006). Firms not required to use the escrow method wired their payments to the state or, in the case of one very small municipal utility, sent a check. The Virginia setup resulted in very rapid settlement. Nearly all payments and allowance account transfers were completed within three days of the completion of the auction.

Every effort should be made to minimize the cost of providing financial assurance. Since these auctions likely will be held on a regular basis, assurance mechanisms, such as using credit ratings and letters of credit, may provide a high level of assurance at very low cost for most bidders. Even insolvent firms can participate using escrow methods, as was the case in the Virginia auction.

5.4 Market Monitoring

Monitoring of behavior in the allowance market would help to address a number of program goals. Monitoring may be effective if it can raise the visibility of behavior that is inconsistent with program goals because it can bring that behavior to the attention of administrators, who may want to make adjustments in the program if necessary. In addition, discovery of such behavior may identify potentially illegal activity. Moreover, raising the visibility of behavior that is inconsistent with program goals may raise the cost of that behavior by threatening to undermine the goodwill and public image of a firm.

There are a variety of potential problems that monitoring could address. Monitoring is a way to guard against potential collusion in the allowance market, either with the intent of realizing gains based on allowance holdings or with the intent of manipulating the price in electricity markets. Monitoring helps build investor confidence in the knowledge that the value of investments will not be eroded by illegal or unsanctioned activity in the market. Also, market monitoring provides an early-warning system for trends in the market that might be important to market administrators, even if the trends do not stem directly from behavior that is inconsistent with program goals.

If monitoring is administratively costly or raises transaction costs, however, then it will undermine the effectiveness of the allowance market by raising costs. Good program design should not allow market monitoring to be intrusive. There is reasonable business interest on the part of firms in limiting the disclosure of information about their strategic investment activities; information about allowance holdings could be a signal of business strategy. If market monitoring leads to the revelation of strategic business information, then it is likely to undermine the effectiveness of the market by limiting participation in the

market. Consequently, program administrators have to balance their interest in information about the workings of the market, even on a day-to-day basis, with a recognition that markets work best if they are unencumbered with unnecessary reporting requirements.

Within this general context concerning the role for market monitoring, it should be clear that market monitoring is not an auction design issue *per se*. However, information from the auction results may contribute to the ability of the market monitor to track behavior in the market. Also, as noted elsewhere, the design of the auction can help to mitigate some of the same problems that market monitoring is intended to address. Therefore, in deciding on auction design one may want to be informed about opportunities for monitoring.

The closest analogies in the United States to the RGGI market are the SO₂ and NO_x markets. Each of these are somewhat larger than the RGGI market is likely to be, with recurring annual values of emission allowances totaling roughly \$2–3 billion, albeit with considerable variance over time. Today there are two organized futures markets for SO₂ and NO_x run by NYMEX and the Chicago Climate Futures Exchange. These futures are a regulated commodity and fully regulated as “designated contract markets” by the Commodity Futures Trading Commission (CFTC). It is noteworthy that although the primary realm of authority for the CFTC is the futures market, the agency can “reach back” into the over-the-counter cash market when there is evidence of cash market transactions that affect the futures market and vice versa. The over-the-counter market, including bilateral transactions and auctions, would not be in the jurisdiction of the CFTC until and unless manipulation in that market affected a regulated futures market. Plans to organize a futures market for RGGI allowances have been announced.

A second agency interested in the performance of emissions markets is the EPA. The agency conducts ongoing analysis looking for a variety of indicators, aggregating allowance holdings by parent company/holding company, to look for concentrations of market power. Markets are examined weekly by looking at trading logs. Attention is paid to the activities of all participants, including non-emitting entities and hedge funds. The EPA’s accounting software, known as the Allowance Tracking System (ATS), is the backbone of both the SO₂ and NO_x markets. The ATS does not require disclosure of trades, but it does enable and require the transfer of allowances among authorized accounts. Allowances must appear in the appropriate account before they can be used for compliance. In practice, the large majority of trades are immediately recorded as transfers and many contracts for trades actually designate the trade as effective when the transfer among accounts is recorded. In no case does the ATS require the disclosure of prices at

which trades occur. However, the annual auction for 2.8% of SO₂ emission allowances provides one transparent price signal. In addition, there are market indices maintained by various entities that are publicly available and are monitored by the EPA.

A third, different oversight role is provided by the Federal Energy Regulatory Commission's (FERC) Oversight Division. The division follows all the markets providing components that contribute to the delivery of electricity services and activities in those markets, including allowance markets. It also monitors the behavior of parties in the electricity markets and various input markets, including emissions markets, on a daily basis. The FERC will have an immediate and ongoing interest in the performance of the RGGI allowance market.

Finally, the three state ISOs serving the electricity industry in the region have a natural interest in monitoring the performance of markets. These organizations monitor and report on the competitive structure, performance, and economic efficiency of the markets, as well as the conduct of market parties, including any attempt to exercise market power or restrict competition.

5.4.1 Guidance on Monitoring in RGGI

The most obvious step that RGGI should take is to coordinate with interested federal and state agencies in the design of monitoring criteria and sharing of information. Secondly, RGGI may want to take one step beyond that required at the federal level by requiring that the authorized account representatives should be obliged to disclose the "beneficial ownership" of any allowance holdings. That is, every participant would have to disclose the party sponsoring or benefiting from the agent's activities in the allowance market if it was other than themselves or their immediate employer. Currently, this is not required of the EPA's ATS. At the least, if there is a suspicion that an auction could exacerbate the possibility of behavior that is inconsistent with program goals, then a requirement of financial prequalification for participation in the auction should include the disclosure of beneficial ownership of any allowance acquisitions. Revelation of this information would allow for the market monitor and interested third parties, including government agencies, to identify allowance holdings that appear in excess of compliance obligations and to use this information as a potential trigger for further investigations.

Third, RGGI might accomplish a lot by simply articulating goals in a Statement of Intent and asking all participants in the auction to acknowledge that statement and agree not to undermine these goals. The goals that might be articulated range from overall

environmental integrity to specific behavior in the allowance market. Fourth, RGGI could establish guidelines that require, or request, allowance holders to report their holdings on an annual basis, although it is not clear how one would enforce such a rule without additional requirements that could have unintended consequences of their own. In evaluating these or other measures, it is important for the auction and for the allowance market that monitoring should strive not to be intrusive, administratively costly, or to require the release of strategic information about normal operating procedures in the electricity market.

Fourth, information from the auction should be, and is likely to be, assimilated into RGGI's allowance tracking system. The identity of winning bidders should be revealed, along with the market-clearing price in the auction. The overall demand (quantity) for allowances in the auction should be revealed, along with the minimum and maximum bids. It may be useful to reveal information on a weighted average basis for quintiles in the auction. This information would serve as a signal if auction participation dips, bringing it to the attention of compliance parties that the auction may represent a favorable opportunity in the future, thereby boosting participation. However, we strongly recommend against the revelation of specific offer prices by individual participants in the auction.

The literature has provided some evidence of tacit collusion under the uniform price format in highly stylized (symmetric) settings. The level of coordination required to sustain such collusive outcomes is unlikely to occur in the RGGI auction with many asymmetric bidders (in terms of number of permits demanded and their values). For this reason, collusion is unlikely to be a problem in the initial RGGI auctions, but it may become more of a problem as bidders learn about others' demands and bidding behavior. RGGI should monitor auction outcomes and be prepared to make adjustments to the auction design if such collusive behavior becomes evident.

6 Auction Performance in Collusion-Enhancing Environments and How to Limit Collusion

When potential bidders collude to coordinate their bidding, it is done with the intent of lowering the price the colluding bidders pay for the goods purchased. As noted earlier, the lost revenue generally is associated with lower efficiency and also with less accurate price discovery, since the clearing price will be lower than would occur in a

competitive market. In addition, collusion may result in outcomes that will be perceived as unfair by other bidders.

Sealed-bid auctions generally are thought to be more resistant to collusion than are multi-round auctions, where repeated signals of value and demand are available to participants (Fabra 2003; Abbink, Brandts et al. 2006; Goeree, Offerman, and Sloof 2006). One way to reduce the impact of collusion in multi-round auctions is to limit the information provided to bidders during the auction to the minimum information needed for participation in the auction (McAfee and McMillan 1996; Klemperer 2002). It is well-understood that some information, if provided to the bidders during the auction, can (and probably will) be used by bidders to implement tacit collusive strategies (Avery 1998; Binmore and Klemperer 2002; Klemperer 2002; Fabra 2003; Dechenaux and Kovenock 2005; Abbink, Brandts et al. 2006). There is a tradeoff, however, since bidders may be able to glean information about changing market information from others' bids during an auction. In the case of RGGI auctions, however, there will be ongoing secondary markets in which allowances are traded, so this within-auction transmission of information is less critical. These results argue strongly for limiting the amount of information made public during the auction. After the auction, the public disclosure of information should be the minimum necessary to ensure the fairness and transparency of the auction. For example, it is not necessary or advisable to make public the actual bids of winning bidders. This information might be used by bidders to coordinate bids in later auctions. The fair operation of the auction will be guaranteed by normal audit procedures.

Another way to improve the performance of multi-round auctions in the presence of possible collusion is to combine them with sealed-bid auctions. Both of these strategies are used in the shot-clock design tested. The bidders know only the current price, not the size of excess demand and not the amounts bid by others. The shootout round provides each colluding party with a opportunity and incentive to renege on collusive agreements and earn extra profits without the other parties to the collusion having a chance to retaliate.

6.1 *Effects of a Loose Cap in Uniform-Price and Discriminatory-Price Auctions*

6.1.1 Motivation

There has been some discussion of the possibility that the cap on CO₂ emissions might be “loose” in the early years of the RGGI program. The actual tightness of the cap, however, is difficult to predict, since there may be some speculative demand in early years

in anticipation of higher allowance prices in the future as the cap is tightened or as economic growth increases demand for electricity. However, extensive energy-conservation measures and economic responses to allowance costs (e.g., fuel switching) might have the opposite effects. The initial Phase 1 experiments were done with a relatively tight cap, so we decided to run some sessions in which the number of allowances to be auctioned was about 10% below the number that would be demanded at a zero price as opposed to the 30% reduction assumed in the standard cases.¹¹ The motivation is that a loose cap creates a less competitive environment, since most bids would be accepted.

In a loose-cap environment, bidders may be more likely to collude tacitly by reducing their bids in a series of auctions, whether these are discriminatory- or uniform-price auctions. Therefore, this experiment provides a type of “stress test” of auction formats in an environment where tacit collusion may develop, which is one of the main objectives of the Phase 2 research. In a uniform-price auction, there even may be a role for the exercise of unilateral market power if the cap is so loose that a single bidder can profit from bidding lower on allowances for marginally profitable capacity units in the hopes of lowering the clearing price on other, more profitable units; this is known in the literature as “demand reduction.” Likewise in a clock auction, bidders might withhold their demand on marginal units to prevent the clearing price from rising and thereby increase the earnings on other units. In a discriminatory-price auction, there is less risk with a loose cap since a higher proportion of bids will be accepted, and the resulting bid reductions may cause revenue to be lower in a multi-unit discriminatory-price auction than in a multi-unit uniform-price auction.¹² This revenue comparison might be reversed if there is demand withholding in a uniform-price auction.

¹¹ These numbers are scaled to the laboratory environment and should not be taken to be the values that actually would occur in the RGGI program. The actual tightness or looseness of the cap in the laboratory also should be evaluated in comparison with the allowance needs of a typical bidder.

¹² Vernon Smith (1967) found that seller revenues were higher in a uniform-price auction than in discriminatory-price auctions in settings with moderate numbers of rejected bids and that this difference is not apparent with high numbers of rejected bids. These were multi-unit auctions, but they are not directly relevant since bidders only were allowed to submit two bids, whereas the RGGI implementation would allow bidders to submit any number of bids for blocks of allowances. Moreover, about a third (8 of 26) of the bids would be rejected even in the treatment with the fewest number of rejected bids, so this is not a “loose cap” of the type that might be observed in the RGGI auctions. The Smith experiment was done in an environment that was motivated by the Treasury Bill auctions, in which the prize values to bidders were identical for all units and were randomly determined (i.e., a random common value).

6.1.2 Procedures

For these experiments, the group size was reduced from 12 (used in all Phase 1 sessions) to 6 bidders. There were six sessions for each of the auction formats: uniform price, discriminatory price, and clock. New sequences of random cost draws were constructed for each of the six “waves,” so that the same set of cost draws was used in the first group of uniform-price, clock, and discriminatory-price auction sessions, a second set of draws was used for the second group of auction sessions, and so on.

6.1.3 Aggregate Results

Revenues and efficiencies are measured as before, as percentages of the maximum levels. Efficiency measures the extent to which the surplus value (price minus cost) is maximized subject to the constraint on total emissions. Efficiency is 100% if the reduction in emissions by five units is accomplished in the least cost manner; that is, the manner that minimizes the economic effect of the cap on emissions. Efficiency in this sense does not require that production be cut back for high users only, since their costs generally are lower. High users required twice as many permits per capacity unit as low users; it is only better to keep these high-user plants in service if the costs are less than half of the costs of the capacity of the low-user plants that would be taken out of production.


Revenue is measured relative to an (unrealistic) maximum in which bidders in a discriminatory-price auction bid full value and, hence, they earn nothing at all, so 100% revenue would indicate zero profits for the bidders.¹³ A more realistic benchmark is obtained by calculating the Walrasian revenue that results if bidders in a uniform-price auction bid full value for all units, since the bids then reveal the demand curve. The resulting cutoff price is determined by the intersection of demand and a supply curve that is vertical (above the reserve price) at the number of units being auctioned. This cutoff is a market-clearing price, and the actual revenue in the auction can be measured as a percentage of the revenue that would result if bidders were to pay the competitive market-clearing price.

¹³ Note, however, that if one auction format were to yield 10% more revenue than another, this would not mean that the high-revenue format would reduce firms’ profits by 10%, since the output price in these experiments is fixed and exogenous. The demand for electricity is relatively inelastic, especially in the short run, so an increase in permit prices caused by a switch in auction formats likely would be passed through to a large extent to consumers.

Table 6.1. Summary Performance Measures for Sessions With a Loose Cap

	Revenue	Efficiency
Uniform Price (6 sessions)	53.5%	98.5%
Clock (6 sessions)	58%	98%
Discriminatory (6 sessions)	64%	99%

The average revenues and efficiencies across all auctions for all of the sessions of each of the three auction types are shown in Table 6.1. The loose-cap environment appears to have no significant effect on efficiency for any of the auction types. The discriminatory-price auction format has a significant revenue advantage, but the difference goes away by auction eight (Figure 6.1).

**Figure 6.1. Average Prices Paid by Auction Format**

The prices paid, averaged over all six sessions in each treatment, are plotted in Figure 6.1, along with the Walrasian price predictions, again averaged over all six waves. This figure shows that prices in discriminatory-price auctions are higher than in the uniform-price and clock auctions, as would be expected from the revenue comparisons

discussed previously. But note that this price difference diminishes steadily and is eliminated by the final auction. The more aggressive bidding in the discriminatory-price auction diminishes over time as bidders come to collude tacitly. In all three formats, prices are considerably below the Walrasian predictions in this loose-cap setting.

6.2 Collusive Environment with Standard Caps

6.2.1 Motivation

A second way to “stress test” an auction design is to introduce the possibility of explicit discussions among bidders prior to the auction. To test how the two sealed-bid auction formats perform in the presence of this type of explicit collusion, we altered several design features to create an environment that was more favorable to collusion than was the case in the Phase 1 experiments. Other studies have shown that group size has a large effect on the ability of participants to exercise market power and keep prices down (or up, as in the case of supply side auctions) (e.g., Bernard et. al. 1998). Therefore, in our collusion treatments, as in the loose-cap experiments, we reduced the group size by half, using 6 instead of the standard 12 participants. In addition, the participants’ costs (and therefore permit values) remained constant from one auction to the next, instead of changing randomly within a certain interval. The idea here was that it would be easier for participants to coordinate bidding strategies when they face the exact same environment from auction to auction. The auctions in these sessions were followed by secondary (spot) markets, so that bidders would be able to obtain needed permits if a collusive effort to bid low in the auction failed as a result of a defection by other bidders from an agreement. Lastly, in half of our collusion treatments we allowed participants to chat with other participants by using instant messaging for one minute prior to each auction. This gave groups an even stronger potential to collude by giving them the opportunity to discuss strategies and make non-binding agreements. We recorded the transcripts from these chat sessions and used them in our data analysis.

6.2.2 Procedures

We ran discriminatory-price and uniform-price auction treatments. Half of the treatments allowed for explicit collusion via an instant messaging “chat” window, and half of the treatments did not allow communication among group members. New sequences of random cost draws were constructed for four “waves,” so that we ran four sessions of both discriminatory-price and uniform-price auctions with communication and four sessions of

both discriminatory-price and uniform-price auctions without communication, for a total of 16 sessions.

6.2.3 Aggregate Results

The results offer strong evidence of tacit collusion under both the discriminatory-price and uniform-price auction types in this environment. One of the ways to ascertain how well participants colluded is to look at the difference between the average price paid for each permit and the Walrasian prediction (“supply equals demand”). In all of the treatments, the average accepted bid in the discriminatory-price auctions and the price in the uniform-price auctions remained below the Walrasian prediction. This collusion appears to be somewhat more successful at lowering price in the discriminatory-price auction than the uniform-price auction but only slightly so (see Panel A of Figure 6.2).




Figure 6.2. Illustration of Tacit Collusion

Allowing bidders to communicate before the auction can facilitate collusion under the right group dynamics. While there were some groups in the chat treatments that were not able to coordinate a strategy to collectively lower prices, there were two groups in particular that colluded more effectively (lowered price to a greater extent) than was achieved with tacit collusion alone.

One group of bidders in the discriminatory-price auction collusion treatment initially agreed that everyone would bid the reserve price of \$2. But when it became clear after several auctions that others were defecting from this agreement, bidders agreed to bid \$2 on just three-fifths of their permits. While the plan was not followed perfectly by

everyone, enough low bids were made that the average accepted bid remained quite low. In fact, for two of the auctions some \$2 bids were accepted.

One group of participants in the uniform-price collusion treatment also colluded quite well. While no explicit strategy was mentioned other than that everyone should bid low, there was a congenial exchange of messages with several jokes, and the majority of players bid less aggressively throughout the session.

On average we did not find a large revenue or efficiency effect of allowing communication prior to the auctions, as can be seen from the averages in Table 6.2. However, it is clear that communication has the potential to facilitate collusion, as it did in the two sessions mentioned above. The second column of Table 6.2 shows that the two sessions with successful explicit collusion had substantially lower revenue (48% and 54%) than other sessions. These results suggest that group dynamics matter more for the effectiveness of explicit collusion than does the type of sealed-bid auction (discriminatory price or uniform price) being used. Collusion does not appear to affect the efficiency of the auction substantially under either auction type.

