# Multiagent System Simulations of Sealed-Bid, English, and Treasury Auctions

by

Alan Mehlenbacher
B.S., University of Michigan, 1968
M.Sc., University of British Columbia, 1970
M.B.A., Simon Fraser University, 1993

A Dissertation Submitted in Partial Fulfillment of the Requirements for the

Degree of

DOCTOR OF PHILOSOPHY

in the Department of Economics

**Multiagent System Simulations of Sealed-Bid, English, and Treasury Auctions**

by

Alan Mehlenbacher
B.S., University of Michigan, 1968
M.Sc., University of British Columbia, 1970
M.B.A., Simon Fraser University, 1993

**Supervisory Committee**

Dr. David Scoones, Department of Economics
**Supervisor**

Dr. Donald Ferguson, Department of Economics
**Departmental Member**

Dr. Linda Welling, Department of Economics
**Departmental Member**

Dr. Tony Marley, Department of Psychology
**Outside Member**

Dr. Jasmina Arifovic, Simon Fraser University, Department of Economics
**External Examiner**

## ABSTRACT

**Supervisory Committee**

Dr. David Scoones, Department of Economics
**Supervisor**

Dr. Donald Ferguson, Department of Economics
**Departmental Member**

Dr. Linda Welling, Department of Economics
**Departmental Member**

Dr. Tony Marley, Department of Psychology
**Outside Member**

Dr. Jasmina Arifovic, Simon Fraser University, Department of Economics
**External Examiner**

I have developed a multiagent system platform that provides a valuable complement to the alternative research methods. The platform facilitates the development of heterogeneous agents in complex environments. The first application of the multiagent system is to the study of sealed-bid auctions with two-dimensional value signals from pure private to pure common value. I find that several auction outcomes are significantly nonlinear across the two-dimensional value signals. As the common value percent increases, profit, revenue, and efficiency all decrease monotonically, but they decrease in different ways. Finally, I find that forcing revelation by the auction winner of the true common value may have beneficial revenue effects when the common-value percent is high and there is a high degree of uncertainty about the common value. The second application of the multiagent system is to the study of English auctions with two-dimensional value signals using agents that learn a signal-averaging factor. I find that signal averaging increases nonlinearly as the common value percent increases, decreases with the number of bidders, and decreases at high common value percents when the common value signal is more uncertain. Using signal averaging, agents increase their profit when the value is more uncertain. The most obvious effect of signal averaging is on reducing the percentage of auctions won by bidders with the highest common value signal. The third application of the multiagent system is to the study of the optimal payment rule in Treasury auctions using Canadian rules. The model encompasses the when-issued, auction, and secondary markets, as well as constraints for primary dealers. I find that the Spanish payment rule is revenue inferior to the Discriminatory payment rule across all market price spreads, but the Average rule is revenue superior. For most market-price spreads, Uniform payment results in less revenue than Discriminatory, but there are many cases in which Vickrey payment produces more revenue.

# Table of Contents

**List of Tables**

**List of Figures**

**Acknowledgements**

I am indebted to David Scoones, Linda Welling, Don Ferguson, and Tony Marley for their valuable support during the course of this research. Specifically, I thank Don Ferguson for first alerting me to the potential for multiagent simulations in Economics and his subsequent teaching and advice on computational modelling in general, Linda Welling for sponsoring my first course in auction theory, David Scoones for ongoing advice and encouragement during the thesis project and for encouraging me to apply to the Social Sciences and Humanities Research Council of Canada for a PhD scholarship, and Tony Marley for his valuable insights from the perspective of computational modelling in Psychology. I would also like to thank David Scoones, Linda Welling, and Don Ferguson for participating with me in numerous joint discussions over the past four years, David Giles for econometric advice and inspiration, and the other members of the Economics Department at the University of Victoria for interesting courses, seminars, and discussions.

# Chapter 1

# Introduction

## 1.1 Motivation

Auctions are complex economic mechanisms that are used for transactions worth trillions of dollars each year throughout the world. Beginning 2000 or more years ago with Babylonian auctions of wives and Roman auctions of property (Smith, 1968), auctions have expanded to include procurement auctions for government goods and services; government asset sales of timber licences, oil leases, telecommunication licences, and treasury securities; commercial sale and procurement of vehicles, flowers, fish, equipment, and wine; and online sales of consumer goods on eBay. The basic concept of an auction is that bidders make decisions about how much they value the auctioned object and then bid in a way that will enable them to obtain the object at a profit. Whether the value of the object is private for each bidder, common to all bidders, or a mixture of private and common values is critical for bid strategies and auction outcomes.

We study auctions to understand how bidders value objects, why they make the bids that they do, how they can improve their bidding, and which auction design results in the most benefit for the seller in a sale auction or the buyer in a procurement auction. The design of an auction can include decisions about how many objects are for sale and whether they are the same or different, how the bids are made (e.g., sealed-bid or open, ascending or descending price), how the payment is calculated (pay-your-bid, pay the average bid, pay the bid of the loser, etc.), and so on. Because of their importance and complexity, auctions are very interesting to study.

## 1.2 Methods

Most studies of auctions use mathematical analysis that involves optimization, order statistics, and supermodularity, and I have read dozens of very interesting papers and books that use these methods (e.g., Krishna, 2002; Milgrom, 2004). However, in order to achieve tractable results in the face of complexity, drastically simplifying assumptions must be made. As freely admitted by the mathematicians themselves (Milgrom, 2004, p. 22), these simplifications put into question the conclusions and predictions produced by mathematical analysis. One possible alternative is to collect auction data and analyze it statistically. I collected five years' worth of highway procurement auction data from Texas, Alberta, and Saskatchewan, and analyzed it using several advanced econometric methods. However, I was disappointed by the limited conclusions that I could draw because this approach is severely constrained by a lack of information about bidders' values. I therefore decided to use a computational agent model of the bidders (Chapter 2) that allows me to endow them with values known to me but not the other bidders, program them with complex auction mechanisms, and run simulations. The agents record data about their bidding decisions, thereby providing me with not only the auction results but also the means whereby the results were produced. This classic approach to science was articulated by Rosenbelueth and Wiener (1945), who made the distinction between formal (mathematical) and material (computational) models. However, while formal mathematical models have a problem with over-simplification, computational modellers must guard against making their model so complex that it confounds interpretation of the results. As Rosenblueth and Wiener pointed out, “The best material model for a cat is another, or preferably the same, cat.”

## 1.3 Learning

The bidding model is a learning model in which the agents learn by repeated experiences with the same auction mechanism. I experimented with several learning methods and then selected Selten’s impulse balance learning method (Ockenfels and Selten, 2005) using criteria explained in Chapter 2. I modified this method to make it suitable for agents learning how to bid in sealed-bid auctions (Chapter 3), and I extended it to agents learning how to average value signals in English auctions (Chapter 4) and to agents learning how to bid for different types of securities in Treasury auctions (Chapter 5). I show in all cases that the learning models result in convergence to steady-state bid prices and bid quantities. These are critical results because it is impossible to interpret bidder profit and auctioneer revenue when there is no convergence in the bidding strategies.

## 1.4 Contributions

The first set of contributions are the development of the multiagent system platform and the adaptation of Selten’s impulse balance learning method to computational models, described above in Sections 1.2 and 1.3 respectively. The next contribution concerns sealed-bid auctions with bidder values that range from private to common. It is recognized that this is an important issue that can drive the results of an auction, and that valuations are usually a mixture of private and common values. However, nearly all theoretical studies consider pure private values, and nearly all human experiment studies consider pure private or pure common values. In the agent environment, I can easily give the agents values that have varying mixtures of private and common values. For sealed-bid auctions (Chapter 3), several auction outcomes are significantly nonlinear across the two-dimensional value signals. As the common value percent increases, profit, revenue, and efficiency all decrease monotonically, but they decrease in different ways. The discovery of these nonlinear relationships is the major contribution of Chapter 3.

The next contribution concerns English auctions, in which bidders' bids are open for the other bidders to observe. A single experimental study (Levin *et alia*, 1996) has shown that bidders modify their bids based on the most recent price at which other bidders drop out. My question was whether or not agents could learn to make this modification and whether or not this modification varied with the mixture of private and common values and/or with the level of uncertainty about the common value. In Chapter 4, I show that the agents do indeed learn to modify their bid strategies, that signal averaging increases nonlinearly as the common value percent increases, and that as the common value signal becomes more uncertain, signal averaging changes.

The final contribution concerns Treasury auctions. In my macroeconomic studies, I learned that the central banks manage the money supply primarily through auctions of bonds and treasury bills. The value of these auctions amounts to several hundred billion dollars in Canada, trillions of dollars in the United States, and trillions of dollars in other countries throughout the world. However, nobody actually knows which payment rule produces the most revenue for the central bank! Moreover, there have been few studies on this because of the complexity introduced by trading in the securities both before and after the auction. This means that most auctions cannot be realistically considered in isolation because they are embedded in an ongoing series of markets or industry dynamics. This message is stressed by Milgrom in his chapter "Auctions in Context" (Milgrom, 2004). Using the multiagent system, I incorporated the before-markets and after-markets as well as the complexities of the Treasury auction itself (Chapter 5). The results are interesting and important. I find that the "discriminatory" payment rule (used by the Bank of Canada) results in less revenue than an "average" payment rule. Whether or not the "discriminatory" payment rule produces more revenue than the "uniform" payment rule (used by the U.S. Treasury) and the "Vickrey" payment rule depends on the price spreads in the markets that occur before and after the auction.

### 1.5 References

Krishna, V., 2002, *Auction Theory*. Academic Press.

Levin, D., Kagel, J.H., and Richard, J.F., 1996. "Revenue Effects and Information Processing in English Common Value Auctions." *American Economic Review*. 86: 442-460.

Milgrom, P., 2004, *Putting Auction Theory to Work*. Cambridge University Press.

Ockenfels, A. and Selten, R., 2005. "Impulse Balance Equilibrium and Feedback in First Price Auctions." *Games and Economic Behavior*. 51: 155-170.

Rosenblueth, A. and N. Wiener, 1945, "The Role of Models in Science," *Philosophy of Scienc*e, 12(4): 316-321.

Smith, V. L., 1968, Review of *Auctions and Auctioneering* by Ralph Cassady, Jr., *The American Economic Review*, 58 (4): 959-963.

# Chapter 2

# Multiagent System Platform for Auction Simulations

## 2.1 Introduction

Multiagent systems have been applied to problems that are dynamic, complex, and distributed, and thus have been used to model the machines in manufacturing and process control, work orders in production scheduling, jobs and departments in business process optimization, planes in air traffic control, treatments and tests in hospital patient scheduling, messages in communication networks, and in many more areas (Weiss, 1999). In economics, agents have been used to model the behavior of, and interactions between, consumers, workers, families, firms, markets, regulatory agencies, and so on (see Tesfatsion, 2003 and 2006), and there have been a few applications of agent systems to auctions (Kim, 2007; Byde, 2002; and Hailu and Schilizzi, 2004). Section 2.2 discusses alternatives to multiagent systems in the analysis of auctions and why the multiagent system method was chosen for the current research.

An agent is a software entity that is autonomous, communicating, and adaptive. Autonomy means that an agent is driven by its own objectives, possesses resources (e.g., information) of its own, is capable of recording information about its environment, and can choose how to react to the environment. An agent is also a communicating software entity. Agents communicate directly with other agents by passing messages. Because each agent is autonomous, an agent must send requests to other agents for things to be done. For example, in this system, agents send messages to coordinate auctions, establish values, send bids, move to a new auctioneer, and so on. The agents are developed using object-oriented design. This means that the system consists of approximately 100 independent programs that are called "classes." Section 2.3 describes the design principles and, together with the Appendix, provides a guide to the classes.

An agent endeavours to improve its state (e.g., profit or revenue) in at least two ways. The first type of learning is reinforcement learning that uses feedback on results of actions to improve the results. The second type of learning is belief learning that involves updating beliefs about the environment, markets, and competitors, which may provide further improvements to the agent. Section 2.4 presents the results of an evaluation of different methods of agent learning.

## 2.2 Alternatives to Multiagent Systems

The major alternatives to using a multiagent system are mathematical theory, lab experiments, econometric models, and computational models.

### *2.2.1 Mathematical Theory*

In a mathematical approach, mathematical machinery is developed (e.g., optimization, order statistics, supermodularity, etc.), simplifying assumptions are made, and results proven using theorems. However, applying these theoretical results to real-world auctions is problematic. For example, Milgrom (2004, p. 22) has identified the following problems: "Academic mechanism design theory relies on stark and exaggerated assumptions to reach theoretical conclusions that can sometimes be fragile. Among these are the assumptions (i) that bidders' beliefs are well formed and describable in terms of probabilities, (ii) that any differences in bidder beliefs reflect differences in their information, (iii) that bidders not only maximize, but also cling confidently to the belief that all other bidders maximize as well."

A more realistic model of industry bidders can be achieved by using artificially intelligent software agents that are designed to optimize adaptively using the information they receive from the seller. This approach directly addresses the problems identified by Milgrom. There are fewer and more flexible simplifying assumptions[1], information to agents can be restricted to own information or expanded to information about other bids, and agents are programmed to maximize within the constraints of the abilities and information they have.

### *2.2.2 Lab Experiments*

One approach to dealing with the limitations of theory has been to perform lab experiments, usually using student subjects (Kagel and Levin, 2002). These experiments have the benefit of bidders that encompass the wide range of human reasoning and feeling, but the disadvantage is the inexperience of the bidders. The subjects, whether students or adults from industry, must learn about the bidding environment from scratch, and this constrains the complexity of the mechanisms that can be studied in the lab. The subjects simply do not have the time to develop the richness of task-specific knowledge that is used again and again in a real-world industry auction (Dyer *et alia*, 1989). Lab experiments are also expensive and time consuming. Because of these constraints, the number of existing publications on human auction experiments is small, and the experiments are limited to relatively simple environments. However, the results provide useful benchmarks to assess the results of the computational models (see Section 2.3.3).

[1] The single major assumption in this approach is the method the agents use to learn bidding strategies.

### *2.2.3 Econometric Models*

There are two types of econometric methods that have been applied to auction data: regression analysis and structural models. For example, regression analysis is applied by De Silva *et alia* (2002, 2003) to bidding data from road construction procurement auctions, by Athey and Levin (2001) to data from U.S. timber auctions, and by Iledare *et alia* (2004) to data of oil lease auctions. The aim of the structural modelling approach is to recover from the auction data distributions of values and bids, in order to then analyze such topics as: whether the values are private, affiliated, or common; the extent of collusion; the impact of entry costs, and so on. Some researchers use parametric distribution functions (Li and Perrigne, 2003; Haile *et alia*, 2003; Li *et alia*, 2000) , but an increasing number of authors are using nonparametric methods (Campo *et alia*, 2003; Hendricks *et alia*, 2003). A thorough overview with several examples is contained in Paarsch and Hong (2006).

The major advantage of the econometric methods is that they use data from real auctions. The most serious disadvantage is that data is very difficult to obtain. In addition, econometric models are restrictive because the econometrician does not know the value estimates of the bidders, and all bidding strategies are based on these valuations. In addition, the structural models assume that bidders use a Bayesian Nash equilibrium bidding strategy, which is a very questionable assumption (Bajari and Hortacsu, 2005).

### *2.2.4 Computational Models*

Another approach is to use a computational method that is not agent-based. Dynamic programming methods have been used to determine optimal bidding strategies for bidders. The use of these methods began with Friedman (1956) and is reviewed in Stark and Mayer (1971). Since a large volume of historical data on competitor bids is required to determine the optimal bidding strategy for a single bidder, the approach is useful for advising bidders in situations in which large volumes of data exists, such as online bidding (Tesauro and Bredin, 2002) and electricity markets (Attaviriyanupap *et alia*, 2005). The main advantage of the dynamic programming approach is that it produces an optimized bidding strategy based on real-world data, but the disadvantage is that such datasets are few and far between.

In summary, there are advantages and disadvantages to each approach. The major advantages of agent computational modelling are that it does not require the simplifying assumptions of mathematical analysis, can model the experienced bidders in complex environments that are beyond the reach of lab experiments, does not require assumptions about values or Bayesian Nash equilibrium required by econometric methods, and does not require large amounts of historical data required by dynamic programming methods.

## 2.3 Design Methods

The object-oriented design methods are described in Section 2.3.1. Section 2.3.2 describes some of the major classes that have been developed for the basic agent functions, auctions, and other applications. In Section 2.3.3, I present the methods that are used to verify the validity of the agent models.

### *2.3.1 Object-Oriented Design*

The multiagent system is designed using object-oriented principles and developed with Java, which is a platform-independent, object-oriented programming language. Two of the main advantages of an object-oriented approach are instantiation and extension. When we develop a Java program, we create a "class" that is an independent program with a specific purpose. This class can be used ("instantiated") one or more times to become an "object" that can then be executed. For example, I program a bidder agent class and then instantiate it many times to produce a large population of bidder agent objects. Each class program consists of properties and methods. For a bidder agent, properties include name, current bid price, and value estimate; and methods include handling a message, moving to a new auction, and adjusting the bid price. All properties are for private use by the class, but these properties may sometimes be set or retrieved by other classes. Some methods are for public access but others are restricted for use only within the object. We can create base classes with common attributes and functions and extend them using more specific attributes and functions. For example, cars, trucks and busses have many common attributes and functions that we would place in a Vehicle class, which is then extended by the classes Car, Truck, and Bus. Then, we can extend the Car class to classes for SUV, Sedan, and so on. In this application, AbstractAgent class is extended by AbstractBidderAgent, which is extended by MultiUnitBidderAgent, which is extended by BankAgent. All of the extensions from the AbstractAgent class are illustrated in Figure 2.1.

### *2.3.2 Classes*

The base multiagent platform is implemented with about 22 Java classes that are shown in Table 2.1. I have previously extended these classes in studies of repeated games with evolving finite automata using about 17 classes, repeated games with probabilistic finite automata using about 9 classes, and a simple trading economy using about 15 classes. The focus of this paper is auction simulations, which have been implemented using about 43 classes that are shown in Table 2.2. Some of the classes are described in the Appendix.

For auctions, there are auctioneer agents (sellers), bidder agents (buyers), and a coordinator agent to implement the important coordination mechanisms (Decker and Lesser, 1995). The basic idea is that each auction format (e.g. single-unit sealed bid, single-unit English, etc.) has an associated auction class to handle the mechanics of fetching bids, choosing a winner, etc., and an associated conversation class that handles the communication between the bidders and the auctioneer. The auctioneer uses the appropriate auction class and the bidder uses the appropriate conversation class. The system supports a wide variety of options for current and future simulations. I can select the auction type (sale or procurement), payment type (first-price, second-price), bid format (sealed, English), numbers (of items being auctioned, auctions, auctioneers, and bidders), value (private value, common value, mixed value), and so on. The major design goal is to provide broad functionality so that different mechanisms can be studied for both single-unit and multi-unit auctions.

### *2.3.3 Verification of Multiagent Models*

Multiagent systems, like other computational methods, have the challenge of verification. In my work, I use four approaches to verification.

First, verification is facilitated in multiagent models by explicitly modelling real world objects and relationships. For example, in the multiagent model of consumer choice in a transportation system, households, persons, and families are modeled with realistic behaviours (Salvini and Miller, 2005) based on observations and data. In my multiagent model, bidder learning is modeled using adjustment rules that are based on results from lab experiments (Ockenfels and Selten, 2005; Neugebauer and Selten, 2006).

Second, verification is strengthened by comparing simulation results to data from lab experiments for the simple cases for which there are such results. This is virtually impossible for very complex auction mechanisms, and in these cases test data itself is generated computationally (Leyton-Brown and Shoham, 2006). For games that are less complex than auctions, there are good opportunities to test learning models against data from lab experiments (Arifovic *et alia*, 2006). The single-unit sealed bid and English auctions that I study are of moderate complexity (Chapter 3 and Chapter 4), and the results can be verified against lab experiments in the simple cases of, for example, pure private values and pure common values. Agreement with this data lends credibility to the validity of the model in the more complex cases.

Third, the model must have as few parameters as possible, and the model must produce results that are stable within a range of the parameters. For example, if a reasonable range of one parameter is [0, 1], the model must be stable within a subset of this range, e.g. [0.3, 0.8]. If there are two or more parameters, then the model must be stable for an intersection of subranges. This is admittedly a subjective process, but it provides a relative measure of confidence in the model if the results are stable over [0.3, 0.8] when the results of another model are stable over [0.5, 0.7].

Fourth, the model must converge for the variables being studied. These convergence results are important, since it is impossible to interpret auction results for bid strategies, profit, revenue, and efficiency when there is no convergence. For example, without convergence the results are different when we stop the simulation in period *t*+10 from the results when stopping the simulation in period *t*. Also, the fact that convergence occurs in less than, say, 100 periods makes it reasonable to infer that the bid strategies of human agents could converge in a realistic number of real-world auctions.

## 2.4 Learning Models

An agent can improve its profit through learning in at least two ways. The first type ("belief learning") is described in Section 2.4.1. Belief learning involves updating beliefs about the environment, markets, and competitors, which may provide further improvements to the agent. Several alternative methods of the second type ("action learning") are described in Section 2.4.2, and Section 2.4.3 presents an evaluation of the alternative methods. I have implemented the two types of learning with about 31 Java classes in three packages (Table 2.3).

### *2.4.1 Belief Learning*

Belief learning is modelled by probabilistic networks (also called Bayesian networks and belief networks), and I developed the Java classes shown in Table 2.3 using the concepts and algorithms in Cowell *et alia* (1999) and Shafer (1996)[2]. Briefly, a probabilistic network is a directed acyclic graph in which nodes represent the random variables, an arrow from node X to node Y means that X has a direct influence on Y, and each dependent node has a conditional probability table. In constructing a probabilistic network, you choose[3] the set of relevant variables that describe the beliefs, add the nodes by adding the "root causes" first, then the variables they influence, and so on until you

---

[2] I developed a compact package of software for agents, but there are several products available that are oriented towards working with large datasets, e.g., Hugin.

[3] Given a large enough dataset, it is possible to for an agent to learn the structure of its Bayesian network (Heckerman, 1998).

reach the leaves which have no direct causal influence on other variables. Finally, you define the conditional probability table for each node, which provides the probability that a given node state will occur, given the states in the preceding nodes. In order for the agent to make inferences from observed facts, the network must be converted into a more compact form called a junction tree. First, the network is moralized, which means that all parents of a node are joined (or "married" and thus becoming "moral"!). Second, the moralized network is triangulated, which means that every polygon larger than a triangle is filled in to produce a network of connected triangles. Third, the triangles are converted into nodes of a junction tree, i.e., a junction tree is network of the triangles. During this process, the conditional probability tables are modified appropriately. Now when the agent observes some change in the environment, the change is propagated to all the nodes of the junction tree and the conditional probability tables are updated. To the agent, this means that its belief system is updated to accommodate the new information.

I performed many computational experiments with agents developing beliefs based on information they compile using the bid distribution classes listed in Table 2.2. These classes provide an agent with distributions of its own results for profit, winning, etc. and the results of other agents (for this, the I3 agents were provided with the identity of other bidders) in order to develop beliefs about relative strength. The bidders then used these beliefs to modify their ongoing bid strategy depending on the specific opponents in each auction. However, I found that using belief learning in this context did not significantly change the overall results for profit, revenue, and efficiency compared to agents who did not use belief learning. This result occurred because the auction-specific strategies are stationary around the ongoing bidding strategy and thus had no effect on the averages. Therefore, in the interests of parsimony, I removed belief learning from the model and have therefore not used it in the current research on auctions. However, I believe that belief learning has potential application in other types of multiagent models, especially in macroeconomic models where expectations play a major role.

### *2.4.2 Action Learning Alternatives*

There is considerable scope for choosing the action learning model for the agents. Alternatives for action learning include simple reinforcement learning, reinforcement learning methods, experience-weighted attraction, learning direction theory, genetic algorithms, and neural networks.

Simple reinforcement learning uses profit to reinforce action weights. Thus, the actions are usually modelled as discrete states that can be weighted, and only one type of information is used (profit). This method has been applied with some success to normal form games (Erev and Roth , 1998) and to auctions by Armantier (2004), Daniel *et alia* (1998), Seale *et alia* (2001), Bower and Bunn (2001), and Nicolaisen *et alia* (2001). I experimented with this simple reinforcement method, but I also extended it using two types of states: the average profit of the bidder and the average profit of the bidder's opponents. In the first, the state is 0 if the bidder's own average profit is negative, and 1 if it is positive. In the second, the state is 0 if the opponents on average are losers (negative average profits), and 1 if the opponents are on average profitable. The action weights occur in pairs, one for each state, that are updated as in the simple reinforcement learning but now depending on the state.

More sophisticated methods of reinforcement learning have not previously been used in auctions, so I performed simulations for common-value first-price sealed-bid auctions using dynamic programming, temporal difference, and Q-Sarsa methods (Sutton and Barto, 1998). The dynamic programming method reinforces actions by both actual profit and expected future profits (based on past profits) as in Sutton and Barto (1998, Chapter 4). I use the states as described above for the extended simple reinforcement learning methods, along with a state transition table containing the probabilities of transition from one state to another. The weights are then updated by combining the profit for the current state with the weights for the states indicated by the state transition table. In the temporal difference method, the agent uses profit to reinforce the current state-action pair as well as the state-action pair that preceded the current action. This approach closely follows Sutton and Barto (1998, Chapter 6). Q-Sarsa learning involves reinforcing the current state-action pair as well as <u>all</u> of the state-action pairs that preceded this action. This involves the use of eligibility traces as described in Sutton and Barto (1998, Chapter 7).

Experience-weighted attraction uses profit for winners and foregone profit for losers to reinforce discrete action states. Camerer has used this method extensively in games (Camerer, 2003; Camerer *et alia*, 2002), and it has been applied to auctions by Rapoport and Amaldoss (2004).

Learning direction theory (Selten, 1998) has been applied as impulse balance learning to auctions by Selten and Buchta (1998), Selten *et alia* (2005), Ockenfels and Selten (2005), and Neugebauer and Selten (2006). The method has also been used to interpret experimental data by Garvin and Kagel (1994) and Kagel and Levin (1999). Impulse balance learning uses foregone profit upon losing as an upward impulse on a continuous bidding strategy and money on the table upon winning as a downward impulse. The downward impulse is weighted by a balance factor that is the ratio of the expected value of the upward impulse to the downward impulse. I augmented this method to include adjustment using actual loss by the winner and foregone loss (the amount the agent would have lost if it had won) by the losers. The agent adjusts the bid strategy for a loser to bid higher depending on the level of foregone profit and bid lower depending upon the level of foregone loss. A winner reduces its bid in proportion to the money on the table if it made a profit and in proportion to the actual loss if it made a loss.

Genetic algorithms have been applied to auctions by Dawid (1999) and Andreoni and Miller (1995), and neural networks have been used by Bengio *et alia* (1999). Genetic algorithms require discrete states, and genetic algorithms and neural networks use only profit to guide the optimization.

### *2.4.3 Action Learning Evaluation*

To guide selection of an appropriate learning method, we need to establish the level of intelligence required. Since the research is motivated by an interest in real-world asset-sale auctions such as those for timber sales, drilling licences, and treasury securities, the agents must simulate experienced real-world auction bidders. A credible learning method for simulating these sophisticated bidders must satisfy four criteria: (1) be a realistic representation of how humans can potentially maximize profit in the auction environment, (2) potentially utilize all available information feedback, (3) handle continuous bids, and (4) be extendable.

Human reasoning cannot be captured with a single computational paradigm but is situational and adaptable and involves a combination of heuristics and rules-of-thumb, together with logic and optimization when required (Dyer *et alia*, 1989; Hutchinson and

Gigerenzer, 2005; Ohtsubo and Rapoport, 2006). In a study using experienced construction executives, Dyer *et alia* (1989, p. 115) concluded that "success in the field thus derives not from conformity to a narrow notion of rationality, but from acquiring and utilizing detailed knowledge of a particular market environment." Genetic algorithms and neural networks are general-purpose methods that require the researcher to fit the reasoning to the algorithm and do not accommodate the specific economic reasoning that goes into developing the various auction strategies. The more straightforward methods like simple reinforcement, experience-weighted attraction, and impulse balance are superior in this regard.

Research in auctions (Dyer and Kagel, 1996; Dyer *et alia*, 1989) demonstrates that bidders acquire and use detailed knowledge in their specific auction environments. Thus, a realistic learning method must accommodate different levels of information and utilize more than just profit. Except for experience weighted attraction and impulse balance, the methods use only profit and are thus too informationally restrictive. Experience weighted attraction uses profit and foregone profit, but impulse balance uses money on the table and foregone profit and can be extended to use profit, loss, and foregone loss.

A further limitation of most of the learning methods is that they are implemented using discrete states. If the discretization is too fine, the implementation is too inefficient; if it is too coarse, the bidding is not realistic enough for meaningful economic conclusions. The impulse balance model is the exception since it deals efficiently with continuous[4] increases or decreases in the bidding strategy.

Finally, the method should be extendable so that other auction mechanisms and context variables can be accommodated in future studies, but only a method like impulse balance can be practically extended in this way. The basic method uses money on the table and foregone profit, and I have extended it to use actual profit, actual loss, foregone loss, and estimates of these impulses when information is restricted.

In summary, the method that comes closest to satisfying the criteria is the augmented impulse balance method. Thus, this method is developed and expanded in subsequent chapters.

## 2.5 Conclusion

The multiagent system approach with agents using modified impulse balance learning has the advantages of not requiring the simplifying assumptions of mathematical theory and of not being constrained in complexity by the limited experience of experimental subjects. Impulse balance learning provides the best foundation for learning in auctions since it is a realistic representation of experienced human bidders, utilizes several types of information feedback, handles continuous bids, and is extendable. Therefore, I modify and extend the impulse balance method in multiagent system simulations of sealed-bid auctions (Chapter 3), English auctions (Chapter 4), and treasury auctions (Chapter 5).

[4] Continuous in this context means real numbers that are not restricted to integers and that are represented by 32 bits.

## 2.6 References

Andreoni, J., Miller, J.H., 1995. "Auctions with artificial adaptive agents." *Games and Economic Behavior*. 19: 39-64.

Arifovic, J., McKelvey, R. Pevnitskaya, S., 2006."An initial implementation of the turing tournament to learning in repeated two person games." *Games and Economic Behavior,* 57: 93-122.

Armantier, O., 2004. "Does observation influence learning?" *Games and Economic Behavior*. 46: 221-239.

Athey, S., Levin, J., 2001."Information and competition in U.S. forest service timber auctions." *Journal of Political Economy*, 109, No. 2, 375–417.

Attaviriyanupap, P., Kita, H., Tanaka, E., Hasegawa, J., 2005."New bidding strategy formulation for day-ahead energy and reserve markets based on evolutionary programming." *Electrical Power and Energy Systems*, 27: 157–167.

Bajari, P. and Hortacsu, A., 2005 "Are structural estimates of auction models reasonable? Evidence from experimental data." *Journal of Political Economy*, 113(4): 703–741.

Bengio, S., Bengio, Y., Robert, J., Belanger, G., 1999. "Stochastic learning of strategic equilibria for auctions." *Neural Computation*. 11: 1199-1209.

Bower, J., Bunn, D., 2001. "Experimental analysis of the efficiency of uniform-price versus discriminatory auctions in the england and wales electricity market." *Games and Economic Behavior*. 25: 561-592.

Byde, A., 2002. "Applying evolutionary game theory to auction mechanism design." HPL-2002-321, Hewlet-Packard Company.

Camerer, C., 2003, *Behavioral game theory: experiments in strategic interaction*. Princeton University Press.

Camerer, C., Ho, T.-H., Chong, J. K., 2002. "Sophisticated experience-weighted attraction learning and strategic teaching in repeated games." *Journal of Economic Theory*. 104: 137-188.

Campo, S., Perrigne, I. , Vuong, Q., 2003."Asymmetry in first-price auctions with affiliated private values." *Journal of Applied Econometrics*, 18: 179-207.

Cowell, R. G., A. P. Dawid, S.L. Lauritzen, Spiegelhalter, D.J. 1999, *Probabilistic Networks and Expert Systems*, Springer.

Daniel, T. E., Seale, D.A., Rapoport, A., 1998. "Strategic play and adaptive learning in the sealed-bid bargaining mechanism." *Journal of Mathematical Psychology*. 42: 133-166.

Dawid, H., 1999. "On the convergence of genetic learning in a double auction market." *Journal of Economic Dynamics and Control*. 23: 1545-1567.

De Silva, D. G., Dunne, T., Kosmopoulou, G., 2002. "Sequential bidding in auctions of construction contracts." *Economics Letters*, 76: 239–244.

De Silva, D. G., Dunne, T., Kosmopoulou, G., 2003. "An empirical analysis of entrant and incumbent bidding in road construction contracts." *The Journal of Industrial Economics*, 51(3): 295–316.

Dyer, D., Kagel, J. H., and Levin, D., 1989. "A comparison of naïve and experienced bidders in common value offer auctions: laboratory analysis." *The Economic Journal,*. 99, 108-115.

Dyer, D., Kagel, J.H., 1996. "Bidding in common value auctions: how the commercial construction industry corrects for the winner's curse." *Management Science*. 42, 1463-1475.

Erev, I., Roth, A.E., 1998. "Predicting how people play games: reinforcement learning in experimental games with unique, mixed strategy equilibria." *American Economic Review*. 88: 848-881.

FIPA, 2002. *FIPA Communicative Act Library Specification*. Foundation for Intelligent Physical Agents, Document SC00037J.

Friedman, L., 1956."A competitive-bidding strategy." *Operations Research*, 4(1):104-112.

Garvin, S., Kagel, J.H., 1994. "Learning in common value auctions: some initial observations." *Journal of Economic Behavior and Organization*. 25: 351-372.

Haile, P. A., Hong, H., Shum, M., 2003."Nonparametric tests for common values in first-price sealed-bid auctions." *Yale University Working Paper*.

Hailu, A., Schilizzi, S., 2004. "Are auctions more efficient than fixed price schemes when bidders learn?" *Australian Journal of Management*. 29: 147-168.

Heckerman, D., 1998."A tutorial on learning with Bayesian networks." In: Jordon, M.I. (Ed), *Learning in Graphical Models,*. Kluwer Academic Publishers, pp. 301-354.

Hendricks, K., Pinkse, J., Porter, R. H., 2003."Empirical implications of equilibrium bidding in first-price, symmetric, common value auctions." *The Review of Economics Studies*, 70, 115-145.

Hendricks, K., Porter, R. H., and Wilson, C. A.,1994."Auctions for oil and gas leases with an informed bidder and a random reservation price." *Econometrica*, 62, No. 6, 1415-1444.

Hutchinson, J. C., Gigerenzer, G., 2005. "Simple heuristics and rules of thumb: Where psychologists and behavioral biologists might meet." *Behavioral Processes*. 69: 97-124.

Iledare, O., Pulsipher, A., Olatubi, W., Mesyanzhinov, D. 2004."an empirical analysis of the determinants and value of high bonus bids for petroleum leases in the U.S. outer continental shelf (OCS)." *Energy Economics*, 26: 239-259.

Kagel, J.H., Levin, D., 1999. "Common value auctions with insider information." *Econometrica*. 67: 1219-1238.

Kagel, J.H., Levin, D., 2002. *Common Value Auctions and the Winner's Curse*. Princeton University Press.

Kim, Y. S., 2007. Maximizing sellers' welfare in online auction by simulating bidders' proxy bidding agents." *Expert Systems with Applications* 32(2): 289-298.

Milgrom, P., 2004. *Putting Auction Theory to Work*. Cambridge University Press.

Leyton-Brown, K. and Y. Shoham, 2006."A test suite for combinatorial auctions." In: Cramton, P., Shoham, Y., and Steinberg, R. (Eds), *Combinatoria Auctions.* The MIT Press, pp. 451-478.

Li, T., Perrigne, I, 2003."Timber sale auctions with random reserve prices." *Review of Economics and Statistics*, 85. 15, No. 1, 189-200.

Li, T., I. Perrigne, Vuong, Q., 2000."Conditional independent private information in OCS wildcat auctions." *Journal of Econometrics*, 98, 129-161.

Neugebauer, T., Selten, R., 2006. "Individual behavior of first-price auctions: the importance of information feedback in computerized experimental markets." *Games and Economic Behavior* , 54, 183-204.

Nicolaisen, J., Petrov, V., Tesfatsion, L., 2001. "Market power and efficiency in a computational electricity market with discriminatory double-auction pricing." *ISU Economic Report No. 52*, Iowa State University.

Ockenfels, A., Selten, R., 2005. "Impulse balance equilibrium and feedback in first price auctions. *Games and Economic Behavior.* 51: 155-170.

Ohtsubo, Y. , Rapoport, A., 2006. "Depth of reasoning in strategic form games." *Journal of Socio-Economics*. 35: 31-47.

Paarsch, H. J., Hong H., 2006, *An Introduction to the Structural Econometrics of Auction Data*. The MIT Press.

Rapoport, A., Amaldoss, W., 2004. "Mixed-strategy play in single-stage first-price all-pay auctions with symmetric bidders." Journal of Economic Behavior and Organization. 54: 585–607.

Salvini, P.A. and E.J. Miller, 2005."ILUTE: An operational prototype of a comprehensive microsimulation model of urban systems", *Networks and Spatial Economics*, 5: 217-234.

Seale, D. A., Daniel, T.E., Rapoport, A., 2001. "The information advantage in two-person bargaining with incomplete information." *Journal of Economic Behavior and Organization*. 44: 177-200.

Selten, R., 1998. "Features of experimentally observed bounded rationality." *European Economic Review*, 42: 413-436.