Table 6.2. Summary Performance Measures for Sessions With and Without Chat

	Revenue Percentages		Efficiency Percentages	
	Session	Overall	Session	Overall
Uniform Chat	71, 48, 64, 65	62	97, 92, 99, 96	96
Discriminatory Chat	64, 68, 67, 54	63	96, 93, 96, 97	96
Uniform No Chat	77, 68, 60, 70	69	95, 98, 94, 97	96
Discriminatory No Chat	63, 72, 63, 63	65	97, 97, 95, 98	97

Allowing participants to communicate prior to the auctions did result in more bids at the reserve price of \$2 than in comparison sessions with no communication. The reserve price was a focal point for many of the proposed bidding strategies discussed by the participants in these experiments. Figure 6.3 shows the relationships between bids and allowance values for the two types of sealed-bid auctions both with and without explicit communication. In these graphs, darker dots indicate that more bids are observed at those points. In both the uniform-price and the discriminatory-price sessions with chat, we see

that the actual bids were more condensed between the values of \$2 and \$3 than occurred in auctions of similar format without the ability to “chat” with others.




Figure 6.3. Bidder Behavior in Uniform-Price and Discriminatory-Price Treatments Both With and Without Communication Opportunities

6.3 Discussion and Extensions

The results of these experiments show that revenues are likely to fall below the Walrasian benchmark with a less competitive market (with fewer bidders, a loose cap, or with explicit collusion), which is quite different from the results of the Phase 1 experiments (with a tighter cap and no chat opportunities), where the more competitive environment resulted in revenues that typically were close to the Walrasian predictions. In the loose-cap setting with uniform-price and clock auctions, the drop in revenue was due in part to demand withholding. Similarly, bidders in the discriminatory-price auction sessions came to realize that low bids involved little risk, especially as bidders reduced bids together in a type of tacit collusion. There is some evidence that bid reductions in the relatively noncompetitive loose-cap environment are more prevalent in uniform-price and clock auctions, although the differences between the discriminatory-price and the other two

auction types were reduced over the course of several auctions conducted in each session. In the sessions with a collusion-enhancing environment and a standard cap, the results indicate that collusion is equally prevalent in the uniform-price and discriminatory-price auction types. In both cases, revenue reductions did not adversely affect efficiency, and any efficiency differences probably would be mitigated by the presence of a spot market.

Finally, we note the possibility of ties among bids. If there is a tie at the auction clearing price level, bids above the auction clearing level should be met first, and secondly the remaining allowances should be divided among the tie bids. Moreover, how the tie is resolved may pose an incentive for collusive bidding. This may be especially true if the equilibrium price in the auction is in the vicinity of the reserve price, in which case the reserve price may provide a focal point for collusive bidding. If ties were resolved through a random assignment of allowances *among bidders* with tied bids rather than through proportional division then there would be an element of risk associated with using the reserve price as a focal point for collusive bidding. Hence, a random assignment has the advantage in that it would lessen the incentive for collusion and help to achieve other auction design criteria. A tie-breaking procedure is only needed if the number of allowances bid for exactly the market clearing price is less than the number available at that price (after allocations to bidders submitting higher bids have been made). *The random tie-breaking procedure should be implemented by assigning a random priority number to each bidder, not to each bid.*

7 Reserve Prices and Price Volatility

7.1 Price Volatility

Because the supply of CO₂ allowances in the RGGI region is fixed, the price of allowances will be more volatile than would the price of a good for which the supply could respond to changes in price. The low short-run elasticity of demand for fossil fuels and electricity adds to the volatility of CO₂-allowance prices. A main reason for the focus on price volatility is that high price volatility is known to dampen the incentives for investment by increasing risk. Because one of the primary justifications for placing a price on CO₂ is to induce investment in non-carbon alternatives, the issue of price volatility merges with concerns over the cost and effectiveness of carbon-reduction policies.¹⁴

¹⁴ Strategies for controlling price volatility are the subject of considerable academic interest (Nordhaus 2007).

Banking in RGGI is one mechanism that helps prevent the price from going over some upper limit, at least in the short run, by acting to smooth expected price variation. The availability of offsets also constrains the price of allowances. The effect of the offset rules is somewhat less clear since it is not known whether offset prices will be higher or lower than the prices in RGGI that trigger the ability to draw on a larger pool of offsets for compliance. There is uncertainty about how allowance prices would evolve in the event that allowance prices are high enough to reach the first offset trigger price but offset prices are far below the trigger price.

Even with these cost-management mechanisms in place, most informed parties generally would agree that reducing investment risk by reducing allowance-price volatility is beneficial for both the carbon market and the electricity market. The research team has been encouraged to give weight to auction features that reduce price volatility.

As we will discuss in detail later in this report, the auction itself should contribute little, if any, to the volatility of prices. Once the allowance market settles into routine operation, we expect the price of allowances at auction to be very close to the price of allowances in the spot market, as discussed in section 8 below. The residual risk in the price of allowances will not be due to the auctioning of allowances but will arise due to shifts in allowance demand caused by changes in external conditions or changes in demand for electricity.¹⁵

In the next section, we will discuss some features of auction design that may be useful in controlling price volatility.

7.2 Reserve Prices

A reserve price is an auction price below which the seller chooses to retain ownership of the item rather than sell it. The most obvious use of a reserve price is to prevent the item from being sold at a price below the seller's opportunity cost. They are especially important where the bidders have very asymmetric willingness to pay for an item or asymmetric information or when participation in the auction is low. Under these conditions, the clearing price may be far below the competitive price unless a reserve is used, and the efficiency of the auction could suffer.

¹⁵ This is true whether the allowances are auctioned or grandfathered. Ultimately, that choice will not affect the price volatility of carbon allowances.

Reserve prices also are very important in reducing the potential damage from collusion because they reduce the profitability of collusion. This is true whether the collusion is tacit or explicit. The importance of reserve prices in limiting collusion is strongly supported in the theoretical literature and in empirical examinations of auction performance (McMillan 1994; Levin and Smith 1996; Binmore and Klemperer 2002; Ausubel and Cramton 2004). The academic literature and numerous notorious examples of failed auctions point to a credible and efficient reserve price as one of the most important aspects of auction design. For example, for the upcoming 700MHz spectrum auction, the FCC has set reserve prices that total more than \$10 billion.

We conclude that the possibility of collusion and the possibility of weak competition among asymmetric bidders make a strong case for establishing a reserve price and committing to a policy that any allowances for which offers do not meet the level of the reserve price not be sold in the current auction. If the reserve price is triggered the reserve price becomes the auction clearing price, and bids at or above that level are accepted.

A reserve price may be publicly announced in advance of the auction, it may be revealed once bidding reaches the reserve level, or it may remain undisclosed by the auctioneer until after the auction is over. Undisclosed reserve prices commonly are used in the standard English clock auction for art, wine, and other valuable commodities (Ashenfelter 1989). They also have been used in the sale of publicly owned assets. Since one of the reasons for having a reserve price is to protect against cases where weak competition may lead to low prices, there are cases where the reserve price is set to a higher value when relatively few bidders are present and lower if the auction is relatively more competitive (Hendricks, Porter et al. 1989).

In regularly repeated auctions by a government, it may be difficult to prevent bidders from learning the reserve price. An agency must have a rule for setting the reserve. Over time, smart bidders will be able to infer the rule for the setting of the reserve and will be able to bid on the basis of this information. Even if the reserve is set with randomness, over time it likely will be possible to infer the distribution of the randomizing method.¹⁶ Thus, many bidders will have a reasonable estimate of the reserve price even though the government agency is operating on the assumption that the reserve price is not known. We conclude that it is not a good strategy to have an undisclosed reserve price, since it cannot

¹⁶ Recent empirical evidence suggests that random reserve prices tend to lower auction revenues without accomplishing any identifiable efficiency objective (Hendricks, Porter et al. 1989).

be assumed that the strategy for selecting the reserve will not become known to the bidders. It is better design to assume that bidders will be able to obtain any information reasonably available and hence make the reserve price public from the outset. In fact, in the English-clock auction, the price used for the first round of the clock generally serves as the reserve price and so is always revealed at the start of the auction. One exception may be the first auction, since bidders would not be able to infer the reserve price from previous auctions. In addition, in the absence of information about the value of allowances from previous auctions and before a robust secondary market emerges, the reserve price might serve as a focal point for bids. Therefore in the first auction there may be some justification for not announcing the reserve price in advance.

How the reserve price is set in the auction interacts with other aspects of program design. One method of setting the reserve price for a RGGI auction would be to set it a level close to but below the expected clearing price for the auction, which is likely to be very close to the current price for allowances in the secondary market. An analysis could be done to produce an estimate of the current market price of allowances, and the reserve price can be set below that amount. For example, a simple rule of thumb for setting the expected price would be to take the average of the spot-market price index in the month preceding the auction and to set the reserve far enough below that amount to account for reasonable short-term variation in prices. If the allowance price were to run up due to an external event such as severe weather or fuel supply disruption, then this strategy would not allow the auction price to lead the market as prices retreat to lower levels over time. For an auction with limited competitiveness, this reserve price may serve as a focal point for coordinated bids and could result in some level of tacit collusion, but the benefits of collusion would be low because of the closeness of the reserve to the market price.

An alternative method for setting the reserve price would be to set it to a level that would maintain a minimum rate of progress in reducing emissions below business as usual, and to maintain the value of investments in new technologies. In this case the reserve price would not be directly linked to market prices, but instead would grow at a constant rate such the rate of interest. Both of these approaches would serve as a short-run way to prevent disruption of the auction due to collusion or weak competition.

In order to maintain the integrity of the reserve price, the allowances may be retired, may be rolled forward to the next auction, or may be placed in a contingency bank.

Retirement: A number of parties to the RGGI auction design discussions have suggested that a low price implies a cap inadvertently set higher than intended when the

original calculations were done to establish the level of the cap. These parties argue that any allowances not meeting the reserve should be retired. The logic of retiring unsold allowances might be applied to a reserve price based on the minimum value to RGGI of a ton of carbon reduced, but it does not apply to a reserve that is set merely to protect against collusion and weak competition.

Rolling forward: Rolling any unsold allowances forward into the next auction is an administratively simple rule, but it may result in an awkward start to the RGGI market if the cap is as slack as some commentators have suggested. In the auction subsequent to the one in which the reserve price was triggered, the new reserve price would be lowered according to the kind of rules we suggest above. Nonetheless, the difficulty is that taking a large block of allowances forward to the next auction, which may itself face relatively slack demand, would result in a large overhang of unsold allowances and a series of auctions with significant proportions of allowances being unsold. This could lead to an impression that the auctions were failing, even though the results primarily were due to a relatively slack cap.

Contingency bank: Allowances rolled into a contingency bank would be released for sale at the next RGGI auction once any RGGI allowance auction closed at a price above the first offset trigger price. This strategy has some distinct advantages over rolling allowances forward. First, it solves the problem of increasing overhang in a slack market. Second, and possibly more importantly, it provides a mechanism for reducing price volatility, which produces gains for all firms using allowances. The offset trigger requires a high price be maintained for RGGI allowances for a year before offsets become available. During that time, prices might continue to escalate. A compliance bank acts much as other banked allowances but is used specifically to blunt rapid spikes in price. Once an auction closes above the offset trigger price, the banked allowances would be released into the subsequent auction, increasing the available supply and reducing the clearing price. In fact, the anticipation that allowance prices might top the offset price at an auction will tend to moderate bid prices since the bidders know that the extra supply is available. Such a strategy could, if a reasonably large bank were to develop, help prevent price bubbles and panics that otherwise might cause considerable concern in electricity markets. It also lowers the profitability of efforts to manipulate markets to raise prices.

This issue is important in the first years of the RGGI allowance market, when the quantity of allowances initially distributed into the market is relatively close to baseline emissions levels. If the market is fully developed, then in the first years of the program the ability to bank allowances protects against the value of allowances falling to zero. This is

because investors and speculators would be willing to buy the allowances during periods of depressed prices and hold the allowances until the initial distribution of allowances declines and prices rise. Hence, in equilibrium, the price in one period should be related to the price in a subsequent period by the opportunity cost of capital. However, at the outset of the program the market may not be in equilibrium as market participants still are learning about how the market will function. If the market is not fully developed in this way, the presence of a reserve price helps to provide stability and provides assurance to those entities that are making efforts to reduce emissions that their emissions reductions have financial value.

The process of rolling allowances from periods of very low prices to periods of extremely high prices would tend to reduce volatility in allowance prices, thereby reducing price risk to generators and their customers.

8 Price Discovery

An efficient market organization maximizes the difference between the value of electricity produced and the cost of production, subject to the cap on emissions imposed by RGGI officials. The (Walrasian) market price of an allowance provides an important signal that allocates electricity production efficiently across independent producers, who may be concerned only with their own earnings. In equilibrium, the production for each emitter is scaled back enough so that the cap is met in the aggregate, and the marginal social cost of additional electricity production (including the implicit cost of emission) is equalized across different producers.

The price of an allowance provides a signal of its scarcity value; that is, a signal of the direct economic cost of a reduction in emissions. Consequently, it is important that the auction provide reasonably accurate price signals. Accurate price discovery in an auction can help establish a market price close to the marginal cost of control. Once the market has reached this equilibrium, then the spot market will provide a continuous summary of current opinions about the current value of allowances and, hence, the current scarcity value of allowances or marginal cost of reducing emissions, for example via fuel switching. This price will adjust daily as expectations change concerning fuel prices, electricity demand, and other factors. As the experience with the SO₂ market has shown, a sealed-bid auction of allowances already traded in a secondary market will closely track the prices in the secondary market (Ellerman, Joskow et al. 2000).

The SO₂ experience and Virginia's NO_x auction experience also clearly demonstrate that an auction need not disrupt the spot market price signal even if the number of allowances sold at auction is much greater than the quantities traded in the spot market on a daily or weekly basis (Ellerman, Joskow et al. 2000). The number of allowances traded on the spot market over a period of days or weeks is not a measure of the "liquidity" of the market. In fact, the opposite is more likely true since a periodic injection of allowances into the market through a sequence of regularly scheduled auctions can lower the perceived risk of illiquid markets, reducing overall price volatility.

When used to sell multiple items, uniform-price auctions give bidders an incentive to shade their bids a bit. This is because a reduction in the bid below a level that actually reflects the bidder's value may result in the auction clearing at a lower price. So, while the bidder may not get as many units as might have been warranted at the clearing price, the bidder will save some money on all of the units that are won at auction. Discriminatory-price auctions are not subject to this incentive to shade bids since lowering the bid on one unit does not affect the price paid on other units. The incentive to shade bids in multi-unit, uniform-price auctions, known as demand reduction, must be weighed against the superior price discovery properties (Ausubel and Cramton 1998; List and Lucking-Reiley 2000). Although discriminatory-price auctions are not subject to this demand-reduction incentive to shade bids, it is worth noting that bids in a discriminatory-price auction typically are well below value, and variations in this strategic bid reduction from one bidder to another also may create inefficiencies.

8.1 Price Discovery: Unanticipated Shift in the Demand for Permits

In order to provide correct price signals concerning the market valuation of permits, well-functioning markets should aggregate information that is dispersed among the participants. Our strategy for the "price discovery" series of experiments was to implement an unanticipated demand shift—that is, an increase in permit values—due to production cost reductions for some bidders but not for others.

The experiment involved equal numbers of high emitters (needing two permits for each unit of production capacity) and low emitters (needing only one permit per unit of capacity). Each session began with a series of three auctions in a baseline condition with 82 permits offered for sale in each auction. After auction three, the costs of low emitters were reduced in a manner that raised the predicted Walrasian permit price from about \$3.50 to about \$5.75, as shown by the dashed line in each panel of Figure 8.1. The low

emitters, who began with high production costs, had some knowledge of the change in market conditions prior to auction four in the sense that each low emitter observed a reduction in their own costs. High emitters in these experiments had relatively low production costs that stayed the same, on average, for all six auctions, so they had no inkling of a demand shift prior to seeing results.




Figure 8.1. Average Price Paid by Auction Type with an Unanticipated Demand Shift

We ran four sessions for each of the three main auction formats with bids submitted once per round: the single-round uniform-price and discriminatory-price auctions and the multi-round clock auction. The prices paid, averaged over all four sessions for each format, are graphed on the left side of Figure 8.1. Overall, the average price deviations are lowest for the uniform-price and clock auctions, with no meaningful difference between the two. It is not surprising that a multi-round auction like the clock auction will pick up the demand shift, since demand is revealed as the clock price is raised. The observation that final clock prices always are below the Walrasian predictions probably is due to tacit collusion, as bidders realize that if they reduce demand they may stop the clock, lowering the prices for the permits that they do purchase, and possibly signaling cooperation that may affect bidding in later auctions.

The single-round, uniform-price auction also tracked the demand shift, although there was again a downward bias in prices relative to the Walrasian predictions, as subjects tended to bid low on some units in an effort to reduce the clearing price (highest rejected bid). Prices in the uniform-price auctions tended to track the demand shift because people

were bidding near value on some of their permits (those with high use values), which is a profitable strategy whether or not others' values have changed.

The discriminatory-price (sealed-bid, single-round) format also had low average deviations but did not pick up the demand shift. The average prices paid for permits by winning bidders are shown on the left side of Figure 8.1 (the thick gray line). Note that these average prices were biased upward in the first three auctions and downward after the demand shift. The upward initial bias is consistent with a tendency (observed in the loose-cap experiment) for auction revenues to be relatively high early in a sequence of discriminatory auctions, but this difference tends to diminish over time. This pattern also is revealed in Figure 8.1 by the downward trend in the average accepted bids for discriminatory auctions in the first auctions.

In addition to the three main auction formats, we ran two sessions each for two additional types of auctions. These were continuous-time auctions in which bidders could raise (but not lower) their bids at any point during a pre-announced time interval of five minutes. Bidders could see which of their bids were “provisionally winning;” that is, among the 82 highest bids at that time.¹⁷ The 82 highest bids at the closing time became the official winning bids. In the continuous discriminatory-price auction, the winning bidders had to pay their own bids for each permit purchased. In contrast, the winning bidders in the continuous-price uniform auction only had to pay the highest rejected bid at the time of the auction close.

The continuous discriminatory-price auction yielded the worst price tracking of any of the five auction types considered, as indicated by the flat gray line on the right side of Figure 8.1. Subjects generally were bidding below their values in the early minutes of these auctions, often near the reserve price level. Some bidders did not even turn in bids in the first three or four minutes. Thus, the remaining bidders would see all of their bids listed as provisionally accepted, even at low bid levels. Then “sniping” in the final 30 seconds of the auction would raise the cutoff prices, and bidders would scramble to leapfrog their bids upward once or twice if they had time. The resulting prices did not increase to the predicted levels, especially after the demand shift.

As with the continuous discriminatory-price auction, bidders in the continuous uniform-price auction could view the status of their bids (provisionally winning or not) and could increase (but not decrease) their bids at any time prior to the end of the auction. The

¹⁷ We decided to consider continuous auctions with a “hard close.”

result of continuous bidding was again a widespread attempt to collude tacitly by bidding at the reserve price on some permits early in the auction, with some bidders not bidding at all until the final seconds. But the uniform-price property allowed the bidders the opportunity to bid aggressively for their most valuable permits in order to ensure some high-value purchases at a price determined by the highest rejected bid. This demand-revelation behavior for high-value units (likely to be purchased) caused the continuous uniform-price format to pick up the *magnitude* of the predicted price increase after the third auction in each sequence, but *levels* of average purchase prices uniformly were too low as a result of signaling and bidding at the reserve price until the final seconds, at which time “sniping” was pervasive.¹⁸

Taken as a group, this demand-shift experiment yields three main conclusions: 1) uniform-price auctions (clock and sealed-bid uniform-price and continuous uniform-price auctions) generate *changes* in purchase prices that are reasonably close to the Walrasian predictions; 2) there is some evidence of tacit collusion causing prices to be too low relative to predictions in most cases, and such tacit collusion is most successful for the multi-round and continuous formats (clock, continuous discriminatory-price and continuous uniform-price auctions) where signaling was possible to some extent, especially in the continuous auctions; and 3) the worst price tracking is for the continuous discriminatory-price auction, where the combined effects of signaling and sniping all but hide the effects of the unanticipated demand shift in auction four of the experiment. Overall, the clock and sealed-bid, uniform-price auctions performed best in this demand-shift environment.

8.2 Relationship Between Auctions and Secondary (Spot) Markets

Most items that are auctioned have value in a secondary resale market (spot market). In general, if the secondary market is not mature and perhaps not well informed about the value of an asset, the auction is expected to contribute to the discovery and realization of the value of the asset. If the auction is well designed and the secondary market is mature, the auction price may be expected to mirror closely the price in the secondary market. What might we expect about the relationship between the auction and the allowance market in RGGI?

¹⁸ One possible solution to the sniping problem is to switch to a “soft close” in which the clock is restarted (e.g., for a minute), as soon as a new bid is received. Such a procedure may create its own problems in terms of the length of the auction.

8.2.1 Experience in Existing Allowance Markets

One can look for evidence in the actual performance of past auctions for emission allowances in the presence of a secondary allowance market. The SO₂ emissions allowance trading program was initiated by Title IV of the 1990 Clean Air Act Amendments. While a majority of allowances are allocated for free to incumbent generators, Title IV specifies that 2.8% of the allowances issued every year should be allocated through a revenue-neutral auction. The proceeds from the auction are returned to industry in proportion to the underlying allocation of the remainder of the allowances.

A primary reason the auction was included in the legislation was to address the concern of independent power producers that new entrants into electricity generation would not have access to allowances if the incumbent generators “hoarded” allowances or for some other reason the secondary market was not liquid (Hausker 1992). In retrospect, however, this concern turned out to be misplaced, as the secondary market has been liquid (if not always exhibiting high volumes). Instead, observers have suggested that the auction performed two valuable if unanticipated functions. One was to prime the pump for trading, so to speak, by forcing a redistribution of some allowances. Second, as a consequence of the pump priming, the auction contributed to price discovery at a time when expectations about compliance costs were varied across the industry (Ellerman, Joskow et al. 2000).

What role will the auction play in revealing willingness to pay for allowances in secondary allowance markets?

An examination of the annual auctions beginning March 1993 shows that the bid schedule of participants was quite steep in 1993, indicating a wide variation in opinions about compliance costs. By 1994, that schedule had flattened out considerably, and by 1995 and beyond, the schedule is almost flat, indicating widespread consensus on the price at which allowances were likely to be sold. The first auction in 1993 achieved a clearing price of \$131 per ton, substantially below previous estimates of compliance costs and the prices of bilateral trades that had been reported in the trade press. In 1994, the spot market-clearing price of \$150 was still 10% lower than the prevailing cost of bilateral transactions. Both these results contributed to a short-term criticism that the auction was not properly reflecting the value of emission allowances. However, by August 1994, the prices reported by the three brokerage firms for allowances traded in the spot market were almost identical to the level established by the 1994 auction. In retrospect, it appears that the auction contributed importantly to price discovery and set the context for an active secondary market (Ellerman et al. 2000, 178–180).

How does the auction price compare with outcomes in the secondary or reserve markets?

By 1995, the secondary market had matured considerably. Figure 8.2 illustrates the pattern of prices in each auction since 1995, along with the spot-market price approximately one month prior to and one month after the auction. In virtually every year, the auction price has been nearly coincident with the spot-market prices in the surrounding months, or it has been in line with a trend in prices. This evidence suggests that the allowance auction has not disrupted price-setting behavior in the spot market and, furthermore, that the auction reflects willingness to pay in a similar manner as does the spot market. Of course, the auction is for a small portion of all allowances, but it is relatively large compared to allowance trading activity in the spot market because most allowances are allocated directly to the firms that use them.




Figure 8.2. SO₂ Auction and Trading Prices

Note: The price for April 2007 was not available at the time of this writing, so March data are used. Market data source: Cantor. "SO₂ Allowance Price Indications: Historic Monthly Bulletins." http://www.noxmarket.com/Environment/?page=USAComp_MarketData-BulletinsHistoric (accessed May 7, 2007). Auction data source: Clean Air Markets. "Annual Auction." EPA. <http://www.epa.gov/airmarkets/trading/auction.html> (accessed May 7, 2007).

In general, what is the effect of the auction on secondary allowance markets, and vice versa?

Largely similar results were achieved with an English-clock auction used to sell 5% of 2004 and 2005 vintage NO_x allowances under Virginia's SIP Call NO_x budget. The auction of 3,710 NO_x allowances was held in June of 2004. Even though the amount of allowances sold was more than 30 times greater than the daily number of trades then occurring in the spot market, the clearing prices for the auction were 5% to 7% higher than the spot-market prices just before the auction. The increase over the recent spot-market price may have involved a small measure of good luck by catching the market before a period of higher prices. The price of NO_x allowances did trend somewhat higher for the months after the auction. However, the sale of a large block of allowances reflected the valuation in the spot market and did not cause a price reduction, as had been forecast by traders in the days before the auction.

8.2.2 Evidence From Experiments

We combined an auction with a spot market in numerous experiments. A consistent relationship emerged that is evident from the experiments and illustrated in the following figure. The first panel of Figure 8.3 compares the auction and spot price for a uniform-price auction format. Over a series of eight sessions, the spot price and the average auction price are close. Both typically are slightly below the Walrasian (competitive market) price, which could reflect a small degree of tacit collusion. The second panel illustrates a similar result for the discriminatory-price auction. In this case, the average accepted bid is reported because bidders pay different prices in this auction format. Again, the average accepted bid is very close to the spot price and to the Walrasian price.




Figure 8.3. Auction Results and Spot Market Prices in Laboratory Experiments

In summary, the empirical evidence from previous allowance markets, including especially the evidence from the SO₂ market, indicates that the auction price and spot-market price track together closely. In that market, the auction price actually played a leading role in identifying the equilibrium value of emission allowances and the spot-market price moved to that value as the market matured. Experimental results corroborate this finding. In the laboratory, the auction and spot-market prices are close and show no regular pattern in the way that they differ. The bottom line is that participants in RGGI should not worry about a potential difference between the auction and spot-market price, and, furthermore, one might expect the auction price to be a leading indicator of equilibrium in the spot market.

9 Hoarding of Allowances

RGGI allowances are in fixed supply. An increase in demand for RGGI allowances will raise their price but will not induce an increased supply, as would be the case in most markets, such as the market for manufactured goods. In addition, RGGI allowances are essential inputs to the production of electricity in the RGGI region. The demand for electricity is known to be inelastic in the short run, so increases in price do not cause large changes in demand in the short run. These characteristic of RGGI allowances have generated discussion about the possibility of large increases in the allowance prices due to increased demand. We already have discussed some aspects of this issue in our discussion of price volatility (section 7.1), where we were implicitly assuming that increases in demand were generated from normal market forces related to the electricity industry in the RGGI region. However, participants in discussions of RGGI auction design have expressed

concern that demand for RGGI allowances may be artificially inflated to levels above those justified by normal market factors within RGGI. Factors that could result in an artificial increase in demand generally have been referred to as “hoarding” of allowances.

The concept of hoarding rarely is made precise. We interpret this to describe the concern that an entity *may obtain and hold allowances in excess of the entity’s own anticipated compliance obligations*. Five possible types of hoarding behavior have been mentioned in RGGI auction discussions:

- Speculation: buy allowances in anticipation of their potential future increase in value so allowances may be sold at a profit.
- Allowance market manipulation: buy allowances not only in anticipation of their potential future increase in value but to directly induce scarcity and raise the price so allowances may be sold at a profit.
- Electricity market interference: buy allowances with the intent of disrupting the electricity market within the RGGI region.
- Competitive advantage (a.k.a. raising rivals’ costs): non-emitting generators buying allowances to raise the cost of generation by emitting generators.
- External compliance: use RGGI allowances as offsets to satisfy voluntary or mandatory CO₂ reductions by sources outside of RGGI.

For the most part, these are not issues of auction design. Rather, these issues arise as a consequence of the structure of the RGGI market. Auctions might contribute to hoarding if somehow auctions made it substantially easier for hoarders to obtain RGGI allowances than would otherwise be possible. However, if there is a liquid allowance market, as most anticipate, then auctions do not provide an opportunity that would not already exist in the allowance market. Our recommendation for a schedule of quarterly auctions, including the advance sale of future vintages, is meant partly to reduce the ability of bidders to “surprise” the allowance market with large, sudden, and unexpected jumps in demand. The schedule of auctions we recommend would result in no more than 12.5% of first-time allowances of a given vintage being sold at a single auction.

RGGI further may wish to follow the procedure used in U.S. Treasury auctions and limit single entities to no more than 33% of the allowances for sale at a given auction.¹⁹

¹⁹ This rule is, like others, cannot be effective against those willing to engage in illegal or unethical behavior. However, it is a relatively low-cost tool for enhancing the competitiveness of the auction.

(Potential buyers possibly could request special permission from RGGI to exceed this limit with sufficient justification.) While unlikely to constrain compliance buyers, this additional constraint would limit a single entity from purchasing more than about 4% of a vintage at a single auction. Rules such as this 33% rule do require some additional enforcement efforts, but these are largely consistent with our recommendations on market monitoring, and they should not be administratively cumbersome. This recommendation would be made more practical if coupled with another recommendation that required allowance account representatives (or at a minimum, parties seeking financial pre-approval for participation in an auction) to disclose the beneficial party for whom they are acquiring allowances if it is other than themselves or their immediate employer. Otherwise, firms could enlist agents and brokers to exceed this purchasing limit. This proprietary information would need to be kept confidential.

Disclosure of beneficial interest as suggested above is an important element. There is one well-known case where a large buyer of U.S. Treasury bonds violated the restrictions by unauthorized use of other entities for purchase with the intent of driving up prices in the market. This resulted in an enforcement action against a large U.S. brokerage firm.²⁰

The uniform-price auction form also provides some protection against disruption of the electricity market, although not against high prices due to increased demand. This protection is due to the uniform-price rule where bidders do not pay their bid but rather the value of the highest rejected bid. If a firm faces a very high cost of not getting some minimum supply of allowances, then it can bid high on non-marginal units and greatly increase the probability of winning on those bids. But the amount paid still will depend on the marginal bid. This is a way of ensuring access to high-value units. Bidders wishing to disrupt the market by buying up available allowances would find it very expensive to purchase these high-value allowances.

Once the use of the auction for surprising the market with a large, sudden spike in demand is restricted, there is little remaining difference between the auction and the spot market with respect to facilitating hoarding. This conclusion is strongly confirmed by our experimental results. In experiments where both auctions and spot markets were present, auction and spot-market prices tracked each other very closely. Any divergences were small and temporary. The possibility of a buyer accumulating allowances for non-

²⁰ See Fuerbringer (1991).

compliance reasons is a property of markets generally. The reason that this issue is of special concern to RGGI is not because of auctioning but rather because the supply of allowances and the demand for electricity both are highly inelastic.

The presence of two price triggers that would allow the use of offsets for compliance shares some similarities with a safety-valve price. The triggers expand the supply of allowances once prices reach unexpectedly high levels. However, the long lead-time between the time prices first rise above trigger levels and the time offsets may be used leaves considerable room for prices to rise well above these levels before offsets become available.

The likely effect of hoarding behavior further is reduced by other design elements of the RGGI market. The ability to bank RGGI allowances is likely to have a very significant effect on hoarding behavior, at least in the first few years. Most observers seem to expect that the RGGI cap will be relatively slack in the first few years of the program. A slack cap implies a relatively low price. At a low price, generators and others will be able to profit by buying allowances now for use in future years. So, as with other emissions-trading programs, traders will anticipate the future tightening of the cap and purchase allowances in the near term, thereby building up a bank. The presence of a substantial privately-held bank of allowances makes it much more difficult to manipulate the market because any attempt to raise the market price will require depleting the bank. In the presence of a large bank, hoarding for market manipulation becomes less profitable and riskier. If RGGI were to choose to implement a contingency bank (as recommended in section 7), this publicly-held bank would serve a complementary function by limiting price increases above the first offset trigger.

The three-year compliance window adds to the risk and expense of any hoarding for market manipulation. Aside from any banked allowances that they might own, generators have three years of auctions and spot-market trading to accumulate the allowances needed for compliance. Buying allowances early in the enforcement cycle will require holding the allowances on the books at the opportunity cost of money for a period of years. Buying late in the cycle is risky because firms already will have purchased their highest value allowances and there may not be much opportunity to drive up prices. Buying and selling within a compliance cycle is likely to be a wash. In other words, with a three-year compliance period, it is hard to squeeze generators by buying allowances. That is not to say that it cannot be done. But it is risky and costly, which drives down the expected net gain of the activity.

9.1 *The Speculation and Market Manipulation Motives*

It is not correct to think of speculators as being able to buy the allowances they want at the current price and thereby drive up prices for others and then turn around and sell the allowances at the new higher price. Simply stated, if buying drives it up, selling drives it down. As a speculator bids in an auction or purchases in the spot market, his purchases will immediately raise the clearing price at auction or the lowest offer to sell in the spot market. For example, if the current spot-market price of allowances is \$1.00, then there is good reason to believe that an auction will clear at a price close to \$1.00 unless something changes. If a firm enters the market with a bid for a large number of allowances at \$1.50, then chances are that it will not get allowances for \$1.00. Rather, the price paid will be between \$1.00 and \$1.50. The clearing price will depend on how many bids are between \$1.00 and \$1.50. If the bid is for a large block, then chances are that the firm will pay close to \$1.50. If demand is slack and there are many bids at or close to \$1.00, then the price will be closer to \$1.00 and the firm will not have succeeded in increasing the price. The reverse logic is true when the firm goes to sell. The firm's participation in the market will drive the price paid to the new equilibrium before the auction clears, not after.

But what if a speculator believes that firms are not buying enough allowances because they have underestimated their need? Then, there is a chance of profiting from the mistakes of others by buying now when prices are low and then selling later when generators have realized their mistake. This effort to profit from the mistakes of others by buying things now that will sell for much more later *is exactly what you want speculators to do*. This type of speculation is socially productive for two reasons. First, it gives people incentive to make better forecasts of future outcomes as doing so will result in lower risk for society. This risk reduction is a welfare-improving investment. Second, when the speculator enters the market, in purchasing the allowances when other people mistakenly believe that they have little value, the speculator will raise the price of allowances and broadcast to the entire market the assessment that the allowances were undervalued at the old price. If the assessment is correct about future demand, the speculator will make a profit, if the assessment is wrong, the speculator will lose money and other participants in the market will profit from his mistake. The difference between this case and the previous one is that demand actually shifts between the purchase and the sale, so the savvy speculator can make a profit by being the better forecaster but also serves a public service by signaling impending scarcity. Eliminating speculative activity of this type would have the likely effect of increasing volatility rather than reducing it. Speculators participate routinely in the energy and generation-capacity markets. Attempting to eliminate their

activity in the RGGI allowance market could give rise to inefficient pricing decisions across these markets.

While it is true that a market manipulator might be able to take advantage of the limited supply of RGGI allowances and the inelastic demand of electricity, these opportunities to profit are both costly and risky for the reasons already discussed. In addition, as the U.S. Treasury sale example demonstrates, there is considerable legal and political risk to engaging in such a strategy.

9.2 *The Market Disruption Motive*

For many of the reasons already stated, the fears that some party may use auctions or the spot market to disrupt the electricity market in the RGGI region seems overblown. An early move to buy allowances will give generators signals that they need to make more aggressive bids on their essential units. A play late in the compliance period is unlikely to be successful because generators most likely will have already purchased their essential units. As already discussed, the uniform-price auction allows bidders to bid high on essential units to ensure winning them at auction but without expecting to pay the high price for those units. Bids necessary to win all of these high-value units immediately would raise suspicions of attempted market manipulation.

9.3 *The Competitive Advantage Motive*

It is possible for non-emitting generators to buy enough allowances to raise the price of allowances. Doing so would raise the cost of generation for emitters. This, in turn, would raise the price of electricity. The increase in price could increase profits of the non-emitting generators enough to offset their costs of purchasing the allowances. But these allowances cannot be sold. To retain the profits, the generator must keep the allowances and not sell them back. The reason for this is that the higher price for allowances will result in a reduction of CO₂ emissions through fuel switching and possibly some limited reductions in the demand for electricity. In either case, there would be lower CO₂ emissions, increasing the net supply of future allowances relative to the need. If the non-emitting utility were to sell the allowances, the allowances would be in surplus relative to the period before the increased price and prices would fall below the previous equilibrium. In turn, electricity prices would fall below the original price, reducing profits of the non-emitting generators. This sequence of events is a consequence of the simple arithmetic of allowance demand.

It is difficult to imagine how a non-emitting generator could justify having a long-term beneficial interest in a large block of allowances that it does not need for compliance. Routine monitoring of RGGI accounts and generator finances would make this form of market manipulation legally and politically risky. Rather than build frictions into the RGGI market to prevent a problem that is not apparent, routine market monitoring such as that suggested in this report likely will prevent the behavior because of the substantial likelihood of getting caught

9.4 The External Compliance Motive

The RGGI states are creating a new asset, the RGGI CO₂ allowance, which may have value outside of RGGI. For example, a corporation wishing to advertise its carbon neutrality could buy RGGI allowances and functionally retire them. The same strategy might be used by a city that has pledged to reduce its carbon footprint. Rather than buy offsets through the voluntary offset market, the city could choose to buy RGGI allowances. The external compliance motive is entirely consistent with the goals of RGGI, but because the world market for such carbon reductions, while relatively small now, ultimately is very large compared to any excess of allowances in RGGI, it is possible that external compliance activities could have an effect on RGGI allowance prices.