Selten, R., Buchta, J., 1998. "Experimental sealed bid first price auctions with directly observed bid functions." In: Budescu, D. V., Erev, I., and Zwick, R. (Eds), *Games and Human Behavior: Essays*. Lawrence Erlbaum Associates, pp. 53-78.

Selten, R., Abbink, K., Cox, R., 2005. "Learning direction theory and the winner's curse." *Experimental Economics* 8: 5-20.

Shafer, G., 1996, *Probabilistic Expert Systems*, SIAM.

Sutton, R. S., Barto, A. G., 1998, *Reinforcement learning: an introduction*, The MIT Press.

Tesauro, G., Bredin, J. L.."Strategic sequential bidding in auctions using dynamic programming." in: *Proceedings of the first international joint conference on Autonomous agents and multiagent systems*, ACM Press.

Tesfatsion, L., 2003. "Agent-based computational economics: modeling economies as complex adaptive systems." Information Sciences 149: 263-269.

Tesfatsion, L. , Judd, K., 2006. *Handbook of Computational Economics: Volume 2 Agent-Based Computational Economics*. Elsevier.

Weiss, G., 1999, *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*. The MIT Press.

## 2.7 Tables

| Table 2.1. Base Packages and Classes | | |
|---|---|---|
| Package | Class | Extends |
| 1. agent | 1. AbstractAgent<br>2. AgentInfo<br>3. Registry | |
| 2. distributions | 4. RandomNumber<br>5. Beta<br>6. Normal<br>7. Uniform | <br>RandomNumber<br>RandomNumber<br>RandomNumber |
| 3. grid | 8. Cell<br>9. Coordinates<br>10. Grid<br>11. Options | |
| 4. gui | 12. BasicMenu<br>13. GuiFrame<br>14. HelpFrame<br>15. InfoPanel | JMenuBar<br>JFrame<br>JFrame<br>JPanel |
| 5. statistics | 16. Moments<br>17. Regression<br>18. TimeSeries | |
| 6. support.filesupport | 19. Tracing | |
| 7. support.guisupport | 20. Console<br>21. MenuCreator<br>22. RadioButtonPanel | <br><br>JPanel |

Table 2.2. Auction Packages and Classes

| Package | Class | Extends |
|---|---|---|
| 1. auction.agent | 1. AbstractAuctioneerAgent | AbstractAgent |
| | 2. AbstractBidderAgent | AbstractAgent |
| | 3. AbstractCoordinatorAgent | AbstractAgent |
| | 4. BankAgent | MultiUnitBidderAgent |
| | 5. BidDistributions | |
| | 6. CentralBankAgent | MultiUnitAuctionerAgent |
| | 7. MultiUnitAuctionerAgent | AbstractAuctioneerAgent |
| | 8. MultiUnitBidderAgent | AbstractBidderAgent |
| | 9. MultiUnitBidDistributions | BidDistributions |
| | 10. SingleUnitAuctioneerAgent | AbstractAuctioneerAgent |
| | 11. SingleUnitBidderAgent | AbstractBidderAgent |
| | 12. SingleUnitBidDistributions | BidDistributions |
| | 13. SingleUnitCoordinatorAgent | AbstractCoordinatorAgent |
| | 14. TreasuryCoordinatorAgent | AbstractCoordinatorAgent |
| 2. auction.bidding | 15. Auction | |
| | 16. AuctionResult | |
| | 17. Bid | |
| | 18. MultiUnit | Auction |
| | 19. MultiUnitEnglish | MultiUnit |
| | 20. MultiUnitSealed | MultiUnit |
| | 21. SecondaryTreasuryMarket | |
| | 22. SingleUnit | Auction |
| | 23. SingleUnitEnglish | SingleUnit |
| | 24. SingleUnitSealed | SingleUnit |
| 3. auction.conversation | 25. MultiUnitConversation | |
| | 26. MultiUnitSealedConversation | MultiUnitConversation |
| | 27. MultiUnitEnglishConversation | MultiUnitConversation |
| | 28. SingleUnitConversation | |
| | 29. SingleUnitSealedConversation | SingleUnitConversation |
| | 30. SingleUnitEnglishConversation | SingleUnitConversation |
| 4. auction.grid | 31. AuctionAgentCell | Cell |
| | 32. AuctionAgentGrid | Grid |
| | 33. AuctionAgentOptions | Options |
| 5. auction.gui | 34. AuctionAgentGuiFrame | GuiFrame |
| | 35. AuctionAgentOptionDialogSingleUnit | JDialog |
| | 36. AuctionAgentOptionDialogTreasury | JDialog |
| | 37. SliderHandlerSingleUnit | |
| | 38. SliderHandlerTreasury | |
| 6. auction.simulation | 39. AveragingImpulseOutput | |
| | 40. BidImpulseOutput | |
| | 41. EfficiencyOutput | |
| | 42. ProfitOutput | |
| | 43. RevenueOutput | |

| Table 2.3. Learning Packages and Classes | | |
|---|---|---|
| Package | Class | Extends |
| 1. learning | 1. Action | |
| | 2. SingleUnitLearning | |
| | 3. Rla | SingleUnitLearning |
| | 4. RLas | SingleUnitLearning |
| | 5. EWA | SingleUnitLearning |
| | 6. DP | SingleUnitLearning |
| | 7. TD | SingleUnitLearning |
| | 8. Q | SingleUnitLearning |
| | 9. IB | SingleUnitLearning |
| | 10. IBA | SingleUnitLearning |
| | 11. SingleUnitImpulse | SingleUnitLearning |
| | 12. MultiUnitAdjustment | |
| | 13. MultiUnitRules | MultiUnitAdjustment |
| 2. probnet.algorithm | 14. CreateJunctionTree | |
| | 15. FindCliques | |
| | 16. InitializePotentials | |
| | 17. Moralize | |
| | 18. PerfectOrder | |
| | 19. Triangulate | |
| 3. probnet.bayesnetwork | 20. ActiveBN | |
| | 21. BayesNetwork | Network |
| | 22. BayesNode | Node |
| | 23. ChainComponent | Node |
| | 24. JunctionTree | Network |
| | 25. JunctionTreeNode | Node |
| | 26. Key | |
| | 27. Network | |
| | 28. Node | |
| | 29. PotentialTable | Table |
| | 30. Separator | Node |
| | 31. Table | |

## 2.8 Figures

Figure 2.1. Simple Class Diagram for Auction Bidder Classes

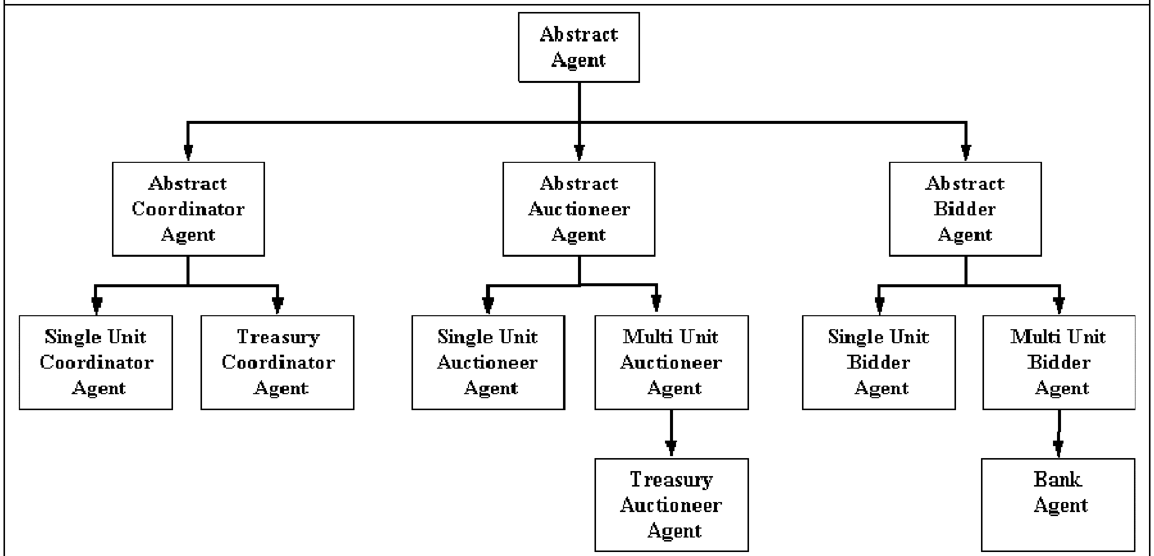

## 2.9 Appendix

This Appendix describes some of the design concepts used in implementing the functionality for Agents, Conversations, and Auctions.

### *Agent Classes*

There is a base AbstractAgent class that provides functions common to all agents. AbstractCoordinatorAgent, AbstractAuctioneerAgent, and AbstractBidderAgent classes extend AbstractAgent and then these in turn are extended for single-unit, multi-unit, and treasury auctions.

A coordinator agent has two major tasks: to create the other agents and coordinate the auctions. For each auction, the coordinator broadcasts a message to every auctioneer to hold an auction and directs the agents to move if there is more than one auctioneer. The coordinator can randomly distribute the bidders equally or unequally to the auctioneers.

An auctioneer agent has three major tasks: execute the auction, notify the bidders, and print results. An auctioneer creates an auction object of the appropriate type (e.g., SingleUnitSealed, MultiUnitSealed, etc.) based on the type of auction that has been set by the experimenter. The auctioneer then uses the auction object to execute the auction, fetch bids, pick winners, and send results to the bidders. For the benefit of the experimenter, the auctioneer agent also prints results for the experimenter using classes in the auction.simulation package.

A bidder agent has three major tasks: learn how to improve bidding, calculate a bid and send it to the Auctioneer using the Bid class (The Bid class holds attributes for a bid: the bidder, value signal, action that led to the bid, and the bid amount plus the resulting ranking, profit, foregone profit, and so on.), and move to a new auctioneer (if the *Bidders* option is "random"). Each bidder agent has a learnBidFactor method that is called when the auction object requests the bidder's participation in an auction. The learnBidFactor method in turn calls one of the learning algorithms (see Section 2.3) to calculate the bid factor. For the benefit of the experimenter, the bidder agent also prints results for the experimenter using the classes in the auction.simulation package.

***Conversation Classes***

The bidder communicates with the auctioneer using protocols encapsulated in conversation classes. The message types are consistent with FIPA Agent Communication Language (FIPA, 2002).

The SingleUnitConversation and MultiUnitCoversation classes tell the bidders to learn and inform them of auction results. They are extended by the classes for sealed-bid and English auctions that retrieve the bids from the bidders. The process involves a single message for sealed-bid auctions, but involves many messages for the English auctions. Starting with a low price, SingleUnitEnglish iterates through a loop: send to active bidders the price and the latest dropout price; remove bidders who reject this price level from the auction; increment the price.

***Auction Classes***

Each auction involves the following four major functions: manage the auction, fetch bids, pick the winner(s), and calculate payment(s). The processes of auction management, picking the winner, and calculating the payment are handled by the SingleUnit and MultiUnit classes. Since the process of fetching bids differs for sealed-bid and English auctions, this function is handled by extensions of these classes.

# Chapter 3

# Multiagent System Simulations of Sealed-Bid Auctions with Two-Dimensional Value Signals

## 3.1 Introduction

This study endows computational agents with a learning model and uses these agents in computational experiments to make three contributions to knowledge about multiagent simulations of sealed-bid auctions.

Several empirical studies have shown that impulse balance learning explains how human bidders in auction experiments adjust their bid price strategies (Selten and Buchta, 1998; Selten *et alia*, 2005; Ockenfels and Selten, 2005; Neugebauer and Selten, 2006; Garvin and Kagel, 1994; Kagel and Levin, 1999). This makes it a promising method to investigate as the learning model in a multiagent system. The first contribution is to adapt Selten's impulse balance learning method for use by agents in a multiagent system.

In real-world auctions (such as those for timber sales, oil leases, spectrum, and services) the item value often has both a private value and a common value component (Goeree and Offerman, 2002). Thus, the second contribution is to determine how profit, revenue, and efficiency change as the common value component increases. There are no lab experiments to indicate whether this change is linear or non-linear. The multiagent simulations show that as the common value percent increases, profit, revenue, and efficiency all decrease monotonically (and often nonlinearly), but they decrease at different rates. Profit curves tend to decrease faster at higher common values, revenue curves tend to decrease more rapidly at low common value percents, and efficiency curves tend to stay high and then decrease rapidly for high percents of common value.

The third contribution is to determine whether it may be worthwhile for a seller (such as a federal or state government) to enforce truthful revelation of the true common value by auction winners. In lab experiments, Kagel and Levin (1999) show that revealing information about the true common value in first-price auctions increased or decreased revenue depending upon the number of bidders and the degree of uncertainty about the common value. The multiagent simulations show that forcing revelation of the true common value may have beneficial revenue effects when there is a higher degree of uncertainty about the common value.

In Section 3.2, I describe the auction model. Section 3.3 provides details of the learning model and its properties of convergence and sensitivity. Section 3.4 compares the results of learning model with results from lab experiments in other studies. Section 3.5 demonstrates the nonlinear variation of revenue and efficiency with the common value percent. Section 3.6 shows the results of requiring the auction winners to reveal the actual common value to the auction losers. Section 3.7 presents conclusions.

## 3.2 Auction Model

The multiagent system platform is described in Chapter 2. In this section, I describe how the system implements values and the value signals for bidders (3.2.1), the levels of information feedback (3.2.2), and the number of periods and bidders (3.2.3).

### *3.2.1 Values and Value Signals*

Before participating in a sealed-bid auction in period *t*, each bidder *i* determines its estimate for the value $v_t^i$ of the item, and this estimate is called a value signal, denoted

$\hat{v}_t^i$.[1] Most auction research has involved a single value signal $\hat{v}_t^i$ that is either pure private ($v_{P,t}^i$) or pure common ($\hat{v}_{C,t}^i$), and these pure signals are called "one-dimensional" value signals. The bidders' value signals are "pure private value" when they base their estimates on their own value for the item, without considering how other bidders might value the item. The value signals are "pure common value" when bidders base their estimates on an estimated future actual value that is common to all bidders, for example a resale price. In the case of pure private values, each bidder will have a different value signal and the estimated value for a bidder is the actual value of the item to that bidder. In the case of pure common values, the actual common value is unknown to the bidders before and during the auction, and is discovered in the markets after the auction only by the winning bidder.

In most real-world situations, a value signal is a mixture of private and common value components. A few researchers (Dasgupta and Maskin, 2000; Jehiel and Moldovanu, 2001; Goeree and Offerman, 2002) have studied these mixed value signals and designated them "multi-dimensional" (or more precisely, "two-dimensional") value signals. For example timber sale auctions and oil leases have a common value component consisting of the volume and market price of the resource and a private value component consisting of firm-specific costs, capacities, and skills (Athey and Haile, 2002; Hendricks *et alia*, 2003; Haile *et alia*, 2003). Similarly, service procurement auctions have a common value component that is the scope of work and a private value component consisting of productivity, wage costs, and overhead costs. Within the

---

[1] The notation is summarized in Table 3.1.

context of a unique mixture of private and common values, the seller establishes the auction rules, the most fundamental of which are the payment rule and the information to be released to the bidders after the auction. In this case, the value signal $\hat{v}_t^i$ is a function of both types of value so that $\hat{v}_t^i = \hat{v}_t^i(v_{P,t}^i, \hat{v}_{C,t}^i)$. Following Goeree and Offerman (2002), I use linear combinations of private values and common value signals to produce mixed value signals that range from pure private value to pure common value. An agent's value signal is $\hat{v}_t^i = (1-\theta_C)v_{P,t}^i + \theta_C \hat{v}_{C,t}^i$, where $\theta_C \in [0,1]$ is the fraction of common value. The actual value, known by the winner, is therefore $v_t^i = (1-\theta_C)v_{P,t}^i + \theta_C v_C$. Two levels of two-dimensional value signals ($\theta_C$ = 0.14 and 0.25) have been investigated in experiments by Goeree and Offerman (2002), but my study is the first to look at the full spectrum of two-dimensional signals and the variation in profit and revenue as well as efficiency.

Values are distributed to the agent bidders in a different way than the distribution to human subjects in lab experiments (Kagel and Levin, 2002). In this study, each bidder agent's private and common value signals, as well as the actual common value, are fixed throughout the auctions. This is an artificial situation, but it has the purpose of identifying the adaptively best bidding strategy for each possible value signal. The alternative, which is used in lab experiments, is to provide each bidder with a random value signal for each auction. This results in each bidder learning an average bidding strategy in response to the full range of value signals. However, since bidding strategies may be different for different value signals, especially in first-price auctions, this average is not very informative.

The experimenter specifies the support $[\underline{S}_P, \overline{S}_P]$ of a distribution of the private value signals $v_{P,t}^i$ and a support $[\underline{S}_C, \overline{S}_C]$ for the common value $v_C$. In most experimental studies and the simulations in this paper $[\underline{S}_C, \overline{S}_C] = [\underline{S}_P, \overline{S}_P]$. There are two methods of providing the bidding agents with value signals from these supports: random and deterministic. In the first method, a bidder's private value $v_{P,t}^i$ is drawn from a distribution (usually the uniform distribution) on the support. Each bidder's common value signal $\hat{v}_{C,t}^i$ is drawn from a distribution on the support centred on the common value $[v_C - \varepsilon, v_C + \varepsilon]$, where the common value is the centre of the support $[\underline{S}_C, \overline{S}_C]$. There is uncertainty among the bidders about what this common value is, and a larger $\varepsilon$ represents more uncertainty. This method is satisfactory for investigating a single point in the two-dimensional value spectrum (i.e. 50% common value, pure common value, etc.) However, for simulations performed across the full two-dimensional spectrum from pure common to pure private value, random draws lead to different value signal profiles at each common value percent. This introduces some unnecessary noise into the results, but in fact does not change the overall results. However, it is preferable to have the same profile across the simulations so that the results are perfectly comparable. Therefore, the second method is a simple algorithm that sets the private and common value signals. Each agent is provided with a unique two-dimensional value signal so that the collection of signals spans the supports. The first method is used for the fixed point simulations and the second is used for the simulations that span the two-dimensional value spectrum.

When using the first method, I use the Uniform distribution of value signals over this support, since this is commonly used in the experiments in Kagel and Levin (2002) and others. I experimented with different distributions (normal, beta(2,2), beta(4,2), and beta(2,4)[2] ) and the results are as expected: the bid price strategies for the symmetric distributions (uniform, normal, and beta(2,2)) were virtually identical and the bid price strategies for the asymmetric distributions (beta(4,2) and beta(2,4)) shift right and left respectively.

### *3.2.2 Information Levels*

The seller must decide how much information should be released to the bidders after the auction, with alternatives ranging from each bidder's own information to information about all bids. Dufwenberg and Gneezy (2002) compare the results from lab experiments for a two-person bargaining game with three incremental levels of information about auction results: no information about others, the winning bid price (semi-full), and all bids (full). Neugebauer and Selten (2006) report the results from lab experiments for first-price sealed-bid auction with three information levels provided in between auctions: no information about others, the winning bid price, and the runner-up bid price. Similarly, in this study I use three levels of information (own, winner, and winner and runner-up) and designate them I1, I2, and I3 respectively.[3]

Bidders do not know other bidders' value signals, nor do they know the actual common value when they do not win. The common value $v_C$ is unknown *ex ante* for all bidders, and only the winning bidders know $v_C$ *ex post*. The actual value known to the winner in a two-dimensional value environment is $v_t^i = (1-\theta_C)v_{P,t}^i + \theta_C v_C$. In this study, I1

---

[2] These distributions are, respectively, more in the middle with tails, more in the middle without tails, more on the high end, and more on the low end.

[3] The notation is defined in Table 3.1 and the information feedback is summarized in Table 3.2.

consists entirely of own information: own value signal $\hat{v}^i$, own bid price $b_t^i$, own ranking $r_t^i$, actual common value upon winning, and own payment $p_t^i$. I2 consists of the own I1 information plus information about the winning bid price $b_t^{(1)}$[4] and the payment $p_t$. I3 consists of the information from levels I1 and I2 plus information about the runner-up bid price $b_t^{(2)}$. The actual value is revealed only to the winner, and it is revealed before the next iteration so that the agent can use the information. However, all bidders know the support $[\underline{S}_C, \overline{S}_C]$ so that the I1 and I2 agents have an estimate for the gap between bids (see Section 3.3.2 and 3.3.3). Since this method is constructed so the agents seek for their optimal bidding strategy for the value signals they have been given, the winning agent does not carry forward its knowledge of the actual common value. One way to interpret this is that it does not know that the common value will stay the same. I have experimented with moving the actual common value randomly from period to period within the $\varepsilon$ neighbourhood of the center of the support, but this has minimal effect on the results.

### *3.2.3 Number of Bidders and Periods*

Four and seven bidders per auction were chosen to be compatible with lab experiments of Kagel *et alia* (1987) and Levin *et alia* (1996). Twenty-five simultaneous seller agents are used when there are four bidders per auction (for a total of 100 agents) and sixteen when there are seven bidders per auction (for a total of 112). These numbers

---

[4] Superscript numbers in parentheses denote order statistics. In a sealed-bid auction, $b_t^{(1)}$ is the highest bid in the auction and $b_t^{(n)}$ is the lowest bid in an auction with n bidders.

are chosen to provide a good mix of bidder agents and to keep the totals approximately equal. Each auction has the same number of bidders and the bidder agents move randomly from seller to seller on a five-by-five or four-by-four torus. This method matches the bidder agents randomly so that each agent has the opportunity to optimize its bidding strategies by bidding against a wide range of values held by the other agents.

Each agent participates in one auction per period. I use 150 periods in order to accommodate learning, but on average the agents converge to a steady state bidding strategy within about 50 auctions (see Figures 3.6 and 3.7).

### 3.3 Learning Model

The first contribution of the study is to determine if Selten's impulse balance learning method is suitable for multiagent simulations. In this section, I describe the impulse balance learning method and then show that it results in an unacceptable amount of negative profit and sensitivity to initial values. A few simple modifications solve both problems and produce a learning method that converges well, is insensitive to the learning rate, and produces results for value-multiplier, profit, revenue, and efficiency that agree closely with results from lab experiments. This demonstrates that a multiagent system with this learning method can be used as a credible alternative to lab experiments, especially where bidding experience is desirable.

There is considerable scope for choosing the learning model for the agents, including reinforcement learning, experience-weighted attraction, impulse balance, and machine learning methods. These methods are reviewed and evaluated in Chapter 2. Modified impulse balance learning provides the best foundation for learning in auctions since it is a realistic representation of experienced human bidders, utilizes all information feedback, handles continuous bids, and is extendable. The impulse balance method uses foregone profit[5] upon losing as an upward impulse on a continuous bidding strategy and money on the table[6] and actual loss upon winning as downward impulses. Several empirical studies have shown that impulse balance learning fits the data for bid adjustments by lab experimental subjects (Selten and Buchta, 1998; Selten *et alia*, 2005; Ockenfels and Selten, 2005; Negebauer and Selten, 2006; Garvin and Kagel, 1994; Kagel and Levin, 1999).

Section 3.3.1 describes Selten's impulse balance learning method. The next two sections describe the adjustment rules for the downward impulses for winners (Section 3.3.2) and the upward impulses for losers (Section 3.3.3) . Section 3.3.4 presents results from using impulse balance learning and an improved learning method: impulse learning with loss aversion (ILA). Section 3.3.5 presents convergence and sensitivity analyses for the ILA method, and Section 3.3.6 compares simulation results to results from lab experiments.

The common value signal supports in this section follow Kagel *et alia* (1989). I use five bidders, $[\underline{S}_{\mathrm{C}}, \overline{S}_{\mathrm{C}}] = [10,\ 30]$, and $\varepsilon = 5$.

### *3.3.1 Impulse Balance Learning*

Ockenfels and Selten (2005) apply impulse balance learning to first-price auctions with private values and Selten *et alia* (2005) apply impulse balance learning to first-price

---

[5] A losing bidder regrets its low bid to the extent that its value signal $\hat{v}_t^i$ is above the winner's payment. This amount is called foregone profit and is denoted $\pi_{F,t}^i = \hat{v}_t^i - p_t$ .

[6] The winning bidder in a first-price auction sacrifices profit unnecessarily to the extent that its bid exceeds the runner-up bid. This is called leaving "money on the table" and is denoted $m_t^i = b_t^i - b_t^{(2)}$ .

auctions with common values. Bids are adjusted using downward $a^i_{-,t}$ or upward $a^i_{+,t}$ adjustments or "impulses" that the agent calculates using profit $\pi^i_t$, foregone profit $\pi^i_{F,t}$, and money on the table $m^i_t$. For profitable winners $a^i_{-,t}$ is money on the table $m^i_t$, and for unprofitable winners it is the loss $\pi^i_t$. For losers, $a^i_{+,t}$ is the foregone profit. A high-value agent wins more frequently than it loses so that typically $E_t\left[a^i_-\right] > E_t\left[a^i_+\right]$ for the high-value agent, and a low-value agent loses more frequently than it wins so that typically $E_t\left[a^i_+\right] > E_t\left[a^i_-\right]$ for a low-value agent. Thus, because the higher-value agent receives more downward impulses than upward impulses, it should put more weight on an upward impulse to compensate for its infrequency. Similarly, a lower-value agent should put more weight on a downward impulse. This is the motivation for the "balance" aspect of the impulse balance method. Each agent *i* determines its balance weight $\lambda^i_t$ as the ratio of its expected value of the upward impulse to the expected value of the downward impulse: $\lambda^i_t = \dfrac{E_t\left[a^i_+\right]}{E_t\left[a^i_-\right]}$. To determine its adjusted bid, the agent weights the impulses by a learning rate $\phi$ and the downward impulse weight $\lambda^i_t$. The bid for period *t+1* is then a revision of the previous bid $b^i_{t+1} = b^i_t + \phi(a^i_{+,t} - \lambda^i_t a^i_{-,t})$. This type of adjustment method does not require assuming that the bidding strategy is a linear function of the bidder's value signal. However, the bid at any time can be expressed as a ratio of the bid to the value estimate, $\gamma^j_t = \dfrac{b^i_t}{\hat{v}^i_t}$, so that we can discuss the value multiplier $\gamma^j_t$ that can be compared with theoretical and experimental results.

### *3.3.2 Downward Impulses for Winners*

A winning agent is assigned a rank of 1, $r_t^i = 1$, and its ordered bid price denoted $b_t^{(1)}$. Similarly, the runner-up has $r_t^i = 2$ with ordered bid price $b_t^{(2)}$, and so on. In calculating its adjustments, the winner considers information ("impulses") about its profit $\pi_t^i$ and, when the payment rule is first price, its money on the table $m_t^i = b_t^{(1)} - b_t^{(2)}$.

***Rule W1***: For all information levels, $r_t^i = 1$, and $\pi_t^i < 0.0$: $a_{-,t}^i = \left|\pi_t^i\right|$.

*Demonstration:* If the agent wins but has a loss of $\left|\pi_t^i\right| = \left|v_t^i - p_t\right|$, it lowers its bid in proportion to the loss in an effort to improve its expected profit. Adjusting for actual loss was found to be a significant factor in bid adjustment by Garvin and Kagel (1994) and Selten *et alia* (2005).

***Rule* W2(I3)**: For I3, $r_t^i = 1$, first-price payment, $\pi_t^i > 0.0$: $a_{-,t}^i = m_t^i$.

*Demonstration:* An agent with I3 can use information about the other bidders, specifically the runner-up, to make a more informed adjustment when it wins. When winning is profitable in a first-price auction, the agent uses the value of the runner-up bid to determine how much it overbid. This overbidding results when the payment rule uses the first-price since the winning bidder's ideal situation is to have bid just slightly above the runner-up bidder. Any amount that the winning bidder bids over the runner-up bidder is called "money on the table" and is denoted $m_t^i = b_t^i - b_t^{(2)}$. For a first-price payment rule, $m_t^i$ is used to adjust the bid down. Money on the table has been shown to be a significant factor in bid adjustment by Selten and Buchta (1998), Selten *et alia* (2005), Ockenfels and Selten (2005), and Negebauer and Selten (2006).

***Rule* W3(I1,I2)**: For I1 and I2, $r_t^i = 1$, first-price payment, $\pi_t^i > 0.0$: $a_{-,t}^i = \hat{m}_t^i$.

*Demonstration:* A profitable agent with I1 and I2 information must use an approximation for money on the table $\hat{m}_t^i$ to determine the adjustment for lowering its bid to improve its profit. The alternative of making no adjustment is not consistent with the impulse balance method, since there would be no downward impulse. Since the agent has information about the number of bidders *n* and the support $[\underline{S}_C, \overline{S}_C]$, it can use this to create an estimate for money on the table. The gap between bids will decrease in proportion to *n*, and since the values are drawn uniformly, an upper bound on an estimate for money on the table is $\frac{\overline{S}_C - \underline{S}_C}{n}$. However, money on the table will be small with large $\pi_t^i$ so a simple estimate for money on the table is $\hat{m}_t^i = \frac{\overline{S}_C - \underline{S}_C}{n} - \pi_t^i$.

### *3.3.3 Upward Impulses for Losers*

***Rule* L1(I2,I3)**: For I2 and I3, $r_t^i > 1$, when $\pi_{F,t}^i \geq 0$, $a_{+,t}^i = \pi_{F,t}^i$

*Demonstration:* If an agent loses, it usually regrets its low bid to the extent that its value signal $\hat{v}_t^i$ is above the winner's payment. This is the concept of foregone profit used by Camerer *et alia* (2002), Selten and Buchta (1998), Selten *et alia* (2005), Ockenfels and Selten (2005), and Negebauer and Selten (2006) with $\pi_{F,t}^i = \hat{v}_t^i - p_t$. An agent with I2 or I3 information knows the payment and so can calculate its foregone profit. When foregone profit is positive the agent increases its bid in proportion to $\pi_{F,t}^i$ since this will improve its probability of winning profitably. If a bidder has a low value signal, the foregone profit will tend to be negative, and the bidder will not increase its bid.

***Rule* L2(I1)**: For I1, $r_t^i > 1$, when $\hat{\pi}_{F,t}^i \geq 0$, $a_{+,t}^i = \hat{\pi}_{F,t}^i$

*Demonstration:* With one exception, I1 agents do not know the payment and must estimate foregone profit $\hat{\pi}_{F,t}^i = \hat{v}_t^i - \hat{p}_t$. The exception is the runner-up bidder, $r_t^i = 2$, in a second-price auction in which $b_t^i = b_t^{(2)} = p_t$ so bidder *i*'s foregone profit is $\hat{\pi}_{F,t}^i = \pi_{F,t}^i = \hat{v}_t^i - b_t^i$. For the other losing bidders, the foregone profit estimate [7]is a fraction of $\hat{v}_t^i - b_t^i$, decreasing with the number of bidders and increasing with the rank. In a second-price auction with $r_t^i > 2$, $\hat{\pi}_{F,t}^i = \left(r_t^i - 2\right)\frac{\hat{v}_t^i - b_t^i}{n}$ and in first-price auction with $r_t^i > 1$ $\hat{\pi}_{F,t}^i = \left(r_t^i - 1\right)\frac{\hat{v}_t^i - b_t^i}{n}$.[8]

### *3.3.4 Negative Profit and Sensitivity to Initial Values*

In this section, I analyze results of simulations and make changes to the impulse balance learning method. The result is a learning method that uses impulses, excludes the balance principle, and includes loss aversion, so a reasonable name for the method is "impulse learning with loss aversion" (ILA).

***Result 1***: Using impulse balance learning in computational experiments results in a high degree of negative profit, i.e. loss, and sensitivity to initial values. To achieve profitability and insensitivity to initial values, I make three changes to the impulse-

---

[7] I am using the term foregone profit for consistency with the other rules, but it is impossible for an I1 agent to estimate the payment and hence the foregone profit. Instead, the agent uses its value gap to calculate the upward adjustment.

[8] This rule is less of a foregone profit and more of a value gap adjustment. With larger n, the gap between agents will be smaller, so the basic adjustment step is inversely proportional to n. Agents with larger $r_t^i$ will need to adjust more than agents that are closer to the winner.

balance method. First, the "balance" part of the method is removed. Second, the loss adjustment in Rule W1 is weighted using a loss aversion factor $L_t^i = E_t\left[\left|\pi_t^i\right| \mid \pi_t^i < 0\right]$ that is the expected value of the magnitude of the losses. Third, when a winning agent has an expected loss ( $L_t^i > 0$ ) and lowers its bid price to reduce its probability of winning, it is counter-productive for the agent to increase its bid price when it successfully reaches the losing state. Thus, a losing agent uses foregone profit to raised its bid price only when $L_t^i = 0$ and the adjustment can be written using the indicator function $\mathbf{1}_{(L_t^i=0)}$, i.e., $a_{+,t}^i = \mathbf{1}_{(L_t^i=0)}\hat{\pi}_{F,t}^i$. In summary, the ILA method is to adjust bids using $b_{t+1}^i = b_t^i + \phi(a_{+,t}^i - a_{-,t}^i)$, where the adjustment rules are:

***Rule* W1**: For all information levels, $r_t^i = 1$, and $\pi_t^i < 0.0$: $a_{-,t}^i = \left(1 + L_t^i\right)\left|\pi_t^i\right|$.

***Rule* W2**: For I3, $r_t^i = 1$, first-price payment, $\pi_t^i > 0.0$: $a_{-,t}^i = m_t^i$.

***Rule* W3**: For I1 and I2, $r_t^i = 1$, first-price payment, $\pi_t^i > 0.0$: $a_{-,t}^i = \hat{m}_t^i$.

***Rule* L1**: For I2 and I3, $r_t^i > 1$, when $\pi_{F,t}^i \geq 0$, $a_{+,t}^i = \mathbf{1}_{(L_t^i=0)}\pi_{F,t}^i$

***Rule* L2**: For I1, $r_t^i > 1$, when $\hat{\pi}_{F,t}^i \geq 0$, $a_{+,t}^i = \mathbf{1}_{(L_t^i=0)}\hat{\pi}_{F,t}^i$

*Discussion*: The results for impulse balance learning in Figure 3.1 show that a large proportion of the bidders (especially those with high value signals) experience losses. This is a much higher level of losses than shown in results from lab experiments and a level of sensitivity to starting values that is undesirable in a computational model. For example, bankruptcies occur in about 6% of the auctions with experienced bidders (Kagel and Richard, 2001). These bankruptcies occurred in two situations: 8% of bidders went bankrupt with a \$10 cash balance with a support of [50, 380], $\varepsilon = 18$, and 7 bidders; 4% of bidders went bankrupt with a $20 cash balance with support of [25, 225], $\varepsilon = 18$, and 4 bidders. Experienced bidders in real-world auctions would likely be skillful enough to avoid losses and bankruptcies altogether, so the goal of the learning model should be a minimal level losses or bankruptcies, at least below the 6% in Kagel and Richard's experiments.

The fact that the high-value bidders are experiencing losses indicates that there is a problem with the learning model for high-value bidders. The balance factor $\lambda_t^i$, which varies with the bidder value, could be expected to deal with this problem but it is not producing satisfactory results. Figure 3.2 shows that the values of $\lambda_t^i$ do vary with the bidder value, and tend to be lower for high-value bidders than for low-value bidders as expected from the discussion in Section 3.3.1 When $\lambda_t^i$ is removed from the model, the results improve slightly as shown in Figure 3.3a. It may still be reasonable to expect the agent to put more weight on a downward impulse than an upward impulse, even though the balance factor may not be the approach that should be used. An agent may obtain improved profits if it weights the downward impulse from negative profit more than the corresponding increase from positive foregone profit. Tversky and Kahneman (1992, Table 3.6) estimate loss aversion factors in the range [0.97, 2.44], but it makes sense in this case of bidders with different value signals to have endogenous loss aversion. Figure 3. 3b shows the results for an endogenous loss aversion where the loss aversion factor is $L_t^i = E_t\left[\left|\pi_t^i\right| \mid \pi_t^i < 0\right]$. Now 35% of the bidders experience losses but the overall average profit is up to -0.15. It also makes no sense for an optimizing bidder to raise its bid after losing, when it has been experiencing losses when it is winning. Thus, I introduce a profit switch $\mathbf{1}_{(L_t^i=0)}$ that the agent uses for its upward impulses. This final modification now raises all agents to non-negative profit as shown in Figure 3.3c, and an overall average profit level of $0.05$.

The overall model implementation is a nonlinear system with the potential of converging or not, or converging to a local optimum instead of a global optimum. As such, it is preferable for the method to be insensitive to initial values (Judd, 1998) and other parameter values. Figure 3.1 shows that the results from impulse balance learning vary significantly with the initial values $0.95 \pm 0.02$, $0.85 \pm 0.02$, and $0.75 \pm 0.02$. However, Figure 3.4 shows that the ILA method is quite insensitive to the initial values. In Figure 3.1, the profit curve for an initial value $0.95 \pm 0.02$ is close to zero for low-value bidders and decreases rapidly to -1.7 for high-value bidders. As the initial value is decreased to $0.85 \pm 0.02$ and then to $0.75 \pm 0.02$, the values for mid-value and high-value bidders increase considerably so that the curve becomes much flatter. In Figure 3.4, the pattern of profit is much more similar across the initial values. The low-value and high-value bidders tend to have profit close to zero, with about twenty mid-value bidders with profits as high as 0.25 for all three initial values.

One of the main methodological differences between the experiments with humans in the various studies cited in this paper and these computational experiments is that here each bidder's private and common value signals are constant throughout the auctions (as explained in Section 3.2.2). The alternative is to provide each agent with a random value signal for each auction. This results in each agent learning an average bid strategy that is the adaptive best response to the full range of value signals. Perhaps the impulse balance method is more suitable to learning an average bid strategy. As shown in Figure 3.5, this is not the case. For varying common value signals, impulse-balance learning results in significant number of agents with high levels of negative profit. However, the ILA method results in most agents achieving positive profits, but with some achieving small negative profits.