RGGI allowances are not usable in the EU ETS, the largest trading system in the world, covering more than 50% of CO₂ emissions in 25 countries. It is unclear whether the greenhouse gas control programs, such as those in California, Illinois, or Florida, will recognize RGGI allowances as offsets. In the short run, the likely source of demand for RGGI allowances will be the voluntary compliance market. This market appears to be growing fairly rapidly. Sources of offsets vary considerably in terms of perceived quality as well as price.

According to some analysts, a large part of the offset market is somewhat idiosyncratic, with offset buyers looking for specific, highly visible projects to sponsor, rather than looking to purchase generic CO₂ reductions. It is simply not known what effect this external market will have on RGGI allowances.

Once again, limiting auction participation is not an effective response to the external compliance demand. The spot market will provide a ready source of allowances for outside buyers. Closing the auction will not change the effect of this external demand on prices to any appreciable extent. This recommendation both is strongly predicted by economic theory and is further confirmed by our laboratory explorations of the impact of a

“world demand” for RGGI allowances. In this sequence of experiments, we set a “world price” for allowances that was above the Walrasian price. Three of the subjects were identified as brokers who had no production capacity and only could buy and sell allowances for their profit. The treatment variable for this set of experiments was the ability of the brokers to participate in the auction. In one set of trials, the brokers were allowed to participate in the auction and the spot market and in the other trials the brokers were allowed only to participate in the spot market but not the auction. In both sets of trials, the spot price for allowances moved quickly to a value close to the world price, and the price at auction moved quickly above the Walrasian price and approached the world price. The restriction of the auctions to producers only did not change the result that world price drives the market price and, hence, the auction price.

The experiments showed that whether allowances could be purchased in auction or not, through exchanges in the spot market, RGGI allowances would become part of a larger pool of carbon assets that have value for voluntary compliance with carbon-reduction commitments. Not enough is known about the voluntary offset market to determine what effect that market will have on RGGI allowance prices. Nor is it known what effect the availability of RGGI allowances will have on the offset market. The current large price spread in the offset market likely reflects the considerable uncertainty over offset quality and value, in addition to uncertainty over the demand and supply of offsets.

9.5 Possible Approaches to Address Hoarding

We already have mentioned some possible strategies that could be used to address hoarding of allowances. Our key conclusion here is that the possibility of hoarding is speculative and that an initial approach of monitoring of the auctions, the spot market, recorded ownership of allowances, and financial records of firms in the RGGI market will provide significant, and probably sufficient, safeguards to prevent hoarding behavior from causing significant problems in the RGGI market. Some of the possible solutions proposed either are likely to be ineffective or may cause more damage than they are likely to prevent.

It is clear from both experiments and theory that limiting auction participation falls in the category of rules that are both ineffective and likely to do more harm than good. By lowering participation rates and restricting participation to firms with a greater ability to tacitly collude, this strategy runs the risk of substantially increasing the risk of collusion in the auction.

There has been some discussion of using a “buy-it-now” rule, where generators would have a chance to buy allowances before the auction at some price higher than the expected clearing price. This is intended to ensure access to allowances at some price. Our choice of the uniform-price auction already provides the substantial equivalent of a buy-it-now rule. Firms can bid aggressively for “must have” allowances, knowing that they will only pay the market-clearing price, which is set to the highest rejected bid. Adding an additional buy-it-now option may give rise to unintended consequences without producing any gain in protection against any likely hoarding behavior.

Another rule that has been suggested by interested parties as a way to address possible hoarding, especially of the external compliance variety, is a three-year limitation on the life of an allowance. This change in the definition of a RGGI allowance would reduce greatly the value of these allowances for outside compliance. Since allowances would have a limited life, they would not satisfy the requirement in most offset programs that emissions reductions be permanent. It would be possible to implement an active trading strategy that would roll-over allowance stocks by selling those allowances about to expire and purchasing new ones. While such a strategy is possible, it would be costly and risky. Trading to roll-over stocks has transaction costs, which would not be insignificant. In addition, buying and selling always exposes the trader to financial risk. Finally, for the party purchasing the voluntary offset, it will be obvious that any commitment to roll-over stocks in perpetuity cannot be enforced effectively.

In addition, this proposal would be difficult to implement given the current flexible compliance horizon in the RGGI program. This proposal also would cause an important additional distinction to arise between the values of different vintages and even between allowances of the same vintage sold on different dates. These difficulties lead us to conclude that a three-year life span of allowances may impose significant costs on the RGGI market. In light of recent analysis of the external compliance market, our conclusion is that the potential costs of this proposal are not justified by the likely impact of hoarding behavior. In the future, a reevaluation of these options may be justified if hoarding behavior appears more likely to harm the RGGI market.

10 Combining Auctions with Free Allocations

To evaluate the effects of no-cost allocations (grandfathering) of a fraction of the permits on the behavior of participants in an auction, we ran matched sessions, each with a series of eight uniform-price auctions followed by spot markets. With no grandfathering, a

total of 60 permits were sold in each auction, and with partial grandfathering, 36 permits (60%) were allocated at no cost and 24 (40%) were auctioned. In each session, there were six low emitters (needing one permit per capacity unit) and six high emitters (needing two permits per capacity unit). The random cost draws were such that the predicted Walrasian price was \$3.50 in all auctions, except for the sixth, in which it was \$3.75.




Figure 10.1. Average Auction and Spot Prices for Uniform-Price Auctions, With and Without Partial Grandfathering

The average auction and spot prices for the two sessions with no grandfathered allocations are shown on the left side of Figure 10.1, and the corresponding averages for the three sessions with partial grandfathering are shown on the right side. As we have observed in other treatments, there is a close correlation between auction and spot prices, and all price series are fairly close to the Walrasian predictions. Although average prices are slightly higher in the sessions with partial grandfathering, the effect is not large relative to the amount of price variability across sessions. Our conclusion is that we do not see an effect in the auction experiments relating to the portion of allowances that are distributed initially through auction and the portion that are distributed at no cost.

Part 3: Recommendations

11 Recommendations for Auction Design

Recommendation 1: Uniform-Price Auction

The RGGI auction should use a uniform-price auction format. The clearing price for the auction is the value of the highest rejected bid. The uniform-price auction format has much to recommend it, including simplicity, relative transparency, and the observed tendency for bidders to ensure purchases of needed allowances by bidding closer to use values. This auction design performed very well in our price discovery experiments. The uniform-price auction also is familiar to the electricity sector, as it is the auction format used in most ISO electricity auctions.

Recommendation 2: Single-Round, Sealed-Bid Format

The RGGI auction should use a single-round, sealed-bid format. The literature suggests that multiple-round auctions can be more conducive to collusion, as they provide participants with opportunities for signaling and detecting when someone has reneged on a collusive agreement. Some have proposed that a continuous auction may be preferable to a single-round auction, but our results indicate that the continuous auction performs less well at promoting price discovery and does more to facilitate collusion. In our preliminary recommendations in Phase 1, we had recommended that the first auction for each vintage be a clock auction (with a final, sealed-bid stage), but further examination suggests that clock auctions perform no better in terms of price discovery than a single-round auction. Ties in the auction should be resolved in a random manner to help guard against collusive bidding.

Recommendation 3: Separate Auctions by Vintage

Separate auctions should be held for different vintages. Since the yearly vintages within a compliance period are not identical assets due to different first years of allowable use and due to the possibility that in exceptional circumstances compliance periods could be extended, they should be sold separately. Equating vintages with three-year compliance periods would simplify the program and would reduce transaction costs, but this would require significant modifications to the proposed compliance period definition.

Recommendation 4: Quarterly Auctions

Auctions should be held quarterly. This schedule of auctions provides the benefits of periodic price discovery and enhanced liquidity without interfering with the performance of a secondary market. Experimental evidence and evidence from other allowance auctions is persuasive that auction and spot-market prices will track each other closely. A regular sequence of auctions for allowances will be built into spot-market participant expectations and is unlikely to cause disruption.

Recommendation 5: Auction Future Allowances in Advance

Future allowances should be made available four years in advance of their vintage. On each of the quarterly auction days, an auction should be held for current vintage-year allowances and an auction should be held for a future vintage. First-quarter auctions would include an auction of allowances from the one-year-ahead vintage, second-quarter auctions would include an auction for the two-year-ahead vintage, and so forth.

Generators have expressed a desire for some degree of certainty regarding future allowance prices and allowance availability to assist in their planning for future investments. They want auctions of allowances of current and future vintages to occur before regional ISO capacity auctions to allow generators to be able to secure the allowances they need to perform future contract obligations. Auctioning future vintages in advance should help with generator planning.

Recommendation 6: Reserve Price

A reserve price should be used in each auction. In general the reserve price should be publicly announced, although in the first auction a reserve price may or may not be announced in advance. A compelling justification for a reserve price can be found in the academic literature and from previous experience with auctions, and the reserve price would help the auction achieve criteria set out in this report.

How the reserve price is set in the auction interacts with other aspects of the program design. Regardless of how the reserve price is set, no bids for allowances should be accepted if the bid price falls below the reserve price.

Recommendation 7: Unsold Allowances

Two options have been identified for what to do with allowances that are not sold in an auction because of insufficient demand or because the reserve price is triggered. One option is that unsold allowances could be rolled into a contingency reserve account. The allowances in the contingency account would not be released for sale until some RGGI auction closes above a specified value, such as the first offset trigger price. Once this condition is met, the contingency reserve allowances would be available for auction on the next quarterly auction date. If the size of the contingency reserve account is limited and that limit is reached, then some unsold allowances could be rolled into the subsequent auction. The contingency reserve account would help to minimize large fluctuations in allowance prices. Price volatility is undesirable from the generators' perspective and could limit incentives for investment in clean technology, so efforts to limit volatility would help to promote the goals of the program. Another option is that all of the unsold allowances could be rolled into the next auction.

Recommendation 8: Lot Size

Lot size at auction should be a minimum of 1,000 allowances. This will reduce administrative costs and bidding costs without placing significant burdens on bidders. The lot size should not be so large that it limits the participation in the auction.

Recommendation 9: Open Auctions to All Qualified Bidders

Auctions should be open to anyone willing and able to meet financial pre-qualification, but no single entity should be able to purchase (or take a beneficial interest in) more than 33% of the allowances for sale in any auction. Open auctions will enhance competition and limit opportunities for collusion. Limiting the share of allowances that a single entity can purchase in any given auction raises the cost of using the auction to corner the market without placing too stringent a restriction on what generators can purchase.

Recommendation 10: Bids are Binding Contracts

Accepted bids should be treated as binding contracts. Bidders must provide strong financial assurance to cover the value of any bids. Financial assurance may include bond ratings, letters of credit, or other instruments of equal quality. Those not able to meet financial qualifications may deposit cash in escrow to cover bids. No bids above financial

assurance levels should be allowed for any bidder. Substantial penalties should be applied to any party not performing their contract obligation to pay the clearing price on all winning bids.

Recommendation 11: Joint and Uniform Auction

There should be a joint and uniform auction for allowances of a given vintage sold from all RGGI states. Allowances should be completely identical, notwithstanding the state of origin. All contract and enforcement terms should be identical for all allowances, notwithstanding the state of origin. (Note: This does not require that all allowances be sold, only that for those sold, they should all be sold through the joint and uniform RGGI auction mechanism.)

A single, uniform auction is recommended for several reasons. Differences in auction design and implementation across states may lead to confusing and irrelevant differences in price signals. States would be tempted to choose the timing of auctions, reserve prices, or other parameters in ways that favor them. In addition, multiple auctions almost certainly will raise the administrative costs of making allowances available to the market and the transaction costs for firms seeking to acquire them.

Recommendation 12: Market Monitoring

RGGI market monitoring efforts should take advantage of existing monitoring activities by federal and state agencies and other interested parties. RGGI should coordinate with the Federal Energy Regulatory Commission, the U.S. Environmental Protection Agency (EPA), the Independent System Operators and the Commodity Futures Trading Commission (CFTC) in designing criteria for detecting market manipulation and for sharing of information regarding the performance of the allowance market and the detection of attempts to manipulate prices.

Recommendation 13: Disclosure of Beneficial Ownership

RGGI should require that the authorized account representatives be obliged to disclose the “beneficial ownership” of any allowance holdings. That is, every participant would have to disclose the party sponsoring or benefiting from the agent’s activities in the allowance market if it was other than themselves or their immediate employer. Currently, this is not required in the EPA’s ATS. This information is proprietary and should be kept confidential.

Recommendation 14: Auction Information Disclosure

Information from the auction that should be publicly disclosed includes the auction clearing price, the identities of winning bidders and the quantity of allowances obtained by each winning bidder. The actual value bid by each auction participant should not be disclosed. Information about losing bidders should not be disclosed.

Recommendation 15: Statement of Intent

RGGI should articulate the auction goals in a “Statement of Intent” and ask all participants in the auction to acknowledge that statement and agree not to undermine these goals. The goals that might be articulated range from overall environmental integrity to specific behavior in the allowance market.

Recommendation 16: Ongoing Evaluation

RGGI should evaluate the performance of the auction on an ongoing basis as part of their administrative oversight of the program.

12 References

- Abbink, K., J. Brandts, and P. Pezanis-Cristou (2006). "Auctions for government securities: A laboratory comparison of uniform, discriminatory and Spanish designs." *Journal of Economic Behavior & Organization* 61(2): 284–303.
- Ashenfelter, O. (1989). "How Auctions Work For Wine And Art." *Journal of Economic Perspectives* 3(3): 23–36.
- Athey, S. and P. A. Haile (2006). Empirical Models of Auctions. National Bureau of Economic Research Working Paper Series. Cambridge, MA: NBER.
- Ausubel, L. (2004). "An Efficient Ascending-Bid Auction for Multiple Object." *American Economic Review* 94(5): 1452–1475.
- Ausubel, L. and P. Cramton (1998). Demand Reduction and Inefficiency in Multi-unit Auctions. Working Paper. College Park: University of Maryland Department of Economics.
- Ausubel, L. and P. Cramton (2004) "Vickrey Auctions With Reserve Pricing." *Economic Theory* 23(3): 493-505.
- Avery, C. (1998). "Strategic Jump Bidding in English Auctions." *The Review of Economic Studies* 65(2): 185–210.
- Bernard, J. C., T. Mount, et al. (1998). "Alternative Auction Institutions for Electric Power Markets." *Agricultural and Resource Economics Review* 27(2): 125–131.
- Binmore, K. and P. Klemperer (2002). "The Biggest Auction Ever: The Sale of the British 3G Telecom Licenses." *The Economic Journal* 112: C74–C76.
- Borgers, T. and C. Dustmann (2005). "Strange bids: Bidding behavior in the united kingdom's third generation spectrum auction." *The Economic Journal* 115: 551–578.
- Burtraw, D. and K. Palmer (2006), "Summary of the Workshop to Support Implementing the Minimum 25% Public Benefit Allocation in the Regional Greenhouse Gas Initiative," RFF Discussion Paper 06-45. Washington, DC: Resources for the Future.

- Burtraw, D., K. Palmer, R. Bharvirkar and A. Paul (2001). "The Effects of Allowance Trading on the Costs of Carbon Emissions Trading," RFF Discussion Paper 01-30. Washington, DC: Resources for the Future.
- Compte, O. and P. Jehiel (forthcoming) "Auctions and Information Acquisition: Sealed bid or Dynamic Formats?" *The Rand Journal of Economics*.
- Cox, J. C., V. L. Smith, et al. (1985). Expected Revenue in Discriminative and Uniform Price Sealed-Bid Auctions. *Research in Experimental Economics* 3: 183–208.
- Dechenaux, E. and Kovenock, D (2005). "Tacit Collusion and Capacity Withholding in Repeated Uniform Price Auctions." Working Paper Series 636. Stockholm: Research Institute of Industrial Economics.
- Ellerman, A. D., P. L. Joskow, R. Schmalensee, J. Montero, E. Bailey, (2000). *Markets for Clean Air: The U.S. Acid Rain Program*. New York: Cambridge University Press.
- Fabra, N. (2003). "Tacit Collusion in Repeated Auctions: Uniform versus Discriminatory." *Journal of Industrial Economics* 51(3): 271–293.
- Fuerbringer, Jonathan, (1991). "Salomon Brothers Admits Violations At Treasury Sales." *New York Times*, August 10.
- Goeree, J. K. and C. A. Holt (2005) "Comparing the FCC's Combinatorial and Non-Combinatorial Simultaneous Multi-Round Auctions: Experimental Design Report," prepared for the Wireless Communications Bureau of the Federal Communications Commission.
- Goeree, J. K. and C. A. Holt (2007) "Hierarchical Package Bidding: A 'Paper and Pencil' Combinatorial Auction," Discussion Paper, Caltech.
- Goeree, J. K., C. A. Holt, and J. O. Ledyard (2006) "An Experimental Comparison of the FCC's Combinatorial and Non-Combinatorial Simultaneous Multiple Round Auctions," prepared for the Wireless Communications Bureau of the Federal Communications Commission.
- Goeree, J. K., C. A. Holt, and J. O. Ledyard (2007) "An Experimental Comparison of the Package Auction Formats," prepared for the Wireless Communications Bureau of the Federal Communications Commission.
- Goeree, J. K. and T. Offerman (2004). "The Amsterdam Auction." *Econometrica* 72(1): 281–294.

- Goeree, J.K., T. Offerman, R. Sloof. (2006). "Demand Reduction and Preemptive Bidding in License Auctions." Working paper.
- Hausker, K. (1992). "The Politics and Economics of Auction Design in the Market for Sulfur Dioxide Pollution." *Journal of Policy Analysis and Management* 11(4): 553–72.
- Hendricks, K., R. H. Porter, et al. (1989). "Random Reservation Prices And Bidding Behavior In OCS Drainage Auctions." *Journal of Law & Economics* 32(2): S83–S106.
- Hendricks, K. and H. J. Paarsch (1995). "A Survey of Recent Empirical Work Concerning Auctions." *The Canadian Journal of Economics / Revue canadienne d'Economie* 28(2): 403–426.
- Holt, C., W. Shobe, et. al. (2006). "An Experimental Basis for Public Policy Initiatives." In *Promoting the General Welfare: American Democracy and the Political Economy of Government Performance*, Gerber and Patashnik, eds. Washington, DC: Brookings.
- Kagel, J. (1987). "Information Impact and Allocation Rules in Auctions with Affiliated Private Values: A Laboratory Study." *Econometrica* 55(6): 1275–1304.
- Klemperer, P. (1999). "Auction Theory: A Guide to the Literature." *Journal of Economic Surveys* 13(3): 227–286.
- Klemperer, P. (2002). "What Really Matters in Auction Design." *The Journal of Economic Perspectives* 16: 169–189.
- Levin, D. and J. L. Smith (1996). "Optimal reservation prices in auctions." *Economic Journal* 106(438): 1271–1283.
- List, J. A. and D. Lucking-Reiley (2000). "Demand reduction in multiunit auctions: Evidence from a sportscard field experiment." *American Economic Review* 90(4): 961–972.
- McAfee, R. P. and J. McMillan (1996). "Analyzing the Airwaves Auction." *The Journal of Economic Perspectives* 10(1): 159–175.
- McMillan, J. (1994). "Selling Spectrum Rights." *Journal of Economic Perspectives* 8(3): 145–62.