### *3.3.5 Convergence and Sensitivity to Learning Rate*

***Result 2***: The ILA method results in value-multipliers that converge in less than 100 periods, and this convergence is independent of the initial values and smoother than the convergence of the impulse-balance method.

*Discussion*: Figure 3.6 shows value-multiplier convergence for the impulse balance and ILA methods. The impulse-balance value multipliers converge to quite different values (0.94, 0.89, and 0.83) for the three initial values, whereas the ILA value multipliers converge to more similar values of (0.92, 0.90, and 0.89). In addition, the pattern of the convergence is much smoother for the ILA method. For the three initial values, convergence requires about 10, 60, and 90 periods. These convergence results are important, since it is impossible to interpret auction results for profit, revenue, and efficiency when there is no convergence. For example, without convergence the results in period 50 are different from the results in period 60, whereas if the results from period 50 to infinity are the same, we can conclude that these are the results of the auction. Also, the fact that convergence occurs in less than 100 periods makes it reasonable to infer that the bid strategies of human agents could converge in a realistic number of real-world auctions.

***Result 3***: The ILA method is insensitive to the learning rate $\phi$.

*Discussion*: Why should we believe that a bidder's downward impulse is all of its loss and not a fraction of the loss, or that a bidder's upward impulse is all of its foregone profit and not a fraction of it? Using different values for the learning rate $\phi$ answers this question. Figure 3.7 shows the value multiplier for first-price auctions using a sample of learning rates in the interval $[0.1, 1.0]$. First, the resulting value multiplier is very close across all of the learning rates (0.92). Second, the pattern of variation is also very similar throughout the range. The difference is that the smaller learning rates tend to produce smoother convergence, with the standard deviation ranging from about 0.0017 for $\phi$=0.1 to 0.0024 for $\phi$=1.0. These results are quite insensitive to initial value, and I use $\phi = 0.5$ and initial values of $0.85 \pm 0.02$ in the simulations as arbitrary choices.

### 3.4 Comparing ILA Results with Lab Experiments

Using simulations with the ILA learning method, I present results and compare them with lab experiment results in Kagel *et alia* (1989), Kagel *et alia* (1995), Kagel and Richard (2001), and Goeree and Offerman (2002). In the lab experiments, the information feedback is equal to or greater than the I3 information level. Thus, in the figures related to this section, we are interested in only the I3 bidders, represented by the dashdot curves. Also, the most frequently-used support for the common value signals is [25, 225] so that is what I use in the computational experiments. For comparability with most of the lab experiment results, I vary the uncertainty using $\varepsilon$= 8, 12, 18, and 27 and use four and seven bidders for both first-price and second-price auctions. The results that are illustrated in the figures are consistent across repeated simulations.

In the following sub-sections I compare results for value multiplier, profit, and efficiency, and these results are summarized in Table 3.3. The value multiplier and profit are straightforward and have been discussed in Sections 3.1 to 3.3, but efficiency requires some explanation. Since efficiency refers to the auction being won by the bidder with the highest value *ex post*, all auctions are equally efficient when the value is 100% common. Thus, in the two-dimensional value environment, efficiency can be considered only when there is some component of private value, i.e., when the common value component is less than 100%. When the common value is 100%, Kagel *et alia* (1989) and Kagel *et alia* (1995) measure efficiency by the percent of auctions won by the bidder with the highest value signal (*ex ante*). This is not really efficiency, but it is still interesting to look at this highest-value-signal winning percent. When the common value percent is less than 100%, private value efficiency as used by Goeree and Offerman (2002) compares the winner's private value with the maximum private value among the bidders, i.e., $\mathcal{E} = \dfrac{v_{P,t}^{winner} - \min\{v_{P,t}^{i}\}}{\max\{v_{P,t}^{i}\} - \min\{v_{P,t}^{i}\}}$. $\mathcal{E} = 1$ when the winner has the highest private value, and $\mathcal{E} = 0$ when the winner is the bidder with the lowest private value. As the common value component increases, the common value signal may undermine the private value efficiency. Consider two bidders, one with a high private value (say corresponding to low costs of production) and one with a low private value. If the low private value bidder has a higher estimate of the common value than the high private value bidder, it may submit a higher bid and win the auction. This leads to private value inefficiency.

### *3.4.1 First-Price Auctions*

**Value Multiplier**: From intuition and theory, we expect that bidders will shade their bids in first-price auctions, i.e., the value multiplier is expected to be less than 1.0. For pure private values, experimental evidence from Kagel and Levin (1993) shows an average value multiplier of 0.92 for first-price auctions, averaged over experiments with five and 10 bidders with I3 information. The simulation results for seven I3 bidders in Figure 3.8 show value multipliers of about 0.90. These pure private-value results are consistent with the data from lab experiments. Kagel and Richard (2001) show value multipliers for first-price common values with $\varepsilon = 18$ of about 0.92 for four I3 bidders and 0.95 for seven I3 bidders in the middle region of the support. Figure 3.8 shows common value multipliers for $\varepsilon = 18$ of about 0.92 for four bidders and of about 0.94 for seven bidders. Thus, both the private-value and common-value results agree very closely with the experimental results.

**Profit**: Data in Kagel *et alia* (1989) for first-price pure common-value auctions averaging about seven bidders show that profit tends to increase with more uncertainty in the common value signal (higher $\varepsilon$). Figure 3.9 shows results for four and seven bidders across the full spectrum of two-dimensional value signals from pure private value (0% common value) to pure common value (100%). For pure common value, profit increases significantly with uncertainty for four bidders (from about 1 to 10), but increases less with uncertainty for seven bidders (from about 0 to about 3). Goeree and Offerman (2002) show profit increasing slightly with more uncertainty in auctions with two-dimensional value signals that are about 14% common value (with six bidders) and 25% common value (with three bidders). At common value percents of 14% and 25%, Figure 3.9 shows that the profit remains the same as the uncertainty increases. Thus, the agent results for the variation of profit with uncertainty are only partially in agreement with the lab experiment results.

Kagel and Richard (2001) find average profit is higher with four bidders than with seven bidders in first-price common-value auctions. Comparing the left column (four bidders) of Figure 3.9 with the right column (seven bidders) shows that profit tends to be higher for four bidders at all levels of uncertainty and all information levels across the full range of common-value percent. Thus, the agent results for the variation of profit with number of bidders are in agreement with the results from the lab experiments.

Figure 3.10 shows the profit results for 12, 25, 50, 100, 200, and 300 bidders for first-price auctions with high ($\varepsilon = 27$) uncertainty. For I2 and I3 bidders, profit further decreases with the increasing number of bidders, resulting in near-zero profits when there are over 200 bidders. The profit of I1 agents continues to decrease significantly below zero as the number of bidders increases. This highlights the importance to bidder profit of being informed about the payment.

**Efficiency**: The efficiency results for first-price auctions with less than pure common value are shown in Figure 3.11, and the results for highest-value-signal winning percent for pure common-value auctions are shown in Table 3.4.

From theory, we expect that efficiency will decrease when there is a common value component (Dasgupta and Maskin, 2000; Jehiel and Moldovanu, 2001). Experiment results from Kagel *et alia* (1989) for about seven bidders show that the highest-value-signal winning percent tends to decrease with more uncertainty in first-price auctions. Table 3.4 shows that for I3 agents the highest-value-signal winning percent in first-price auctions decreases with increased uncertainty for both four bidders (from 94% to 28%) and seven bidders (from 95% to 22%), which is consistent with the results from the lab experiments.

Goeree and Offerman (2002) show that efficiency is lower with more uncertainty for common value percents of 14% (with six bidders) and 25% (with three bidders). Figure 3.11 shows that for I3 bidders, efficiency tends to stay the same or decrease slightly with increased uncertainty at both 14% and 25% common value for both four and seven bidders, which is consistent with the results from the lab experiments.

### *3.4.2 Second-Price Auctions*

**Value Multiplier**: From intuition and theory for second-price auctions, we expect that bidders will bid their values when the value is private and shade their bids when the value is common. For private-value second-price auctions, experimental evidence from Kagel and Levin (1993) shows an average value multiplier of 1.02 for second-price auctions, averaged over experiments with five and ten bidders (assuming an average of about seven bidders). The results in Figure 3.12 for seven bidders show value multipliers of about 0.99 for private-value second-price auctions. For common-value second-price auctions, regressions in Kagel *et alia* (1995) show value multipliers of about 0.97, 0.96, 0.94, and 0.92 for $\varepsilon$=8, 12, 18, and 27 respectively. The results in Figure 3.12 for seven bidders show a similar magnitude and pattern of value multipliers, namely about 0.98, 0.97, 0.96, and 0.95.

**Profit**: For second-price common-value auctions, Kagel *et alia* (1995) find that profit increases with $\varepsilon$ for four bidders, but decreases with $\varepsilon$ for seven bidders. The results shown in Figure 3.13 are consistent with their results for four bidders (profit increases with $\varepsilon$), but not for seven bidders (no change in profit with $\varepsilon$). Figure 3.13 also shows that profit decreases slightly with an increase in the number of bidders, for all information levels and across the full range of common value percent. This is the same computational result that was obtained for private values. It is also consistent with the experimental results from Kagel *et alia* (1995) who found that profits were higher for four bidders than for seven.

**Efficiency**: The efficiency for second-price auctions with less than pure common value is shown in Figure 3.14, and the highest-value-signal winning percent for pure common-value auctions is shown in Table 3.4. Kagel *et alia* (1995) found that the highest-value-signal winning percent was lower in second-price auctions than in first-price auctions, when the level of uncertainty is $\varepsilon = 27$. However, Table 3.4 shows that the highest-value-signal winning percent is higher in second-price auctions. Second-price auctions are more efficient than first-price auctions for the agents because they are bidding closer to their values. Since this is what is expected from theory, the agents are bidding more like optimizing agents than like the inexperienced agents in the experiments. The results of Kagel *et alia* (1995) also show that highest-value-signal winning percent in second-price auctions is slightly lower for seven bidders than for four bidders. Table 3.4 shows that the agents produce similar results for levels of uncertainty above $\varepsilon = 8$.

## 3.5 Variation of Profit, Revenue, and Efficiency with Common Value Percent

The second contribution of this study is to determine how the auction results change as the common value component increases, and specifically whether the change is linear. In the figures used in the previous section, it is obvious that the results across the two-dimensional value signal are usually not linear. As the common value percent increases, profit, revenue, and efficiency all decrease monotonically, but they decrease in different ways. The seller endeavors to choose the payment rule and information level that maximizes its revenue, maximizes efficiency, or maximizes both. Therefore, the seller is interested in whether the different payment rules and information levels produce different levels of revenue and efficiency, or whether they are equivalent. In this section, I discuss the results for all three information levels.

### *3.5.1 Profit*

***Result 4***: Profit curves decrease nonlinearly for first-price auctions and linearly for second-price auctions. In first-price auctions the nonlinearity usually involves decreasing faster at higher common value percents.

*Discussion*: See Figure 3.9 for first-price auctions and Figure 3.13 for second-price auctions. The main difference between learning in first-price auctions and learning in second-price auctions is that the agents use money on the table in first-price auctions but not in second-price auctions. When the value signal is dominated by private value (i.e., a low common value percent), the bid reduction from money on the table keeps the profit high. As the common value component increases, the contributions from money on the table to profit become dominated by the effects of the common value signal.

There is also some interesting variation with the level of information feedback. The curves tend to be the same for I1, I2, and I3 information levels at lower levels of uncertainty, but as uncertainty about the common value increases profit is higher for the more informed I3 bidders.

### *3.5.2 Revenue*

See Figures 3.15 and 3.16 for revenue results for first-price and second-price auctions, respectively. Revenue tends to decrease with increasing common value.

***Result 5***: Revenue curves decrease nonlinearly for first-price auctions and linearly for second-price auctions. In first-price auctions the nonlinearity usually involves decreasing faster at lower common value percents.

***Result 6***: In most cases, the seller receives less revenue when it provides bidders with more information feedback.

*Discussion*: The figures show that in most cases the seller receives less revenue when the bidders have I3 information. However, for first-price auctions (Figure 3.15) with lower levels of uncertainty ($\varepsilon \leq 18$), I3 agents provide higher revenue at high common value percents (30% to 90%) than the I1 and I2 agents, although this effect diminishes with more bidders. Once again, the major difference between I3 agents and the I1 and I2 agents is that the former can calculate money on the table while the latter can only estimate it. The estimate becomes less reliable as the uncertainty increases so that I3 agents are better able to keep their bid strategies profitable, taking more of the surplus and yielding lower revenue for the seller.

### ***3.5.3 Efficiency***

See Figure 3.11 for first-price auctions and Figure 3.14 for second-price auctions. As the common value component increases, the common value signal disrupts the private value efficiency. Second-price auctions tend to be more efficient than first-price auctions because the agents bid closer to their values.

***Result 7***: Private-value efficiency curves tend to stay high at low percent common value and then decrease rapidly for higher common value percents.

*Discussion*: The nonlinearity is especially pronounced in second-price auctions (Figure 3.14) where efficiency remains close to 1.0 until relatively high levels of common value percent, and then decreases rapidly to efficiency as low as 0.8.

## 3.6 Revelation of Common Value to Losers

The third contribution of this study is to determine whether it may be worthwhile for a seller (such as a federal or state government) to enforce truthful revelation of the true common value by auction winners. In the experiments studied so far, and in nearly all real-world auctions, losing bidders do not know the actual common value. The winner discovers the true common values after the auction. For example, in timber sale auctions the winners learn the true quantity and value of timber; in highway procurement auctions, the winner discovers the true scope of the project; in oil lease licences, the winner discovers the true quantity of oil; and so on. These values are carefully guarded company secrets (Baldwin *et alia*, 1997) and are not intentionally revealed to other bidders. However, some experimental work has studied the effects on bidding of revealing some information about the common value to all bidders (Kagel and Levin (1999) for first-price auctions and Kagel *et alia* (1995) for second-price auctions). This raises the question of whether a buyer or seller, say the government operating procurement or asset-sale auctions, should require the auction winners to reveal the common value that they discover after winning. I know of some attempts to do this in Canadian federal government procurement auctions in which the government asks bidders to reveal their costs. Of course the costs provided are not truthful! If it were worthwhile, the government could rationally decide to invest in implementing regulations and enforcement of truthful revelation of the winner's value. Given this information, losing agents would use it in their calculation of the amount to increase their bid (in Rules L1 and L2). To obtain a computational answer to this question, I perform experiments with value revelation and observe the revenue. Figure 3.17 shows for I3 bidders the differences between the revenue with revealed common value and the revenue without revelation. The results that are illustrated in the figures are consistent across repeated simulations.

***Result 8***: For first-price auctions, when the common value percent is high (>60%) and there is a high degree of uncertainty in the common value signal ($\varepsilon > 12$, revealing information about the common value increases revenue for I3 information.

*Discussion*: Experiments by Kagel and Levin (1999) for first-price common-value auctions with $\varepsilon = 27$ show that revealing information about the true common value increased revenue (+2.75) for four bidders but decreased revenue (-0.88) for seven bidders. For the same conditions, the top row of Figure 3.17 shows that revenue increases (+7) for four bidders and increases less (+3) for seven bidders.[9] These results are consistent with the experimental data in that revenue increases more for four bidders than for seven, but is inconsistent in the direction of change for the seven bidders. Figure 17 also shows that the revenue effects are smaller as uncertainty decreases. For 100% common value, the benefits become negligible when $\varepsilon = 8$. For values with less common value percent, common value revelation sometimes has a negative impact on revenue. Thus, for the auction designer the percent common value and the degree of uncertainty about the common value all affect the impact of value revelation on revenue.

---

9 Since total revenue is approximately 125, the revenue increase of 7 is approximately 5%.

***Result 9***: For second-price auctions, revealing information about the common value significantly increases revenue for I3 bidders, especially when there is a high degree of uncertainty in the common value signal.

*Discussion*: Experiments by Kagel *et alia* (1995) show that revealing information about the common value in second-price auctions increased revenue (+0.31) for four-five bidders but decreased revenue (-2.5) for six-seven bidders. The bottom row of Figure 3.17 shows that revenue increases (+4) for four bidders and increases less (+2) for seven bidders. Again, the results are in partial agreement with the experimental results.

## 3.7 Conclusion

I find that Selten's impulse balance method can be adapted for use in multiagent simulations of auctions with values that have some common value component (Result 1 in Section 3.3.4). The resulting ILA (impulse learning with loss aversion) method converges within 100 periods and is insensitive to the learning rate (Results 2 and 3 in Section 3.3.5).

I use the ILA method in multiagent simulations for first-price and second-price payment rules, three different information levels, and two-dimensional value signals that vary from pure private value to pure common value. The results are compared to data from lab experiments in other studies (summarized in Table 3.3), and the agent results for the value multiplier, profit, and efficiency are usually consistent with results from lab experiments. These consistencies support the real-world validity in this context of using multiagent simulations with learning agents.

For values in between pure private and pure common value, curves for profit, revenue, and efficiency are nonlinear especially when the payment rule is first price. The profit curves tend to decrease nonlinearly for first-price auctions and linearly for second-price auctions (Result 4 in Section 3.5.1). The nonlinear revenue curves tend to decrease more rapidly at low common value percents (Result 5 in Section 3.5.2). The very nonlinear efficiency curves tend to stay high and then decrease rapidly for common value percents (Result 7 in Section 3.5.3). In addition, revenue in most cases decrease with increasing information feedback to the bidders (Result 6 in Section 3.5.2).

Simulations also show that forcing revelation of the true common value may have beneficial revenue effects when the common-value percent is high and there is a high degree of uncertainty about the common value (Results 8 and 9 in Section 3.6).

Using multiagent simulations has provided some insights into single-unit sealed-bid auction performance for different levels of information feedback across different levels of common value. The next paper will expand the approach to analyze English auctions.

### 3.8 References

Athey, S., Haile, P., 2002. "Identification of Standard Auction Models." *Econometrica*. 70: 2107-2140.

Byde, A., 2002. "Applying Evolutionary Game Theory to Auction Mechanism Design." HPL-2002-321, Hewlet-Packard Company.

Camerer, C., Ho, T. H., Chong, J. K., 2002. "Sophisticated Experience-Weighted Attraction Learning and Strategic Teaching in Repeated Games." *Journal of Economic Theory*. 104: 137-188.

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

Decker, K.S., Lesser, V. R., 1995. "Designing a Family of Coordination Algorithms." U. Mass Computer Science Technical Report 94-14.

Dufwenberg, M., Gneezy, U., 2002. "Information disclosure in auctions: an experiment." *Journal of Economic Behavior and Organization*. 48: 431-444.

Garvin, S., Kagel, J.H., 1994. "Learning in Common-value auctions: Some Initial Observations." *Journal of Economic Behavior and Organization*. 25: 351-372.

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

Haile, P. A., Hong, H., Shum, M., 2003. "Nonparametric Tests for Common Values in First-Price Sealed-Bid auctions." *NBER Working Paper No. 10105*.

Hailu, A., Schilizzi, S., 2004. "Are Auctions More Efficient Than Fixed Price Schemes When Bidders Learn?" *Australian Journal of Management*. 29: 147-168.

Hendricks, K., Pinkse, J., Porter, R.H., 2003. "Empirical Implications of Equilibrium Bidding in First-Price, Symmetric, Common-value auctions." *Review of Economic Studies*. 70: 115-145.

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

Judd, K.L., 1998, *Numerical Methods in Economics*. The MIT Press.

Kagel, J.H., Levin, D., 1993. "Independent Private Value Auctions: Bidder Behavior in First-, Second- and Third-Price Auctions with Varying Numbers of Bidders." *Economic Journal*. 108: 868-879.

Kagel, J.H., Levin, D., 1999. "Common-value auctions with Insider Information." *Econometrica*. 67: 1219-1238.

Kagel, J.H., Levin, D., 2002. *Common-value auctions and the Winner's Curse*. Princeton University Press.

Kagel, J.H., Levin, D., Harstad, R.M, 1995. "Comparative Static Effects of Number of Bidders and Public Information on Behavior in Second-Price Common-value auctions." *International Journal of Game Theory*. 24: 293-319.

Kagel, J.H., Levin, D., Battalio, R.C, Meyer, D.J., 1989. "First-Price Common-value auctions: Bidder Behavior and the Winner's Curse." *Economic Inquiry*. 27: 241-258.

Kagel, J.H., Richard, J.F. , 2001. "Super-Experienced Bidders in First-Price Common-Value Auctions: Rules of Thumb, Nash Equilibrium Bidding, the Winner's Curse." *Review of Economic Statistics*. 83: 408-419.

Kim, Y. S., 2007. "Maximizing sellers' welfare in online auction by simulating bidders' proxy bidding agents." *Expert Systems with Applications*. 32: 289-298.

Neugebauer, T., Selten, R., 2006. "Individual Behavior of first-price auctions: The importance of information feedback in computerized experimental markets." *Games and Economic Behavior*. 54: 183-204.

Ockenfels, A., Selten, R., 2005. "Impulse Balance Equilibrium and Feedback in First Price Auctions." *Games and Economic Behavior*. 51: 155-170.

Selten, R., 1998. "Features of Experimentally Observed Bounded Rationality." *European Economic Review*. 42: 413-436.

Selten, R., Abbink, K., Cox, R., 2005. "Learning Direction Theory and the Winner's Curse." *Experimental Economics*. 8: 5-20.

Selten, R., Buchta, J., 1998. "Experimental Sealed Bid First Price Auctions with Directly Observed Bid Functions." In: Budescu, D. V., Erve, I., Zwick, R. (Eds), *Games and Human Behavior: Essays*. Lawrence Erlbaum Associates, pp. 53-78.

Tesfatsion, L., 2003. "Agent-based computational economics: modeling economies as complex adaptive systems." *Information Sciences*. 149: 263-269.

Tesfatsion, L., Judd, K.L., 2006, *Handbook of Computational Economics : Agent-Based Computational Economics*. North Holland.

Tversky, A., Kahneman, D.,1992, "Advances in Prospect Theory: Cumulative Representation of Uncertainty." *Journal of Risk and Uncertainty*. 5: 297-323.

## 3.9 Tables

| Table 3.1. Sealed-Bid Model Notation Summary | |
|---|---|
| Symbol | Description |
| $b_t^i$ | Bid price of bidder i in auction t. |
| $b_t^{(1)},...,b_t^{(n)}$ | Ordered bid prices in a sealed-bid auction where $b_t^{(1)}$ is the highest bid. |
| $\varepsilon$ | Radius of the support for the common value signal. . |
| $\gamma_t^i$ | Value multiplier: $b_t^i = \gamma_t^i \hat{v}_t^i$ . |
| $\lambda_t^i$ | The balance weight in the impulse balance learning method. |
| $m_t^i$ | Money left on the table by a profitable winner for first-price payment: $m_t^i = b_t^{(1)} - b_t^{(2)}$ . |
| $p_t$ | Payment by winner in auction t. |
| $\phi$ | Learning rate of bidder i at period t. |
| $p_t^i$ | Payment made by bidder i, given that it wins. |
| $\pi_t^i$ | Profit of bidder i in auction t. |
| $\pi_{F,t}^i$ | Foregone profit of bidder i in auction t. |
| $r_t^i$ | Ranking of bidder i in auction t. The winner is $r_t^i = 1$ , the runner-up $r_t^i = 2$ , etc. |
| $\overline{S}_{\mathrm{P}},\underline{S}_{\mathrm{P}},\overline{S}_{\mathrm{C}},\underline{S}_{\mathrm{C}}$ | Upper and lower bounds of supports for the private value signal and common value. |
| $\theta_C$ | Common value component of the value signal. |
| $\hat{v}_t^i$ | Value signal of bidder i in auction t. |
| $v_t^i$ | Actual value, revealed only to winner. |

| Table 3.2. Information Levels (incremental) | | |
|---|---|---|
| Level | Description | Feedback |
| I1 | Number of bidders | n |
| | Value signal support | $[\underline{S}_C, \overline{S}_C]$ |
| | Value Signal: Own | $\hat{v}_t^i$ |
| | Bid Price: Own | $b_t^i$ |
| | Ranking: Own | $r_t^i$ |
| | Payment: Own | $p_t^i \mid r_t^i = 1$ |
| | Value: Own | $v_t^i \mid r_t^i = 1$ |
| I2 | Bid Price: Winner | $b_t^{(1)}$ |
| | Payment | $p_t$ |
| I3 | Bid Price: Runner-up | $b_t^{(2)}$ |

| Table 3.3. Summary of Comparison with Lab Experiments (Section 3.3.6) | | | | | | | |
|---|---|---|---|---|---|---|---|
| | | Private Value | | 14%, 25% Common Value | | 100% Common Value | |
| Payment | Result | Lab | Agent | Lab | Agent | Lab | Agent |
| First Price | Value Multiplier | 0.92 (1) | 0.90 | | | 0.92, 0.95 (2) | 0.92, 0.94 |
| | Profit | | | Increases with ε (4) | Constant with ε | Increases with ε and n (2,3) | Increases with ε and n |
| | Efficiency | | | Decreases with ε (4) | Decreases slightly with ε | Decreases with ε (3) | Decreases with ε |
| Second Price | Value Multiplier | 1.02 (1) | 0.99 | | | 0.97, 0.96, 0.94, 0.92 (5) | 0.98, 0.97, 0.96, 0.95 |
| | Profit | | | | | Increases with ε for n=4 Decreases with ε for n=7 Decreases with n (5) | Increases with ε for n=4 No change with ε for n=7 Decreases with n |
| | Efficiency | | | | | Decreases with n (5) | Decreases with n |

1 Kagel and Levin (1993)
2 Kagel and Richard (2001)
3 Kagel *et alia* (1989)
4 Goeree and Offerman (2002)
5 Kagel *et alia* (1995)

Table 3.4. Highest-Value-Signal Winning Percent for Pure Common Value Information Level I3

| Payment | Uncertainty $\varepsilon$ | Four Bidders | Seven Bidders |
|---|---|---|---|
| First Price | 8 | 93% | 95% |
| | 12 | 88% | 72% |
| | 18 | 44% | 30% |
| | 27 | 28% | 22% |
| Second Price | 8 | 94% | 94% |
| | 12 | 81% | 92% |
| | 18 | 61% | 72% |
| | 27 | 44% | 46% |

## 3.10 Figures

Figure 3.1. Impulse Balance Learning: Profit by Common Value Signal
Common Value, First Price, I3

| Initial value multiplier $0.95 \pm 0.02$ | Initial value multiplier $0.85 \pm 0.02$ | Initial value multiplier $0.75 \pm 0.02$ |
|---|---|---|
| 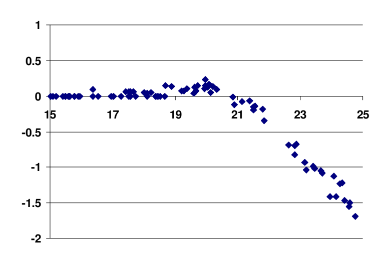 | 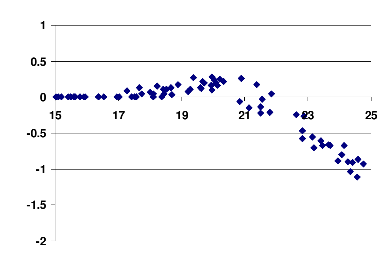 | 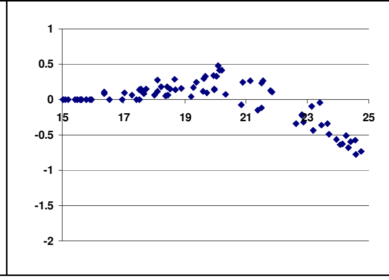 |

Figure 3.2. Impulse Balance Learning: $\lambda_t^i$ by Common Value Signal
Common Value, First Price, I3

| Initial value multiplier $0.95 \pm 0.02$ | Initial value multiplier $0.85 \pm 0.02$ | Initial value multiplier $0.75 \pm 0.02$ |
|---|---|---|
| 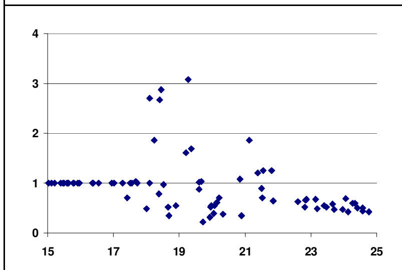 | 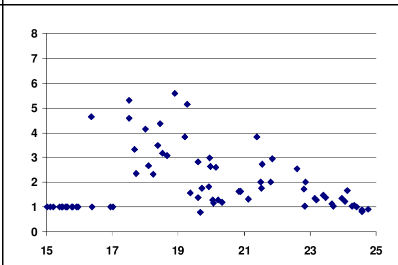 | 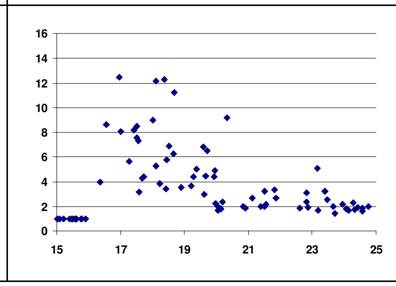 |

Figure 3.3. Learning Alternatives: Profit by Common Value Signal
Common Value, First Price, I3, Initial value multiplier $0.95 \pm 0.02$

| 3a. No Balance | 3b No Balance, Loss Aversion | 3c. No Balance, Loss Aversion, Profit Switch |
|---|---|---|
| 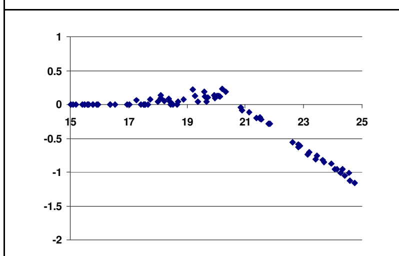 | 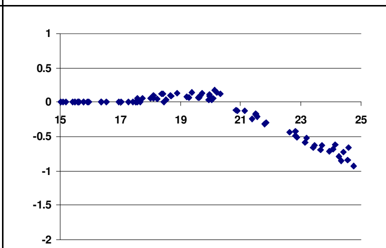 | 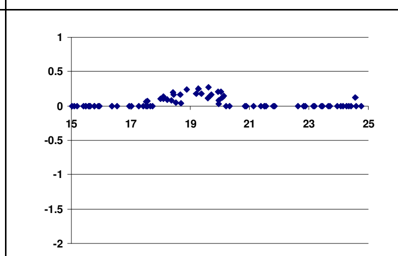 |

| Figure 3.4. ILA learning: Profit by Common Value Signal<br>Common Value, First Price, I3 | | |
|---|---|---|
| Initial value multiplier $0.95 \pm 0.02$ | Initial value multiplier $0.85 \pm 0.02$ | Initial value multiplier $0.75 \pm 0.02$ |
| 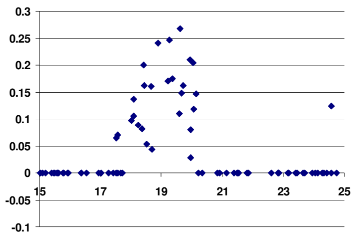 | 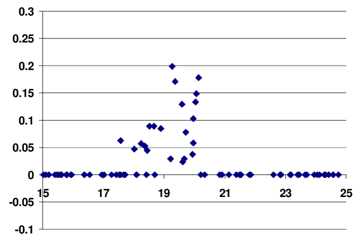 | 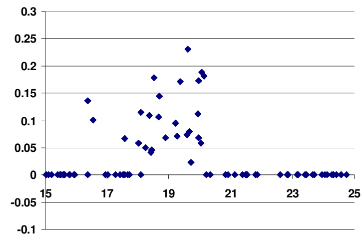 |

| Figure 3.5. Bidders with Varying Value Signals<br>Profit by Common Value Signal<br>Initial value multiplier $0.95 \pm 0.02$<br>Common Value, First Price, I3 | |
|---|---|
| Impulse Balance | ILA |
| 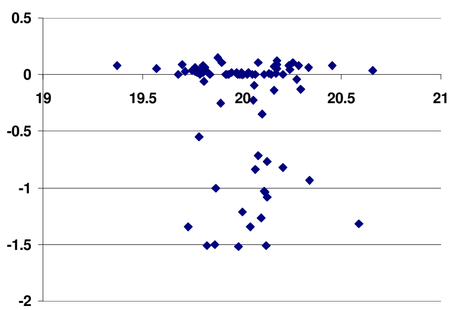 | 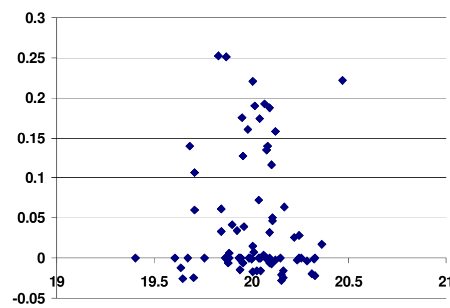 |

Figure 3.6. Convergence of Value Multiplier
Value Multiplier by Period
Common Value, First Price, I3 Information

| Impulse Balance | ILA |
|---|---|
| Initial Value multiplier: $0.95 \pm 0.02$ | Initial Value multiplier: $0.95 \pm 0.02$ |
| 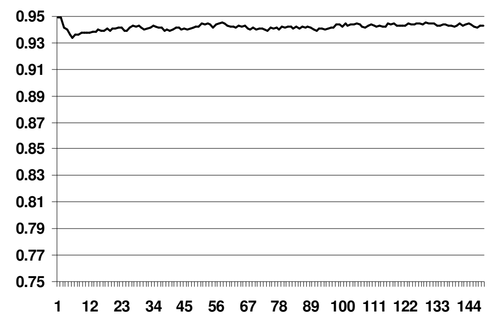 | 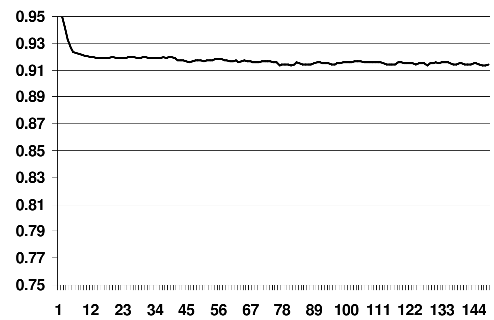 |
| Initial Value multiplier: $0.85 \pm 0.02$ | Initial Value multiplier: $0.85 \pm 0.02$ |
| 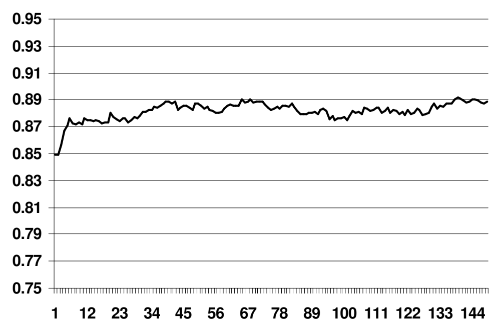 | 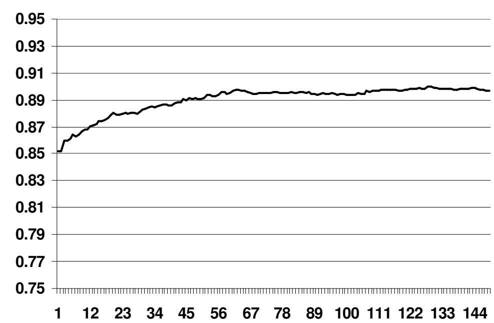 |
| Initial Value multiplier: $0.75 \pm 0.02$ | Initial Value multiplier: $0.75 \pm 0.02$ |
| 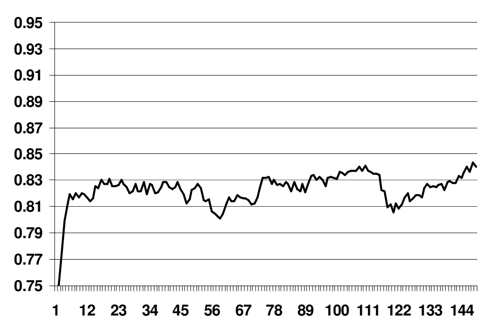 | 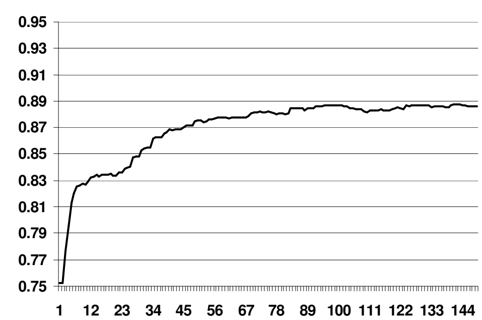 |

Figure 3.7. Sensitivity to Learning Rate
Value Multiplier by Period
Common Value, First Price, I3 Information
(I1: Solid, I2: Dot, I3: DashDot)