- Miller, G. J. and C. L. Plott (1985). "Revenue-Generating Properties of Sealed-Bid Auctions: an Experimental Analysis of One-Price and Discriminative Processes." *Research in Experimental Economics* 3: 159–182.
- Misiolek, W.S. and H.W. Elder (1989). "Exclusionary Manipulation of Markets for Pollution Rights." *Journal of Environmental Economics and Management* 16: 156–166.
- Nordhaus, W. D. (2007). "To Tax or Not to Tax: Alternative Approaches to Slowing Global Warming." *Review of Environmental Economics and Policy* 1(1): 26–44.
- Porter, D., W. Shobe, V. Smith, A. Winn (2007). "The Design, Testing, and Implementation of Virginia's NOx Allowance Auction." *Journal of Economic Behavior and Organization*.
- Rassenti, S., V. Smith, B. Wilson. (2002). "Using Experiments to Inform the Privatization/Deregulation Movement in Electricity." *Cato Journal* 21(3): 29.
- Smith, V., (1967). "Experimental Studies of Discrimination Versus Competition in Sealed Bid Auction Markets." *Journal of Business*. 40(1): 56-84.
- Staropoli, C. and C. Jullien (2006). "Using Laboratory Experiments to Design Efficiency Market Institutions: The Case of Wholesale Electricity Markets." *Annals of Public and Cooperative Economics* 77(2): 23.
- Vickrey, W. (1961) "Counterspeculation, Auctions, and Competitive Sealed Tenders." *Journal of Finance* 16: 8–37.
- Whitford, A.B. (2007) "Designing Markets: Why Competitive Bidding and Auctions in Government Often Fail to Deliver." *Policy Studies Journal* 35(1): 64–85.

13 Appendix A: Statement of Work Map

Statement of Work Questions	Discussion in Text
1. How often and when should RGGI allowance auctions be held?	5.2
2. How large should each auction be?	5.2
3. Since the allowances have 3-year time stamps and they might be auctioned in sequence, shall the first auction precede the announcement of other allocation decisions?	5.2
4. How do staggered implementations by RGGI states of auctions affect the efficiency and fairness criteria laid out below?	5.2
5. Are there negative consequences to interactions with other allocation methods such as grandfathering, future allocations, or federal auctions? What are the rules needed to address these interactions?	10
6. What can or should be done to prevent the hoarding of allowances?	9
7. Does open-ended participation in the auction by “non-compliance” entities have an adverse impact on auction performance? If so, what policies, rules or strategies should be used to mitigate these impacts?	9
8. What role will the auction play in revealing willingness to pay for allowances in secondary allowance markets?	8.2.1
9. How might the timing of the auction help best to disseminate information about costs and contribute to price stability?	5.2
10. In general, what is the effect of the auction on secondary allowance markets, and vice versa?	8.2.1
11. How does the auction price compare with outcomes in the secondary or reserve markets?	8.2.1
12. What is the role of a reserve price?	7.2
13. Should the reserve price be publicly announced? Why/why not?	7.2
14. What shall be the designation for allowances that are not sold if the reservation price is triggered?	7.2
15. Shall such allowances contribute to a “strategic allowance reserve” that carries forward to insure against potential future allowance shortfalls?	7.2
16. How much information about the identity and offer prices of bidders should be shared?	7.2
17. Shall offers of losing bidders be revealed? Why/why not?	7.2
18. What is the interplay of asymmetries among bidders, such as size of the firm, capitalization, etc., and how do these affect perceptions of fairness?	6

14 Appendix B: Annotated Bibliography

Abbink, K., J. Brandts, P. Pezanis-Cristou (2006). "Auctions for government securities: A laboratory comparison of uniform, discriminatory and Spanish designs." *Journal of Economic Behavior & Organization* 61(2): 284-303.

This paper studies an auction format used by the Spanish government to sell bonds. There has been a large debate regarding the use of discriminatory versus uniform price auctions to sell multi-units such as in the treasury auctions. However, the ranking is generally model-dependent and we do not have a final answer regarding which institution is superior.

This paper takes a look at a third option that has been used in practice but for which little or nothing is known theoretically. It compares such third option, the "Spanish auction" (using experimental data), with the other two commonly used formats. The "Spanish auction" is somewhat of a hybrid format in that the bidders pay the market clearing price if their bid is above that value and they pay their bid otherwise. One important assumption that is made is that the bidders have a common valuation for the units for sale. This is argued to be a plausible assumption in the treasury auctions due to the fact that bidders buy to resell in the same market. Although the auction is in common value, the results show that the winner's curse does not appear to be relevant. It is suggested that this may be due to the fact that the units demand was taken to be lower than in comparable experiments. The authors argue that this was done to fit more closely with the treasury auction application. The main result provided is that the "Spanish auction" performs closely to the uniform price auction and that both formats outperform the discriminatory design in terms of revenues. Additionally, these formats display less price variability.

The paper fails to explain why such format should be preferred to the better known uniform price auction. Given the absence of theoretical work on this auction, more work seems to be needed to assess its eventual benefits. A clear ranking could perhaps emerge if preference asymmetries were introduced.

Athey, S. and P. Haile (forthcoming). "Non parametric approaches to auctions."

This is a very important and comprehensive survey that presents the currently available econometric techniques used in empirical auction literature. This survey also reviews the main results.

Athey, S. and P. A. Haile (2006). Empirical Models of Auctions. National Bureau of Economic Research Working Paper Series. Cambridge, MA, NBER.

This paper offers a survey of empirical models to estimate behavior in auctions. It also identifies some of the main findings using recently developed techniques. Several other citations are offered and discussed in detail.

Hendricks, Pinkse and Porter (2003) analyzed oil lease auctions, which are a well-known example of a common value problem. The price of the resource is determined by an external market, so it is common to all bidders. This gives rise to the possibility of a winner's curse, because whoever bids high enough to win the auction apparently assumes a higher resource availability than others and hence has reason to doubt their own expectations over the resource availability. Winning bidders tend to be the ones who have over-estimated the resource value. The authors find that subtle inferences are economically important and are incorporated in bidding strategies. The authors also find that the magnitude of the curse is significant. It appears larger when there is greater anticipated competition, which follows from the assumption of symmetric pure common values.

Haile, Hong and Shum (2003) examine common and private values in first-price auctions. A hypothesis is that as the number of competing bidders increases so will the winner's curse. However, they find this may not be important in the case of timber contracts.

Haile and Tamer (2003) analyze ascending auctions, which they characterize as a dynamic game with a rich strategy space. They also analyze the role of reserve prices. They show that bidders make an inference based on the reserve price policy. Actual reserve prices in timber auctions are likely to be below optimal levels, but raising them would have only a small effect on expected revenues and on the probability of a sale. Hence, one might conclude that it is more important to have

reserve prices than to worry too much about their levels, at least with respect to maximizing revenues.

In another study of timber auctions, Athey, Levin and Seira (2004) look at variation in auction format between ascending and first-price auctions to assess competitiveness and the widely-believed notion that ascending auctions are more susceptible to collusion. The Revenue Equivalence theorem (Vickrey, 1961) implies that if bidders are risk-neutral, have independent and identically distributed values, and bid competitively, the two auction formats yield the same winner, same expected revenue and the same bidder participation. However if these assumptions are relaxed, then auction format becomes relevant. The authors go on to cite Maskin and Riley (2000) who find that first-price auctions lead to inefficient entry and bidding. The effect on revenue is ambiguous. The authors conduct new empirical work to find that bidding behavior in the timber industry is less aggressive in ascending auctions, suggesting collusion. In some cases they cannot reject the hypothesis that the bids in ascending auctions are equal to those predicted for the first-price auction, but in another case they find the ascending auction is less competitive. The setting here is when bidders are face-to-face, which provides some opportunity for signals. They find similar welfare effects for a fixed number of participants. When the number of participants is endogenous, the sealed bid auction increases revenue, suggesting in this case that auction format is important with respect to collusion.

Jofre-Bonet and Pesendorfer (2003) consider the role of capacity constraints that make winning an auction affect valuation in future auctions. A study of highway construction reveals asymmetries in bidding strategies based on point in time, which may depend on their performance in a previous auction. The winner in one auction affects the bids in subsequent auctions, given capacity constraints.

Hortacsu (2002) looks at whether to use discriminatory or uniform price auction formats for treasury bills, drawing on evidence from Turkey. Friedman (1960) considers which format will raise the most revenue. Bidding one's true marginal valuation is not an equilibrium strategy in either auction. In the discriminatory auction "truthful" bidding would lead to zero surplus to any bidder. In a uniform-

price auction, bidders also have an incentive to shade bids below marginal valuations, since a bidder's own bid may set the price for all infra-marginal units. Which format yields the most revenue depends on the primitives of the problem (Ausubel and Cramton, 2002). Hortacsu finds that a uniform auction would not enhance revenues in the case he empirically considered.

Athey, S. and J. Levin (2001). "Information and Competition in U.S. Forest Service Timber Auctions." *The Journal of Political Economy* 109(2): 375-417.

This paper looks at scale sale auctions used for the allocation of timber by the forest service. The forest service provides as public information estimates on the amount of the different types of timber present in the forest. Bidders are then allowed to inspect the forest and make up their own estimates. Bidders are asked to produce a unit price for each type of timber. The winner is the one that produces the highest expected value for the overall units on sale based on the announced quantities. The ultimate payment is based on the actual number of trees extracted, which can only be observed ex-post.

We should expect bidders' private information to be present and to bias the bids. If the bidder estimate is different from the one announced, the payment he expects to make if he were to bid without taking into account his own private estimate is different from the value of his bid. However, a rational bidder should exploit his private information to bid more aggressively on those timbers for which his estimate is lower than the one announced (this increases the bid by less than the expected payment). The authors empirically test the actual bidding observed in these auctions, and they indeed find that the bidding behavior underlines the existence of private information. An alternative selling method that could be used is to ask a for lump sum payment. The reason why this method is not chosen may be due to the excess risk left on the winning bidder.

Ausubel, L. (2004). "An Efficient Ascending-Bid Auction for Multiple Object." *American Economic Review* 94(5): 1452-1475.

This paper provides us with an ascending price auction for the allocation of multiple homogeneous objects that display remarkably good characteristics. As a matter of fact, such format inherits many of the advantages of the single unit English Auction. In particular, if values are private, sincere bidding is in equilibrium and leads to an efficient allocation (in this case as in his static Vickrey

counterpart). If values are affiliated, the Ausubel auction remains efficient, while the static Vickrey auction is not. This replicates the relationship between English Auction and the SPA for the multi-object scenario.

The auction works as follows. The auctioneer calls an initial starting price and the bidders report how many units they are willing to buy at such price. As long as there is excess demand the price is raised by the auctioneer. The allocation rule is the following. Suppose that at price, p , the residual demand for bidders other than i is equal to $n-1$, where n is the number of units for sale. Then we say that bidder i has "clinched" one unit. The price he pays for such unit is the one for which the other bidders residual demand drops to $n-1$. The auction stops when all units have been "clinched".

Ausubel, L. and P. Crampton (1998). Demand Reduction and Inefficiency in Multi-unit Auctions.

This paper points out that the sealed bid uniform price auction used for the allocation of multiple units is generally inefficient due to a phenomenon known as demand reduction. Essentially, a large bidder has a stronger incentive to shade his bid than a small bidder, and this may cause him to lose some units for which his value is ex-post higher than the one of the small bidder. The reason is simple to understand and it is analogous to the reason why a monopolist sells a less-than-efficient quantity. Recall that in a uniform price auction the bidder pays the market clearing price. Such price with positive probability is determined by one of the prices posted by the large bidder. As the price he pays is the same for all units, it may payoff for him to "risk" to win less units but induce a lower price on all of them.

Avery, C. (1998). "Strategic Jump Bidding in English Auctions." *The Review of Economic Studies* 65(2): 185-210.

This paper uses the Milgrom- Weber (1982) affiliated value model to study an open auction where bidders are allowed to raise the price discontinuously (unlike in the standard model). It shows that bidders can exploit such possibility to implement a form of implicit collusion by signaling their type in a first stage to understand who among them is the strongest bidder. Once they have signaled such information, in the second stage they use a less aggressive strategy if they are considered weak and

a more aggressive one if they are strong. The use of such asymmetric strategies decreases the seller's revenues.

The model gives rise to a multiplicity of signaling equilibriums. It may therefore be difficult in practical terms to predict which one bidders will coordinate to (and there could be mis-coordination problems). Further, it is not very clear why a weak bidder should commit to a less aggressive strategy in the second stage compared to the one of the standard symmetric equilibrium.

Back, K. and J. Zender (1993). "Auctions of Divisible Goods: On the Rationale for the Treasury Experiment." *The Review of Financial Studies* 6(4): 733-764.

This paper looks at multi-unit auctions and compares the performance of the sealed bid uniform price auction with that of the sealed bid discriminatory auction. The main difference between the two formats is that in the first one the winning bidders pay the market clearing price, while in the second one they pay their own bid. Notice the analogy with the Second Price Auction and the First Price Auction here. However, one of the points of the paper is exactly to argue that the main insights on the single object framework (FPA vs. SPA) cannot directly be replicated in their multi-units counterparts. The paper highlights the pros and cons of these two formats.

Bergemann, D. and S. Morris (2005) Robust Mechanism Design. *Econometrica* Volume, 1771-1813 DOI:

This is one of the most important recent theoretical contributions in mechanism design; auction theory is one of the leading applications. This paper is motivated by the emerging concern that some of the theoretical results derived by the literature to date are sensitive to the details of the format or the assumptions regarding the information that the bidders (or the seller) know or can report. In particular, often strong common knowledge assumptions are made regarding the distribution from which bidders' private information are drawn. This paper is a first important step that calls attention to auction designs that are more "robust" when considering the specific details of the environment for which the auctions are designed.

Bergemann, D. and J. Valinaki (2006). *Information in Mechanism Design*.

This is a good survey regarding the emerging literature on information acquisition in auctions (and more generally mechanism design). Standard auction models

typically assume that the information that bidders hold is exogenous. This paper reviews those contributions where information is taken as endogenous, where the bidders (and/or the seller) can decide to acquire (typically costly) or release information. This is very important because in many action applications bidders invest large amounts of money to improve their information in order to be successful in the auction. Different mechanisms provide different incentives for information acquisition, which in turn affects revenues and efficiency. For example, one crucial distinction is whether the format is open or sealed. While the former typically allows bidders to gather information during the auction, the latter allows only information acquisition before the auction begins. It is important to understand which formats perform better when information is endogenous. Such an important question has only recently begun to be addressed. Some of the main results are reported in this survey.

Bernard, J. C., T. Mount, et al. (1998). "Alternative Auction Institutions for Electric Power Markets." *Agricultural and Resource Economics Review* 27(2): 125-131.

Many electric power restructuring proposals in the northeast currently include single-sided auction mechanisms for the wholesale generation market. In this study, researchers used laboratory experiments to examine the performance of two uniform price auctions: last-accepted-offer (LAO) and first-rejected-offer (FRO). With both of these auction types the offers are submitted in sealed bids and are ranked from lowest to highest. A single buyer then purchases the cheapest units until supply is equal to demand. At that point, the buyer pays the same price for all units; either the lowest-rejected-offer or highest-accepted-offer, according to which auction type is being employed.

Researchers ran experiments with group sizes of two, four, and six subjects, reflecting different degrees of market competition. They found that for both auction types, the smaller the group, the higher the uniform price. Overall, they found that the LAO slightly outperformed the FRO under the same cost and demand conditions. However, group size was a much stronger determinant of price than the auction type, indicating that market power could be a concern.

Bikhchandani, S. and C.-f. Huang (1989). "Auctions with Resale Markets: An Exploratory Model of Treasury Bill Markets." *The Review of Financial Studies* 2(3): 311-339.

This may well be the earliest contribution to the small but emerging literature on auction with resale. The primary market (auction market) is assumed to be a common value market. Bidders in the secondary market (resale market) can observe the bidding in the first market and derive information from that. Hence, this creates an information linkage between the two markets that needs to be carefully evaluated.

The authors examine two auction formats: the uniform price auction and the discriminatory auction. They show that if the winning bids are announced by the seller, the uniform price auction may not have an equilibrium. However, if a uniform price auction has a symmetric equilibrium, it yields more revenues than a discriminatory auction symmetric equilibrium.

Binmore, K. and P. Klemperer (2002). "The Biggest Auction Ever: The Sale of the British 3G Telecom Licenses." *The Economic Journal* 112: C74-C76.

The authors review the British third-generation (3G) mobile-phone license auction that concluded in April 2000. The auction raised \$34 billion, equivalent to 2.5% of British GNP. The authors discuss the lessons learned, and the merits of using an auction compared to "beauty contests" for administrative allocation of licenses according to various qualifying criteria.

It is important to note the differences between radio spectrum auctions and allowance auctions. In many ways the radio spectrum auction issues are much more complicated. The path was cleared for auctions in telecommunications by the US Federal Communications Commission (FCC) use of a simultaneous ascending auction design in 1994. That auction raised about \$20 billion. The UK continued to use administrative "beauty contests" through the 1990s for its 2G phones. The central virtue of an auction is that it is the method that is expected to allocate resources to those who can use them most valuably. The authors offer several citations (e.g. Milgrom 2000) that the secondary market will not be as efficient. They also indicate that an auction approach will lead to less litigation than an administrative approach, with references to experiences in Spain and Sweden.

While there may be good grounds for direct allocation, the regulatory will have to answer “Why subsidize this industry rather than others?”

A major concern of the UK auction was to promote entry since there were a small number of bidders chasing a small number of licenses. Where entry is important, an ascending price auction is not ideal. The reason is that one powerful bidder can effectively threaten to raise their bid as long as necessary, and thereby defer entry. Sealed bid auctions would be better at promoting entry because they give entrants a better chance of winning against strong incumbents. However, they do not give bidders the opportunity to gather information about the business plans of their rivals, or to update their expectations if in the case of a common (associated) value auction. A potential fix to this is labeled an Anglo-Dutch auction, which resembles the shot-clock approach used in our experiments. This approach encourages entry by closing with a sealed bid, but allows for discovery in the early stages of the auction with an ascending clock. Another way that the UK design attempted to limit collusion was to limit the number of licenses that could be purchased to one.