Learning Rate 0.1

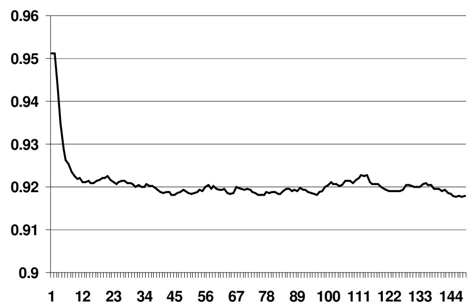


Learning Rate 0.3

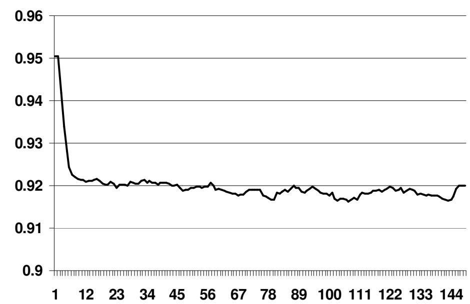


Learning Rate 0.5


Learning Rate 0.7

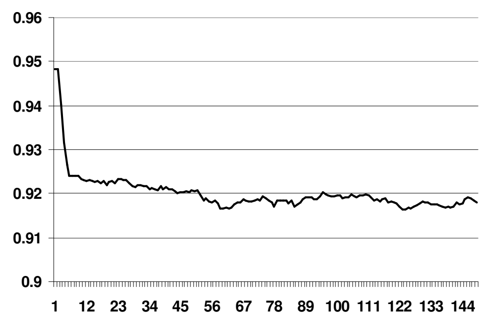


Learning Rate 0.9

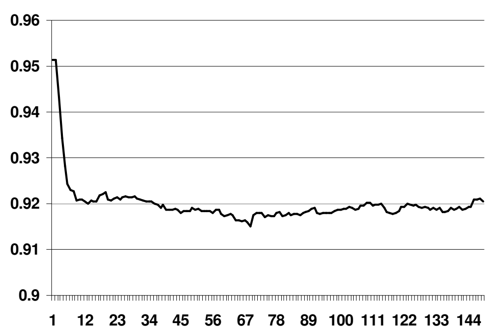


Learning Rate 1.0

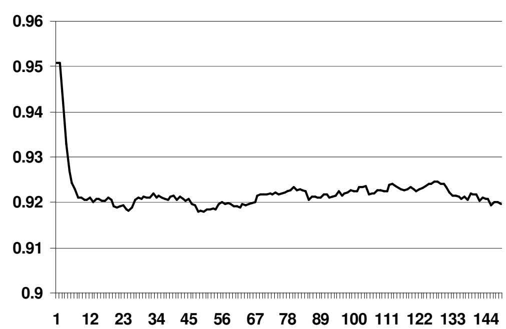

| Figure 3.8. Value Multiplier: First-price auctions<br>Value Multiplier by Common Value Percent<br>(I1: Solid, I2: Dot, I3: DashDot) | |
|---|---|
| Four Bidders | Seven Bidders |
| $\varepsilon = 8$ | $\varepsilon = 8$ |
| 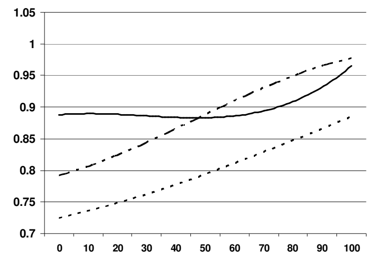 | 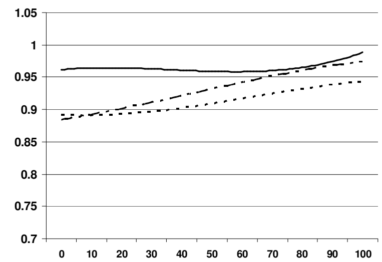 |
| $\varepsilon = 12$ | $\varepsilon = 12$ |
| 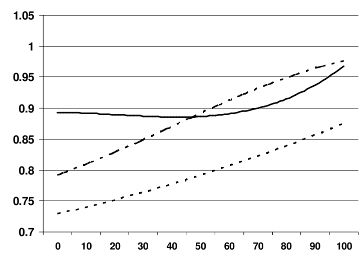 | 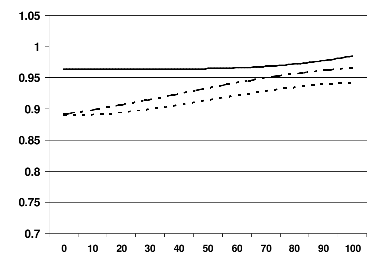 |
| $\varepsilon = 18$ | $\varepsilon = 18$ |
| 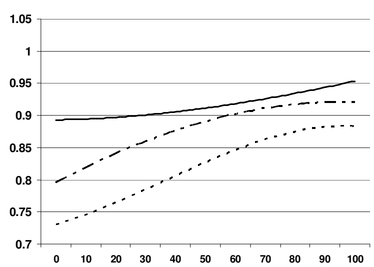 | 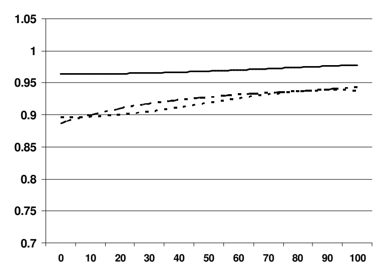 |
| $\varepsilon = 27$ | $\varepsilon = 27$ |
| 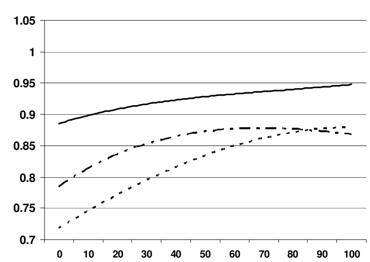 | 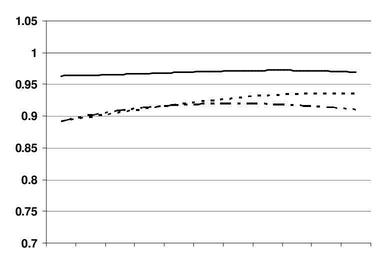 |

| Figure 3.9. Profit: First-price auctions<br>Profit by Common Value Percent<br>(I1: Solid, I2: Dot, I3: DashDot) | |
|---|---|
| Four Bidders | Seven Bidders |
| $\varepsilon = 8$ | $\varepsilon = 8$ |
| 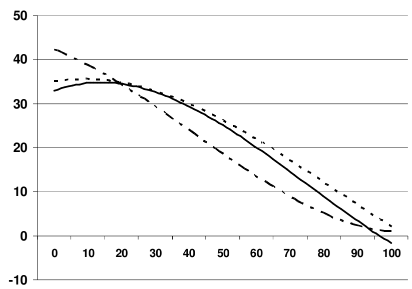 | 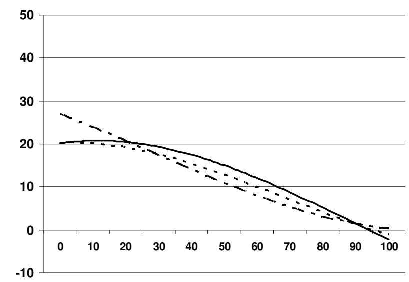 |
| $\varepsilon = 12$ | $\varepsilon = 12$ |
| 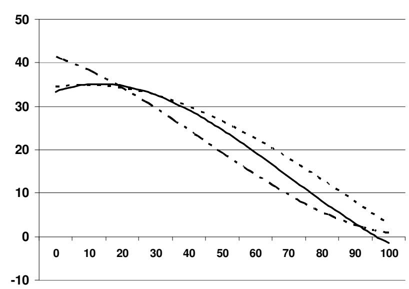 | 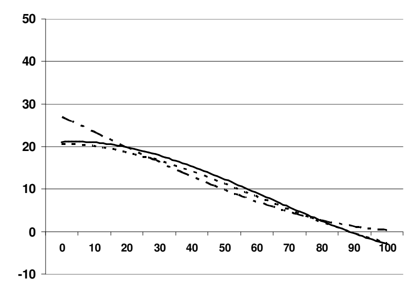 |
| $\varepsilon = 18$ | $\varepsilon = 18$ |
| 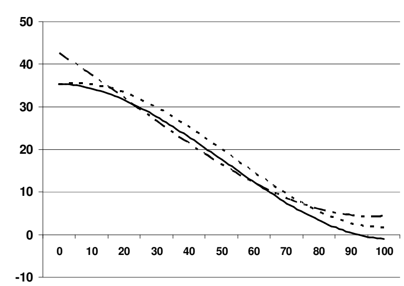 | 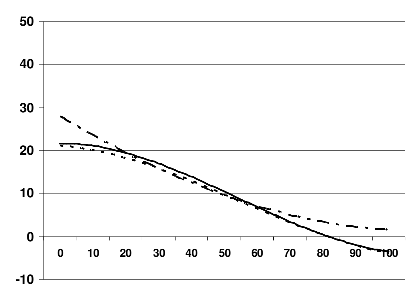 |
| $\varepsilon = 27$ | $\varepsilon = 27$ |
| 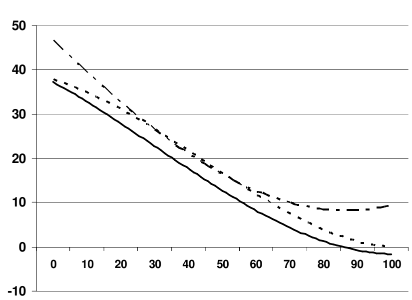 | 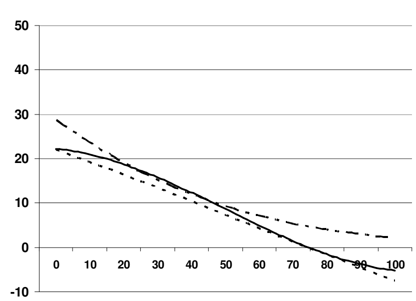 |

Figure 3.10. Profit with Many Bidders
Profit by Common Value Percent
First-price Auctions, $\varepsilon = 27$
(I1: Solid, I2: Dot, I3: DashDot)

12 Bidders

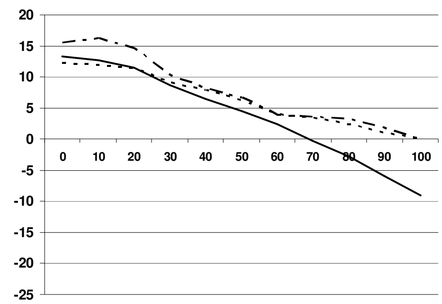


25 Bidders

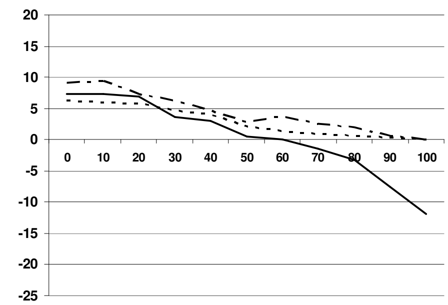


50 Bidders

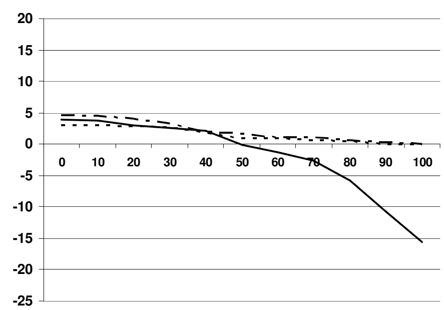


100 Bidders

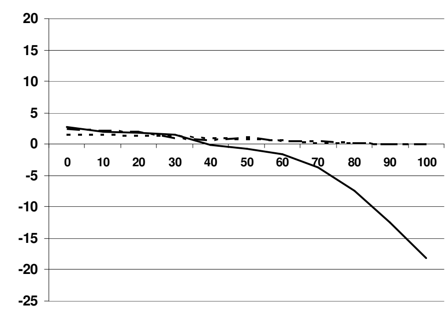


200 Bidders

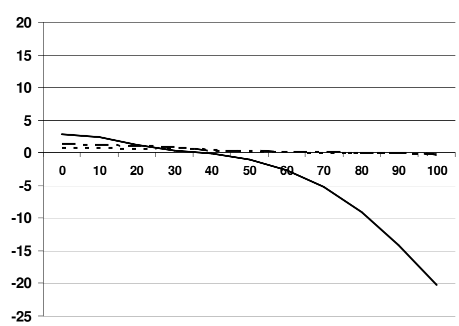


300 Bidders

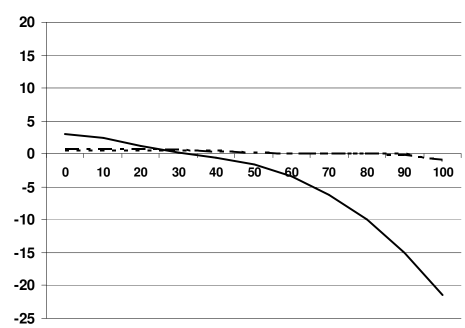

<table>
<tr><td colspan="2">Figure 3.11. Efficiency: First-price auctions<br>Efficiency by Common Value Percent<br>(I1: Solid, I2: Dot, I3: DashDot)</td></tr>
<tr><td>Four Bidders</td><td>Seven Bidders</td></tr>
<tr><td>$\varepsilon = 8$</td><td>$\varepsilon = 8$</td></tr>
<tr><td>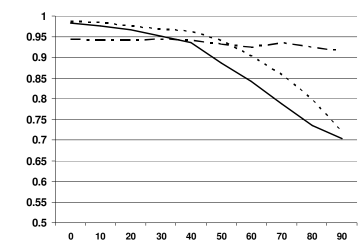</td><td>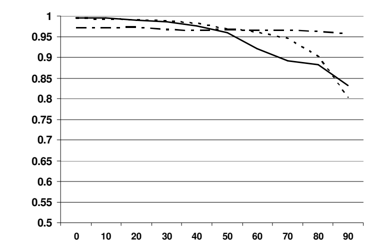</td></tr>
<tr><td>$\varepsilon = 12$</td><td>$\varepsilon = 12$</td></tr>
<tr><td>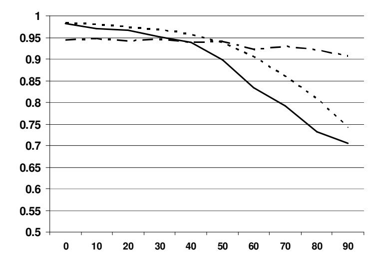</td><td>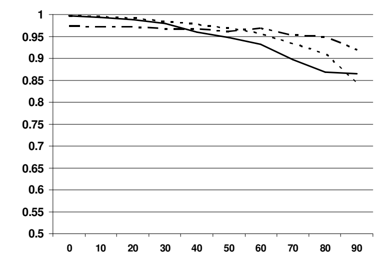</td></tr>
<tr><td>$\varepsilon = 18$</td><td>$\varepsilon = 18$</td></tr>
<tr><td>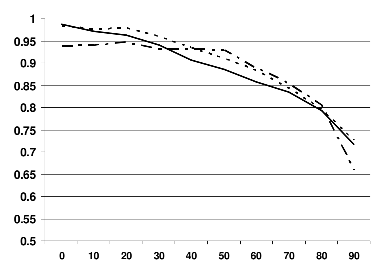</td><td>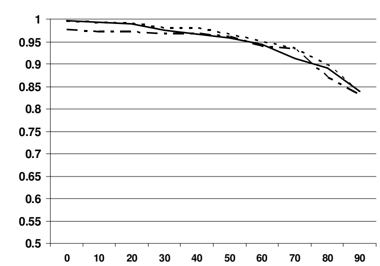</td></tr>
<tr><td>$\varepsilon = 27$</td><td>$\varepsilon = 27$</td></tr>
<tr><td>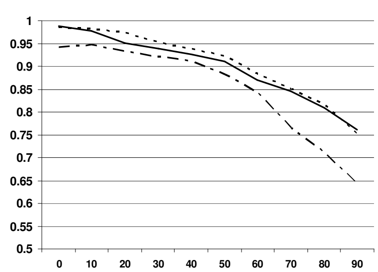</td><td>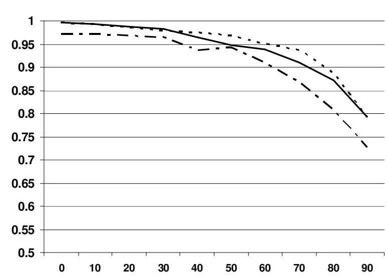</td></tr>
</table>

Figure 3.12. Value Multiplier: Second-price auctions
Value Multiplier by Common Value Percent
(I1: Solid, I2: Dot, I3: DashDot)

Four Bidders

$\varepsilon = 8$

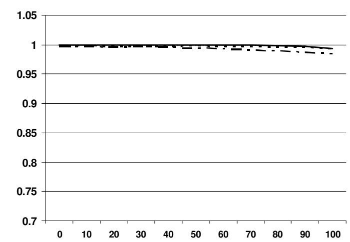


$\varepsilon = 12$

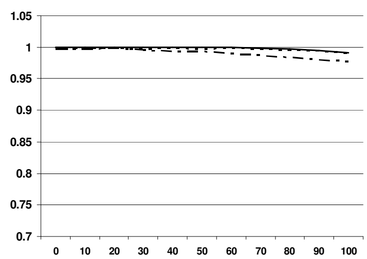


$\varepsilon = 18$

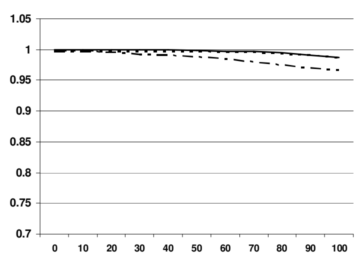


$\varepsilon = 27$

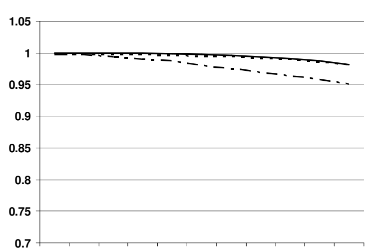


Seven Bidders

$\varepsilon = 8$

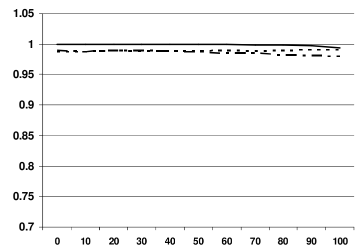


$\varepsilon = 12$

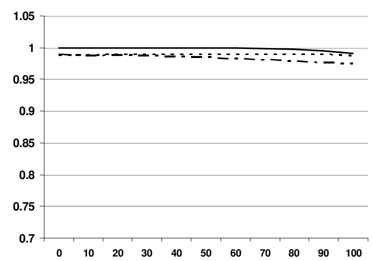


$\varepsilon = 18$

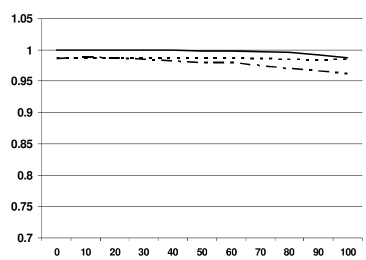


$\varepsilon = 27$

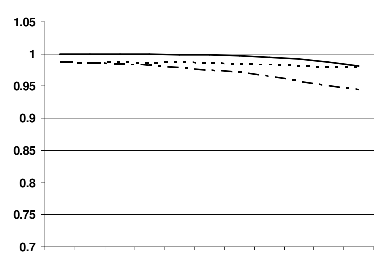

Figure 3.13. Profit: Second-price auctions
Profit by Common Value Percent
(I1: Solid, I2: Dot, I3: DashDot)

| Four Bidders | Seven Bidders |
|---|---|
| $\varepsilon = 8$ | $\varepsilon = 8$ |
| 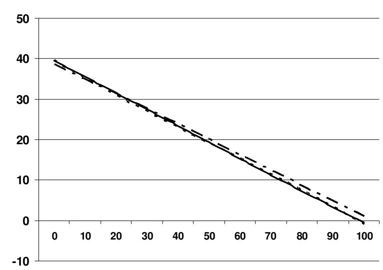 | 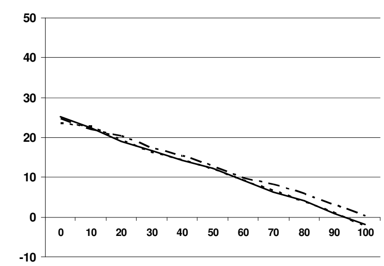 |
| $\varepsilon = 12$ | $\varepsilon = 12$ |
| 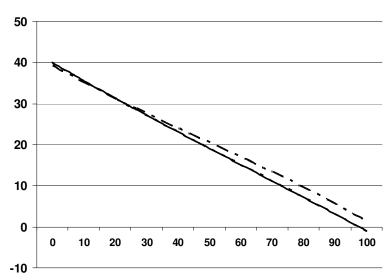 | 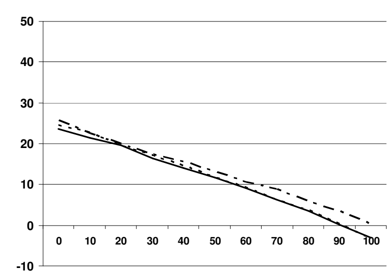 |
| $\varepsilon = 18$ | $\varepsilon = 18$ |
| 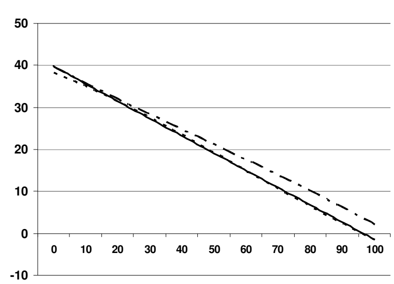 | 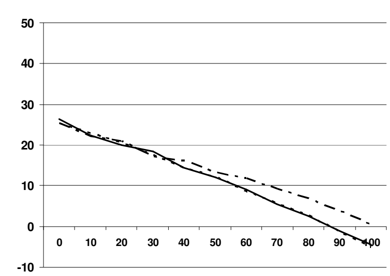 |
| $\varepsilon = 27$ | $\varepsilon = 27$ |
| 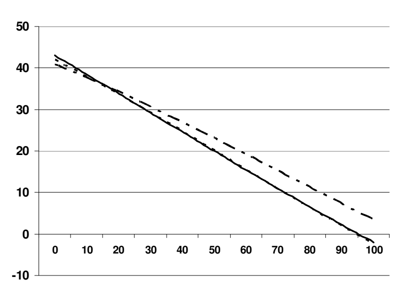 | 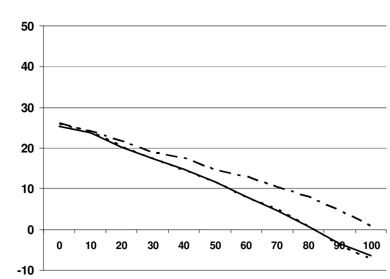 |

Figure 3.14. Efficiency: Second-price auctions
Efficiency by Common Value Percent
(I1: Solid, I2: Dot, I3: DashDot)

| Four Bidders | Seven Bidders |
|---|---|
| $\varepsilon = 8$ | $\varepsilon = 8$ |
| 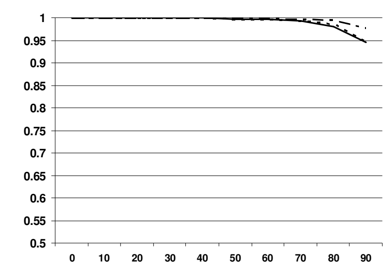 | 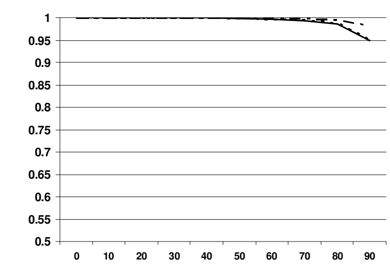 |
| $\varepsilon = 12$ | $\varepsilon = 12$ |
| 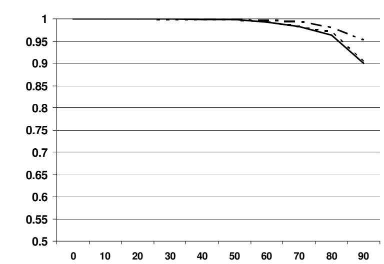 | 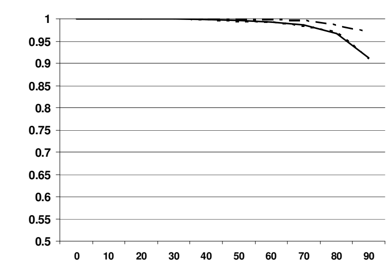 |
| $\varepsilon = 18$ | $\varepsilon = 18$ |
| 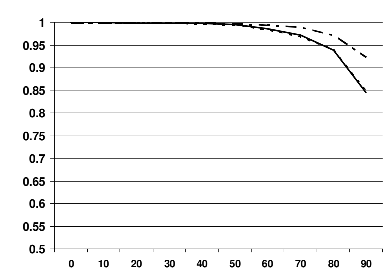 | 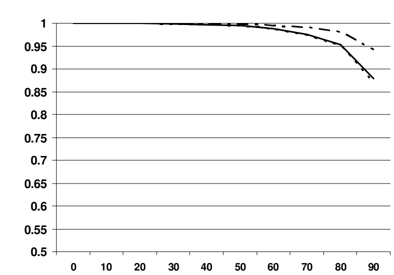 |
| $\varepsilon = 27$ | $\varepsilon = 27$ |
| 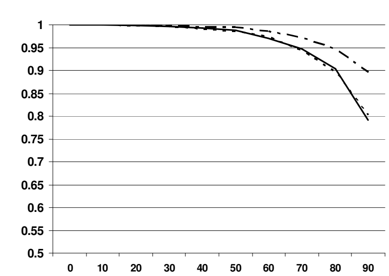 | 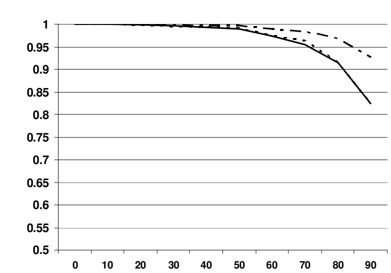 |

Figure 3.15. Revenue: First-Price Auctions
Revenue by Common Value Percent
(I1: Solid, I2: Dot, I3: DashDot)

| Four Bidders | Seven Bidders |
|---|---|
| $\varepsilon=8$ | $\varepsilon=8$ |
| 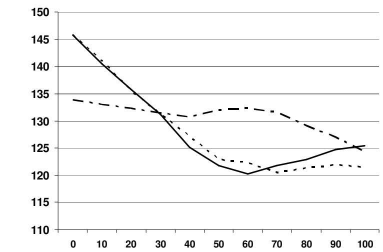 | 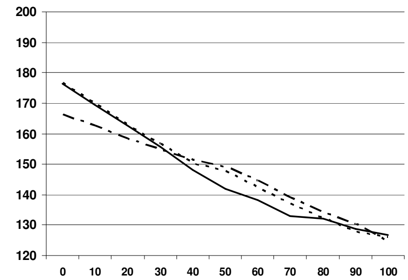 |
| $\varepsilon=12$ | $\varepsilon=12$ |
| 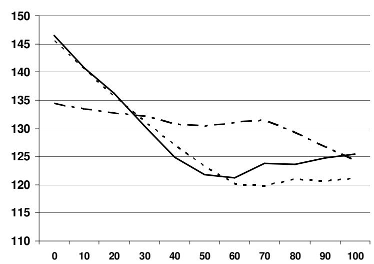 | 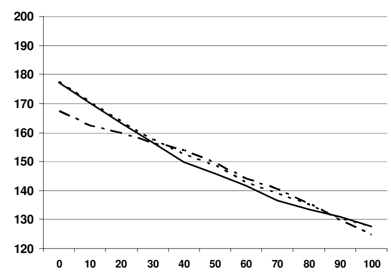 |
| $\varepsilon=18$ | $\varepsilon=18$ |
| 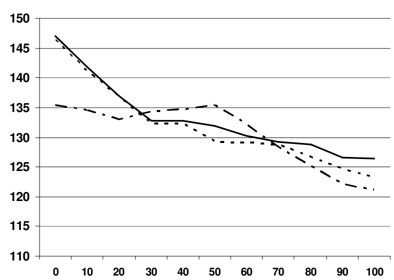 | 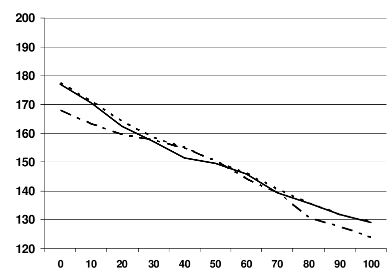 |
| $\varepsilon=27$ | $\varepsilon=27$ |
| 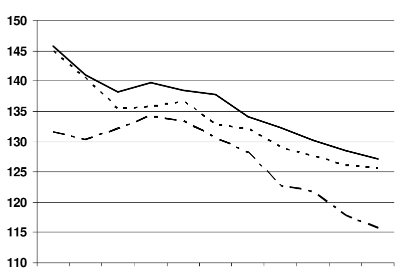 | 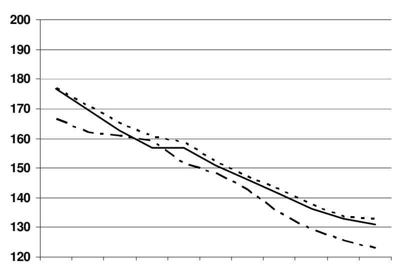 |

<table>
<tr><td colspan="2">Figure 3.16. Revenue: Second-Price Auctions<br>Revenue by Common Value Percent<br>(I1: Solid, I2: Dot, I3: DashDot)</td></tr>
<tr><td>Four Bidders</td><td>Seven Bidders</td></tr>
<tr><td>$\varepsilon = 8$</td><td>$\varepsilon = 8$</td></tr>
<tr><td>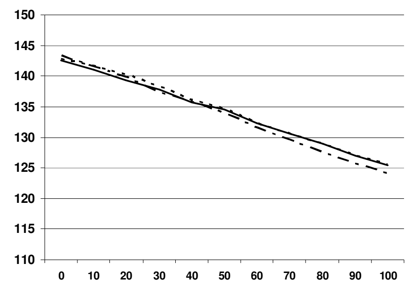
</td><td>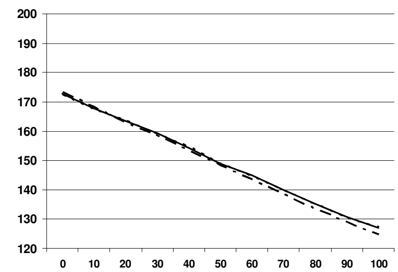
</td></tr>
<tr><td>$\varepsilon = 12$</td><td>$\varepsilon = 12$</td></tr>
<tr><td>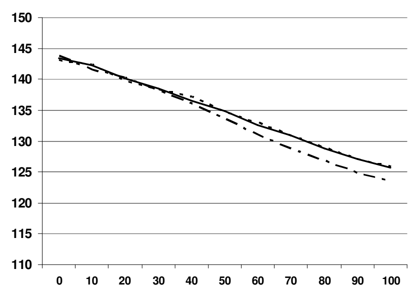
</td><td>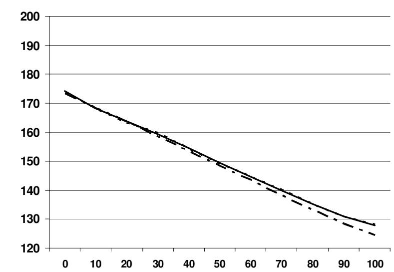
</td></tr>
<tr><td>$\varepsilon = 18$</td><td>$\varepsilon = 18$</td></tr>
<tr><td>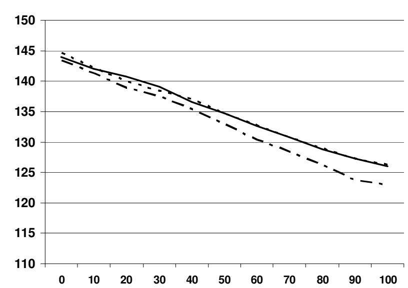
</td><td>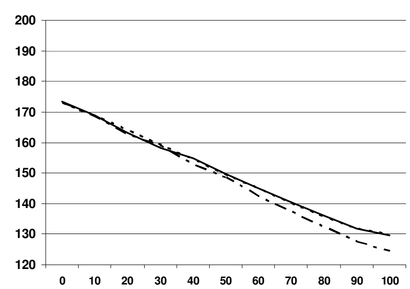
</td></tr>
<tr><td>$\varepsilon = 27$</td><td>$\varepsilon = 27$</td></tr>
<tr><td>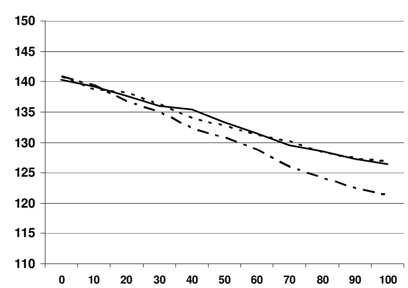
</td><td>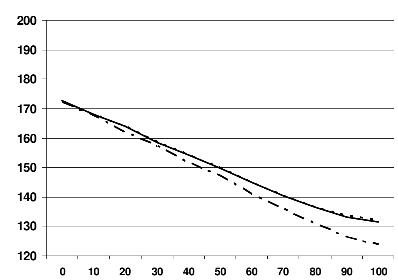
</td></tr>
</table>

Figure 3.17. Revenue Effects of Revealed Common Value
Revenue Difference from No Information, by Common Value Percent
I3 Information
(Solid: $\varepsilon = 8$; Dot: $\varepsilon = 12$; Dash-Dot: $\varepsilon = 18$; Dash-Dot-Dot: $\varepsilon = 27$)

First-Price, Four Bidders

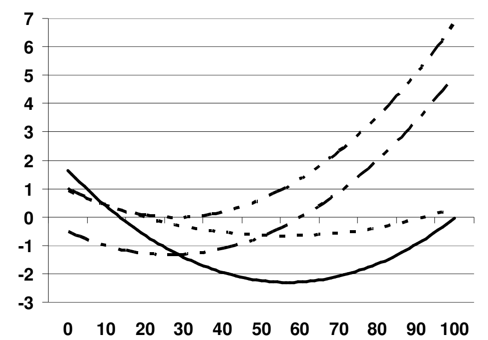


First-Price, Seven Bidders

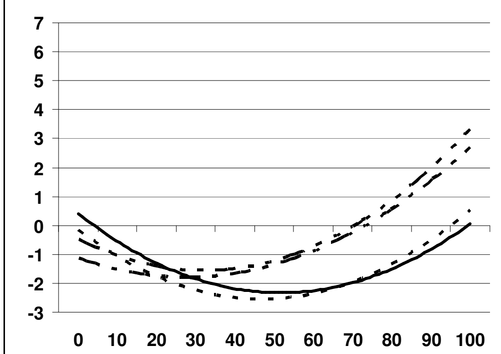


Second-Price, Four Bidders

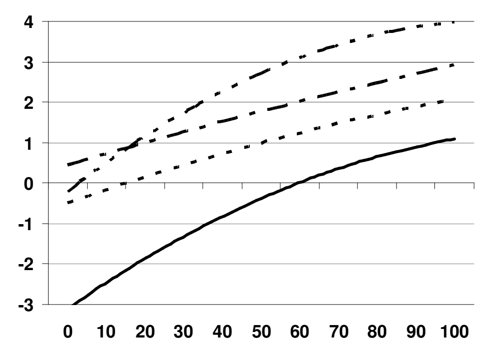


Second-Price, Seven Bidders

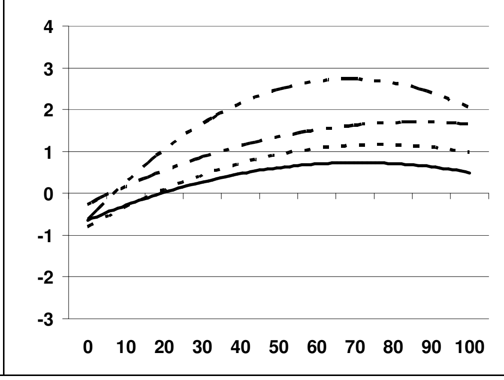

# Chapter 4

# Multiagent System Simulations of Signal Averaging in English Auctions with Two-Dimensional Value Signals

## 4.1 Introduction

In each step *s* of an English ascending dropout auction, the auctioneer raises the price in small increments, and bidders drop out when the price exceeds the price they are willing to pay. Thus, each bidder has information about prices at which other bidders drop out. Empirical analysis of experimental data for common value auctions (Levin *et alia*, 1996) has shown that bidders at all experience levels modify their bids based on the most recent price at which other bidders drop out, called the dropout price. A bidder modifies its bid because the information provided by the dropout price results in the bidder revising its estimation of the value (its "value signal") of the object being auctioned. The bidder combines the information from the dropout price with its original value signal as a weighted average to produce a revised value signal. The bidder continues to do this signal averaging until the auction price exceeds its revised value signal at which point it drops out.

This study makes three contributions. First, is this signal averaging something that computational agents can learn to do effectively? Second, if so, how does this averaging vary across the percent of common value in the multi-dimensional value signal? Third, how does signal averaging affect profit, revenue, and efficiency? I find that agents learn to signal average, but it takes many periods; that signal averaging increases nonlinearly as the common value percent increases; and that the major effect of signal averaging is on the percentage of auctions won by bidders with the highest common value signal.

Section 4.2 describes the English auction format. Section 4.3 provides details of the agent learning for the value multiplier and for the signal-averaging factor. The section also analyzes the convergence and sensitivity properties of the learning method. Section 4.4 analyzes the signal-averaging results, including how they compare with the results from lab experiments, how signal averaging varies with the common value percent, and the effects of signal averaging on the value multiplier. Section 4.5 demonstrates how signal averaging affects profit, revenue, and efficiency. Section 4.6 presents conclusions.

## 4.2 English Auction Model

The multiagent system platform is described in Chapter 2, and Chapter 3 describes how the system implements values and the value signals for bidders, the levels of information feedback, and the number of periods and bidders. This section describes the format of the English auction and how bidders average their original value signals with information conveyed during the auction.

### *4.2.1 Ascending Dropout Process*

Following Milgrom and Weber (1982) and Levin *et alia* (1996), I use the "ascending dropout" or "button auction" format for the English auction. The seller begins the bidding at a low offer price and then raises it in small increments at each step $s$. If this offer price $b_{t,s}$ exceeds a bidder's threshold price $b^i_{t,s}$, the bidder drops out of the auction. Each bidder is provided with the information about whether a dropout has occurred at $b_{t,s}$. Thus, each bidder has information about prices at which other bidders drop out. Empirical analysis of experimental data for common value auctions (Levin *et alia*, 1996) has shown that bidders at all experience levels modify their bids based on the most recent price at which other bidders drop out, denoted $b_{t,s}^{d}$.

### *4.2.2 Signal Averaging*

Assuming a linear bidding strategy for bidder *i* of the form $b_{t,s+1}^{i} = \gamma_t^i \hat{v}_{t,s+1}^{i}$ (where $b_{t,s+1}^{i}$ is the bid price at step s+1, $\gamma_t^i$ is a learned Value Multiplier, and $\hat{v}_{t,s+1}^{i}$ is bidder *i*'s revised value estimate at step *s*+1), the dropout value contains information about other bidders' values. Agent i's revised value estimate $\hat{v}_{t,s+1}^{i}$ in step s+*1* is a weighted average of its previous value estimate and the latest dropout price $\hat{v}_{t,s+1}^{i} = \left(1-\delta_t^i\right)\hat{v}_{t,s}^{i} + \delta_t^i b_{t,s}^{d}$, where $\delta_t^i$ is a learned signal-averaging factor (see Section 4.3.2) and $\hat{v}_{t,0}^{i} = \hat{v}_t^i$ is the agent's original value estimate before it observes any dropout bids. The agent's revised price threshold is then $b_{t,r+1}^{i} = \gamma_t^i \hat{v}_{t,r+1}^{i}$, where $\gamma_t^i$ is a learned Value Multiplier (see Section 4.3.1). The agent can either assume that $b_{t,s}^{d}$ represents the dropout agent's value (which is what is assumed here), or it can assume that the agent has a value multiplier similar to its own. In the latter case, the revised value would be $\hat{v}_{t,s+1}^{i} = \left(1-\delta_t^i\right)\hat{v}_{t,s}^{i} + \delta_t^i \frac{b_{t,s}^{d}}{\gamma_t^i}$, but this modification does not change the results significantly because the value multiplier is close to 1.0 (see Figures 4.7 and 4.8).

## 4.3 Learning Model

Each agent uses information feedback at the I1, I2, or I3 level to learn the most profitable value multiplier $\gamma_t^i$ and signal-averaging factor $\delta_t^i$. The agents use impulse learning with loss aversion (ILA) method for learning $\gamma_t^i$ . The ILA method was developed in Chapter 3 and is based on Selten's impulse balance method (Selten and Buchta, 1998; Selten *et alia*, 2005; Ockenfels and Selten, 2005; Neugebauer and Selten, 2006). In the impulse balance method, the agents use profit and foregone profit that are balanced using a factor that depends on the frequency of winning and losing. The ILA method replaces the balance factor with a loss aversion factor and a loss indicator function. The loss aversion factor is $L_t^i = E_t\left[\left|\pi_t^i\right| \mid \pi_t^i < 0\right]$, i.e., the expected value of the magnitude of the losses. The indicator function $\mathbf{1}_{(L_t^i=0)}$ obtains the value 1 when $L_t^i = 0$. Since a losing agent uses foregone profit to raise its bid price only when $L_t^i = 0$, the adjustment can be written $a_{+,t}^i = \mathbf{1}_{(L_t^i=0)}\hat{\pi}_{F,t}^i$ . In this study the ILA method is extended to learning the value multiplier in English auctions (Section 4.3.1) and the signal-averaging factor (Section 4.3.2).

### *4.3.1 Learning the Value Multiplier*

There are two differences in learning the Value Multiplier in English auctions as opposed to sealed-bid auctions. First, money on the table (the difference between the winning bid and the runner-up bid) is not relevant since in English auctions the winner pays the bid of the runner-up bidder. Second, the calculation of foregone profit is different for English auctions than for the sealed-bid auctions. Let $s = D$ be the step in which agent *i* drops out so then $\hat{v}_{t,D}^i$ is its revised value at the end of its participation in the auction in period *t*. It makes more sense for the agent to use this revised value signal instead of its initial value signal to calculate foregone profit. Thus, instead of $\pi_{F,t}^i = \hat{v}_t^i - p_t$ , we have $\pi_{F,t}^i = \hat{v}_{t,D}^i - p_t$ .

The value multiplier $\gamma_t^i$ is adjusted using downward $a_{-,t}^i$ or upward $a_{+,t}^i$ impulses that are calculated from information feedback. Profitable winners make no adjustment since their strategy has produced a good result, but for unprofitable winners $a_{-,t}^i$ is the loss. As in Chapter 3, the loss adjustment is weighted using a loss aversion factor $L_t^i = E_t\left[\left|\pi_t^i\right| \mid \pi_t^i < 0\right]$ that is the expected value of the magnitude of the losses, and once an agent has a loss ( $L_t^i > 0$ ), it ceases to process upward impulses when it loses. For losers, $a_{+,t}^i$ is the foregone profit.

The value multiplier is updated using $\gamma_{t+1}^i = \gamma_t^i + \phi(a_{+,t}^i - a_{-,t}^i)$, where $\phi$ is a learning rate. The adjustment rules for the value multiplier are:

***Rule*** **VM1**: For all information levels, $r_t^i = 1$, and $\pi_t^i < 0.0$: $a_{-,t}^i = \left(1 + L_t^i\right)\dfrac{\left|\pi_t^i\right|}{\hat{v}_t^i}$.

***Rule*** **VM2**: For I2 and I3, $r_t^i < 1$, when $\pi_{F,t}^i \geq 0$, $a_{+,t}^i = \mathbf{1}_{(L_t^i = 0)} \dfrac{\pi_{F,t}^i}{\hat{v}_t^i}$

***Rule*** **VM3**: For I1, $r_t^i < 1$, when $\hat{\pi}_{F,t}^i \geq 0$, $a_{+,t}^i = \mathbf{1}_{(L_t^i = 0)} \dfrac{\hat{\pi}_{F,t}^i}{\hat{v}_t^i}$

### *4.3.2 Learning the Signal-Averaging Factor*

Each agent adapts its signal-averaging factor $\delta_t^i$ using an impulse $c_t^i$ that raises it or lowers it. A winning agent learns the true value $v^i$, so it can use $\hat{v}_{t,D}^i$, $\hat{v}_t^i$, and $v^i$ to produce the impulse. The signal-averaging factor is adjusted using $\delta_{t+1}^i = \delta_t^i + \varphi c_t^i$, where $\varphi$ is the learning rate for signal averaging. The revision varies depending on whether it has been averaging its value estimate up ( $\hat{v}_{t,D}^i > \hat{v}_t^i$ ) or down ( $\hat{v}_{t,D}^i < \hat{v}_t^i$ ). When the agent wins, it learns the true value and can assess the benefit of signal averaging by comparing $\hat{v}^i_{t,D}$ to $v^i$. When an agent loses, it does not have information about the actual value and thus has no means of assessing and adjusting its signal averaging.

***Rule* SA.1**: If $\hat{v}^i_{t,D} > \hat{v}^i_t$, $c^i_t = \frac{v^i - \hat{v}^i_{t,D}}{v^i}$

*Demonstration*: If a winning agent has been averaging up, the revised value is greater than the original value signal ($\hat{v}^i_{t,D} > \hat{v}^i_t$). The agent is averaging too little if the revised value is still less than the true value $(\hat{v}^i_{t,D} < v^i)$. In order to achieve a better estimate of $v^i$, the agent will raise $\delta^i_t$ using $c^i_t = \frac{v^i - \hat{v}^i_{t,D}}{v^i}$. On the other hand, the agent is averaging up too much if the revised value is greater than the true value ($\hat{v}^i_{t,D} > v^i$) so in this case it lowers $\delta^i_t$ using $c^i_t = \frac{v^i - \hat{v}^i_{t,D}}{v^i}$.

***Rule* SA.2**: If $\hat{v}^i_{t,D} < \hat{v}^i_t$, $c^i_t = \frac{\hat{v}^i_{t,D} - v^i}{v^i}$

*Demonstration*: If the agent has been averaging down, the agent is averaging too little if the revised value is greater than the true value ($v^i < \hat{v}^i_{t,D}$), so it raises $\delta^i_t$ using $c^i_t = \frac{\hat{v}^i_{t,D} - v^i}{v^i}$. On the other hand, the agent is averaging down too much if $\hat{v}^i_{t,D} < v^i$, so it lowers $\delta^i_t$ using $c^i_t = \frac{\hat{v}^i_{t,D} - v^i}{v^i}$.

***Rule* SA.3**: If $\hat{v}^i_{t,D} = \hat{v}^i_t$, $c^i_t = \frac{\left|\hat{v}^i_t - v^i\right|}{v^i}$

*Demonstration*: If there has been no signal averaging and therefore no effect on the agent's value estimate, the agent has not been averaging enough. In this case, it will raise $\delta_t^i$ using the magnitude of the difference between the true value and the value signal. This rule is not applied by an agent that has learned to reduce $\delta_t^i$ to zero.

#### *4.3.3 Sensitivity and Convergence*

The learning method for signal averaging described in Sections 4.3.1 and 4.3.2 is insensitive to learning rates. However, the method is sensitive to the initial value, but this is understandable and explained below. Results are shown for n = 4 and $\varepsilon = 8$, and are similar for n = 7 and for other values of $\varepsilon$.

***Result 1***: Signal averaging learning is insensitive to learning rates in the range $1 < \varphi \leq 3$ and $\delta_t^i$ converges to about 0.4 in about 400 periods.

*Discussion*: Figure 4.1 shows that learning rates less that 1.0 require over 600 periods for convergence. For all of the learning rates greater than 1.0, the average signal-averaging factor converges to about 0.4 is achieved within the 600 periods. As the rate increases from 1.0 to 3.0, the number of periods required for convergence decreases gradually from about 600 to about 400. Higher learning rates converge to slightly higher values, without much improvement in the rate of convergence. For example, when $\varphi = 7$ it takes about 300 periods to converge to a value for $\delta_t^i$ of about 0.49. Such a large number of learning periods implies that $\delta_t^i$ is a difficult factor to learn. This is discussed more in Section 4.6. In the simulations, I use $\varphi = 2$.

***Result 2***: For agents that have a significant level of learning, the learning is insensitive to the starting value.

*Discussion*: Figure 4.2 illustrates convergence for different initial values of $\delta_t^i$.

When the initial value is 1.0, the average signal-averaging factor is about 0.75. As the initial value decreases, the final value decreases until it is about 0.4 for the initial value of 0.0. The explanation is that agents receive impulses to adjust their signal-averaging factor only when they win. Thus, agents that win at very low frequency do not have much opportunity to adjust their factors and remain stuck at the initial value. The agents that win at very low frequency will, of course, tend to be the lower valued bidders. This is illustrated in Figure 4.3. When the initial value is 0.0, the signal-averaging factors of most of the low-valued bidders remain at or near 0.0. When the initial value is 1.0, the signal-averaging factors of most of the low-valued bidders remain at or near 1.0. These low-valued signal-signal-averaging factors are shown by the box. However, Figure 4.3 also shows that learning of signal averaging has similar results (indicated by the triangle) for the higher-valued bidders. This is illustrated further in Figure 4.4 that shows the convergence for the initial values for the higher valued bidders only. For the agents that have a significant level of learning, the learning is insensitive to the starting value, since they all end up converging to a value near 0.9. This works out fine for the agents, since signal-averaging is more critical to the higher-valued agents because it helps them to avoid the winners' curse.

## 4.4 Results of Signal Averaging

In this section, I compare the simulation results to data from lab experiments, analyze the variation of signal averaging with the common value percent, and comment on the effect of signal averaging on the value multiplier $\gamma_t^j$. The results that are illustrated in the figures are consistent across repeated simulations.

### *4.4.1 Comparing Signal Averaging Results with Lab Experiments*

Levin *et alia* (1996) present results for signal-averaging factors for four and seven bidders with I3 information averaged over $\varepsilon$= 18 and 30. Regressions of bids with value signals and the recent dropout price show coefficients for the dropout price averaging about 0.78 for four bidders and 0.69 for seven bidders. Running similar regressions on data produced by the agents, results in coefficients for the dropout price of 0.55 (0.009) for four bidders and 0.31 (0.003) for seven bidders. Thus, the agents on average are placing a lower weight on dropout bids than the lab experimental subjects, but in both cases signal averaging decreases with the number of bidders.

Levin *et alia* (1996) show for four experienced bidders that profit increased with $\varepsilon$ (18 and 30) from an average of about 1 to 7. Figure 4.9 shows that for four agents profit increases with $\varepsilon$ from about 4 to about 7, which is similar to the results from the lab experiments.

### *4.4.2 Variation of Signal Averaging with Common Value Percent*

For compatibility with Chapter 3 and with the lab experiments of Levin *et alia* (1996), the computational experiments use a value support of [25, 225] and common value uncertainties of $\varepsilon$= 8, 12, 18, and 27, and bidder numbers of four and seven. Figure 4.5 shows the signal-averaging factor by common value percent.

***Result 3***: Signal averaging increases as the common value percent increases.

*Discussion*: The most obvious result in Figure 4.5 is that signal averaging increases as the common value percent increases. This is reasonable since signal-averaging is not useful to an agent when its value is private, but it is useful when the agent can benefit from other agents' information about the common value, as signaled by the dropout bids. It is interesting to see that in most instances (except for four bidders with $\varepsilon = 27$) the pattern of increase is very nonlinear. There is very little signal averaging up until the common value percent is between 60 to 80, and then it increases rapidly for the higher percents. In other words, signal averaging is not used in linear proportion to the common value percent. When the common value component is less than half, the agents place most of their weight on their own value estimates. This is due to the fact that, up to a common value percent of 60%, the upward and downward impulses are small in magnitude (e.g., about 0.2 or less) and about equal. For common value 70% and over, the magnitude of the upward impulses grows larger than that of the downward impulses.

***Result 4***: Signal averaging decreases as the number of bidders increases.

*Discussion*: Figure 4.5 also shows that signal averaging is less when there are seven bidders than when there are four. The detailed data for the agent learning shows that for seven bidders, there are about half as many upward impulses and they average about half the size as they do for four bidders. At the same time, there are almost twice as many downward impulses and they are about three times the magnitude as for four bidders.

***Result 5***: As the common value signal becomes more uncertain (higher $\varepsilon$), signal averaging decreases for high common value percents and increases slightly at the middle percentages.

*Discussion*: Figure 4.5 shows that the curves tend to flatten out with increasing uncertainty. Increasing uncertainty in all agents' common value signals results in less signal improvement with averaging, so the averaging is not reinforced as much when there is less uncertainty.

### *4.4.3 Effects of Signal Averaging on the Value Multiplier*

Figure 4.6 shows one case of the value multiplier convergence. The value multiplier is initialized at 0.93 and rises to about .975 by period 50. Signal averaging learning starts at period 51 and the value multiplier decreases to 0.97 before rising to a final value between 0.975 and 0.980.

Figures 4.7 and 4.8 show the value multiplier $\gamma_t^j$ without signal averaging (on the left) and with signal averaging (on the right). For four bidders, the value multipliers tend to increase slightly with common value percent (e.g. from 0.94 to 0.96) and stay flat or decrease slightly for seven bidders (e.g. 0.96 to 0.94). For sealed-bid auctions (Chapter 3), the value multiplier increased with the common value percent for first-price auctions (e.g. from 0.8 to 0.95) and decreased slightly for second-price auctions.(e.g. from 1.0 to 0.97). Thus, as expected, the English value multipliers are more similar to the second-price value multipliers than to the first-price value multipliers.

There appear to be no significant effects from signal averaging on the value multiplier when the common value uncertainty is low, but the multiplier decreases when $\varepsilon > 12$. This is demonstrated more clearly in Figure 4.9, which shows that when $\varepsilon > 12$ bidders not only reduce their higher values using signal averaging but also increase their bid shading.

## 4.5 Effects of Signal Averaging on Profit, Revenue, and Efficiency

This section analyzes the effects of signal averaging on bidder profit, seller revenue, and efficiency.

### *4.5.1 Profit*

Figures 4.10 and 4.11 show that the variation of profit with common value percent is very close to linear and is decreasing. In all cases, I3 agents are profitable, but I1 and I2 agents are often unprofitable when the value is 100% common value.

***Result 6***: When bidders use signal averaging, profit is slightly higher at high percents of common value. This difference is higher for bidders with I2 and I3 information.

*Discussion*: Figure 4.12 shows increases in profit with signal averaging for four bidders for all information levels, and these differences increase as $\varepsilon$ increases. For seven bidders, there is no profit improvement for I1 bidders with signal averaging, some improvement for I3, and even more improvement for I2.

Signal averaging causes an increase in average profit because the agents with the highest common value signals win less often than without signal averaging. This reduction in the frequency of winning with the highest signal reduces the probability of losses and thus abates the winner's curse. Result 3 showed that signal averaging is higher at high common value percents. This in turn results in agents with high common value signals lowering their value estimates so that they win less than without signal averaging. This winner's curse abatement effect is illustrated in Figures 4.13 and 4.14, which show the fraction of auctions won by the bidder with the highest common value signal (before signal averaging). There are some interesting observations to make about these figures. When value signals are less than pure common value, the private value estimate has an increasingly dampening effect on the winner's curse[1]. Then, when the value reaches 100% common value, the fraction of auctions won by the bidder with the highest common value signal jumps back up, in many cases to 1.0. With signal averaging, the fraction of auctions won by the bidder with the highest original common value signal is not only lower across the full range of two-dimensional signals, the jump at pure common value is dampened considerably, especially for bidders with I2 and I3 information.

### *4.5.2 Revenue*

Figure 4.15 shows revenue for four bidders, and Figure 4.16 for seven bidders. Revenue decreases almost linearly with increased percent of common value. The most obvious revenue variation is with information level, with more information (I3) always resulting in slightly less revenue than with less information (I1 and I2), and the difference increases with the common value percent.

***Result 7***: When bidders use signal averaging, revenue is slightly lower at high percents of common value. This difference is greater for bidders with I2 and I3 information.

*Discussion*: Figure 4.17 shows that there are some small revenue decreases from signal averaging, reflecting the profit increases described in the previous section. There are decreases in revenue with signal averaging for four bidders for all information levels, and these differences increase as $\varepsilon$ increases. For seven bidders, there is no revenue decrease for I1 bidders with signal averaging, some decrease for I3, and even less

---

[1] A bidder with a common value signal that is higher than the actual common value is more likely to win the auction than a lower-valued bidder. Because the winner loses money in this case, the winner curses.

revenue for I2.

### *4.5.3 Efficiency*

In this section, I use the definition of "private value" efficiency described in Goeree and Offerman (2002) and used in Chapter 3. Figure 4.18 shows efficiency for four bidders and Figure 4.19 for seven bidders. There is a marked nonlinear relationship of efficiency with common value percent. The negative effect of common value percent on efficiency increases at an increasing rate as the percent increases. This is similar to that reported for first-price and second-price auctions in Chapter 3.

***Result 8***: When four I3 bidders use signal averaging, private value efficiency decreases slightly at high common value percents.

*Discussion*: Figure 4.20 shows that signal averaging has some small effects on private value efficiency when the common value percent is over 60%. The most significant effect is for I3 bidders when uncertainty is high, i.e., $\varepsilon \geq 12$.

## 4.6 Conclusion

I have found that agents learn to signal average and learn to increase signal averaging as the common value percentage increases. Signal averaging decreases as the number of bidders increases. The uncertainty of the common value signal has an effect on signal-averaging and the results of signal averaging. As the common value signal becomes more uncertain, signal averaging changes. It decreases for high common value percents and increases slightly at the middle percentages. When bidders use signal averaging, profit is slightly higher when the common value component is high. This difference increases as the uncertainty about the common value increases. For I3 bidders using signal averaging, profit increases with an increase in $\varepsilon$, but this does not occur for

I1 and I2 information levels. The most obvious impact of signal averaging is on the frequency at which the bidder with the highest common value signal wins the auction.

The problem with learning to average value signals is that it takes so long because bidders have information on the actual value only when they win. We know that real-world bidders will average signals (Levin *et alia*, 1996), but they will not have the opportunity to learn to signal average as well as they could if they had 300 repeated auctions.

Using multiagent simulations has provided some insights into single-unit auctions, both sealed-bid auctions (Chapter 3) and English auctions. Chapter 5 will expand the method from single-unit auctions to the multi-unit Treasury auctions.

## 4.7 References

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

Kagel, J.H., Harstad, R.M., and Levin, D., 1987. "Information Impact and Allocation Rules in Auctions with Affiliated Private Values: A Laboratory Study." *Econometrica*. 55(6): 1275-1304.

Levin, D., Kagel, J.H., and Richard, J.F., 1996. "Revenue Effects and Information Processing in English Common Value Auctions." *American Economic Review*. 86: 442-460.

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

Neugebauer, T., Selten, R., 2006. "Individual Behavior of first-price auctions: The importance of information feedback in computerized experimental markets." *Games and Economic Behavior*. 54: 183-204.

Ockenfels, A. and Selten, R., 2005. "Impulse Balance Equilibrium and Feedback in First Price Auctions." *Games and Economic Behavior*. 51: 155-170.

Selten, R., Abbink, K., and Cox, R., 2005. "Learning Direction Theory and the Winner's Curse." *Experimental Economics*. 8: 5-20.

Selten, R. and Buchta, J., 1998. "Experimental Sealed Bid First Price Auctions with Directly Observed Bid Functions." In: Budescu, D. V., Erve, I. and Zwick, R. (Eds), *Games and Human Behavior: Essays*. Lawrence Erlbaum Associates, pp. 53-78.

## 4.8 Tables

Table 4.1. English Auction Model Notation Summary

| Symbol | Description |
|---|---|
| $b_t^i$ | Bid level of bidder i in auction t. |
| $b_{t,s}^i$ | Bid of agent i at step s of auction period t. |
| $b_{t,r}^d$ | Latest dropout price |
| $\delta_t^i$ | Signal-Averaging Factor |
| $\varepsilon$ | Radius of the support for the common value signal. |
| $\gamma_t^i$ | Value multiplier or bid factor. |
| $p_t$ | Payment by winner in auction t. |
| $\phi$ | Learning rate for the value multiplier $\gamma_t^i$ |
| $\varphi$ | Learning rate for the signal-averaging factor $\delta_t^i$ |
| $p_t^i$ | Payment made by bidder i, given that it wins. |
| $\pi_t^i$ | Profit of bidder i in auction t. |
| $\pi_{F,t}^i$ | Foregone profit of bidder i in auction t. |
| s, D | s is step of bidding in an English auction; D is the step in which a bidder drops out |
| $r_t^i$ | Ranking of bidder i in auction t. |
| $\overline{S}_{\mathrm{P}}, \underline{S}_{\mathrm{P}}, \overline{S}_{\mathrm{C}}, \underline{S}_{\mathrm{C}}$ | Upper and lower bounds of the supports for the private value signal and common value. |
| $\theta_C$ | Common value component of the value signal. |
| $\hat{v}_t^i$ | Value signal of bidder i in auction t. |
| $\hat{v}_{t,s}^i$ | Revised value signal of bidder i at step r of auction t. |
| $v_t^i$ | Actual value, revealed only to winner. |

| Table 4.2. Information Levels (incremental) | | |
|---|---|---|
| Level | Description | Feedback |
| IL1 | Number of bidders | n |
| | Value Signal: Own | $\hat{v}_t^i$ |
| | Bid: Own | $b_t^i$ |
| | Ranking: Own | $r_t^i$ |
| | Payment: Own | $p_t^i \mid r_t^i = 1$ |
| | Value: Own | $v_t^i \mid r_t^i = 1$ |
| IL2 | Bid: Winner | $b_t^{(1)}$ |
| | Payment | $p_t$ |
| IL3 | Bid: Runner-up | $b_t^{(2)}$ |

## 4.9 Figures

Figure 4.1. Sensitivity of Signal-Averaging Factor to Learning Rate
I3 Information, 100% Common Value, n = 4, $\varepsilon = 8$
Signal-Averaging Factor by Period

Learning Rate 0.5

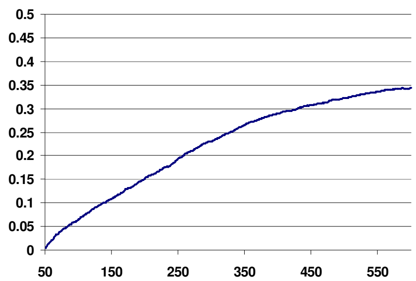


Learning Rate 1.0

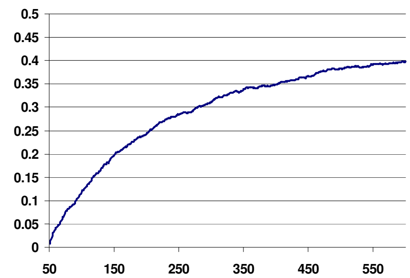


Learning Rate 1.5

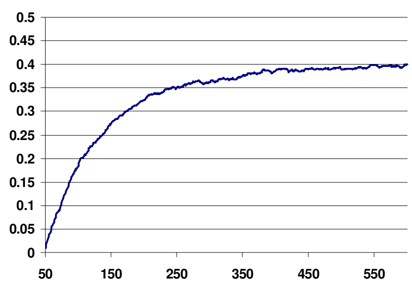


Learning Rate 2.0

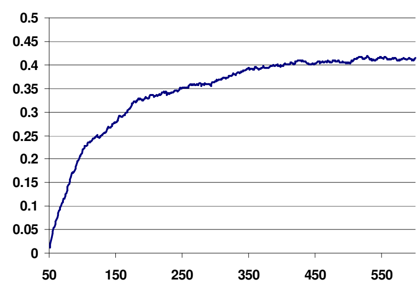


Learning Rate 2.5

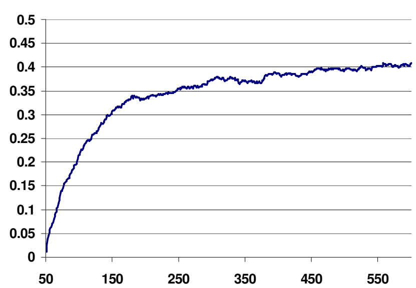


Learning Rate 3.0

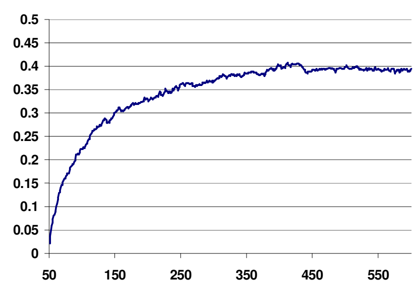

Figure 4.2. Convergence of Signal-Averaging Factor
I3 Information, 100% Common Value, n = 4, $\varepsilon = 8$
Signal-Averaging Factor by Period

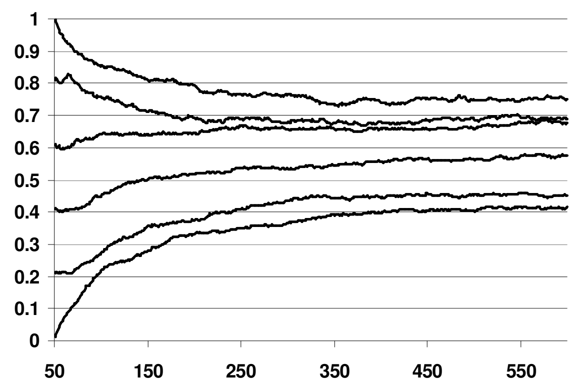


Figure 4.3. Distribution of Signal-averaging Factor with Signal Value
I3 Information, 100% Common Value, n = 4, $\varepsilon = 8$
(I1: Solid, I2: Dot, I3: DashDot)

| Initial Value 0.0 | Initial Value 1.0 |
|---|---|

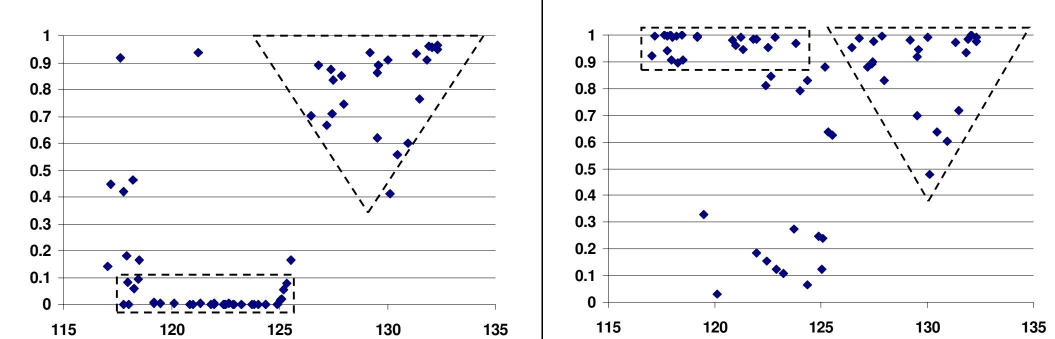

Figure 4.4. Convergence of Signal-Averaging Factor for Different Initial Values
Restricted to high value bidders
Signal-Averaging Factor by Period
I3 Information, 100% Common Value, n = 4, $\varepsilon = 8$

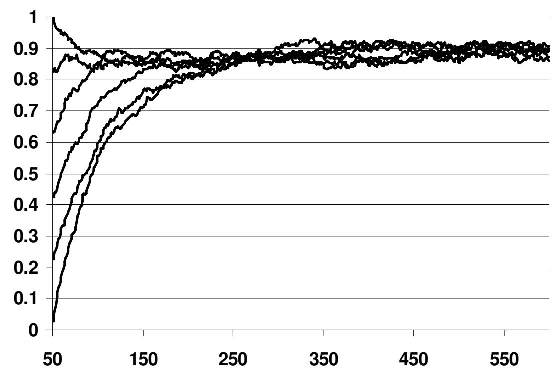

Figure 4.5. Signal-Averaging Factor
Signal-Averaging Factor by Common Value Percent
(I1: Solid, I2: Dot, I3: DashDot)

<table>
<tr><th>Four Bidders</th><th>Seven Bidders</th></tr>
<tr><td>$\varepsilon = 8$</td><td>$\varepsilon = 8$</td></tr>
<tr><td>
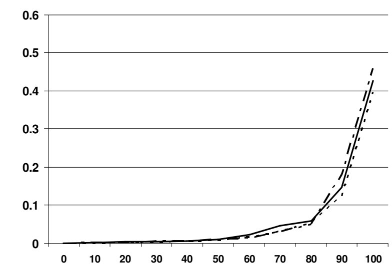

</td><td>
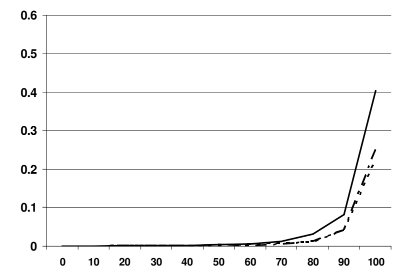

</td></tr>
<tr><td>$\varepsilon = 12$</td><td>$\varepsilon = 12$</td></tr>
<tr><td>
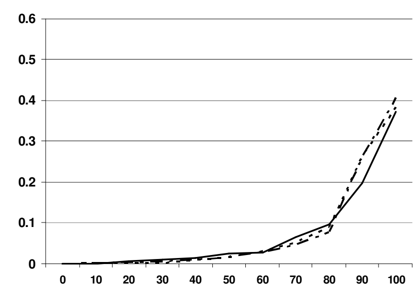

</td><td>
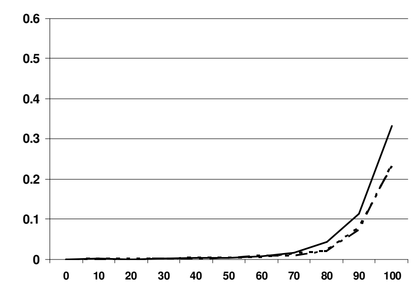

</td></tr>
<tr><td>$\varepsilon = 18$</td><td>$\varepsilon = 18$</td></tr>
<tr><td>
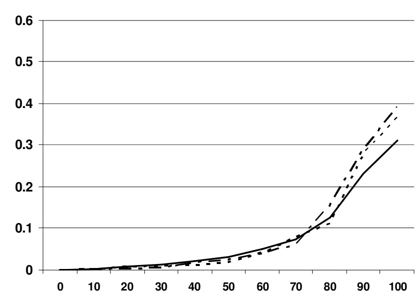

</td><td>
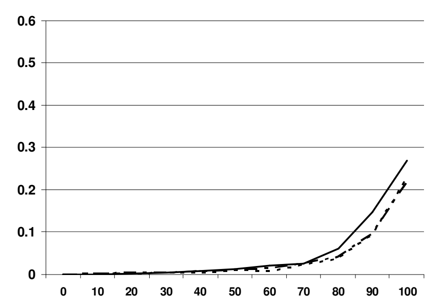

</td></tr>
<tr><td>$\varepsilon = 27$</td><td>$\varepsilon = 27$</td></tr>
<tr><td>
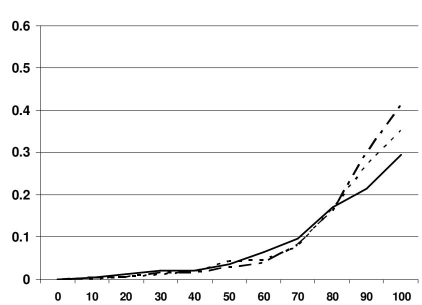

</td><td>
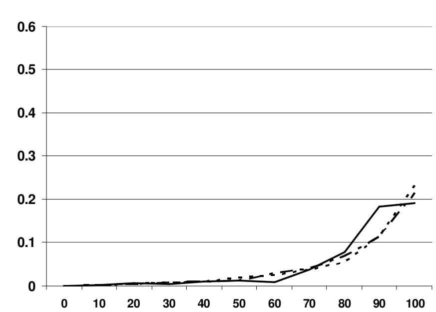

</td></tr>
</table>

Figure 4.6. Convergence of Value Multiplier
Value Multiplier by Common Value Percent
Common Value, First Price, I3 Information, $\varepsilon = 8$

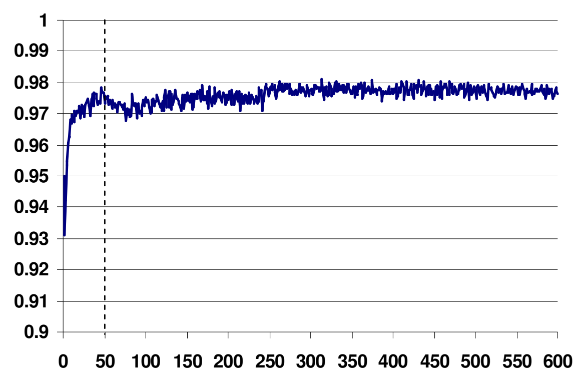

Figure 4.7. Value Multiplier: Four Bidders
Value Multiplier by Common Value Percent
(I1: Solid, I2: Dot, I3: DashDot)

| Without Signal Averaging | With Signal Averaging |
|---|---|
| $\varepsilon = 8$ | $\varepsilon = 8$ |
| 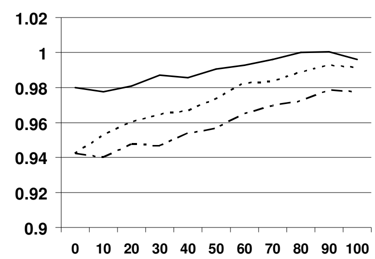 | 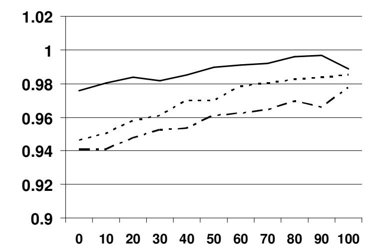 |
| $\varepsilon = 12$ | $\varepsilon = 12$ |
| 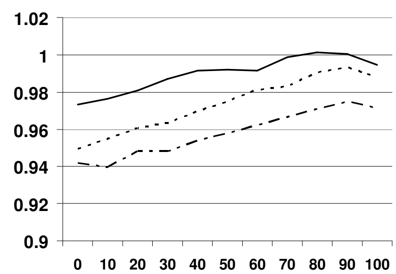 | 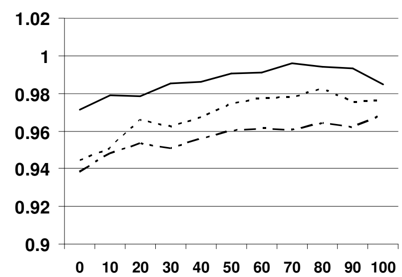 |
| $\varepsilon = 18$ | $\varepsilon = 18$ |
| 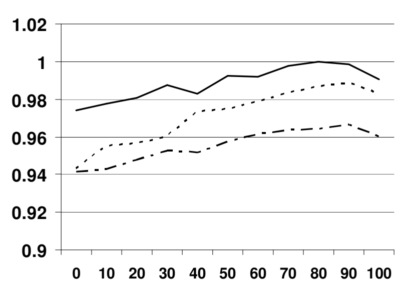 | 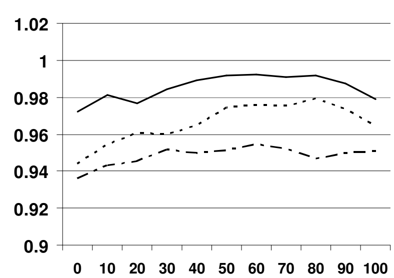 |
| $\varepsilon = 27$ | $\varepsilon = 27$ |
| 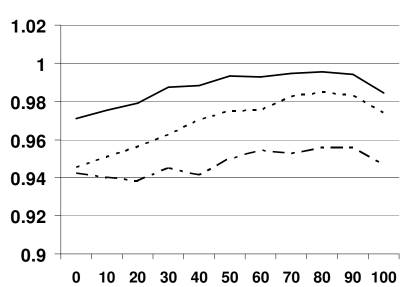 | 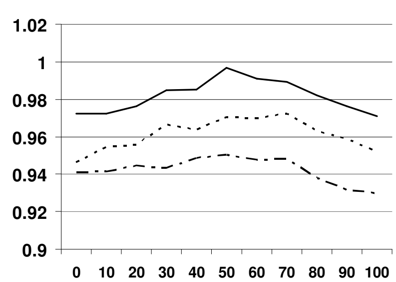 |