The authors stress that any reserve price should be a clear commitment not to sell if the bids do not meet the price. If the bidders expect the government would subsequently resell at a lower price then the bidders will behave strategically to push the price down.

The authors also note the significant effect on non-economists of having the opportunity to play in the experiment. “By contrast, mathematical equations have very little persuasive power.”

Several potential mistakes did not surface in the UK, but there are lessons for other auctions. One chief problem was the inadequacy of the deposits that bidders were required to put down. This also has been noted in emissions allowance auctions. Plus, the longer the time in clearing the auction (the UK auction ran for several weeks) the more likely that external events may change values and cause bidders to retract previous bids.

Finally the authors emphasize that the UK auction should be copied, but that auctions should never be copied without attention to local circumstances. The really bad mistake is to take an auction design off the shelf. There is no “one size fits all.”

Borgers, T. and C. Dustmann (2005). "Strange bids: Bidding behavior in the united kingdom's third generation spectrum auction." *The Economic Journal* 115: 551-578.

This paper analyzes the actual bidding behavior in the UK third generation spectrum auction of the year 2000. The authors argue that even though the auction generated high revenues, the actual bidding behavior is far from understood using auction theory.

The format used was a simultaneous ascending price auction. The authors assume as a theoretical benchmark that bidders had private values and that they were bidding straightforwardly. By this they mean that a bidder would be active on the license for which he would hold the current highest surplus. If bidders were to follow such strategy, an efficient outcome would result.

However, several unexplained deviations from this strategy were observed. The authors suggest some possible reasons, such as the presence of financial constraints and allocative externalities.

Bose, S. and G. Delta (2007). "Exclusive Versus Non-exclusive Dealing in Auctions with Resale." *Economic Theory* 31: 1-17.

The main point of this paper is to show that when a seller has the option to open the market to the final consumers or to sell exclusively to a reseller, such latter option may be superior if the reseller can access a big enough share of the market.

The intuition for the result is that if final consumers are allowed to participate in the market, winning the auction may bring bad news to the reseller as it means that the value attached to the object by the final consumers is low. In equilibrium this fact is anticipated by the reseller that hence bids less aggressively, thus inducing less revenues.

The seller may then be better off contracting exclusively with the reseller when the uncertainty about the market is still unresolved.

Bulow, J. and J. Roberts (1989). "The Simple Economics of Optimal Auctions." *Journal of Political Economy* 97(5): 1060-1090.

This paper makes the Myerson optimal auction "accessible for the crowds." Whereas the Myerson paper takes the abstract (and powerful) mechanism design point of view as a starting point, these authors show that the final results can be interpreted in the more familiar language of a price-discriminating monopolist. In particular, the optimal reserve prices and bidding credits (used by the FCC, for instance) are derived by considering standard monopoly maximization problems.

Cassady, R. (1967). *Auctions and Auctioneering*, Berkeley: University of California Press.

This is an old but quite comprehensive study that provides a detailed description of many auction mechanisms used in practice. It is a source very often cited by the theoretical papers.

Che, Y.-K. and I. Gale (1998). "Standard Auctions with Financially Constrained Bidders." *The Review of Economic Studies* 65(1): 1-21.

Standard models assume that bidders do not suffer from any financial constraint so that in principle they can bid any sum of money. This paper looks at the standard private value model but relaxes the assumption that bidders do not suffer from financial constraints; in particular, it assumes that any bidder has a certain budget and that he holds private information regarding it. The main result is that the First Price Auction (FPA) outperforms the Second Price Auction (SPA) both in revenues and efficiency. The intuition as to why revenue equivalence breaks in such direction is that in the FPA bidders shade their bids in equilibrium so that the added constraint is less likely to bind than in the SPA (where instead bidders would be willing to bid up to their value).

Compte, O. and J. Philippe (2007) *Auctions and Information Acquisition: Sealed bid or Dynamic Formats*. *The Rand Journal of Economics* Volume, DOI:

This paper highlights a possible virtue of open formats versus sealed formats: the former allows a bidder to observe the strength of competition during the auction provide better incentives for acquiring finer information, which in turn boosts both revenues and efficiency. The setting they observe is the standard one of private value, amended to allow for (costly) information acquisition during the auction. Some bidders are assumed to be informed about their exact valuation, while others know only the expected value of it and by paying some cost can find out the exact one.

In a sealed format a bidder needs to make his decision prior to the start, when little is known about the strength of the actual competition. Conversely, in an open format a bidder can postpone his decision and observe how many bidders are left. If the realized level of competition turns out to be low, he will invest. Competition may turn out to be weaker than expected in which case he would not invest in the sealed format but you would in the open format.

Cox, J., B. Roberson, et al. (1982). "Theory and Behavior of single object auctions." *Research in Experimental Economics*.

This article is one of the early contributions in testing auction theory in experimental labs. Theory dictates that in the independent private value model, the First Price Auction is strategically equivalent to the Dutch auction and the Second Price Auction is equivalent to the English auction. Thus the same bids should be observed for equivalent formats. Moreover, revenue equivalence states that the same revenues should be generated in all four formats. The experimental evidence contradicts the theoretical strategic equivalence within the former two and latter two mechanisms. In particular, this paper shows that the First Price Auction and the Dutch auction are not isomorphic; the difference is mainly driven by bidders bidding below the equilibrium prediction in the Dutch auction.

Cox, J. C., V. L. Smith, et al. (1985). *Expected Revenue in Discriminative and Uniform Price Sealed-Bid Auctions*. *Research in Experimental Economics*. V. L. Smith. Greenwich, CT, JAI Press Inc. 3: 183-208.

Researchers used experiments to test the predictions of Nash equilibrium bidding theory for expected revenue in uniform and discriminative price auctions. Theory predicts that the expected revenues from uniform price auctions and discriminative auctions where all bidders are risk neutral should be equal. Where some bidders are

risk averse, the revenues from discriminative price auctions should be greater than those from uniform price auctions: $E(RU) = E(RDN) < E(RDA)$. They found that the bidding behavior of more experienced and graduate student subjects resulted in revenues markedly below that predicted by Nash equilibrium bidding theory. This was true in both uniform and discriminative price auctions. They suggested that there was a kind of “tacit cooperation” between subjects.

Their results also corroborated those of Miller and Plott (1985):

“Increasing the absolute value of the slope of the demand curve and controlling for height, we increase the revenue generated in the discriminative auction relative to the uniform. Also, increasing the height of the demand curve at the $Q+1$ highest resale value and controlling for slope, we increase the revenue generated in the uniform price auction relative to that of the discriminative auction” (p.205).

Crémer, J. and R. P. McLean (1988). "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions." *Econometrica* 56(6): 1247-1257.

This paper together with the previous Crémer and McLean (1985) represents a very important theoretical contribution to the auction literature. It shows that if the private information that bidders receive is not drawn independently of the other bidders information (as normally assumed), but rather is correlated, then the seller can construct a mechanism that extracts all bidders surplus (notice that this implies that the mechanism is also efficient). The result is theoretically very remarkable as it holds very generally even when the correlation is very weak. In such case, however, it requires the use of lotteries involving potentially extremely high payments that in practice no bidder would be willing to accept (or would simply default).

Dasgupta, P. and E. Maskin (2000). "Efficient Auctions." *Quarterly Journal of Economics* 115(2): 341-388.

Sometimes the main objective of a seller is not to maximize revenues, but rather to achieve an efficient outcome, whereby “efficient” means that he would like to allocate the object to the bidder who values it the most. This could be the case, for instance, when the seller is a public authority allocating some scarce resource. This paper extends the Vickrey-Clark-Groves mechanism to a setting where bidders valuations are interdependent and provides a mechanism that can allocate the object

efficiently (under the assumption that bidders' private information is unidimensional and that some technical conditions necessary for efficiency are satisfied). The paper presents an important theoretical result. However, such mechanism is not observed in practice as it requires bidders to report to the seller an unrealistic amount of information (essentially, each bidder should state his value contingent on each possible realization of his opponents private information). For a similar mechanism that still achieves efficiency but requires less information to be reported, see Perry and Reny (2002).

Environmental Resources Management (2005). EU ETS: Planning for Auction or Sale. G. Cook, L. Solsbery, P. C. Cramton and L. M. Ausubel.

This consultancy study was done for the UK to help them decide how to liquidate surplus allowances in their New Entrant Reserve. A draft report was developed that considered four circumstances listed below, based on criteria used to assess the suitability of each method. Stakeholders were invited to respond in writing. The feedback was favorable, but some concerns were identified. A value was placed on continuity. Concern about the EU-wide policy was considered to be secondary. There was a general preference for auctions on the basis of transparency, openness of participation and continuity (in view of future phase disposals). The ascending clock auctions were preferred, particularly among service providers and electricity generating companies. Other themes that were valued included easy and low cost participation, and adequate time for participant preparation.

Four proposals that were considered included:

- Liquid EU ETS market and low surplus volume -> market order
- Liquid EU ETS market and low moderate volume -> sequence of market orders
- Illiquid EU ETS market and low surplus volume -> uniform-price sealed-bid auction
- High surplus volume or illiquid EU ETS market with moderate surplus volume -> ascending clock auction

The consultants concluded with a recommendation of an ascending clock auction as the default method, to be used if either of the following conditions hold one month before the scheduled auction:

1. Volume to dispose is more than 5% of average daily volume.
2. Volume to dispose times the average sale price is more than 2 million pound.

Otherwise a sequence of market orders should be used. A market order is the offer to sell at the current market price. In contrast, a limit order is the offer to sell at a fixed price.

That is, in summary, an ascending auction should be used unless the market is sufficiently liquid and the volume to dispose is small. Except for 'small' volumes, auctions are thought to perform better than sales.

Two main considerations were economic efficiency and good value for the taxpayer. Further, the auction should be simple with low transaction cost.

Three factors are important in considering the auction or sale theory:

1. The product is homogeneous and divisible.
2. The product is actively traded in secondary markets.
3. The market for allowances is not concentrated.

Fabra, N., N.-H. von der Fehr, et al. (2006). "Designing Electricity Auctions." *RAND Journal of Economics* 37(1): 23-46.

The authors develop a theoretical model to characterize bidding behavior and market outcomes for discriminatory and uniform electricity auctions. They use a basic duopoly model with 2 single unit suppliers with asymmetric capacities and marginal costs. Initially they face a known, perfectly inelastic demand curve. The model is then extended to consider multiple-step bid functions, downward-sloping demands, an oligopoly case, and long-lived bids.

They found that uniform price auctions result in higher average prices than discriminatory auctions. In terms of which kind of auction is more efficient, the results are more nebulous—it depends on which equilibrium strategy is played. Long-lived bids do better than short-lived bids in the uniform auction suggesting that simplified bid formats are preferable. Having a responsive demand curve directly improves allocative efficiency and it increases supply security. Reserve price affects prices when the price cap is binding and it indirectly affects price through increased competition.

Garratt, R. and T. Troge (2006). "Speculation in standard auctions with resale." *Econometrica* 74(3): 753-769.

This paper looks at the single object private value model when a resale market is introduced, and a speculator which attaches zero value to the object is known to be present in the auction. In an auction with resale, a speculator with no private value for the object may want to win the object as the resale stage essentially introduces a common value component to the primary action. The winner regardless of his value can extract up to the difference (if positive) between the price paid and the maximum valuation of his opponents. The authors look at how standard formats are affected. The Second Price and the English auction conserve their efficient equilibrium where all bidders bid their value and the speculator cannot win. However, they also display a continuum of equilibria where the speculator wins with positive probability and makes positive profits. The First Price Auction and Dutch auction have a unique equilibrium where sometimes the speculator wins. Interestingly, the speculator makes zero expected profits.

Goeree, J. K. and e. al. (2006). "Using first-price auctions to sell heterogeneous licenses." *International Journal of Industrial Organization* 24(3): 555-581.

This paper experimentally tests the performance of three different First Price Auctions versus the simultaneous ascending one. It is motivated by the theoretical insight that a first price component should increase competition and thus revenues. The three formats are the simultaneous FPA, the sequential FPA and the descending auction. The authors look at different settings: ex-ante symmetric private values, asymmetric values, and finally a setting where bidders are uncertain about their exact value.

The main findings are as follows. The ascending auction is superior on efficiency grounds. However, all FPA are on average superior in terms of revenues and display less variability. The sequential FPA is the format that generates more revenues if the licenses are sold from highest to lowest quality. Weak bidders are worse off in such case. Finally, the simultaneous FPA is the less affected by the winner's curse.

Goeree, J. K. and T. Offerman (2002). "Efficiency in Auctions with Private and Common Values: An Experimental Study." *The American Economic Review* 92(3): 625-643.

This paper tests experimentally the performance of the First Price Auction in an environment with private and common values that in general cannot yield an efficient allocation. The observed inefficiency in the experiments is not far from the one predicted by the Nash equilibrium. The authors also show that increased competition increases both revenues and efficiency as it forces the bidders to put more weight on their private value signal rather than the common one. Providing information on the common component increases efficiency, but not as much as theory would predict.

Goeree, J. K. and T. Offerman (2004). "The Amsterdam Auction." *Econometrica* 72(1): 281-294.

Often a seller faces a pool of potential bidders that are known to be ex ante asymmetric, with some bidders being stronger than others. Myerson (1981) shows that the optimal mechanism should favor the weak bidders, allocating to them more often than efficiency would require. In practice, the optimal mechanism is difficult to implement. This paper shows that in the presence of strong asymmetries the Amsterdam auction performs better than the English auction and close to the theoretical optimal. The experimental results provided confirm this fact. The Amsterdam auction is used for the sale of real estate in Amsterdam. It is a two stage mechanism. The first part is an ascending auction that selects two bidders and sets a reservation for the second part, which is a sealed bid part. The interesting fact is that in the second stage both bidders are offered a premium proportional to the incremental value they offer relative to the reserve price. Such a premium encourages entry from the weak bidders.

Gupta, M. and B. Lebrun (1999). "First Price Auctions with Resale." *Economics Letters* 64(2): 181-185.

This paper provides a characterization of the equilibrium for the First Price Auction with asymmetric bidders (more precisely for asymmetric c.d.f, but identical signal supports) when a resale stage is added. The result is provided for the two bidder case only.

Hafalir, I. and V. Krishna (forthcoming). "Asymmetric Auctions with Resale." *American Economic Review*.

This paper looks at a private value setting where bidders' valuations may be drawn from asymmetric distributions. It compares First Price Auction and Second Price Auction in a setting where a resale stage is added after the end of the auction.

Asymmetric equilibria in FPA are typically hard to derive. Interestingly, adding an apparent further complication makes the construction of equilibria easier.

Moreover, while without resale the ranking in terms of revenues of FPA and SPA (under asymmetries) is in general ambiguous, with resale the authors are able to show that the FPA dominates the SPA. It is to be stressed that such result holds for two bidders and has not been extended to more than two bidders.

Hernando-Veciana, A. a. M., Fabio (2007). *Second Best Efficiency in Auctions*.

This paper characterizes the constrained efficient (or second best) allocation when the conditions to implement the first best are not satisfied. This is done for a setting analogous to Myerson (1981). Looking for second best efficiency is a very relevant issue as, as shown by the authors, the conditions that are necessary for full efficiency fail in many applications of interest such as under the presence of an insider or of allocative externalities. Interestingly, the second best allocation may involve no sale with positive probability. Another result provided is that for the case of two bidders the English auction implements the second best. The author can construct an equilibrium of the English auction that is constrained efficiency with more than two bidders, but they show that such equilibrium is in general not robust in the sense that may involve the use of weakly dominated strategies.

Holt, C. (1980). "Competitive Bidding for Contracts under Alternative Auction Procedures." *The Journal of Political Economy* 88(3).

This paper is one of the first to study the impact of bidders risk aversion on the auction format performance (the results are stated for a procurement auction but they analogously hold for a standard auction). The main result is that under risk aversion the First Price Auction (FPA) outperforms the open auction. The intuition why the revenue equivalence breaks is simple. If bidders are averse to risk they are

willing to give up part of their potential profits to increase their chances of winning. This fact drives the winning price up. Such adjustment is possible in the FPA, unlike in the open auction, as there is the risk that neutral equilibrium bidders shade their value.

For a study of the optimal format under risk aversion, see Maskin and Riley (1984).

Jehiel, P. and B. Moldovanu (2001). "Efficient Design with Interdependent Valuations." *Econometrica* 69(5): 1237-1259.

One of the important contributions of this paper is to provide a negative (or impossibility) result. It shows that if the information that bidders receive is multi-dimensional, there is no hope (in general) of achieving the efficient allocation, i.e. to allocate to the bidder with the highest valuation.

This opens an important problem as there are many circumstances where bidders hold multi-dimensional information, a very natural setting being for instance the one of multi-objects auctions.

The conditions that are necessary for efficiency are provided both for the multi-dimensional and the uni-dimensional case.

Jehiel, P. e. a. (1996). "How (Not) to Sell Nuclear Weapons." *The American Economic Review* 86(4): 814-829.

This paper looks at the optimal selling procedure in the presence of allocative externalities.

Allocative externalities are often relevant when the object sold is a resource that is necessary for the bidders in a downstream market where they compete. The identity of the bidder that gets the object may affect the willingness to pay of each bidder.

The authors show that this implies that the seller can extract some surplus from the agents that have not been allocated with the object. Also, if externalities exceed the value for the object the seller may optimally retain the object.

See on this topic also "auctions with downstream interaction among buyers" by Jehiel and Moldovanu (2000).

Joskow, P. L., R. Schmalensee, et al. (1998). "The Market for Sulfur Dioxide Emissions." *The American Economic Review* 88(4): 669-685.

The Joskow article argues that the theoretical and experimental findings by Cason are not actually born out in practice in the EPA SO₂ allowance auctions. The most plausible explanation is that the EPA auctions compose a small part, only about 2.8%, of the overall SO₂ allowance trading market. Hence the participants in the auction, both buyers and sellers, have a real functioning trading market that they can fall back on. As the authors put it, "The development of the outside market significantly tightened the 'opportunity cost bounds' on the behavior of auction participants".

The authors looked at both spot auctions and six and seven-year advance auctions, but the results were similar for both types. The main evidence that they offer in support of the thesis is that as time went and the SO₂ trading market got more and more robustly established with reliable prices, deviations from market prices in the auctions diminished. In 1993, the first year allowances were auctioned, several buyers put in low-ball bids as the market price was not yet established. By 1995 the buyers' bids were much flatter and only went down to 10% below the best available estimate of market price. Additionally, in 1993 the lowest winning bid (the market clearing price) was 20.6% below the average winning bid in the spot auction, but by 1997 it was only 3.4% below. The authors concede that the outcome could be different if the EPA auction were the only way to acquire allowances.