Figure 4.8. Value Multiplier: Seven Bidders
Value Multiplier by Common Value Percent
(I1: Solid, I2: Dot, I3: DashDot)

| Without Signal Averaging | With Signal Averaging |
|---|---|
| $\varepsilon = 8$ | $\varepsilon = 8$ |
| 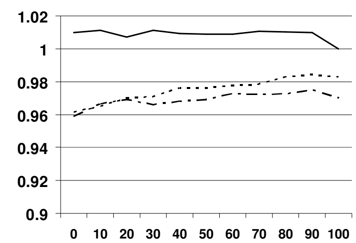 | 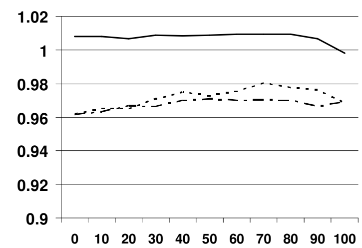 |
| $\varepsilon = 12$ | $\varepsilon = 12$ |
| 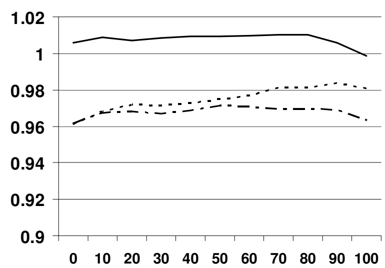 | 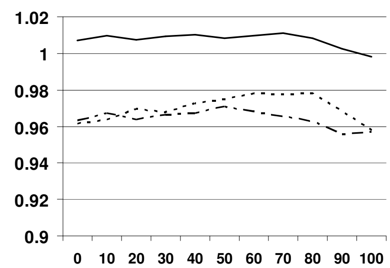 |
| $\varepsilon = 18$ | $\varepsilon = 18$ |
| 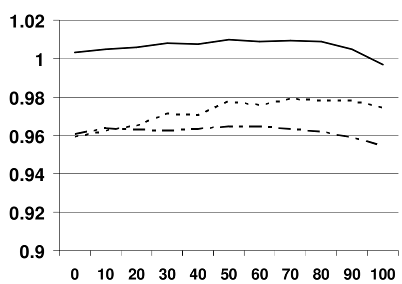 | 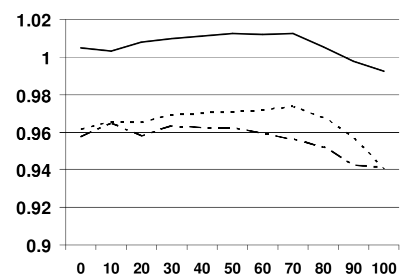 |
| $\varepsilon = 27$ | $\varepsilon = 27$ |
| 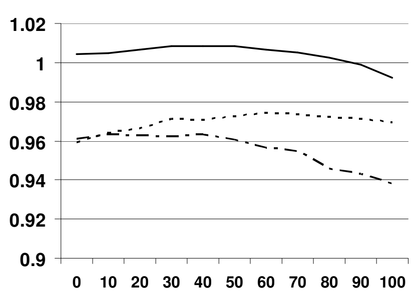 | 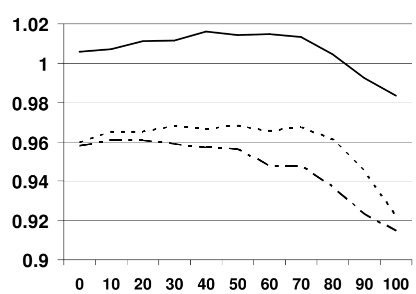 |

Figure 4.9. Value Multiplier Changes with Signal Averaging
by Common Value Percent
(I1: Solid, I2: Dot, I3: DashDot)

**Four Bidders**

$\varepsilon = 8$

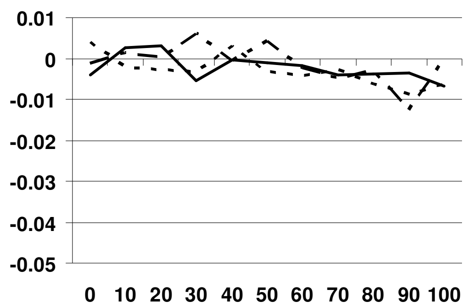


$\varepsilon = 12$

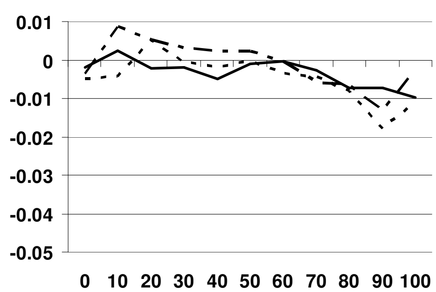


$\varepsilon = 18$

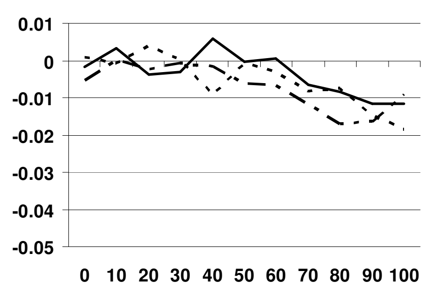


$\varepsilon = 27$

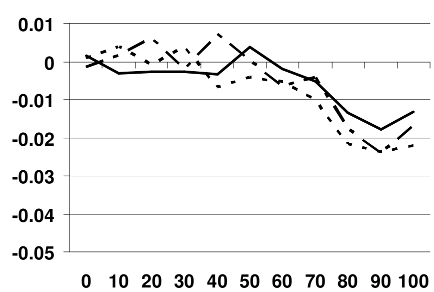


**Seven Bidders**

$\varepsilon = 8$

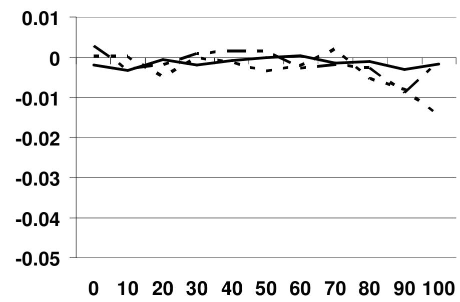


$\varepsilon = 12$

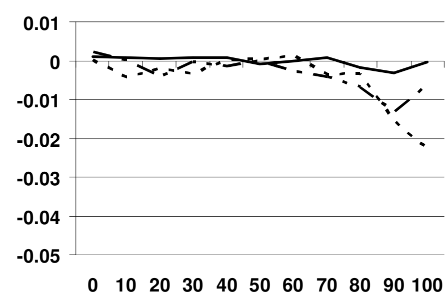


$\varepsilon = 18$

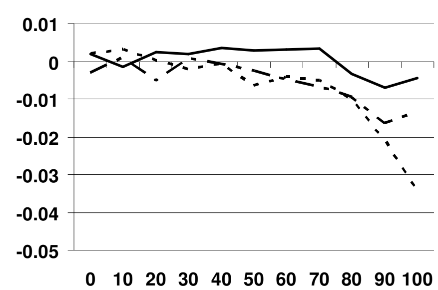


$\varepsilon = 27$

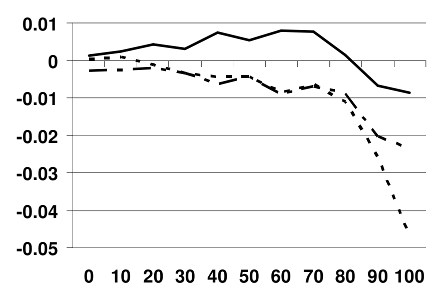

Figure 4.10. Profit: Four Bidders
Profit by Common Value Percent
(I1: Solid, I2: Dot, I3: DashDot)

| Without Signal Averaging | With Signal Averaging |
|---|---|
| $\varepsilon = 8$ | $\varepsilon = 8$ |
| 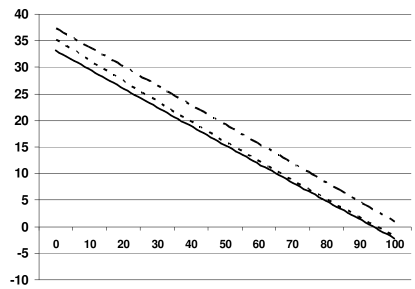 | 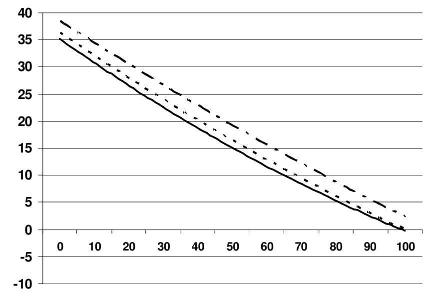 |
| $\varepsilon = 12$ | $\varepsilon = 12$ |
| 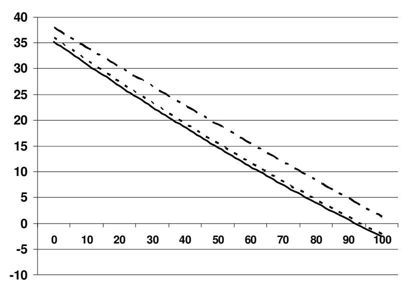 | 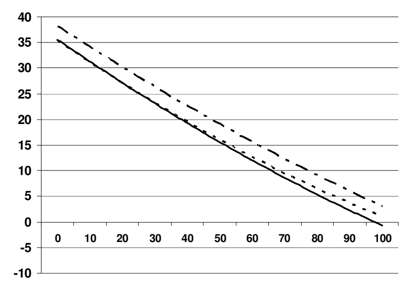 |
| $\varepsilon = 18$ | $\varepsilon = 18$ |
| 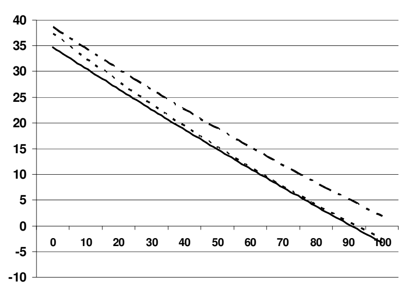 | 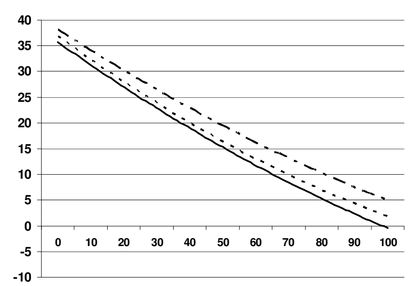 |
| $\varepsilon = 27$ | $\varepsilon = 27$ |
| 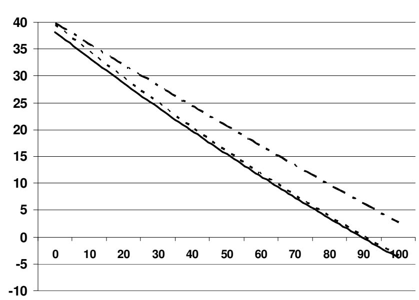 | 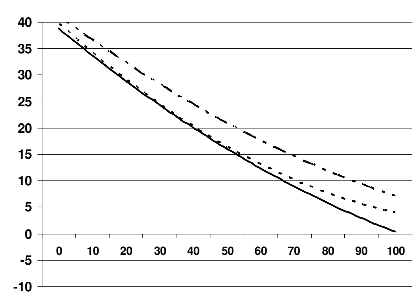 |

Figure 4.11. Profit: Seven Bidders
Profit by Common Value Percent
(I1: Solid, I2: Dot, I3: DashDot)

| Without Signal Averaging | With Signal Averaging |
|---|---|
| $\varepsilon = 8$ | $\varepsilon = 8$ |
| 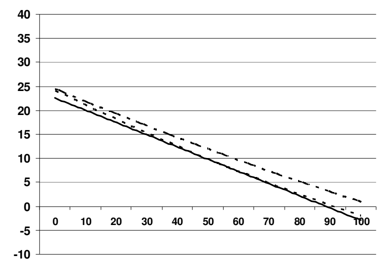 | 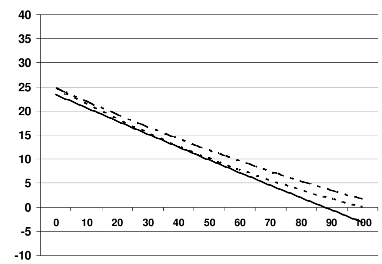 |
| $\varepsilon = 12$ | $\varepsilon = 12$ |
| 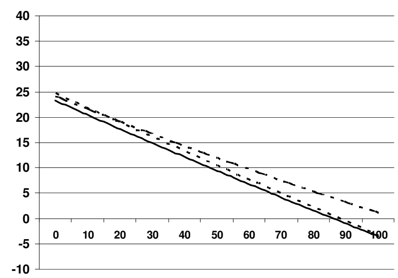 | 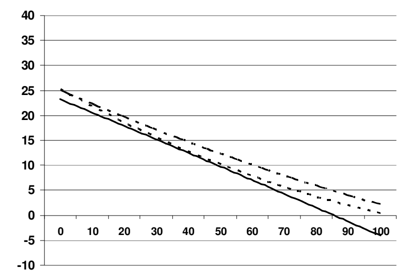 |
| $\varepsilon = 18$ | $\varepsilon = 18$ |
| 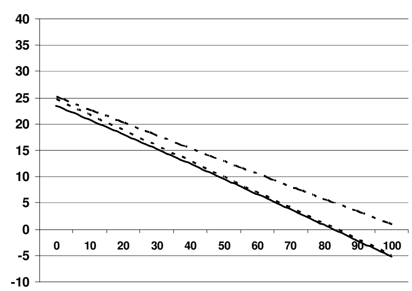 | 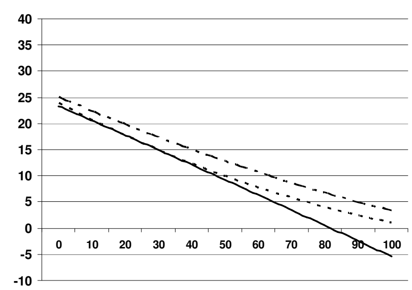 |
| $\varepsilon = 27$ | $\varepsilon = 27$ |
| 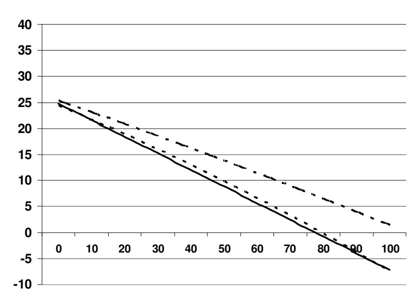 | 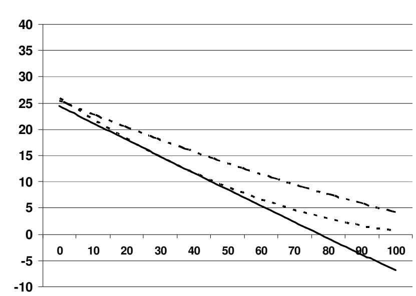 |

Figure 4.12. Profit Changes with Signal Averaging
by Common Value Percent
(I1: Solid, I2: Dot, I3: DashDot)

| Four Bidders | Seven Bidders |
| --- | --- |
| $\varepsilon=8$ | $\varepsilon=8$ |
| 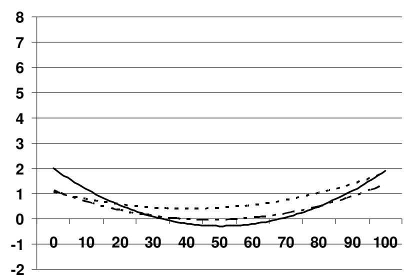 | 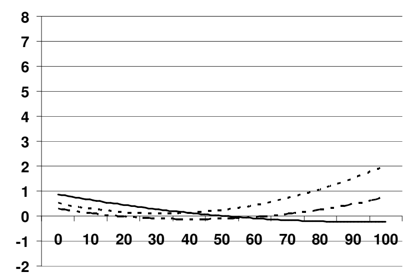 |
| $\varepsilon=12$ | $\varepsilon=12$ |
| 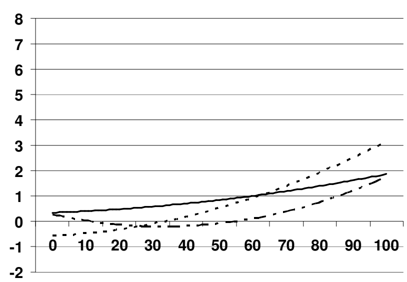 | 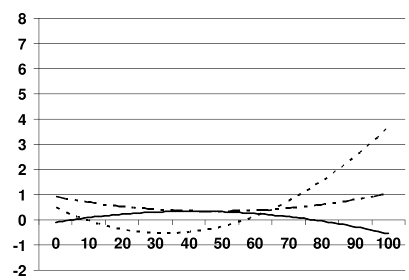 |
| $\varepsilon=18$ | $\varepsilon=18$ |
| 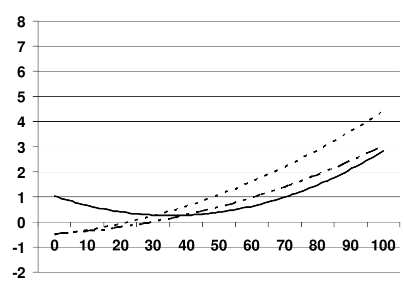 | 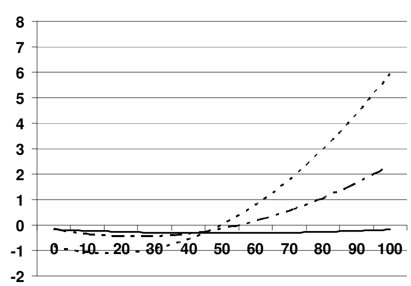 |
| $\varepsilon=27$ | $\varepsilon=27$ |
| 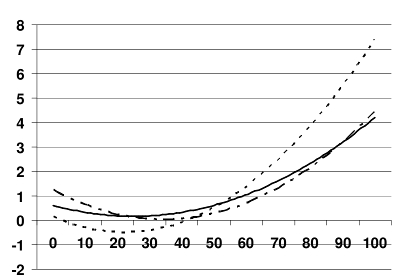 | 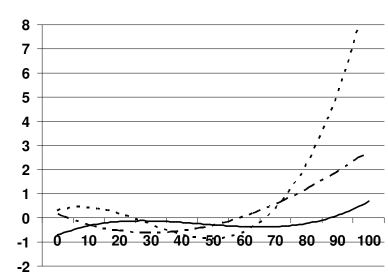 |

Figure 4.13. Winner's Curse Abatement: Four Bidders
Fraction of Auctions Won by Highest Common Value Signal by Common Value Percent
(I1: Solid, I2: Dot, I3: DashDot)

| Without Signal Averaging | With Signal Averaging |
|---|---|
| $\varepsilon = 8$ | $\varepsilon = 8$ |

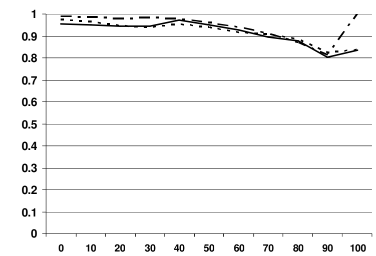


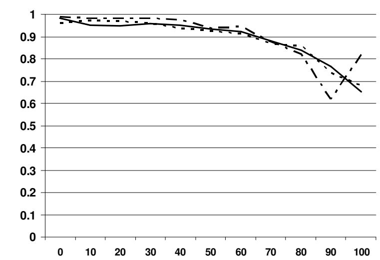


| | |
|---|---|
| $\varepsilon = 12$ | $\varepsilon = 12$ |

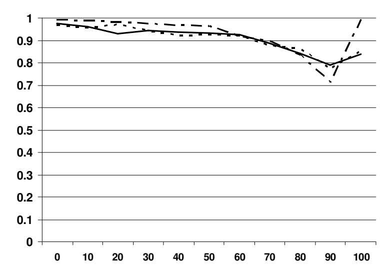


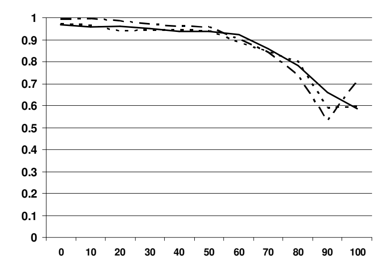


| | |
|---|---|
| $\varepsilon = 18$ | $\varepsilon = 18$ |

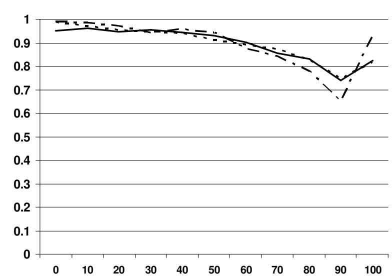


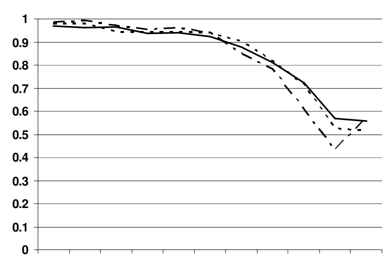


| | |
|---|---|
| $\varepsilon = 27$ | $\varepsilon = 27$ |

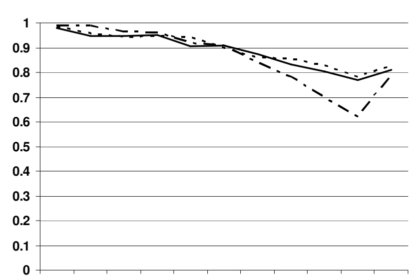


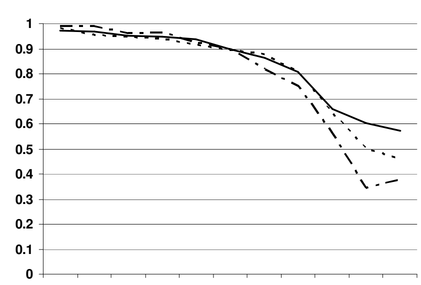

Figure 4.14. Winner's Curse Abatement: Seven Bidders
Fraction of Auctions Won by Highest Common Value Signal by Common Value Percent
(I1: Solid, I2: Dot, I3: DashDot)

| Without Signal Averaging | With Signal Averaging |
|---|---|
| $\varepsilon = 8$ | $\varepsilon = 8$ |
| 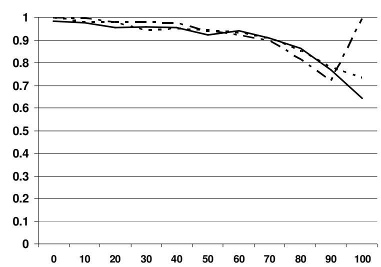 | 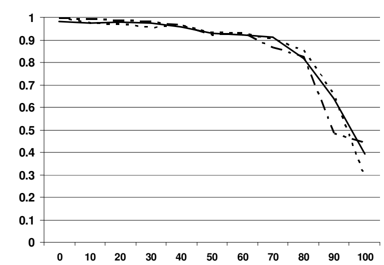 |
| $\varepsilon = 12$ | $\varepsilon = 12$ |
| 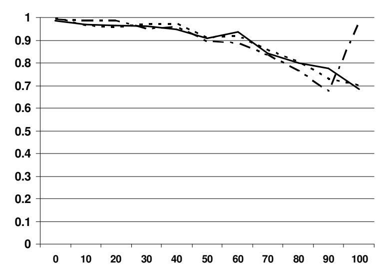 | 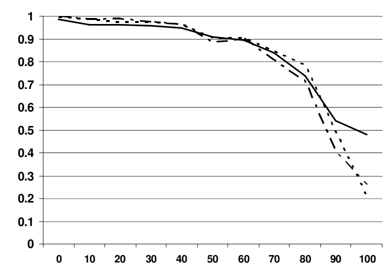 |
| $\varepsilon = 18$ | $\varepsilon = 18$ |
| 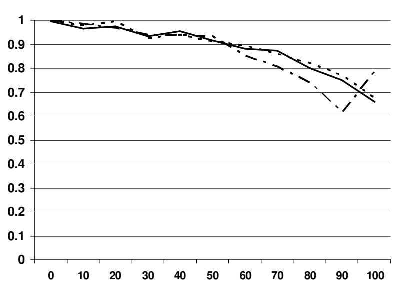 | 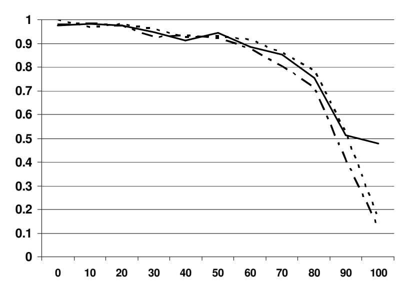 |
| $\varepsilon = 27$ | $\varepsilon = 27$ |
| 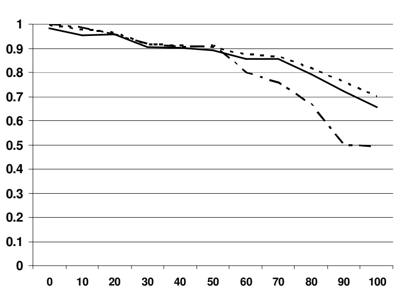 | 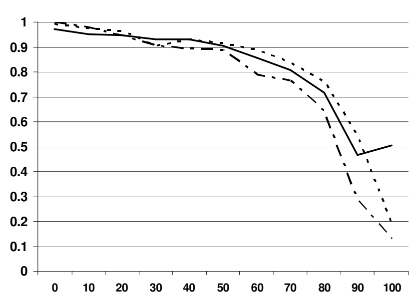 |

Figure 4.15. Revenue: Four Bidders
Revenue by Common Value Percent
(I1: Solid, I2: Dot, I3: DashDot)

**Without Signal Averaging**

$\varepsilon = 8$

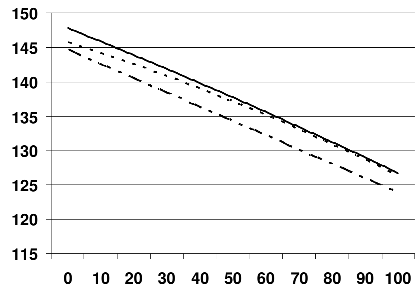


$\varepsilon = 12$

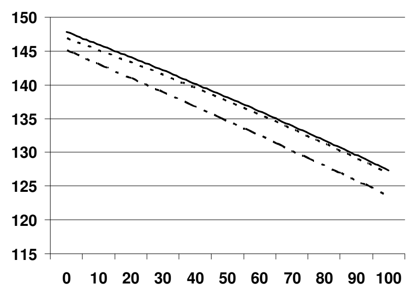


$\varepsilon = 18$

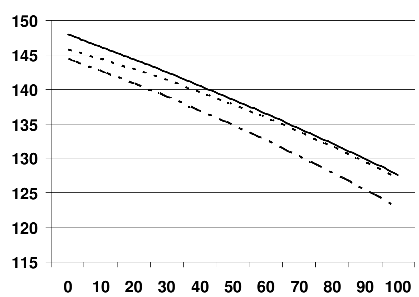


$\varepsilon = 27$

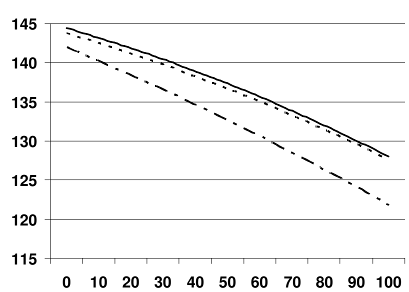


**With Signal Averaging**

$\varepsilon = 8$

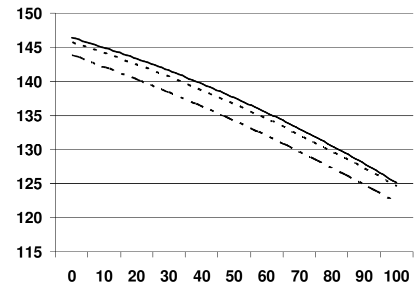


$\varepsilon = 12$

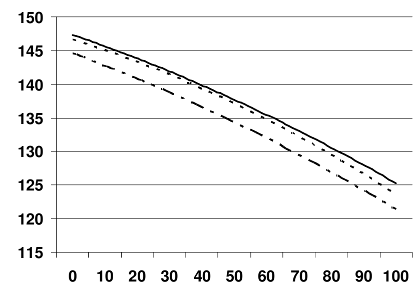


$\varepsilon = 18$

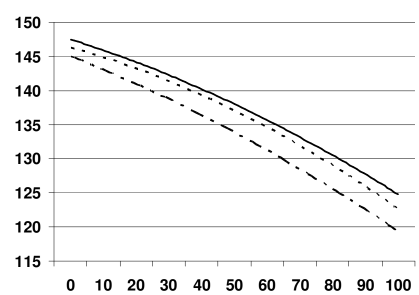


$\varepsilon = 27$

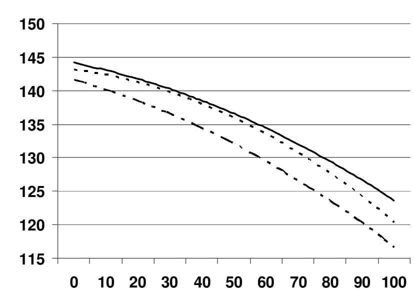

Figure 4.16. Revenue: Seven Bidders
Revenue by Common Value Percent
(I1: Solid, I2: Dot, I3: DashDot)

| Without Signal Averaging | With Signal Averaging |
|---|---|
| $\varepsilon = 8$ | $\varepsilon = 8$ |
| 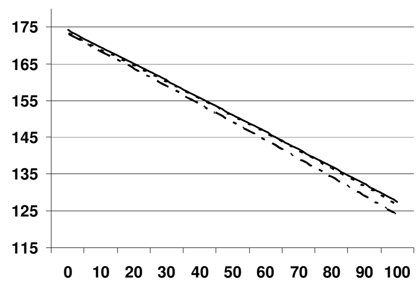 | 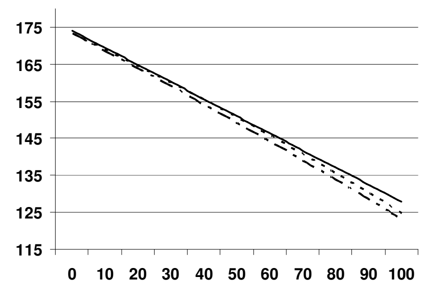 |
| $\varepsilon = 12$ | $\varepsilon = 12$ |
| 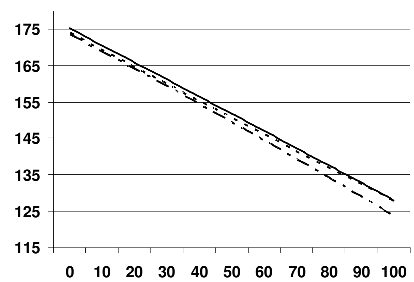 | 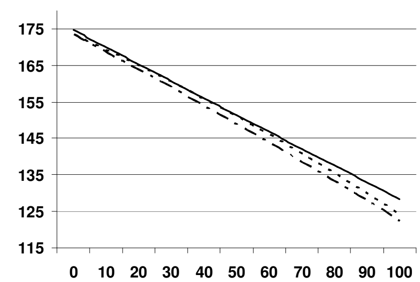 |
| $\varepsilon = 18$ | $\varepsilon = 18$ |
| 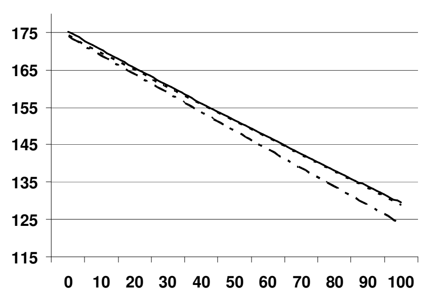 | 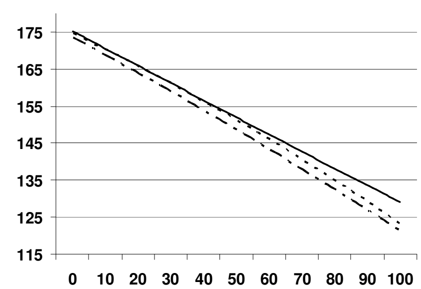 |
| $\varepsilon = 27$ | $\varepsilon = 27$ |
| 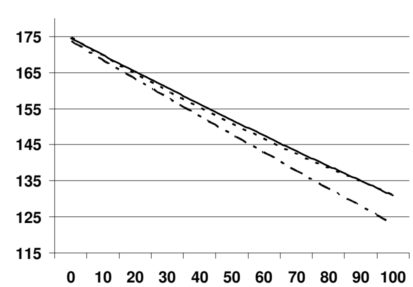 | 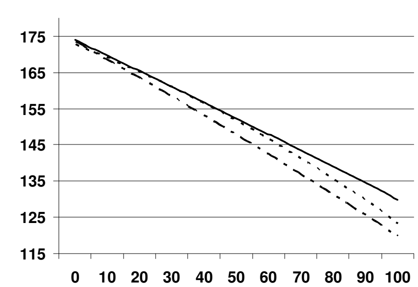 |

<table>
<tr><th colspan="2">Figure 4.17. Revenue Changes with Signal Averaging<br>by Common Value Percent<br>(I1: Solid, I2: Dot, I3: DashDot)</th></tr>
<tr><th>Four Bidders</th><th>Seven Bidders</th></tr>
<tr><td>$\varepsilon = 8$</td><td>$\varepsilon = 8$</td></tr>
<tr><td>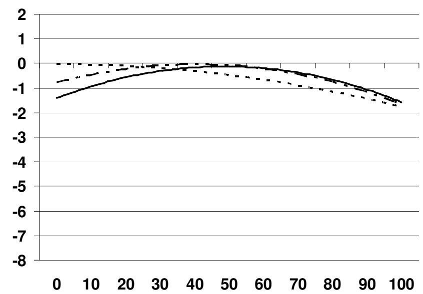</td><td>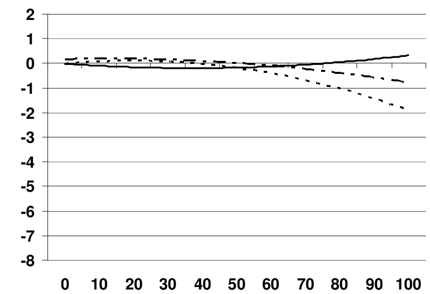</td></tr>
<tr><td>$\varepsilon = 12$</td><td>$\varepsilon = 12$</td></tr>
<tr><td>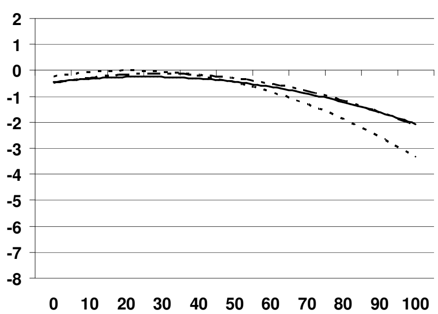</td><td>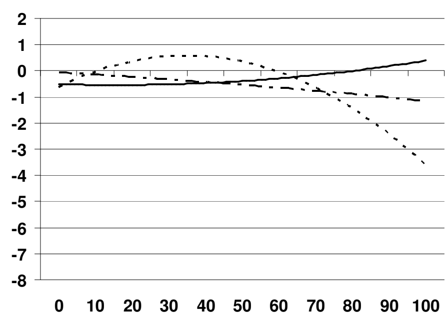</td></tr>
<tr><td>$\varepsilon = 18$</td><td>$\varepsilon = 18$</td></tr>
<tr><td>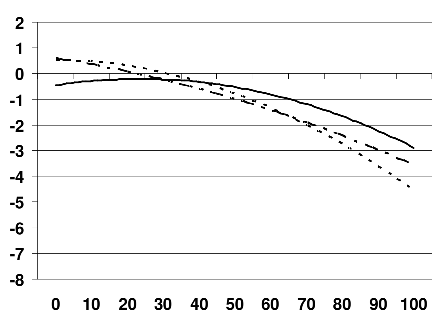</td><td>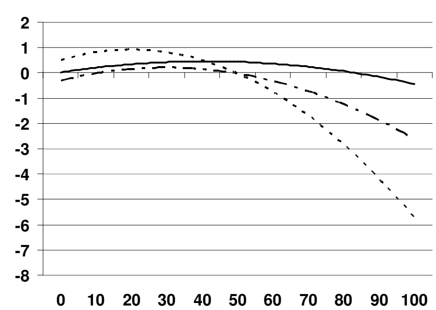</td></tr>
<tr><td>$\varepsilon = 27$</td><td>$\varepsilon = 27$</td></tr>
<tr><td>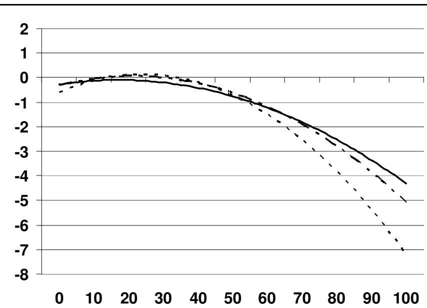</td><td>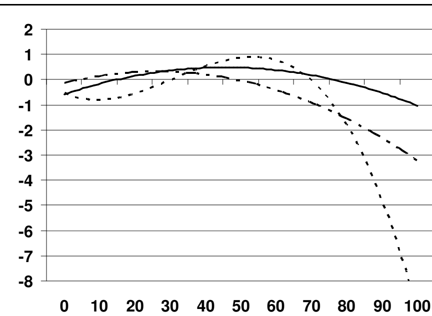</td></tr>
</table>

Figure 4.18. Private-Value Efficiency: Four Bidders
Private-Value Efficiency by Common Value Percent
(I1: Solid, I2: Dot, I3: DashDot)

| Without Signal Averaging | With Signal Averaging |
|---|---|
| $\varepsilon = 8$ | $\varepsilon = 8$ |
| 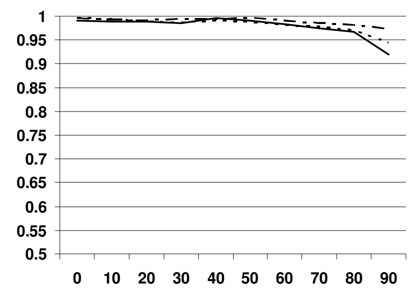 | 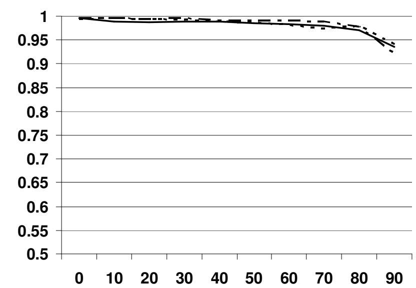 |
| $\varepsilon = 12$ | $\varepsilon = 12$ |
| 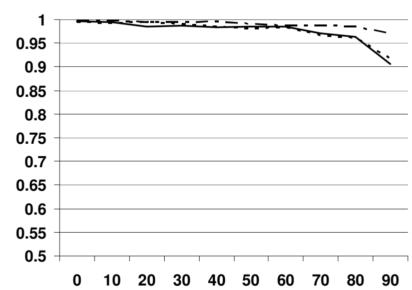 | 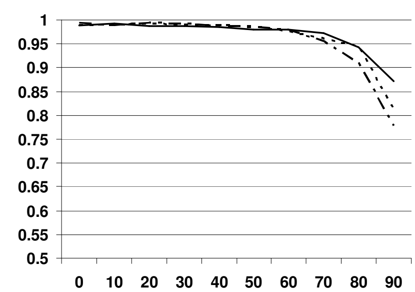 |
| $\varepsilon = 18$ | $\varepsilon = 18$ |
| 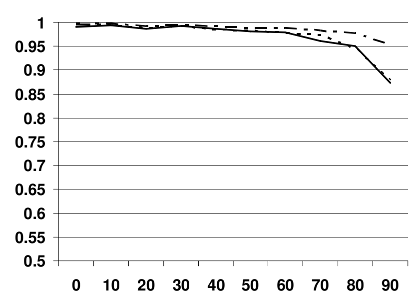 | 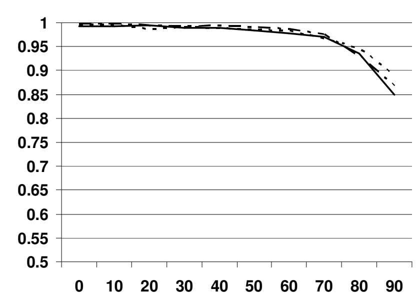 |
| $\varepsilon = 27$ | $\varepsilon = 27$ |
| 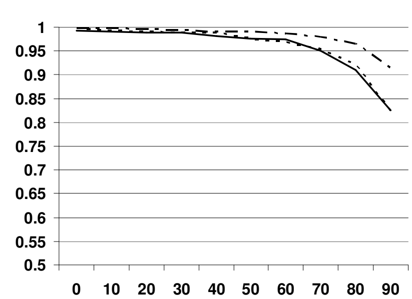 | 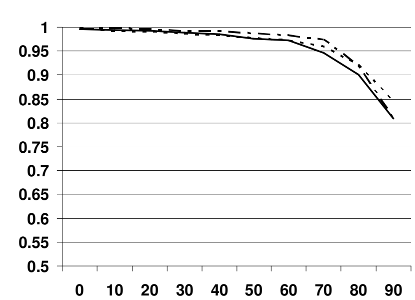 |

| Figure 4.19. Private-Value Efficiency: Seven Bidders<br>Private-Value Efficiency by Common Value Percent<br>(I1: Solid, I2: Dot, I3: DashDot) | |
|---|---|
| Without Signal Averaging | With Signal Averaging |
| $\varepsilon = 8$ | $\varepsilon = 8$ |
| 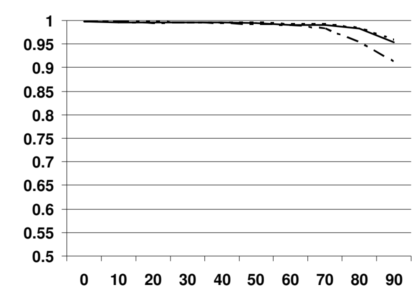 | 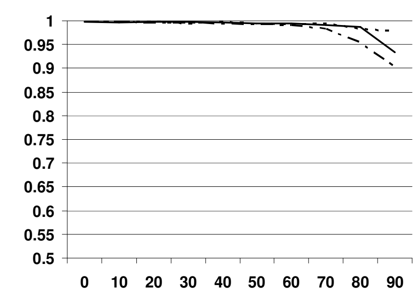 |
| $\varepsilon = 12$ | $\varepsilon = 12$ |
| 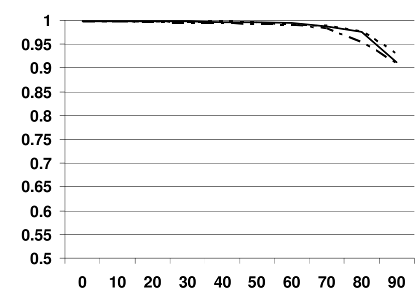 | 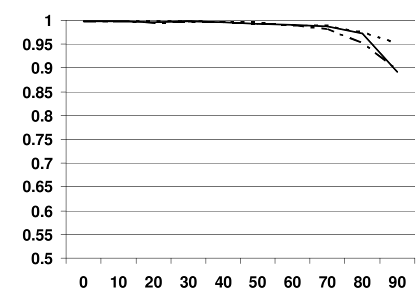 |
| $\varepsilon = 18$ | $\varepsilon = 18$ |
| 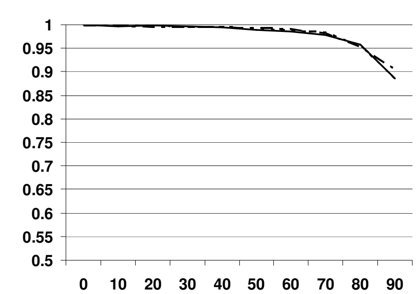 | 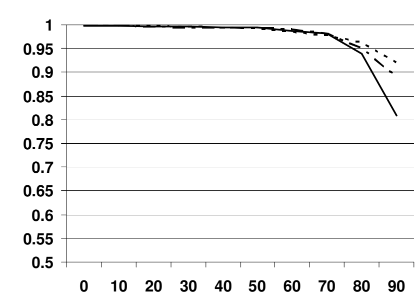 |
| $\varepsilon = 27$ | $\varepsilon = 27$ |
| 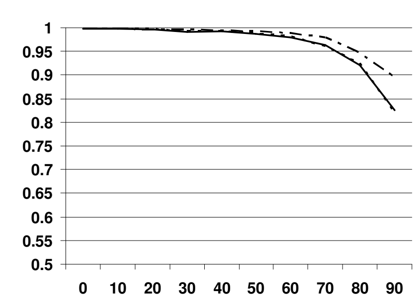 | 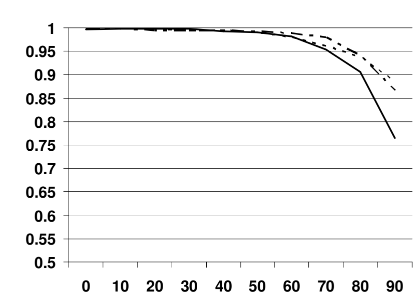 |

<table>
<tr><td colspan="2">Figure 4.20. Private Value Efficiency Changes with Signal Averaging<br>by Common Value Percent<br>(I1: Solid, I2: Dot, I3: DashDot)</td></tr>
<tr><td>Four Bidders</td><td>Seven Bidders</td></tr>
<tr><td>$\varepsilon = 8$</td><td>$\varepsilon = 8$</td></tr>
<tr><td>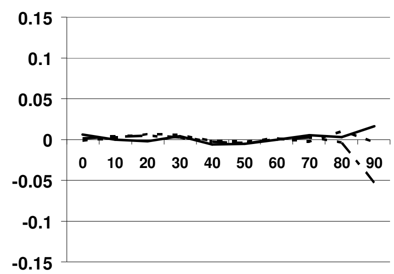
</td><td>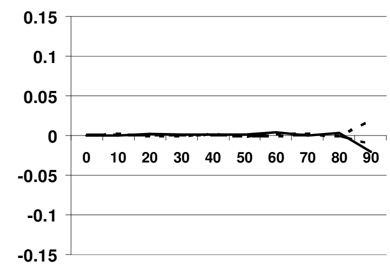
</td></tr>
<tr><td>$\varepsilon = 12$</td><td>$\varepsilon = 12$</td></tr>
<tr><td>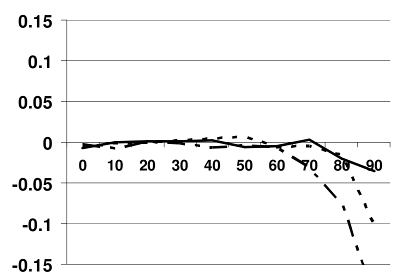
</td><td>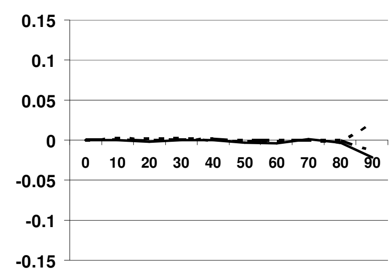
</td></tr>
<tr><td>$\varepsilon = 18$</td><td>$\varepsilon = 18$</td></tr>
<tr><td>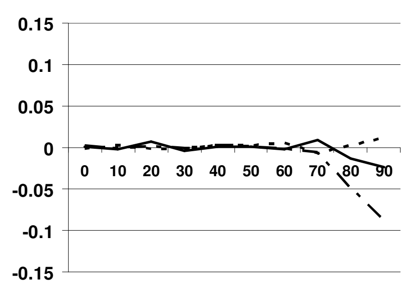
</td><td>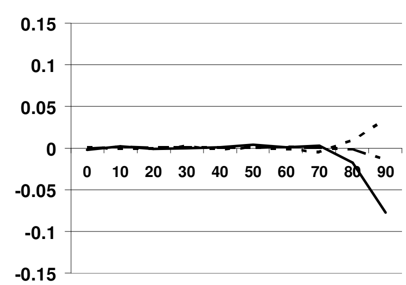
</td></tr>
<tr><td>$\varepsilon = 27$</td><td>$\varepsilon = 27$</td></tr>
<tr><td>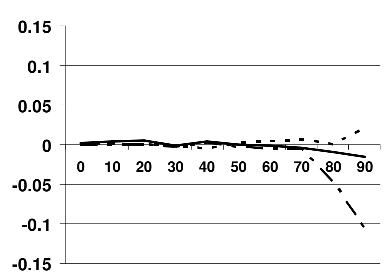
</td><td>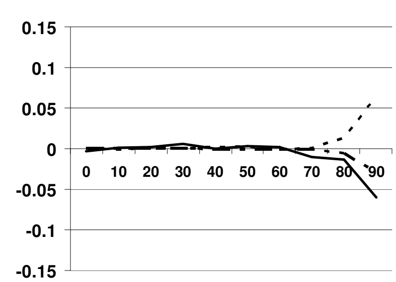
</td></tr>
</table>

# Chapter 5

# Multiagent System Simulations of Treasury Auctions

## 5.1 Introduction

Auctions for treasury securities are used in about 42 countries (Bartolini and Cottarelli, 1997), including Canada. The Bank of Canada and many other central banks manage government debt and the money supply by buying or selling T-bills and bonds of several denominations in treasury auctions. During the period 1998-2002 the U.S. Treasury held more than 800 auctions with a total nominal value of $12.7 trillion (Nyborg and Strebulaev, 2004), and the 2005 and 2006 totals were both about $3.4 trillion (U.S. Department of Treasury). During the last three fiscal years in Canada, the treasury auction volumes have been approximately $175 billion (Government of Canada, 2005, 2006, 2007). Following the auction, there is an inter-dealer secondary market for securities. D'Souza and Gaa (2004) find that average daily trading volumes in the secondary market were $22.4 billion for Canada and $433.5 billion for the U.S. in 2003.

Arnone and Iden (2003) surveyed 39 countries in 2001 and found that 29 had primary dealers and 10 do not. Among the countries with primary dealers, the average number of primary dealers was 14.5. For example, in 2001 Canada had 11 primary dealers for bonds and 9 for T-bills, with bi-annual status reviews. Currently, Canada has 12 primary dealers for bonds and 10 primary dealers for T-bills. Sareen (2005) reports that one Canadian primary dealer explicitly quit in 2001, that dealers will be promoted to primary dealers if they bid for too much, and primary dealers will be demoted if they bid for too little. All in all the number of primary dealers has increased by a couple, but basically remains quite constant. I treat the number of primary dealers and regular dealers as exogenous treatment variables, setting each to 10. The designation of dealers could in principle be made endogenous by using the same formula that the Bank of Canada uses to calculate the bidding limits, but this formula is confidential.

The objective of the Bank of Canada is to sell its entire issue of securities in order to finance the projected government expenditures. Failure to sell this amount will have adverse effects on the government spending ability and would also have a bad signalling effect for investors that the country's debt is not a good investment. So, part of the primary dealer contract (Bank of Canada, 2005) is to ensure that each issuance is sold in full. Each primary dealer has a minimum quantity constraint, but regular dealers have no minimum constraints. In addition to minimum constraints, maximum bid constraints were set because the Bank of Canada was concerned that the auction process could be used to implement a squeeze in the secondary market (Bank of Canada, 1998).

The target overnight interest rate is the middle of the 0.5%-wide operating band. The bank rate is the top of the band. Up until 1996, the bank rate was the average yield at the Bank of Canada's weekly auction of 3-month treasury bills plus 0.25%, so the rate changed weekly with every auction. However, since 1996, the Bank of Canada has decoupled the bank rate from the auction, and the rates are managed through a daily process that is facilitated by the 1999 implementation of an electronic payments system, the Large Value Transfer System (LVTS). Having announced a target interest rate, if the rate being used for trades is less than the target rate, the Bank of Canada uses Sale and Repurchase Agreements (SRA's) to raise the rate. If it is more than the target rate, the Bank of Canada uses Special Purchase and Resale Agreements (SPRA's) to lower the rate. In Europe, this process is usually accomplished with an auction, but in Canada the trades are completed bilaterally between the Bank of Canada and the primary dealers, each having a predetermined limit for both types. The weekly average value of traded SRA's and SPRA's for the period between April 2005 and March 2007 was $140 million (Belisle, 2007). Although there is no public information about returns in these transactions, since they are one-day transactions the return is likely no more than about 0.01%, which amounts to an average weekly profit of only $14,000 to be distributed among the ten or so primary dealers. Thus, the benefit of being a primary dealer is obviously not profit, but there may be some reputational benefits that result in a large client network (Sareen, 2005).

Most of the countries with Treasury auctions use the discriminatory payment rule. A survey by Bartolini and Cottarelli (1997) showed that 39 of 42 countries used the Discriminatory payment rule.[1] Briefly, with Discriminatory payment each bidder pays its bid price, and with Uniform payment all bidders pay the same cutoff price. There has been a long-standing question for Treasury auctions of whether Discriminatory payment or Uniform payment will result in more revenue. Friedman (1963) proposed that the Uniform payment rule would result in more revenue and less collusion, and the U.S. switched from Discriminatory payment to Uniform payment in 1992. The econometric study by Malvey and Archibald (1998) shows slightly higher revenue for the Uniform payment, but they compare revenue by comparing the auction yield to the yield in the when-issued market.

[1] The payment rules are described in Section 5.2.4.3.

The treasury auctions in Canada and some other countries are preceded by when-issued markets and followed by secondary markets,[2] both of which may affect the bidding strategies in the treasury auctions themselves. Nyborg and Strebulaev (2004) criticize Malvey and Archibald's use of the when-issued yield as a measure of revenue, by arguing that the type of auction can influence the yield in the when-issued market. This interaction between the auction payment format and when-issued markets was studied by Nyborg and Sundaresan (1996) who found significant impact of the when-issued market on the auction results. This is, of course, another reason for including the when-issued and secondary markets in the model of the treasury auction.

I develop a model that includes the when-issued market, auction, secondary market, multi-unit bidding in the auction, ten primary dealers and ten regular dealers, and constraints on bidding. Although I am using the Canadian treasury auction as the basis for this study, the results have broader applicability to any auction with before markets, after markets, and bidder constraints. Previous models (see Table 5.1) have been limited to one or two markets and used single-unit bidding in the auction. They have not included the constraints on the primary dealers, and some of the broader models (Chatterjea and Jarrow, 1998 and Nyborg and Strebulaev, 2004) strongly restrict the number of primary dealers who bid or win in the auction.

---

[2] The when-issued and secondary markets are described in Sections 2.3 and 2.5 respectively. The term "when issued" is a short form of "when, as, and if issued" and the market begins after the issue is announced, usually a week before the auction. The "secondary" market refers to the resale market for securities that were purchased in the auction.

This model is used to answer the following research question: Which payment rule provides the most revenue to the central bank and does the answer depend upon the spread of the prices in the when-issued and secondary markets? To answer the research question, I have extended the more general multiagent system that is described in Chapter 2. The agents learn their bidding strategies using a modified impulse balance learning method similar to the method in Chapter 3 for single-unit auctions. In this model, agents attempt to maximize their profits by adjusting bidding strategies in repeated auctions that are the same except for the strategies of other bidders that are also being adjusted.

I find that the Spanish payment rule is revenue inferior to the Discriminatory payment rule across all market price spreads, but the Average rule is revenue superior. For most market-price spreads, Uniform payment results in less revenue than Discriminatory, but there are many cases in which Vickrey payment produces more revenue.

The market model is presented in Section 5.2 and the agent learning model in Section 5.3. Section 5.4 describes the learning adjustment rules in detail, but more details are also provided in the Appendix. Section 5.5 analyzes the model's sensitivity to parameter values, demonstrates convergence, and discusses the endogenous agent variation. Section 5.6 presents the main results of the simulations that compare revenue across the payment rules and market price spreads. Section 5.7 presents conclusions.

## 5.2 Market Model

This section first provides a brief overview of the market model illustrated in Figure 5.1. Section 5.2.1 explains the market prices in the when-issued and secondary markets, which are used as values by the bidders in the auction. This leads to a detailed discussion of the three markets in the model: when-issued market (Section 5.2.2), auction (Section 5.2.3), and secondary market (Section 5.2.4). The model notation is summarized in Table 5.2.

Figure 5.1 illustrates an overview of the three markets that occur sequentially in every period *t*. There are three security categories: securities that are sold short in the when-issued market ("security S"), bought long to sell to short-selling dealers in the secondary market ("security D"), and bought to sell to customers in the secondary market ("security C"). A dealer is said to be selling short when it sells securities that it does not yet own and buying long when it does not yet have customers to sell to.

The first market is the when-issued market in which a dealer *i* sells $q_{i,t}^{S}$ securities at the market price $P_S$, where the quantity $q_{i,t}^{S}$ is endogenous but the market price $P_S$ is exogenous. In the auction, the dealer will attempt to buy the securities to cover its short sale quantity $q_{i,t}^{S}$. The when-issued market price $P_S$, together with the secondary market prices price $P_C$ and $P_D$ are exogenous in the current version of the model. To make them endogenous, I would create a model of the when-issued market and a model of the secondary market that dynamically extends across time periods. This is a possibility for future work.

The second market is the treasury auction, which consists of dealers submitting sealed bids to the central bank. Dealers who sell short in the when-issued market will be trying to cover their position by bidding for $q_{i,t}^{S}$ securities at bid price $b_{i,t}^{S}$. If their bid price is high enough they will be allocated $x_{i,t}^{S}$ securities at a payment $p_{i,t}^{S}$. Some dealers may bid for $q_{i,t}^{D}$ securities at bid price $b_{i,t}^{D}$ to sell to short sellers who are unable to cover their position in the auction (i.e., $x^S_{i,t} < q^S_{i,t}$). If their bid price is high enough they will be allocated $x^D_{i,t}$ securities at a payment $p^D_{i,t}$. Bidders may also bid for securities to sell directly to customers in the secondary market. They bid for $q^C_{i,t}$ securities at bid price $b^C_{i,t}$, and are allocated $x^C_{i,t}$ securities at a payment $p^C_{i,t}$. The bid prices, quantities, allocations, and payments are all endogenous.

The third market is the secondary market in which long-buying dealers sell to short-selling dealers at price $P_D$ and all dealers sell to customers at price $P_C$, where $P_D$ and $P_C$ are exogenous. The long dealers sell as much as they can to the short sellers at price $P_D$ and sell the remainder $\delta^D_{i,t}$ to customers at the lower price $P_C$. Any dealer may have been allocated $x^C_{i,t}$ securities to sell directly to customers, which they sell at price $P_C$. The amount of allocated $x^D_{i,t}$ that long-buyers sell to short-sellers is designated $\tau^D_{i,t}$ and the amount of allocated $x^C_{i,t}$ that is sold to customers is $\tau^C_{i,t} = x^C_{i,t}$. The amount of $x^D_{i,t}$ that is not sold to short-sellers is designated $\delta^D_{i,t}$, and this amount is sold to customers.

### *5.2.1 Market Prices*

Participants in treasury auctions submit bids that consist of one or more pairs of quantity and yield. For ease of understanding, I replace yield with the equivalent price per $100 as presented in the auction results from the Bank of Canada. Since the purpose of the model is to understand how the when-issued market, secondary market, and auction mechanism affect the auction outcomes, the choice variables of the model are the auction bid prices and quantities; the prices in the when-issued and secondary markets are exogenous parameters.

I use a set of three posted prices: the price charged to customers by short sellers in the when-issued market $P_S$, the price charged to short sellers by dealers in the secondary market $P_D$, and the price charged by all dealers to customers in the secondary market $P_C$. A dealer sells securities to customers in the when-issued market with the intention of buying them in the auction, a process called "short selling." It competes in the auction with dealers who want to take advantage of the short-sellers' contractual commitments to their customers by buying all of the securities in the auction, a process called "long buying." If a short seller fails to cover its position in the auction, it is forced to pay the high price $P_D$ in the secondary market, called a "squeeze." In the secondary market, a high squeeze price $P_D$ occurs in reality (Merrick *et alia*, 2005) and is used in other models such as Chatterjea and Jarrow (1998) and Nyborg and Strebulaev (2004).

I set the non-auction prices using the relative magnitudes $P_D > P_S > P_C$. The first inequality is straightforward, since a dealer pays the high price $P_D$ in the secondary market if it is squeezed as a result of a short position not being covered by auction winnings. According to Nyborg and Sundaresam (1996) the when-issued market provides price and quantity insurance so the inequality $P_S > P_C$ can be justified by when-issued customers willing to pay a premium rather than risk not obtaining their quantity in the secondary market. The ordering $P_D > P_S > P_C$ is consistent with prices in previous models (Chatterjea and Jarrow, 1998; Nyborg and Strebulaev, 2004).

The market prices are parameters in the simulations. First, I assume that the base yield is 4.25% for a 12-month T-Bill. Then the associated price is approximately \$95.90

and I set this as the short-sale price $P_S$ in the when-issued market.[3] Then, for simulation purposes I define variables $\varepsilon_D$ and $\varepsilon_C$ by $P_D = P_S + \varepsilon_D$ and $P_C = P_S - \varepsilon_C$, and vary $\varepsilon_D$ and $\varepsilon_C$. Second, I assume that this price interval [$P_S - \varepsilon_C$, $P_S + \varepsilon_D$] is no larger than the price equivalent of about 50 basis points or 0.5%. For the purpose of the simulations, I use $P_S = \$95.90$ and both epsilons less than \$0.20, i.e., $\varepsilon_D$ and $\varepsilon_C$ both less than the price equivalent of 25 basis points. Thus, the simulations are run to analyze prices that are in the ranges $\$95.90 < P_D < \$96.10$ and $\$95.70 < P_C < \$95.90$.

### *5.2.2 When-Issued Market*

Bikhachandani *et alia* (2000) explain that dealers will sell short to customers in the when-issued market since this does not provide information to other dealers. They do not buy long because they would have to do this with other dealers, and this would convey too much information about order flow. In general, a sizable portion of the issue is sold in the when-issued market. Hortacsu and Sareen (2005) found that in 2002 Canadian Treasury auctions 54% of the total volume traded in the when-issued and secondary markets consisted of customer orders. Each dealer receives customer orders and fills these orders by short selling the quantity $q_{i,t}^S$ at the price $P_S$ for when-issued revenue of $P_S q_{i,t}^S$. When-issued trading volume is endogenous in the model. The simulations involve agents adjusting their $q_{i,t}^S$ in order to maximize profits.

---

[3] A one-year T-bill with face value of \$100 and a yield of 4.25% has a price of 100/(1+.0425) = \$95.92.

### *5.2.3 Auction Market*

#### 5.2.3.1 Bidding

In the auction, each dealer endeavours to maximize overall profit by choosing prices and quantities. Each bid consists of up to three price-quantity pairs, i.e., it is a three point demand schedule. Table 5.3 shows bids from a Government of Canada auction reported in Lu and Yang (2003). Most bids have three yield-quantity pairs, but some have two or four points. For the sake of simplicity, in this model all bids have three price-quantity pairs for the quantities to cover the short selling (S) from the when-issued market, to allow a dealer squeeze (D) in the secondary market, and to sell to customers (C) in the secondary market. Thus, each dealer *i* submits a set of 3 bids consisting of a chosen price that is associated with these quantities: $\beta_{i,t} = \left\{(b_{i,t}^S, q_{i,t}^S), (b_{i,t}^D, q_{i,t}^D), (b_{i,t}^C, q_{i,t}^C)\right\}$.

Each dealer will hope that its bidding strategy covers its short position since it is otherwise subject to being squeezed in the secondary market. The dealer will be squeezed if its auction winnings are less than its short position in the amount $\Delta_{i,t}^S = q_{i,t}^S - x_{i,t}^S$. Given short-selling revenue $P_S q_{i,t}^S$ in the when-issued market, a dealer makes the following decisions in the auction:

1. a bid price decision $b_{i,t}^S$ to maximize profit from short selling, resulting in an allocation of $x_{i,t}^S$.
2. a decision about bid price and quantity $(b_{i,t}^D, q_{i,t}^D)$ in expectation of profits from selling these to the short-selling dealers in the secondary market, resulting in an allocation of $x_{i,t}^D$. I assume that any D securities that are not sold to short-sellers are sold to customers, with the former quantity denoted $\tau_{i,t}^D$ and the latter $\delta_{i,t}^D$ so that $\tau_{i,t}^D + \delta_{i,t}^D = x_{i,t}^D$.

3. a decision about $(b_{i,t}^C, q_{i,t}^C)$, resulting in an allocation of $x_{i,t}^C$, and where I assume that all securities can be sold to customers in the secondary market.

Depending upon the auction payment rule, the payments paid by dealer *i* for its allocations $x_{i,t}^S, x_{i,t}^D$, and $x_{i,t}^C$ are $p_{i,t}^S, p_{i,t}^D$, and $p_{i,t}^C$. Thus, I can write the *ex ante* profit decision for the dealer in the auction as

$$\max_{b_{i,t}^S, b_{i,t}^D, q_{i,t}^D, b_{i,t}^C, q_{i,t}^C} \left\{ \begin{aligned} &P_S q_{i,t}^S - E[p_{i,t}^S x_{i,t}^S] - P_D(q_{i,t}^S - E[x_{i,t}^S]) \\ &+E[\tau_{i,t}^D](P_D - E[p_{i,t}^D]) + E[\delta_{i,t}^D](P_C - E[p_{i,t}^D]) \\ &+E[x_{i,t}^C](P_C - E[p_{i,t}^C]) \end{aligned} \right\}$$

Each primary dealer has a minimum quantity constraint of $\underline{\theta}_i$ = 50% of its maximum bid limit, but regular dealers have no minimum constraints. The aggregated maximum is $\overline{\theta}_i$ = 40% of the issue amount $Q$ for primary dealers and $\overline{\theta}_i$ = 20% of the issue amount for regular dealers. Then, for agent *i* the maximum total bid quantity is $q_i^{\max} = \overline{\theta}_i Q$, the minimum is $q_i^{\min} = \underline{\theta}_i Q$. After every quantity adjustment the bidders must ensure that $q_i^{\min} \le q_{i,t}^D + q_{i,t}^C + q_{i,t}^S \le q_i^{\max}$. For example, if the current set of bid quantities sums to less than $q_i^{\min}$, the agent must increase the bid quantities.

#### 5.2.3.2 Cutoff Price and Allocations

Bids are allocated starting with the highest price and continuing down until the cutoff price $p_t^{cut}$ is reached, where $p_t^{cut}$ is the highest price such that the total bid quantity is greater than or equal to the issued amount. If the total quantity bid at $p_t^{cut}$ is greater than the issued amount, the bids at $p_t^{cut}$ are only partially allocated. Also, if there is more than one bidder at $p_t^{cut}$, the remaining quantity is shared among the bidders according to the quantities in their bids. For each bidder, the allocated quantities are $x_{i,t}^S$, $x_{i,t}^D$, and $x_{i,t}^C$.

#### 5.2.4.3 Payment rules

Now we can easily substitute different payment rules, and the simulations in this paper use five payment rules: Discriminatory, Uniform, Average, Spanish, and Vickrey. Let $p_{i,t}^j$ be the auction price for security type $j \in$ {S, D, C}. The Discriminatory payment rule is pay-your-bid, so the payments are $p_{i,t}^S = b_{i,t}^S$, $p_{i,t}^D = b_{i,t}^D$, and $p_{i,t}^C = b_{i,t}^C$. For a Uniform cutoff price payment rule, all securities are priced at the cutoff price so $p_{i,t}^S = p_{i,t}^D = p_{i,t}^C = p_t^{cut}$, and for an Average payment rule, all securities are priced at the average price so that $p_{i,t}^S = p_{i,t}^D = p_{i,t}^C = p_t^{avg}$. Abbink *et alia* (2006) describe the Spanish auction that uses a combination of Average and Discriminatory approaches. Winning bids that are above the weighted average winning bid pay the same price, namely the weighted average winning bid. Winning bids that are below the weighted average winning bid pay the bid price, as in the Discriminatory method. Vickrey payment requires that a dealer who wins *k* units in a multi-unit auction pays the *k* highest losing bids of the other bidders. Ausubel (2004, 2006) has recently proposed some new auction designs that use the English auction format, but this paper focuses on the sealed-bid format currently used by most of the central banks.

### *5.2.4 Secondary Market*

As illustrated in Figure 5.1, the secondary market potentially consists of long-buying dealers selling what they can to short-selling dealers and selling the residual to customers, and of all dealers selling to customers. The secondary market for dealers that have acquired C securities in the auction is straightforward. The acquired $x_{i,t}^{C}$ securities in the auction are all sold at price $P_C$ in the secondary market. The situation for securities S and D is more complicated. If a short-selling dealer does not cover its short position $q_{i,t}^{S}$ with its allocation $x_{i,t}^{S}$ in the auction, it must buy $\Delta_{i,t}^{S} = (q_{i,t}^{S} - x_{i,t}^{S})$ at a high price $P_D$ in the secondary market. Similarly, the long-buying dealer will try to sell its allocation $x_{i,t}^{D}$ to the short sellers. Let $\sum^{D}$ denote the total amount of security D that is available for sale by the long buyers, and $\sum^{S}$ denote the total amount of security S that needs to be bought in the secondary market by the short sellers. Then $\sum^{D} = \sum_{i=1}^{n} x_{i,t}^{D}$ and $\sum^{S} = \sum_{i=1}^{n} \Delta_{i,t}^{S}$. While the agents are learning, there may be an excess on either side of the market, i.e., we can have $\sum^{D} > \sum^{S}$ or $\sum^{S} > \sum^{D}$. In the first case, there is an excess of security D since the squeezing dealers have overbought in the auction, but the squeezed short-sellers can buy what they need to cover their short sells. In the second case, there is an excess of security S since the squeezing dealers have under bought in the auction and the short-sellers cannot cover their positions. Neither one of these situations should be an equilibrium by the end of the simulation since the agents will make quantity adjustments (see Section 5.4).

There are at least three ways to model the secondary market transactions. The first method simply allocates the securities proportionally. Denoting the actual cleared amounts $\tau_{i,t}^{D}$ and $\tau_{i,t}^{S}$, in the case of an excess of security D ($\sum^{D} > \sum^{S}$) we have $\tau_{i,t}^{D} = \frac{\sum^{S}}{\sum^{D}} x_{i,t}^{D}$ and $\tau_{i,t}^{S} = x_{i,t}^{S}$; in the case of an excess of security S ($\sum^{S} > \sum^{D}$) we have

$\tau_{i,t}^{S} = \frac{\sum^{D}}{\sum^{S}} x_{i,t}^{S}$ and $\tau_{i,t}^{D} = x_{i,t}^{D}$. Using this method, however, results in instability in the learning model. Agents can oscillate back and forth between over allocation and under allocation by a very small quantity that in turn results in oscillation between profitability and unprofitability. The learning model is stable when secondary market trades are based on market power, as done by Nyborg and Strebulaev (2004). There are two ways to implement this. One method is to allocate secondary market sales to the largest security holders first. Given an ordering of the quantities for both the short sellers (S) and long buyers (D), match the largest long buyer with the largest short seller and if there are long holdings left the largents long sells to the next largest short, and so on. Once the largest long buyer runs out of holdings to sell, the next largest long buyer takes over and continues to sell down the list to the shorts in decreasing order. Another method is to make the realistic assumption that the primary dealers have more market power than the regular dealers. In this case, primary dealer short sellers and primary dealer long buyers are given first priority in secondary market transactions. This second method is the one used in the simulations.

In the case of an excess of security D in the secondary market, a long-buying agent may up holding $\delta_{i,t}^{D} = (x_{i,t}^{D} - \tau_{i,t}^{D})$ securities that it can sell to customers at the lower price $P_C$. The overall profit for security D is therefore $\pi_{i,t}^{D} = P_D \tau_{i,t}^{D} + P_C \delta_{i,t}^{D} - p_{i,t}^{D} x_{i,t}^{D}$. With an excess of security S in the secondary market, a short selling agent may not be able to clear a portion of its position $\delta_{i,t}^{S} = (\Delta_{i,t}^{S} - \tau_{i,t}^{S})$. The short-selling agent cannot hold these securities because they belong to customers, and the short-seller must cancel the orders resulting in a refund of $-P_S \left( \Delta_{i,t}^{S} - \tau_{i,t}^{S} \right)$. The overall profit for security S is thus

$\pi_{i,t}^{S} = P_S q_{i,t}^{S} - \left( p_{i,t}^{S} x_{i,t}^{S} + P_D \tau_{i,t}^{S} \right) - P_S \left( \Delta_{i,t}^{S} - \tau_{i,t}^{S} \right)$. In summary, the profit from short selling is the revenue from the when-issued market minus the sum of the amounts paid in the auction and secondary market, minus the revenue from cancelled orders. The profit from long buying is the revenue from the secondary market trading with the short sellers and customers minus the amount paid in the auction.

## 5.3 Learning Model

This section starts with a description of the linear bid adjustment strategy (Section 5.3.1). The discussion of information levels in Section 5.3.2 explains why the agents use a limited set of private information. Section 5.3.3 explains the general concepts of learning how to adjust bid strategies in a multi-market environment.

### *5.3.1 Linear Adjustments*

When an auction is considered in isolation (i.e. before and after markets are not considered), an agent sets a bid price in terms of its value estimate in the auction by using feedback on profit, foregone profit, and money on the table (as in the models of Chapter 3 and Chapter 4). However, the market model described in Section 5.2 uses three markets and three securities so that an agent's value is a mix of the three market prices $P_S$, $P_D$, and $P_C$. An agent adjusts both bid prices and bid quantities for the three securities using simple linear adjustments. Several empirical studies on impulse-balance learning have shown that linear adjustments to bidding strategies explain how bidders adjust their bid strategies (Selten and Buchta, 1998; Selten *et alia*, 2005; Ockenfels and Selten, 2005; Negebauer and Selten, 2006; Garvin and Kagel, 1994; Kagel and Levin, 1999). I have previously adapted the empirical impulse balance method (Ockenfels and Selten, 2005) for use in computational models for single-unit sealed-bid (Chapter 3) and English auctions (Chapter 4). In this model, it is extended further to the multi-unit case.

The bid price is adjusted using $b_{i,t+1}^{j} = b_{i,t}^{j} + \phi a_{i,t}^{j}$, where $\phi$ is the price learning rate and $a_{i,t}^{j}$ is the period *t* price adjustment for agent *i*'s security *j*. Similarly, the quantity is adjusted using $q_{i,t+1}^{j} = q_{i,t}^{j} + \chi z_{i,t}^{j}$, where $z_{i,t}^{j}$ is a period *t* quantity adjustment for agent *i*'s security *j* and $\chi$ is the quantity learning rate.