Kagel, J. a. a. (1986). "The Winner's Curse and Public Information in Common Value Auctions." *The American Economic Review* 76(5): 894-920.

This paper experimentally tests the common value model. Theory acknowledges a winner's curse problem to which in equilibrium bidders should take account and adjust their bids.

Essentially, winning the object tells the bidder that his estimate was the highest. If there are many bidders this implies that such estimate is likely to be too high.

The authors use experienced bidders that should be less sensitive to this problem.

Nonetheless, they find that bidders bid more aggressively if the number of bidders is higher. This is the reverse of the equilibrium prediction as with more bidders the winner's curse is stronger and one should adjust downwards his bid.

They also show that providing some public information to the bidders in this context decreases revenues. This is not surprising as it helps a bidder not to fall in the winners' curse trap.

Kagel, J. e. a. (1987). "Information Impact and Allocation Rules in Auctions with Affiliated Private Values: A Laboratory Study." *Econometrica* 55(6).

This paper tests the auction theory predictions for the affiliated private value model where each bidder value is private but a higher realization of it makes it more likely that the opponent's value is also higher.

One contribution is to show that the strategic theoretical equivalence of English Auction and Second price auction fails. This is due to some bidders bidding over their valuation in the Second Price Auction, which then generates more revenues.

It also looks at the effect of providing some public information regarding the opponents' values. The experiment confirms that more information raises revenues, but the effect is less relevant than the one theoretically predicted.

Kiesling, L. and B. J. Wilson (2007). "An Experimental Analysis of the Effects of Automated Mitigation Procedures on Investment and Prices in Wholesale Electricity Markets." *Journal of Regulatory and Experimental Economics* 31(3): 313-334.

Many people have argued that market monitoring is an important component of a deregulated electricity market. The Federal Energy Regulatory Commission's Orders 888 and 2000 allow ISOs and RTOs to implement market mitigation including AMP. The New York ISO has used market mitigation since 1999. AMP

uses a historically-based price cap on suppliers' offers to screen and mitigate offers that exceed the pre-defined threshold. The screen and mitigated prices vary over time and across different individual generators.

In this study, researchers used experiments to test the effects of an automated mitigation procedure (AMP) on whole sale electricity prices and the capacity investment behavior of suppliers. They found that: 1) Investment in new capacity is the only variable that reduces long-run prices; 2) The use of AMP does not affect investment in new capacity nor does it affect the long-run price of electricity relative to markets without AMP; 3) Subjects were still able to manipulate the market by finding prices higher than the reference price, but not high enough to trigger AMP.

Klemperer, P. (2002). "What Really Matters in Auction Design." *The Journal of Economic Perspectives* 16: 169-189.

Klemperer analyzes auction format with special attention to telecommunications as an operating example. He argues that the key concerns about auction design can be learned from elementary economics and strategic behavior. He examines collusion, entry-detering behavior and predatory behavior. He suggests that most literature focuses on issues that are second-order for practical design, that is a fixed number of noncooperative bidders, and it emphasizes effects such as risk aversion, correlation of information, budget constraints and complementarities. However, these may be more important issues for allowance markets, compared to telecommunications markets, which have fewer bidders and fewer goods (licenses) to be auctioned. He finds that ascending and uniform-price auctions are both very vulnerable to collusion and efforts to deter entry. He nominates a final sealed-bid stage into an otherwise ascending auction to create an Anglo-Dutch auction to address this. This design resembles the "shot clock" auction design we have modeled.

The concern about tacit or explicit collusion has been important in the multiunit (simultaneous) ascending auctions. When there are limited buyers, the ascending clock provides information that helps collusion. This was seen in the German experience with the spectrum auction in 1999. There is also evidence of collusion

in US markets, because of repetition in different geographic areas, which provides the ability to retaliate.

A frequently repeated auction, sometimes called a repeated stationary auction, is particularly vulnerable to collusion. Electricity markets are a good example. In the case of allowances, a lesson here might be that absent other good reasons to have frequent auctions, fewer auctions would help deter collusion.

Ascending auctions are especially bad at attracting bidders (Bulow and Klemperer, 1996). There is a strong presumption that the firm that values the item the most will be the ultimate winner, so competitors are not enticed to enter. Other auction forms can have similar problems if there are great asymmetries among bidders and entry costs are large. An example is the UK 1991 sealed-bid auction of television franchises.

The winner's curse can depress bidding in some ascending auctions, when bidders have close to common values for the item being auctioned and when there is some uncertainty about its actual value. The winner's curse affects weak firms more than strong ones. Hence the advantaged bidder usually wins, and can pay a low price. This is illustrated by the 1995 auction in Los Angeles for mobile-phone licenses. A strong bidder also has the incentive to establish a reputation for aggressiveness. If this effect is present in an allowance auction, it suggests that smaller firms will defer to the secondary market or to brokers, and the difference in price between the auction and secondary market will earn rents to large firms that bid in the auction.

Many poor experiences in ascending price auctions were aggravated by the failure to set a proper reserve price. Inadequate reserve prices increase the incentive for predation and may encourage collusion. Not only are serious reserve prices opposed by bidders, but often by politicians who fear the embarrassment of not selling the item.

Similarly, sealed-bid auctions can also be embarrassing. One example was discovered by the BSCH (Spain's biggest bank) when they won an auction for the Sao Paulo bank Banespa at three times the bid of the runner up. So firms may oppose first-price auctions. But the converse second-price auction would be embarrassing for the auctioneer in this case, as occurred in New Zealand, and this could be remedied by a reserve price.

Loopholes and special strategic opportunities can plague auctions. Test-bedding is essential to discover these loopholes. The Turkey auction of telecom licenses sequentially is an example of this. In the US bidders have won spectrum auctions but have defaulted on their commitments after long delays. In these cases the default penalties were small and bidders are bidding for options on prizes rather than the prizes themselves.

Ascending auctions are subject to rule breaking by bidders, because it allows cheat time.

Auction design may be less important when there is a large number of bidders for whom entry is easy. The Treasury auctions are an example. Experiments with different kinds of auctions have led to inconclusive results.

The author offers solutions, one of which is to make ascending auctions more robust. An ascending auction may succeed in allocating to the bidder who values an item most. It also helps bidders learn about the market by inference of the value to others. To avoid signaling, bidders can be forced to bid round numbers. Keeping secret the numbers of bidders remaining makes collusion harder.

In sealed bid auctions firms are unable to retaliate and collusion is difficult. However the advantaged bidder will probably win, but it must make its single offer in the face of uncertainty about its rival's bids. Weaker bidders hence have some chance. They are more attractive to entrants. Also, the winner's curse is less severe in the case of common values. However, by giving some chance of victory to

weaker bidders, the sealed bid auction is less likely to lead to efficient outcomes. Also, bidders need good private information about value. Hence pay-your-bid discriminatory auctions may discourage bidders. The entry problem is less serious when small bidders can buy from intermediaries, such as brokers, who can aggregate smaller bidders demands and bid in their place as occurs in auctions of Treasury bills.

The Anglo-Dutch (or shot clock in our terminology) is a suggested remedy bringing the best of both auction types. An auction with similar features is the OpenBook auction for corporate bonds. eBay auctions also have this feature, with an ascending price that rushes toward closure so bidders have one last opportunity to bid their best and final offer. This approach will repel collusion, and encourage entry. But it also is more likely to sell to the highest valued buyer as will an ascending auction.

The author notes in closing that most auctions work well, and even cases where there have been problems probably are better outcomes than the administrative “beauty contest” alternative.

One size does not fit all. Auction formats should be tested.

Kline, J. J. and F. M. Menezes (1999). "A simple analysis of the US emission permits auctions." *Economics Letters* 65(2): 183-189.

This paper focuses exclusively on a stylized version of the EPA SO₂ auction method and uses it to prove two propositions under complete information. The propositions are: “that there are either inefficient equilibria (where no goods are exchanged) or efficient equilibria (where all possible gains from trades are realized). The efficient equilibria have the property that all trades occur at a uniform price.”

Two examples are also provided where the participants are under incomplete information. The first case results in the sellers shading their bids up when both buyers and sellers are behaving strategically. The other case results in both buyers

and sellers shading their bids down. It is an important finding that under some circumstances sellers will inflate their asking price when behaving strategically, which is counterintuitive to what one might expect.

List, J. A. and D. Lucking-Reiley (2000). "Demand reduction in multiunit auctions: Evidence from a sportscard field experiment." *American Economic Review* 90(4): 961-972.

This paper uses a field experiment to test the theory of Multi-units auctions.

Field experiments differ from lab experiments in that bidders are confronted with a "real" auction, where if they win they are awarded the object for sale (rather than some induced monetary payoff). Also unlike lab experiments, there is little common knowledge among bidders (bidders for instance typically do not know the underlying distributions of other bidders' values). The drawback is that there is less control of the variables at play.

In this experiment the objects for sale are sportscards. Some units are auctioned via a uniform price auction, while others via a Vickrey auction. The specific setting is one with two units on sale and two bidders. In some sessions experienced bidders are used.

The main results are the following. As predicted by the theory, the demand reduction phenomena is present much more strongly in the uniform price auction, which displays many more zero bids for the second unit. Due to the demand reduction phenomena, the uniform price auction is less efficient. Contrary to the theory prediction that bidders should bid their value for the first unit, some overbidding is observed in the uniform price auction.

Overall, the differences in terms of revenues are small so that, given the higher efficiency it displays, the Vickrey auction seems to be the preferable mechanism (at least for the setting proposed).

Lucking-Reiley, D. (2000). "Auctions on the Internet: What's Being Auctioned, and How?" *Journal of Industrial Economics* 48(3).

This is an interesting survey of selling mechanisms used on the internet. Some of the main issues coming from internet auctions can be addressed using standard auction theory. Some new features of internet auctions, however, bring new and interesting challenges some of which are highlighted in this work.

Mandell, S. (2005). "The choice of multiple or single auctions in emissions trading." *Climate Policy* 5(1): 97-107.

The main issue addressed by this article is the frequency of CO₂ permit auctions. One of the auction frequencies considered is the 'single-auction approach', in which a single auction is held at the beginning of a commitment period to sell the entire volume of allowances for that period. The alternative is the 'multiple-auction approach' in which several auctions are used throughout the commitment period to sell the volume of allowances. The primary contribution of this article is to address auction frequency in the context of the 'winner's curse'.

The author makes some arguments for why an auction is preferable to any system of free allowance allocation. Two assumptions underpin the discussion of auction frequency. First, the market for CO₂ permits is 'small'. Second, any CO₂ allowance auction is run as an ascending clock auction. The author acknowledges that ascending auctions for multi-unit goods may yield inefficient prices when large bidders choose to shade their bids, but this concern vanishes when the number of bidders is 'large' or the secondary market is competitive. To the extent that either of these is true, the bid shading problem is more a question of wealth distribution than efficient allocation.

A literature review reveals two pre-existing ideas on auction frequency. The first is that higher frequency allows firms a shorter planning horizon when bidding into each auction. This benefit of higher frequency auctions is undermined by an efficient secondary market for allowances. The second benefit of high frequency auctions is cash-flow management, but it is undermined by a perfect market for capital.

A perfectly competitive secondary market for CO₂ allowances removes the benefits of a multiple-auctions approach to initial allowance allocation. However, the

market will not be perfectly competitive if it provides too little price information (e.g. prices are confidential) or it is illiquid.

The ‘winner’s curse’ may arise in an auction for a good in which all bidders will value the good equally, but with uncertainty. Consider the bidders’ expectation of the good’s value to be i.i.d. with mean at the true value. The winner of the auction will, upon winning the auction, expect that his valuation of the good must have been too high since all others valued it lower. Formally, “the bidder’s expected value of the good prior to the bidding process is larger than the expected value conditional on winning.” The bidders’ recognition of the winner’s curse will cause them to adjust their bids downwards. The key is uncertainty. With no uncertainty, the winner’s curse vanishes. As uncertainty grows, so too does the impact of the winner’s curse.

To put the winner’s curse back into the context of CO₂ allowance auction frequency, we must reconsider the secondary market for allowances. If it is perfectly competitive, then there is no uncertainty about allowance valuation and the winner’s curse does not exist. To the extent that the secondary market is imperfect, the frequency of allowance auctions can affect the information available to the bidders. Increasing auction frequency improves market information and minimizes the effect of the winner’s curse.

Allowance auction frequency has two other notable impacts. More frequent auctions yield higher overall transaction costs and are more vulnerable to collusion. Higher transaction costs are an obvious result of frequent auctions. Greater vulnerability to collusion results from the opportunity to follow through with threats of punishment more quickly.

High frequency allowance auctions can offset the problems presented by the ‘winner’s curse’, but at the cost of higher transaction costs and greater vulnerability to collusion. Since the winner’s curse only emerges in the presence of an imperfect secondary market for allowances, the viability of frequent auctions only emerges in the same case. The author assumes that secondary markets for CO₂ allowances will

evolve over time. He therefore asserts that, "A plausible policy recommendation... would then be to use frequent auctions during the early years of the scheme with the intention of decreasing the frequency in future stages..."

Margolis, M. and J. F. Shogren (2004). "Implementing the efficient auction: initial results from the lab." *Economics Letters* 84(1): 141-147.

Vickrey's second-price auction (for 1 good) is demand revealing and efficient in theory. As such, it is not directly relevant to an allowance auction. A variant to Vickrey's auction has bidders with affiliated values, i.e. the value to one bidder depends in part on information available only to some other bidder. This auction is neither demand revealing nor efficient. Dasgupta and Maskin (2000) constructed a generalized version of the Vickrey auction, called an efficient auction, in which each bidder expresses his bid as a function of other bidders' signals. This auction is theoretically efficient even when the bidders have affiliated values. This may have some relevance to allowance auctions as bidders in these markets may have affiliated values.

This paper experimentally assesses the efficiency of an efficient auction with inexperienced bidders. The finding is that bids are systematic, but yield a bid curve flatter than the truthful one.

Maskin, E. (2003). *Auctions and Efficiency*, Cambridge University Press.

This is a nice survey regarding the important issue of allocating efficiently an object by means of an auction or some other mechanism. It presents the main results available in the literature (up to his publication) and pinpoints some of the questions that remain to be answered.

Maskin, E. and J. Riley (2000). "Asymmetric Auctions." *The Review of Economic Studies* 67(3).

The vast majority of models in auction theory assume that bidders are perfectly symmetric. This is an interesting theoretical benchmark to begin with, but cannot adequately represent many practical applications, especially if the asymmetries are expected to be strong.

The problem with asymmetric models of auctions is that deriving the equilibrium bidding behavior is often too a difficult task.

This paper looks at the private value model and separately introduces three different and very specific types of asymmetries regarding the distribution from which the bidders' private information is drawn. It shows that the ranking of First Price Auction and open auction depends critically on the type of asymmetry taken into consideration. It thus provides some policy recommendations for a seller who has some knowledge of the type of asymmetry bidders display (note that Revenue Equivalence does not hold as bidders' beliefs are asymmetric).

McAfee, P. and J. McMillan (1987). "Auctions with a stochastic number of bidders." *Journal of Economic Theory* 43(1): 1-19.

Most of the models in auction theory assume that the number of bidders participating in the auction is fixed and known by all participants. This assumption is not realistic in many applications, however. This is one of the first papers to assume that the number of bidders is stochastic.

It shows that if bidders are risk averse, the revenues for the seller are higher when a bidder perceives the number of his opponents as uncertain. It therefore suggests that, when possible, the seller should conceal the number of participants.

McMillan, J. (1994). "Selling Spectrum Rights." *Journal of Economic Perspectives* 8(3): 145-62.

This paper nicely introduces the reader to the most important issues in the practical design of auctions taking as a leading example the experience of the 1993 FCC spectrum license auction in the US. It points out the key elements to consider in the design of such a complex auction and highlights why auction theory has been so successful in providing useful insight to practitioners.

The author was himself hired by the FCC for the designed of the mentioned auction. Many other leading auction theorists were hired by the phone company bidding in the auction.

Milgrom, P. and R. Weber (1982). "A Theory of Auctions and Competitive Bidding." 50(5): 1089-1122.

This paper introduces "common values," where all bidders derive the same value from the object being sold but no bidder knows its exact value (e.g. bidding for the rights to drill for oil on a certain tract). The authors derive the theoretical equilibria of the first, second, and open ascending (English) auctions, and show that the English auction yields more revenue than a second-price, which in turn yields more revenue than a first price auction. This continues to be one of the most influential papers in auction theory, and one of the most cited.

Miller, G. J. and C. L. Plott (1985). Revenue-Generating Properties of Sealed-Bid Auctions: an Experimental Analysis of One-Price and Discriminative Processes. *Research in Experimental Economics*. V. L. Smith. Greenwich, CT, JAI Press Inc. 3: 159-182.

Researchers performed an auction experiment comparing the performance between discriminative and one-price auctions under several induced demand structures. Of particular interest was the relative revenue generating capacity of the two auction types. It was found that under steeper demand conditions, the discriminative auction generated more revenue; while under flatter demand conditions, the one-price auction generated more revenue.

The demand conditions were categorized as perfectly flat (PF), flat (F), moderately steep (M), and steep (S). In each case, the equilibrium price was the same value. Each experimental session had multiple individual auctions and the induced demand shifted twice during every session except for one, shorter session.

The results are summarized as follows:

- The revenue-generating capacity of each type of auction is sensitive to the slope of the demand function.

- With steeper demand conditions discriminative price auctions generate more revenue than one-price auctions.

- Under flatter demand conditions, one-price auctions raise more revenue than discriminative auctions.

- For almost all units auctioned, bids were higher under one-price conditions.
- The variance of the bids was higher (especially in later periods) in the one-price conditions in all but the five paired periods with flat demand.
- In the one-price auctions, the price was near the competitive equilibrium and the degree of demand revelation increased to nearly perfect revelation.
- In the discriminative auctions:
 - The lowest accepted bid converged on the equilibrium price.
 - The accepted bids tended toward the equilibrium price over time.
 - The bids on the extramarginal units converged on the limit price (full revelation).
 - When there is a moderately sloped demand curve, average price may stay systematically and significantly above the equilibrium price even though an increasing number of individual bids are tending toward equilibrium price.

Muller, R. A., S. Mestelman, et al. (2002). "Can double auctions control monopoly and monopsony power in emissions trading markets?" *Journal of Environmental Economics and Management* 44(1): 70-92.

This paper experimentally tests whether the double auction market can somehow mitigate the effects of market power. The setting the authors try to replicate is the one of the trading of pollution emissions rights.

Market power there is likely to be relevant with a country such as the USA having a predominant role as a buyer. Different market structures are used. In particular both the monopolist benchmark case and the monopsony are covered.

The main finding is that traders are able to exploit their market power to increase their profits compared to the competitive benchmark. This is done achieving some degree of price discrimination. Thus the efficiency level is not severely affected and it is not far from the competitive one.

From a policy point view, the second conclusion is to be seen with favor, the first one though may be politically undesirable.

Myerson, R. (1981). "Optimal Auction Design." (6): 58-73.

This paper is one of the milestone contributions in the theory of auctions. It solves the seller revenue maximization problem subject to the bidders incentive compatibility constraints, thus providing the optimal allocation and the corresponding payment scheme. In solving such problem, a corollary result is the important revelation principle that allows one to restrict attention to direct mechanisms.

The study yields important practical insights. In particular, it shows that the optimal mechanism involves the use of a reserve price. It also shows that if bidders are asymmetric the optimal mechanism should favor the weaker bidders.

Finally, it provides a formal proof of the revenue equivalence theorem, anticipated, but not formally proven, by Vickrey in his 1961 classic (with this respect see also Riley and Samuelson (1981)).

Paarsch, H. (1992). "Deciding between the common and private value paradigms in empirical models of auctions." 51(1-2): 191-215.

This is one of the early important contributions to empirical auctions. The theory of auctions yields different predictions depending on whether the object for sale is assumed to fit the private value paradigm or the common one.

It is extremely important that the seller is aware of which of the private or common value component is more relevant. This might be difficult to assess in general.

This paper tries to use the theoretical bidding behavior to assess the actual bidding one and empirically establish whether the particular market they look at fits better the private value paradigm or the common value one.

Pagnozzi, M. (2007). Should speculators be welcomed in auctions?

This paper looks at the effect of introducing resale in the sale of multiple-objects via a uniform price auction. The possibility of being able to buy the object in the

resale market is shown to affect the incentive for demand reduction (see Ausubel and Crampton (1998)) making it stronger and thus suppressing revenues.

It then looks at the effect of introducing the presence of speculators that attach no value to the object but may be willing to pay a positive price for the object in order to sell it in the resale market. The presence of speculators has the positive effect of increasing competition. This fact needs to be traded-off with the demand reduction effect.

Pagnozzi shows that sometimes a strong bidder should optimally choose an accommodating strategy and let the speculator win some units. This decreases revenues. Such accommodating behavior is not always optimal. It is shown that the effect is in general ambiguous and depends on how clustered bidders valuations are.

Palfrey, T. (1983). "Bundling Decisions by a Multiproduct Monopolist with Incomplete Information." *Econometrica* 51(2): 463-483.

This paper is one of the earliest contributions on multi-object auctions. It compares the monopolist choice of selling different objects via separate auctions versus bundling them all and selling them via a single unit auction.

An important underlying assumption is that the value for the bundle equals the sum of the values for the single objects. The main result is that with few buyers the seller should optimally bundle the objects. This makes all bidders worse off. Conversely, if there are many buyers, he should sell the objects separately. In such case bidders with high demand are worse off in that they would prefer to be proposed the bundle, while low demand bidders are better off.

For more recent contributions on bundling, see Armstrong (2000) and Jehiel et al. (2007).

Persico, N. (2000). "Information Acquisition in Auctions." *Econometrica* 68(1): 135-148.

This paper studies the impact of information acquisition (prior to the auction) on the revenue performance of two auction formats: the First Price Auction (FPA) and

the Second Price Auction (SPA) (the paper has also more general results but the auction environment represents his main application).

The setting studied is one in which bidders' information is affiliated (a form of positive correlation). Milgrom-Weber (1982) in this setting rank the SPA as superior to the FPA in the absence of information acquisition. Persico shows that the incentive to acquire information is greater in the FPA and that sometimes such effect is sufficient to reverse the ranking.

The basic intuition is rather simple. If values are correlated getting a more precise signal enables a bidder to have a better estimate of his opponents values. This allows him to leave on the table only the amount that is really needed to win. Such information is less valuable in the SPA as there the amount paid is independent of ones own bid.

Plott, C. R. (1983). "Externalities and Corrective Policies in Environmental markets." *The Economic Journal* 93(369): 106-127.

This paper compared three policy measures using laboratory experiments. The policies were a tax, standard and pollution license. The experiments examined price behavior, efficiency and distributional consequences. In the absence of policies, subjects ignore the externality in their private market behavior. The key result is that experiments confirm economic theory in a variety of ways, including specifically the internalization of social costs when taxes and pollution licenses are used. The most efficient policy was the pollution license, and second most efficient was the tax. The paper allocates the emissions licenses in a somewhat random way so that the secondary market plays an important role in their use.

Porter, D., W. Shobe, et al. (2007 forthcoming). "The Design, Testing, and Implementation of Virginia's NOx Allowance Auction." *Journal of Economic Behavior and Organization*.

Researchers compared the performance of three different auction designs for Virginia's NOx Allowance Auction.

The Allowances: Researchers explored the dynamics of auctioning two different allowance vintages (2004 and 2005). The two allowance vintages were asymmetric substitutes because allowances from 2004 could be used for compliance in 2005, but the reverse did not hold. In addition, by law, the number of allowances carried over from 2004 to 2005 could not exceed 10% of the regional budget for 2005. Thus, if the number of banked allowances exceeded the limit, firms could only use a fixed proportion of them for compliance. For example, if regional banking in 2004 was 15% of the 2005 budget, firms could only use two thirds (10%/15%) of their banked allowances for 2005 compliance.

The Auctions Considered: Experimenters looked at 3 different auction types: Combinatorial Sealed Bid (CSB), Sequential English Clock (SEC), and Combinatorial English Clock (CEC).

In the CSB auctions, bidders submitted bids in the format (p4, Q4| p5, Q5), where p4 is the price per allowance that the bidder was willing to pay for up to Q4 2004 allowances, and p5 was the price per allowance that the bidder is willing to pay for up to Q5 2005 allowances. Units were allocated to the highest bidders and they paid the price that they bid.

With the SEC auctions, one year's allowances were sold during one English Clock Auction (ECA) and the other year's allowances were sold in a separate (ECA) shortly following the first.

For the CEC auction researchers ran two clock auctions simultaneously, one for each vintage. Bidders could substitute between vintages as long as: when switching from a lower to a higher priced vintage, the quantity of the higher priced vintage was limited to the quantity currently demanded at the lower price; while when switching from a higher to a lower priced vintage, the budget for the of the lower price vintage could exceed the amount of money committed to the current bid.

Results: The CEC auction outperformed the other two auction types. In terms of efficiency, SEC and CSB were indistinguishable. In terms of revenue, both clock auction types did better in an elastic environment; while in an inelastic environment, revenue was reduced in across all mechanisms with a stronger effect on the clock auctions than CSB.

Rassenti, S., V. Smith, et al. (2002). "Using Experiments to Inform the Privatization/Deregulation Movement in Electricity." *Cato Journal* 21(3): 29.

This survey article summarizes a body of work done by the authors and others in the late 1980s and 1990s related to the question of efficiency gains from restructuring or liberalizing electricity markets. A main focus of the Rassenti and Smith papers summarized here was testing the feasibility of trading electricity in decentralized markets where offers and bids are submitted to a centralized computer program which uses an algorithm that maximizes the gains from trade, given the physical limitations of the electricity grid and associated electricity losses, to decide who sells and who buys and the market clearing prices at each node on the grid. Prior to this work, the conventional wisdom (Joskow and Schmalensee 1983) was that decentralized bidding might not be appropriate in electricity markets because of the externalities imposed by the grid. Experiments showed that decentralized markets would achieve 90 – 100% efficiency as a result of simultaneous determination of allocations, which means that each agent bears the opportunity cost imposed on others by its actions at the margin.

The second question that was addressed in this series of experiments was regarding the importance of allowing demand side bidding on the efficiency of this decentralized market approach. They found that adding demand side bidding brought prices in shoulder and off-peak periods much closer to competitive prices and reduced the volatility of electricity prices across days.

Rassenti, S., V. Smith, et al. (2003). "Controlling Market Power and Price Spikes in Electricity Networks: Demand-side bidding." *Proceedings of the National Academy of Sciences* 100(5): 7.

This article looks at the effect of ownership concentration and demand side bidding on price realization and efficiency of electricity spot markets using experimental methods. The exercise uses a simple representation of electricity markets and the transmission grid with a three node radial network, negligible transmission losses,

no obligation to serve on the part of load serving entities and no reserve market. The ownership of generation assets in the “market power” treatment is designed in such a way that particular generators can increase their profits by increasing their bids or withholding capacity. In another treatment, the “no power” treatment, generation ownership is redefined in such a way that the ability to exercise market power is no longer possible. In these experiments, subjects who were suppliers submit offers to supply powers that are expressed as a step function that indicates the amount of power they are willing to produce and sell at each price. With no demand side bidding, the demand side of the market is represented by a willingness to pay schedule within the software. With demand side bidding, real load serving agents participate as buyers in the market. The market clearing price at each node in the transmission grid is found by identifying the combination of bids that maximizes the gains from trade in electricity.

The main results of this article are that the market power treatment results in substantially higher prices in shoulder and off-peak periods, but adding demand side bidding neutralizes market power. In the no-power treatment, demand side bidding reduces prices to close to the 100% efficient levels.

Robinson, M. (1985). "Collusion and the Choice of Auction." *The RAND Journal of Economics* 16(1): 141-145.

This short paper clearly points out that when the possibility that bidders may form a collusive cartel is a concern, the use of an open format may not be ideal. In fact, such a format intrinsically provides bidders with a better institution to credibly implement a collusive agreement compared to a sealed format such as the First Price Auction (FPA).

The insight is rather simple. The open format allows the other members of the cartel to observe a bidder not respecting the agreement and allows them to react and punish such bidder. Unless the same auction is repeated many times, the possibility of punishing a deviator is not allowed by the sealed format, which in turn discourages the formation of such agreement.

For a more formal study of collusion in open auctions, see Graham and Marshall (1987). For a study of collusion in FPA, see McAfee and McMillan (1992).

Staropoli, C. and C. Jullien (2006). "Using Laboratory Experiments to Design Efficiency Market Institutions: The Case of Wholesale Electricity Markets." *Annals of Public and Cooperative Economics* 77(2): 23.

This paper surveys the main contributions of experimental economics to the design of wholesale electricity markets. It says nothing of allowance markets. Experiments on market design should be taken as a complement to theoretical market design.

The first part of the paper addresses experimental contributions toward the understanding of 'the general architecture of the market' (Wilson, 1999). The second part addresses the details of the auctions rules in the insights garnered from experimentation.

General Architecture

The work of Vernon Smith of the U of Arizona in the mid-1980s found that experimental markets figuring energy sales and purchases expressed as 'offers to sell' and bids to buy', with simultaneous determination of allocations and physical constraints imposed by the grid, are feasible and efficient. This peaked interest and led to more studies.

Transmission constraints

The three main issues that arise from the possibility of transmission constraints are monopoly power of the owner of wires, local market power of generators, and the allocation of ownership rights to use the network.

The vertical disintegration of a utility that own generation capacity and transmission capability modifies incentives and may lead to a distortion of information on congestion to manipulate the expectations of producers. Beckerman

et al (2000) conducted an experiment to assess who, among the supply side, demand side, and transmission owners, can capture the rents from such distortions. They also addressed how the distribution of rents varies under alternative auction rules. The experiment uses a uniform price double auction mechanism. The two variants of the auction mechanism are ‘both-sides rule’ that gives the opportunity to any market participant to accept any offers on each side of the market before the market is called, and ‘other-side rule’ in which in order to have an offer accepted, each participant must meet the terms of the unaccepted bid or ask on the other side of the market. In theory the transmission owners will capture congestion rents. Under experimentation, generators capture some of the congestion rents and demanders are unaffected by who receives the rents. Also, ‘both-sides rule’ is more efficient than ‘other-side rule’. Staropoli et al. conclude that “this experiment suggests that in this environment, no incentives are given to transmission line owners to invest in new transmission capacities as they do not capture the rents in the system.”

Transmission constraints can create local market power in which generation on one part of the grid can have consequences for generation scarcity on other parts of the grid. Those in small areas of generation scarcity may exercise market power to yield inefficiently high prices. Zimmerman et al. (1999) experimentally showed that using a uniform price sealed-bid auction with two competing generators in the area of generation scarcity will yield market prices close to duopoly levels.

The ownership rights on the transmission grid matter, especially at points of congestion. There are two types of property rights, as defined by Kench (2004): financial and physical. Financial rights entitle the owner of a wire to collect congestion rents across it. Physical rights authorizes owners to send power through a line. Kench experimented with a model market governed by a continuous double auction, where both the demand-side and the supply-side are active. He finds that the assignment of either type of right yields greater efficiency than a failure to assign either right to anyone. At points of congestion, the assignment of physical rights yields a more efficient equilibrium than the assignment of financial rights.

Demand-side participation

Demand for electricity has long been considered quasi-inelastic, but real demand response in electricity markets could bring discipline to generators. Rassenti et al (2002) tested this and concluded that active participation on the demand side “neutralizes market power and price spikes on peak in the laboratory as well as it lowers prices”.

Auction Details

Electricity markets are inherently complicated. The types of auctions that might govern them vary along four dimensions: single unit vs. multiple unit, uniform pricing vs. discriminatory pricing, sealed-bid vs. sequential-bid, and repeated vs. non-repeated. The theoretical analytics of the entire four-dimensional matrix of auction options remain incomplete. The experimental analyses of these addresses two of the dimensions: uniform vs. discriminatory pricing and sealed vs. sequential bids.

Sealed-bid vs. Sequential-bid

Bernard et al. (1998) compares two uniform price auction rules: last accepted offer (LAO) as a sealed-bid auction, first rejected offer (FRO) as a sequential-bid auction. In a single unit auction FRO is incentive compatible whereas LAO is not. In a multiple unit auction the incentive compatibility of FRO is lost.

Denton et al. (2001) show that when generators are endowed with identical portfolios, i.e. the auction acts like a single unit auction, a sealed-bid offer (SBO) rule is significantly more efficient than a real-time uniform price double auction (UPDA) rule.

Uniform price vs. Discriminatory price

Many different studies have shown that discriminatory auctions do not perform as well as uniform price auctions in electricity markets. Specifically, Hahn and Van Boenig (1990) showed that SBO beats the split-saving rule (SSR) in terms of price outcome deviations from the competitive equilibrium. Olson et al. used experimentation to compare a day-ahead sealed-bid auction (SB) and a

simultaneous continuous double auction (CDA). The SB was slightly more efficient than the CDA. Moreover, the CDA equilibrium yielded higher prices than the SB auction. So under a CDA institution, the consumers of electricity are nearly as efficient as under an SB institution, but more poorly so.

Vickrey, W. (1961) "Counterspeculation, Auctions, and Competitive Sealed Tenders." *Journal of Finance* 16(8-37).

This is the pioneer work in auction theory and as such is among the most cited ones. It introduces the independent private value model under which each bidder's valuation is independent of the information held by his opponents. It provides the equilibrium bidding behavior for open ascending price auction (English Auction), Second Price Auction and First Price Auction (for the case of uniform distributions).

The Second Price Auction is first introduced in this paper. Vickrey is the first to note the importance of making the winner's payment independent of his bid in order to induce truthful bidding and an efficient outcome. This is the main characteristic of the SPA and provides a fundamental lesson for auction theorists.

He is also the first to point out that in independent private values the revenues generated by all those formats mentioned above coincide. Such a remarkable result, known as revenue equivalence theorem, is proven more generally (and formally, showing under which assumptions holds true) in subsequent papers.

Wolfram, C. (1998). "Strategic Bidding in a Multiunit Auction: An Empirical Analysis of Bids to Supply." *Rand Journal of Economics* 29(4): 703-725.

This is an empirical analysis of bidding behavior using data from the daily electricity auction in England and Wales. Wolfram finds that companies do strategically manipulate their bids in order to raise the price that they are paid for inframarginal capacity. Theory predicts that a firm's incentive to strategically raise their bid goes up as the number of inframarginal units that they are bidding on goes up; but that this effect is tempered by the fact that bidding high reduces the chances of that bid being marginal. Wolfram found that a bigger company did have higher markups because it had more inframarginal bids. However, the larger the individual plant that the firm was bidding on, the less likely they were to markup because the

losses from a larger plant not being included was greater (evidence of the moderating effect). Bids are higher for a given unit if more of the units are likely to run before that one is available for supply. There were only small differences in the operating costs between the small and large firm, so not much efficiency was lost by strategic bidding.

Zheng, C. (2002). "Optimal Auction with Resale." *Econometrica* 70(6): 2197-2224.

Myerson (1981) solves the seller revenue maximization problem and determines the optimal allocation from the seller point of view. This paper adds the often realistic assumption that the winning bidder may try to resell the object to some of the losing bidders.

The paper defines conditions under which Myerson's (1981) original optimal allocation can still be achieved when resale cannot be prohibited. This is done for the two bidder case. For the generic n bidder case the result can be proven only for some special cases.

Zimmerman, R., J. Bernard, et al. (1999). *Energy Auctions and Market Power: An Experimental Examination*. 32nd Hawaii International Conference on System Sciences.

The research reported in this paper has two parts. The first looks at the performance of different auction types with markets of different sizes and numbers of competitors. The auction types analyzed include the last accepted offer (LAO) version of the uniform price auction, the first rejected offer (FRO) form of the uniform price auction and a multiple unit Vickrey auction, under which winners pay the opportunity cost they impose by being in the auction. These auction types were tested in settings with 2, 4 and 6 subjects, each offering to supply electricity. Demand was assumed to be perfectly inelastic. All auctions included a reservation price. Optimal prices in these auctions were defined as the equilibrium prices that would result if all participants offered all of their capacity at its marginal cost. The optimal price depends on the auction type, with slightly lower offers potentially prevailing in the last accepted offer auction than in the first rejected offer auction, but does not vary with number of participants in the auction. However, the optimal price does vary with the number of participants in the Vickrey auction. The results show that group size is a much more important determinant of price outcome than auction type with the price in a two party auction being nearly double the

competitive level. In general, the LAO auction type, the form often used in wholesale electricity spot markets, performs slightly better than the FAO type. The LAO and FRO auctions had similar effects on getting bidders to reveal their true costs but with the Vickrey auction low cost units tended to bid under cost.

The second issue is the effect of a transmission network with a single auction type, the LAO. In these experiments, the authors constructed a network with 30 nodes and transmission constraints between regions that gave two of the six generators market power in a particular part of the grid as given transmission constrains those generators must operate to meet demand. In the experiments they found that in most sessions the generators with market power were able to coordinate their bidding strategies and exploit that market power. In one session this took a while and in others, including those that involved professional electricity traders instead of students, it happened right away. In one session, prices remained near competitive levels throughout the 75 round auction. The authors also demonstrate that voltage limits and reactive power requirements can create opportunities for market power on their network.