### *5.3.2 Information Feedback*

When the agent wins a full or partial allocation of a security, it can determine its profit and foregone profit from its private information about its payment. When the agent has no allocation, when the payment rule is Discriminatory it can calculate foregone profit using its bid price $b_{i,t}^{j}$. The Bank of Canada posts public information after each auction: the maximum bid price $b_{t}^{high}$, average price $b_{t}^{avg}$, and low (cutoff) price $b_{t}^{cut}$. Using this information, the agent can also determine its foregone profit when it has no allocation for a security and the payment rule is non-discriminatory. For the Uniform payment rule, the foregone profit is estimated using payment $b_{t}^{cut}$; for the Average payment rule it is $b_{t}^{avg}$; and for the Spanish payment rule the payment used is $b_{i,t}^{j}$ since with no allocation $b_{i,t}^{j}$ will be less than $b_{t}^{avg}$. For the Vickrey payment rule, the losing bids are not known so the agent uses the upper bound on the payment, $b_{t}^{cut}$.

For the following reasons, I assume that the agents do not use the public information feedback to modify their bidding strategies. An agent might reason that it may not want to reduce its bid price below $b_{t}^{cut}$, since this might adversely affect its probability of winning, and may not want to increase its bid price above $b_t^{high}$, since this might unnecessarily reduce its profitability. This would result in a constraint on the bid price strategy of $b_{i,t+1}^j = \min\left\{b_t^{high}, \max\left\{b_t^{cut}, b_{i,t}^j\left(1+\phi^j a_{i,t}^j\right)\right\}\right\}$. However, simulations not reported here demonstrate that this constraint on adjustments results in negative profit, and thus agents in fact learn to not apply this constraint. The reason for the constraint's poor outcomes is the complexity of the interplay between allocation and clearing, especially for security S, and the narrow price band within which the adjustments must occur. The constraint interferes with the agent's ability to make the fine adjustments in bid price that are required to maximize profits in this complex environment.

### *5.3.3 Bid Adjustment Concepts*

For each security, the agent applies different adjustments depending upon its allocation in the auction and the amount of clearing in the secondary market. Security S (sold short in the when-issued market) is distinguished by being bought in the secondary market, whereas securities D (sold to dealers) and C (sold to customers) are sold. Security C is distinguished by being treated as fully sold to customers in the secondary market, whereas sales of D are matched with buys of S, with leftovers sold to customers.

Figure 5.2 illustrates the general concepts of the learning model. The general approach to the bid price is to raise it when the agent can make a profit with more allocation, and lower it when more allocation would lead to a loss. Figure 5.2 shows that, for each security, the agent generally increases its bid price when the foregone profit is greater than zero and decreases its bid price when the foregone profit is less than zero. The agent's intended effect of raising its bid price is to increase the probability of achieving an allocation in the auction, and the intended effect of lowering its bid price is to increase its expected profit, usually through indirect effects on the payment.

The general approach to the bid quantity is to raise it when there is a profitable full allocation and to lower it when there is an actual loss or an allocation or clearing shortfall. Figure 5.2 shows that generally the quantity is increased when there is a full allocation ( $\Delta_{i,t}^{j} = 0$ ). For securities D and S, quantity is usually decreased when there is a clearing shortfall ( $\delta_{i,t}^{D} > 0$, $\delta_{i,t}^{S} > 0$ ). For security C, the main condition for decreasing quantity is a partial or nil allocation ( $\Delta_{i,t}^{C} > 0$ ). The agent's intended effect of raising its bid quantity is to increase overall profits, and the intended effect of lowering its bid quantity is to reduce current or potential losses. For quantity adjustments, the agent raises its bid quantity when the current quantity has been fully allocated in the auction, fully cleared the secondary market, and the resulting profit is positive.

Each time the bidder makes quantity adjustments, it checks to make sure that the constraint $\underline{q}_i \leq q_{i,t}^{D} + q_{i,t}^{C} + q_{i,t}^{S} \leq \overline{q}_i$ is satisfied. If the maximum constraint is exceeded it, the bidder lowers each security quantity in proportion to achieve the upper bound. If the minimum constraint is subceeded, the agent has more of a challenge in determining how to raise its bid quantities. The challenge comes when the agent needs to significantly reduce the quantity of a bid for Security S because it is unprofitable even with full allocation (Rule S2) of its when-issued quantity. Since a straightforward balancing is likely to raise this quantity right back up again, the agent conditions on this rule and does not scale the quantity up for this security and instead scales up security D, C, or both.

## 5.4 Learning Adjustment Rules

The adjustment rules are described in detail the following sections and summarized in Tables 5.4 to 5.8. In these tables a downward arrow denotes a downward adjustment, and an upward arrow denotes an upward adjustment. The adjustments for bid price and quantity are shown as functions $b(\bullet)$ and $q(\bullet)$ respectively, with the former being a function of profit $\pi$ or foregone profit $\pi^F$ and the latter being a function of 1, the allocation shortfall $\Delta$, the clearing shortfall $\delta$, or the bid quantity $q$. Table 5.4 shows adjustment rules when there is full allocation of securities S and C and full allocation and full clearing of security D. Table 5.5 shows the adjustment rules for security S when the agent has only a partial or nil allocation and thus must try to cover the shortfall in the secondary market. The adjustments depend upon whether the bidder achieves a full or partial/nil clearing of this shortfall in the secondary market. Tables 5.6 to 5.9 show the adjustment rules for security D when there is only partial/nil clearing with full allocation and when there is partial/nil allocation. Table 5.10 shows the adjustments for security C for partial allocation and nil allocation. Recall that security C always clears the secondary market. Tables 5.11 to 5.13 list the adjustment rules for securities S, D, and C respectively.

### *5.4.1 Full Allocation and Full Clearing of S, D, and C*

The main objective of bidding for security S is to achieve a full allocation in the auction. In this case the agent avoids paying a high squeeze price in the secondary market and avoids the risk of cancelling the unfilled customer orders. The objective of an agent bidding for security D is to achieve a full allocation that it also fully clears by selling to short sellers in the secondary market. Security C is set up to be fully cleared in the secondary market so an agent's only concern is to achieve a full allocation in the auction. As shown in Table 5.4, the agent will adjust its bid based on its profitability and as shown in the Appendix the maximum profits for the three security categories are $\pi^S_{i,t} = (P_S - p^S_{i,t})q^S_{i,t}$. $\pi^D_{i,t} = (P_D - p^D_{i,t})q^D_{i,t}$, and $\pi^C_{i,t} = (P_C - p^C_{i,t})q^C_{i,t}$.

**Rule j1 (j = S, D, C)**: When profit is non-negative, the agent maintains its price since raising it risks unprofitability and lowering it risks the full allocation. However, the agent may be able to improve its overall profit by increasing its quantity. Since too large an increase may result in unallocated units (that require purchase in the secondary market) , the agent's increment is limited to the minimal quantity of 1.

**Rule j2 (j = S, D, C)**: When the agent is losing money, it lowers its probability of allocation by its bid price in proportion to the loss $\pi^S_{i,t}$. On the other hand, the loss is occurring with the most ideal results (full allocation or full allocation and full clearing) so a loss results only when $p^j_{i,t} > P_j$, which is usually always an irreversible state. The agent has no potential benefit from bidding for this security, so it lowers its bid quantity in proportion to the current quantity $q^j_{i,t}$.

### *5.4.2 Security S*

If the agent has a partial (or nil) allocation of security S in the auction ($0 \le x^S_{i,t} < q^S_{i,t}$), but is fully cleared in the secondary market ($\tau^S_{i,t} = \Delta^S_{i,t} \to \delta^S_{i,t} = 0$), the adjustments depend upon the level of profit and foregone profit. Table 5.5 shows the adjustment rules for full clearing of partial allocation (S3 to S5) and for partial clearing of partial allocation (S6 and S7). All of the adjustment rules for security S are listed in Table 5.11. The derivations of the profits and foregone profits are in the Appendix.

**5.4.2.1 Partial/Nil Allocation with Full Clearing**

**Rule S3**: If profit and foregone profit are both non-negative ($\pi_{i,t}^{S} \geq 0$ and $\pi_{i,t}^{F,S} \geq 0$), the agent anticipates further profits if it can achieve a full allocation. Thus, the agent raises its bid price in proportion to the foregone profit in an attempt to achieve a full allocation in the auction. There is no reason for the agent to raise or lower its quantity until it tries to achieve full allocation.

**Rule S4**: If foregone profit is non-negative, but actual profit is negative ($\pi_{i,t}^{S} < 0$ and $\pi_{i,t}^{F,S} \geq 0$), the agent can anticipate profits if it can achieve a full allocation. The agent raises its bid price in proportion to the foregone profit in an attempt to achieve a full allocation in the auction. However, since a full allocation may not be achievable, the agent reduces its current actual loss by decreasing its bid quantity by the marginal quantity of 1.

**Rule S5**: A partially-allocated, fully-cleared agent that has a negative foregone profit lowers its price since it does not want to achieve an allocation that will result in a negative profit. This lowering is in proportion to the negative foregone profit. As shown in the Appendix, foregone profit is $(P_D - p_{i,t}^{S})\Delta_{i,t}^{S}$, and it is negative when $p_{i,t}^{S} > P_D$. This state is more likely to occur for the Average and Spanish payment rules that tend to have higher payments and, since the agents in general have will have been increasing their bid prices to achieve full allocations, this state tends to be a steady state. Thus, the agent will never be able to achieve an actual profit by being allocated an additional $\Delta_{i,t}^{S}$ units so it reduces its bid quantity in proportion to $\Delta_{i,t}^{S}$. The result is that the bid quantities will be shifted to more potentially profitable securities.

### 5.4.2.2 Partial/Nil Allocation with Partial/Nil Clearing

**Rule S6:** For partial (or nil) allocation ($0 \le x_{i,t}^S < q_{i,t}^S$) with partial clearing ($0 < \tau_{i,t}^S < x_{i,t}^S \to \delta_{i,t}^S > 0$), the agent is short-selling securities that it is not able to cover in the auction or in the secondary market. When the agent receives no allocation in the auction ($x_{i,t}^S = 0$), the agent attempts to cover the entire bid quantity in the secondary market. When there is a partial allocation, the agent knows the payment and can calculate the foregone profit $\pi_{i,t}^{F,S}$. When it achieves no allocation, the agent estimates foregone profit $\hat{\pi}_{i,t}^{F,S}$, as explained in the Appendix. To the extent that the possibility of profit exists ($\pi_{i,t}^{F,S} \ge 0$ or $\hat{\pi}_{i,t}^{F,S} \ge 0$), the agent raises its bid price to try to achieve a full allocation in the auction, which is the preferred result. It will also want to reduce the bid quantity to reduce its exposure in the secondary market. Since the agent knows that $\delta_{i,t}^S$ units have not cleared the secondary market, it reduces its bid quantity in proportion to $\delta_{i,t}^S$.

**Rule S7**: If the partially-allocated, partially cleared agent has a negative foregone profit or negative estimated foregone profit, it lowers its bid price in order to reduce its allocation. From the Appendix, foregone profit is $(P_D - p_{i,t}^S)\Delta_{i,t}^S - (P_D - P_S)\delta_{i,t}^S$. This will be negative when $p_{i,t}^S \le P_D$ and $(P_D - p_{i,t}^S)\Delta_{i,t}^S < (P_D - P_S)\delta_{i,t}^S$ or when $p_{i,t}^S > P_D$. In order to shift its bid quantities to potentially more profitable securities, the agent reduces its bid quantity using $\delta_{i,t}^S$ in the former case and using $\delta_{i,t}^S + \Delta_{i,t}^S$ in the latter case. With an indicator function $\mathbf{1}_{(\bullet)}$, the adjustment can be written $\delta_{i,t}^S + \Delta_{i,t}^S \mathbf{1}_{(p_{i,t}^S > P_D)}$.

### *5.4.3 Security D*

The objective of an agent bidding for security type D is to sell the shares to short sellers in the secondary market at price $P_D$. After buying $x_{i,t}^D$ securities at a price of $p_{i,t}^D$ in the auction, the agent earns revenue in the secondary market by selling $\tau_{i,t}^D$ securities at the price $P_D$ and selling $\delta_{i,t}^D = x_{i,t}^D - \tau_{i,t}^D$ securities at price $P_C$. Adjustment rules are summarized in Table 5.6 for partial and nil clearing of full allocation (D3 to D4), in Table 5.7 for partial allocation with some clearing (D5 to D10), in Table 5.8 for partial and nil allocation (D11 to D13), and in Table 5.9 for nil allocation (D14 and D15). All of the adjustment rules for security D are listed in Table 5.12. The derivations of the profits and foregone profits are in the Appendix.

#### 5.4.3.1 Full Allocation with Partial/Nil Clearing

When there is full allocation with partial/nil clearing, the foregone profit is never negative, but profit can be negative (see the Appendix).

**Rule D3**: Full allocations of D clear the secondary market at price $P_D$ only if there is enough demand by the short-sellers, i.e., only if there is enough of security S in the secondary market. For full allocation with partial clearing ($\delta_{i,t}^D > 0$) and nonnegative actual profit, the agent lowers its bid quantity in proportion to the clearing shortfall $\delta_{i,t}^D$. Since it has a full allocation that is profitable, it does not change its bid price.

**Rule D4**: When actual profit is negative, the agent lowers its bid quantity in proportion to the clearing shortfall $\delta_{i,t}^D$ and also lowers its bid price in proportion to its loss in order to reduce its probability of an unprofitable allocation. For non-discriminatory payment rules, this will also influence the payment downwards.

### 5.4.3.2 Partial Allocation with Full Clearing

For partial allocation with full clearing, the profit is $(P_D - p_{i,t}^D)x_{i,t}^D$ and the foregone profit is $(P_D - p_{i,t}^D)\Delta_{i,t}^D$ (see the Appendix).

**Rule D5**: For partial allocation of security D with full clearing, non-negative foregone profit, and non-negative profit, the agent tries to increase its allocation by raising the bid price in proportion to the foregone profit. It does not increase the quantity since it is not fully allocated.

**Rule D6**: Profit is negative when $p_{i,t}^D > P_D$, and the agent will want to reduce its probability of receiving an allocation by lowering its bid price in proportion to the loss. It is also desirable to decrease the portion of its total bid quantity that is devoted to the currently unprofitable security D. Since the loss is $(P_D - p_{i,t}^D)x_{i,t}^D$, the agent reduces its bid quantity in proportion to $x_{i,t}^D$.

**Rule D7**: When the foregone profit $(P_D - p_{i,t}^D)\Delta_{i,t}^D$ is negative, the agent has no potential for profit. Thus, it reduces its probability of allocation by lowering the bid price in proportion to the foregone profit and lowers its bid quantity in proportion to its allocation shortfall $\Delta_{i,t}^D$.

### 5.4.3.3 Partial Allocation with Partial Clearing

In the Appendix, I show that for this case the profit is $P_D\tau_{i,t}^D + P_C\delta_{i,t}^D - p_{i,t}^D x_{i,t}^D$ and the foregone profit $P_D(q_{i,t}^D - \tau_{i,t}^D) - p_{i,t}^D\Delta_{i,t}^D$.

**Rule D8**: For partial allocation of security D with partial clearing (at price $P_D$), non-negative foregone profit, and non-negative profit, the agent can try to increase its profit by increasing its probability allocation by raising its bid price in proportion to the foregone profit. Should it increase or decrease its bid quantity? If it increases its bid quantity by 1 and its allocation increases by 1, it will clear at $P_D$ or $P_C$. Since it is currently partially clearing at $P_D$, it is almost certain that this additional unit will clear at $P_C$. Thus, if the agent's allocation increases by 1, the expense will increase by $p_{i,t}^D$ and the revenue will increase by $P_C$. If $P_C > p_{i,t}^D$, this will increase the profit, but if $P_C < p_{i,t}^D$ it will decrease the profit. Similarly, if the agent's allocation decreases by 1, the expense decreases by $p_{i,t}^D$, the revenue decreases by $P_C$, and the profit will increase if $P_C < p_{i,t}^D$. Therefore, the agent should increase its bid quantity by 1 if $P_C > p_{i,t}^D$ and lower its bid quantity by 1 if $P_C < p_{i,t}^D$. This bid quantity adjustment can be written $\mathbf{1}_{(p_{i,t}^D < P_C)} - \mathbf{1}_{(p_{i,t}^D > P_C)}$.

**Rule D9**: When profit is negative ($P_D \tau_{i,t}^D + P_C \delta_{i,t}^D < p_{i,t}^D x_{i,t}^D$), the agent will want to reduce its probability of allocation by lowering its bid price in proportion to the loss. When $P_D \tau_{i,t}^D + P_C \delta_{i,t}^D < p_{i,t}^D x_{i,t}^D$, $P_C < p_{i,t}^D$ so lowering the bid quantity by $x_{i,t}^D$ reduces the loss. However, the major objective of lowering the bid quantity by $x_{i,t}^D$ is to decrease the portion of its total bid quantity that is devoted to this unprofitable security D.

**Rule D10**: When the foregone profit, $P_D(q_{i,t}^D - \tau_{i,t}^D) - p_{i,t}^D \Delta_{i,t}^D$, is negative, $P_D(q_{i,t}^D - \tau_{i,t}^D) < p_{i,t}^D \Delta_{i,t}^D$ and the agent has no potential for profit. Thus, it reduces its probability of allocation by lowering the bid price in proportion to the negative foregone profit and reduces the amount it bids on D by lowering the bid quantity in proportion to $\Delta_{i,t}^D$.

### 5.4.3.4 Partial Allocation with Nil Clearing

For this case, the Appendix shows that profit is $(P_C - p_{i,t}^D)x_{i,t}^D$ and the foregone profit

$P_D q_{i,t}^D - P_C x_{i,t}^D - p_{i,t}^D \Delta_{i,t}^D$.

**Rule D11**: For partial allocation of security D with no clearing (at price $P_D$), non-negative foregone profit, and non-negative profit, the agent increases its probability of allocation by raising its bid price in proportion to the foregone profit. Since $P_C > p_{i,t}^D$ it can also increase its potential profit by raising its bid quantity by 1. If its allocation increases by 1, its profit increases by $P_C - p_{i,t}^D$. .

**Rule D12**: When profit is negative, the agent reduces its probability of allocation by lowering its bid price in proportion to the loss. Negative profit occurs when $P_C < p_{i,t}^D$ so also lowering the bid quantity by $x_{i,t}^D$ reduces the loss. However, the major objective of lowering the bid quantity by $x_{i,t}^D$ is to decrease the portion of its total bid quantity that is devoted to this unprofitable security D.

**Rule D13**: When the foregone profit, $P_D q_{i,t}^D - P_C x_{i,t}^D - p_{i,t}^D \Delta_{i,t}^D$, is negative, $P_D q_{i,t}^D - P_C x_{i,t}^D < p_{i,t}^D \Delta_{i,t}^D$ and the agent has no potential for profit. Thus, it reduces its probability of allocation by lowering the bid price in proportion to the negative foregone profit. With no potential for profit, the agent also reduces the quantity that it bids on security D by lowering its bid quantity in proportion to $\Delta_{i,t}^D$.

### 5.4.3.5 Nil Allocation

In the case of no allocation, the agent's profit for security D is zero and an estimate of the foregone profit is $(0.5P_D + 0.5P_C - p_{i,t}^D)q_{i,t}^D$.

**Rule D14**: When an agent has no allocation for security D ( $x_{i,t}^{D} = 0$ ), if the estimated foregone profit is non-negative, it raises its bid price to increase its probability of an allocation.

**Rule D15**: When the agent has no allocation of D and estimates a negative foregone profit, since it already has no allocation there is no reason to decrease the bid price to lower the probability of allocation. However, a decrease in the bid price may indirectly influence the payment to decrease for non-discriminatory payment rules. Thus, the agent should decrease its bid price in proportion to the foregone loss. It should also decrease the quantity it bids on this security in to $q_{i,t}^{D}$.

### *5.4.4 Security C*

Since security C is set up to be fully cleared in the secondary market, i.e., the allocation $x_{i,t}^{C}$ in the auction clears the secondary market at price $P_C$ , an agent's objective is to be fully allocated in the auction and to make a profit. Table 5.10 shows the adjustment rules for partial allocation (C3 and C4) and nil allocation (C5 and C6). All of the adjustment rules for security C are listed in Table 5.13.

**Rule C3**: In the case of partial allocation ( $0 < x_{i,t}^{C} < q_{i,t}^{C}$ ) and non-negative foregone profit, the agent raises its bid price in proportion to the foregone profit in an attempt to achieve full allocation.

**Rule C4**: In the case of partial allocation with negative foregone profit, the agent lowers its bid price to reduce its probability of allocation, and reduce its bid quantity. . Foregone profit $(P_C - p_{i,t}^{C})\Delta_{i,t}^{C}$ is negative when $p_{i,t}^{C} > P_C$ , and the agent reduces its quantity in proportion to $\Delta_{i,t}^{C}$.

**Rule C5**: When an agent has no allocation for security C and the foregone profit is non-negative, the agent increases its bid price in order to increase its probability of allocation.

**Rule C6**: When the agent has no allocation of C and estimates a negative foregone profit, it reduces its bid quantity. Foregone profit is $(P_C - p_{i,t}^C)q_{i,t}^C$, so when the agent has no profit potential it reduces its bid quantity in proportion to $\Delta_{i,t}^C = q_{i,t}^C$.

## 5.5 Sensitivity, Convergence, and Agent Variation

The simulations use an issue size (in millions of Canadian dollars) $Q$ = 2,500, which is representative of typical Bank of Canada auctions (Hortascu and Sareen, 2005). There is a single auctioneer representing the central bank, and there are twenty bank agent bidders consisting of ten primary dealers and ten regular dealers. In this section, I discuss three important modelling considerations: Are there parameter ranges within which the model results are stable; does the model converge to steady state bidding strategies and steady state market trading; and do the agents vary in their bidding strategies and results?

### *5.5.1 Sensitivity to Parameters*

There are four parameters that must be set: the initial values for bid price and bid quantity and the learning rates for price and quantity. Because of the maximum and minimum bid quantities, the model results are very insensitive to variations to the initial quantities. In the simulations, each agent's initial bid quantity for each of the three securities is an equal portion of its maximum bid limit. For the other three parameters, sensitivity analysis helps to determine the values to use in the simulations. To examine the sensitivity of revenue to the initial bid price-value percent (the initial price expressed as a percentage of the value), bid price learning rate $\phi$, and the bid quantity learning rate $\chi$, I ran simulations for price spreads from $\varepsilon_D = \varepsilon_C = 0.02$ to $\varepsilon_D = \varepsilon_C = 0.20$, five payment rules, initial price-value percents on [95.5%, 1.0%], price learning rate $\phi$ on [1, 10], and quantity learning rate $\chi$ on [0.1, 1.0].

When the price spread is small, e.g., $\varepsilon_D = \varepsilon_C < 0.10$, the revenue results are insensitive to the parameter values. Sensitivity increases as the price spread increases and I therefore present results for the worst case, i.e., $\varepsilon_D = \varepsilon_C = .20$. For purposes of illustration, Figure 5.3 illustrates the revenue topography within the parameter ranges of initial price-value percent $\in [99.0,\ 99.9]$, price learning rate $\phi \in [1.0,\ 5.5]$, and quantity learning rate $\chi \in [0.1,\ 0.95]$. The sensitivity is assessed at a variation of 0.1%, i.e., a variation of \$2.5 million on total revenue of approximately \$2.5 billion. The light areas in the figures represent parameter values for which the revenue variation is less than 0.1%, and the darker areas represent parameter values for which revenue varies from the light area by more than 0.1%.

**Result 1**: The model's revenue results are insensitive across the five payment rules when the parameters are restricted to initial price-value percent $\in [97.5,\ 99.5]$, price learning rate $\phi \in [3.5,\ 9]$, and quantity learning rate $\chi \in [0.1,\ 0.7]$.

*Discussion*: The first observation about the results in Figure 5.3 is that the Uniform and Spanish payment rules are the least sensitive to changes in the parameter values. The second observation is that the model is insensitive within a wide range of parameter values. Specifically, the left column of Figure 5.3 shows that the price learning rate should be in the range $\phi \in [3.5,\ 9]$ and that the initial price-value percent should be in the range [97.5, 99.5]. The middle column shows the same range for the initial price-value percent and that the quantity learning rate should be in the range $\chi \in [0.1,\ 0.7]$. The third column shows that the quantity learning rate should be in the range $\chi \in [0.1,\ 0.7]$ and the price learning rate in the range $\phi \in [3.5,\ 9]$. Thus, to ensure consistent results across the five payment rules, the parameters should be within these ranges. For the simulations in the following sections, I use price-value percent = 98.5, $\phi = 4$, and $\chi = 0.1$.

### *5.5.2 Convergence to Steady State*

Convergence is illustrated in Figure 5.4, which presents charts for the convergence for bid price, bid quantity, allocation, and clearing for the three securities. Note that the lines in the charts that represent the securities S, D, and C are distinct except for the charts for Clearing. In the Clearing charts, the dark gray line is hiding a black line behind it, i.e., the dark gray line represents the amount of security S that must be purchased in the secondary market from sellers of security D. In other words, this is the "squeeze."

**Result 2**: The model results in convergence of bid strategies and market trading. The payment rule with the fastest convergence is the Discriminatory payment rule, followed in order by Spanish, Uniform, Average, and Vickrey.

*Discussion*: The model results in convergence of bid price, bid quantity, auction allocation, and secondary market clearing. Convergence requires about 150 periods for the Discriminatory and Spanish payment rules, about 200 periods for Uniform, and 300 for Average and Vickrey. The agents find it easiest to learn how to bid with the Discriminatory payment rule because the payment is a straightforward pay-your-bid rule, and Spanish is similar because bidders below the average pay their bid. The most difficult to learn payment rules are Vickrey, which depends upon losers' bids, and

Average, which forces bidders to pay more than their bid when they are below the average.

**Result 3**: After converging to steady state, agents are profitable since short-sellers learn to avoid being squeezed in the secondary market.

*Discussion*: In the Clearing charts in Figure 5.4, the dark gray line represents the amount of security S that is being squeezed in the secondary market. All of the agents learn to avoid being squeezed after about 100 periods. Being squeezed is unprofitable for the short sellers and they all learn to adjust their bidding strategies to eventually avoid the loss. The Allocation charts show that the agents all still sell short in the when-issued market (black lines), but they sell only the amount that they have learned they can cover in the auction. When-issued selling is slightly higher for Discriminatory, Average, and Vickrey payment rules than it is for the Uniform and Spanish payment rules. As agents learn to avoid being squeezed, they also learn to avoid buying long in the auction in order to sell in the secondary market. The long-buying agents learn quickly when the source of security S dries up in the secondary market.

### *5.5.3 Endogenous Agent Variation*

Figure 5.5 illustrates the endogenous agent variation in bid price and Figure 5.6 illustrates the endogenous agent variation in bid quantity, auction allocation, and secondary market clearing.

**Result 4**: The bid price variation is greatest for the Uniform payment rule.

*Discussion*: First, compare the bid prices for Discriminatory and Uniform payment rules. Figure 5.5 shows that bid prices are much more variable for the Uniform payment rule than for the Discriminatory. These results are consistent with the empirical results in

Goldreich (2006) who analyzed US data for 105 Discriminatory and 178 Uniform auctions and found that there is a wider dispersion of bid prices for the Uniform payment rule than for Discriminatory. They are also consistent with experimental results in Abbink *et alia* (2006) who find a much higher dispersion of bid prices for Uniform than Discriminatory. Second, Figure 5.5 shows that the bid prices for Spanish auctions are slightly more variable than for Discriminatory, which is also consistent with the slightly higher dispersion for Spanish auctions found by the experiments of Abbink *et alia* (2006).

**Result 5**: Bidding is more dispersed among primary and secondary dealers for Discriminatory and Spanish payment rules than for the other payment rules. When-issued selling varies among the payment rules, but it is always done by the primary dealers.

*Discussion*: The left column of Figure 5.6 shows examples of variation for the bid quantities. The first observation is that the bidding is dominated by the primary dealers (Agents 1 to 10), which is expected because of their minimum bidding constraints. The second observation is that regular bidders (Agents 11 to 20) bid much higher quantities (almost all Security C) for the Discriminatory and Spanish payment rules. Thus, pay-your-bid payments are more conducive to auction participation. Third, for all payment rules, the primary bidders do all of the short-selling in the when-issued market and all of the resulting bidding for security S in the auction. Fourth, primary and regular dealers bid small quantities for security D that they sell to customers in the secondary market for the price of security C.

**Result 6**: The pattern across agents of auction allocation and secondary clearing is similar for Discriminatory and Spanish payments, and these are somewhat different from Uniform and Average, and very different from Vickrey. Most auction allocation is for security S. All secondary clearing is for security D or C at the price of security C.

*Discussion*: In Figure 5.6, the middle column shows the variation among the bidders for the auction allocations and the right column shows the variation for the clearing in the secondary market. First, note that for all of the payment rules, the primary dealers win allocations of security S to completely cover their short sales in the when-issued market. Second, note the variation in allocation and clearing among the payment rules. For Discriminatory and Spanish auctions, all dealers (primary and regular) win small allocations of security C that is cleared in the secondary market. For Uniform auctions, primary dealers win larger allocations of security C and regular dealers win small allocations of security D, both of which are sold in the secondary market at the price of security C. In Average and Vickrey auctions, the dealers win allocations of security D that clear the secondary market at the security C price.

## 5.6 Comparisons Across Market Price Spreads

In this section, I compare the auction revenue results and the activity in the when-issued and secondary markets for the five payment rules across the spectrum of market price spreads, i.e., values of $\varepsilon_D$ and $\varepsilon_C$ ranging from 0.02 to 0.20. The results are illustrated in Figure 5.7 to Figure 5.10, and these results are consistent across repetitions of the simulations.

### *5.6.1 Revenue*

**Result 7**: Discriminatory payment results in revenue that varies the least across the market price spreads, followed by Average, Vickrey, Uniform, and Spanish.

*Discussion*: Figure 5.7 shows the levels of revenue for the 100 combinations of $\varepsilon_D$ and $\varepsilon_C$. Generally, for each payment rule, revenue is lower when the market price spreads are higher. Revenue for Discriminatory payment is quite uniform across the spreads, with gradual monotonic decrease as the spreads increase. The results for Average payment are similar but more irregular. The revenue results for Spanish payment show a steep, regular monotonic decrease as $\varepsilon_C$ increases. For each value of $\varepsilon_D$, the decrease is the same, i.e., it just varies with $\varepsilon_C$. In contrast, for Uniform payment a similar decrease with $\varepsilon_C$ occurs except when $\varepsilon_D$ is small. For $\varepsilon_D \leq 0.10$, the revenue does not decrease with $\varepsilon_C$. Similarly for Vickrey payment, revenue does not decrease for $\varepsilon_D \leq 0.10$, but for higher values of $\varepsilon_D$ the results are irregularly higher or lower.

**Result 8**: The Spanish payment rule is revenue inferior to the Discriminatory across all market price spreads, but the Average rule is revenue superior. For both Uniform and Vickrey payment rules, the revenue comparison with Discriminatory depends on the market price spreads.

*Discussion*: Figure 5.8 shows the differences between the revenue for the four non-discriminatory payment rules and the revenue for the discriminatory payment rule. The Spanish payment rule is revenue inferior to the Discriminatory across all market price spreads, but the Average rule is revenue superior. Of course, the Average rule would probably not be accepted by bidders because they sometimes would pay more than their bid, whereas for all other payment rules they pay their bid or less.

For both Uniform and Vickrey payment rules, the revenue comparison with Discriminatory depends on the market price spreads. The Vickrey payment rule results in more revenue for low values of $\varepsilon_C$ with high values of $\varepsilon_D$ and for low values of $\varepsilon_D$ with high values of $\varepsilon_C$. For most cases, the Uniform payment rule results in less revenue. However, when $\varepsilon_D \leq 0.10$, there are several cases when Uniform payment produces more revenue than Discriminatory. Whether or not these differences are statistically significant is indicated in Table 5.14 by a t-statistic greater than 2.756 ($\alpha = 0.01$). The cases for which the Uniform revenue exceeds the Discriminatory revenue are all statistically significant, since there is little variation across the simulations. Whether these differences, which are \$1.5 million or less, are practically significant would be up to the central bank. Daripa (2001) develops a model that predicts that the Discriminatory and Uniform payment rules will result in the same revenue. I ran simulations without the when-issued and secondary trading, and indeed the revenue for all payment rules is equivalent. However, as seen in Figure 5.8, this result does not hold when the when-issued and secondary markets are included. The experiments of Abbink *et alia* (2006) result in higher revenue for the Spanish and Uniform auctions than the Discriminatory, but again these experiments do not include when-issued and secondary markets.

### *5.6.2 When-Issued and Secondary Markets*

It is interesting to examine the corresponding quantities that are traded in the when-issued and secondary markets. Figure 5.9 shows the volume of short selling in the when-issued market for the 100 combinations of $\varepsilon_D$ and $\varepsilon_C$. When-issued trading is consistently high for the Discriminatory and Average payment rules across all of the market price spreads, although it tends to decrease for higher values of $\varepsilon_D$. When-issued trading is also high for Uniform payment when $\varepsilon_D < 0.1$ and for Vickrey payment when both $\varepsilon_D < 0.1$ and $\varepsilon_C < 0.1$. For the Spanish payment rule, when-issued trading is lower for all market-price spreads, diminishing substantially as $\varepsilon_D$ increases.

Why is a smaller amount of when-issued trading with higher $\varepsilon_D$ a consistent result across the payment rules? With narrow price spreads, especially when $\varepsilon_D$ is small, it is more likely that a dealer bidding for S securities can outbid the price for D securities. This makes a squeeze less likely and thus makes it likely that when-issued selling will be higher. With larger price spreads, especially when $\varepsilon_D$ is large, bid prices for security D are more likely to exceed those for security S so that long-buyers are more likely to an allocation and short-sellers are not. This makes a squeeze more likely for short-sellers and causes agents to learn to restrict their short-selling in the when-issued market.

Figure 5.10 shows the volume of selling in the secondary market. There is no secondary market trading in security S and D at price $P_D$, and all secondary market trading is for securities C and D at price $P_C$. Spanish payment, which had the lowest when-issued trading, has the consistently highest secondary market trading across the market price spreads, with some diminution occurring when $\varepsilon_D < 0.06$. Discriminatory and Average payment rules have consistently low levels of secondary market trading, while Uniform and Vickery auctions have unevenly high trading volumes when $\varepsilon_D > 0.06$.

## 5.7 Conclusion

Using a multiagent system with learning agents has produced a simulation tool that produces consistent results that are stable across a range of initial values and learning rates. This provides a method to analyze the bidding behaviour and revenue results for a variety of payment rules across a full spectrum of market price spreads. Most previous studies have focused on possible revenue improvements from adopting Uniform payment over Discriminatory payment, and results have been inconclusive (Swierzbinski and Borgers, 2004). Taking into account the effects of a when-issued and secondary market on the auction bidding strategies, this study shows that for most market-price spreads Discriminatory payment used by the Bank of Canada results in higher revenue than Uniform payment used by the United States. Similarly, Vickrey payment is revenue inferior to Discriminatory payment for most market-price spreads. Spanish payment is revenue inferior to Discriminatory payment across all market-price spreads, but Average payment is revenue superior across all market-price spreads.

The agents in this model are all able to avoid unprofitable short-selling and long-buying by learning through repetitions of the same environment. In future work it will be interesting to gradually introduce extensions such as varying the issue size and market price spreads with *t*, creating a secondary market that dynamically extends across time periods, making the market prices endogenous, and introducing constraints that correspond to the U.S. rules.

## 5.8 References


Abbink, K., Brandts, J. , Pezanis-Christou, P., 2006, "Auctions for government securities: a laboratory comparison of Uniform, Discriminatory, Spanish designs," *Journal of Economic Behavior and Organization*, 61: 284-303 .

Arnone, M., Iden, G., 2003, "Primary dealers in government securities: policy issues and selected countries' experiences," IMF Working Paper 03/45, 64 pages.

Ausubel, L., 2004, "An efficient ascending-bid auction for multiple objects," *American Economic Review*, 94(5), 1452-1475),

Ausubel, L., 2006, "An efficient dynamic auction for heterogeneous commodities," *American Economic Review,* 96(3): 602-629.

Bank of Canada, 1998, "Revised rules pertaining to auctions of Government of Canada securities and the Bank of Canada's surveillance of the auction process." Final Report, 27 pages.

Bank of Canada, 2005, "Terms of participation in auctions for government securities distributors," Effective 13 December 2005.

Bartolini, L., Cottarelli, C., 1997. "Designing effective auctions for treasury securities." *Current Issues in Economics and Finance 3*, Federal Reserve Bank of New York, pp. 1–6.

Belisle, C., 2007, Bank of Canada, Personal Communication.

Bikhchandani, S., Edsparr, P., Huang, C., 2000, "The treasury bill auction and the when issued market: some evidence," *UCLA Anderson Graduate School of Management Working Paper*, 17 pages.

Chatterjea, A., Jarrow, R. A. , 1998, "Market manipulation, price bubbles, a model of the U.S. Treasury securities auction market," *Journal of Financial and Quantitative Analysis*, 33(2), 255-289.

Daripa, A., 2001, "A theory of treasury auctions," *Journal of International Money and Finance*, 20, 743-767.

D'Souza, C., Gaa, C., 2004, "The effect of economic news on bond market liquidity," Bank of Canada Working Paper 2004-16, 22 pages.

Goldreich, D., 2007, "Underpayment in Discriminatory and Uniform-price treasury auctions," *Journal of Financial and Quantitative Analysis*, 42(2): 443-466.

Government of Canada, Department of Finance, 2005, "Debt management strategy 2005-2006," 33 pages.

Government of Canada, Department of Finance, 2006, "Debt management strategy 2006-2007," 35 pages.

Government of Canada, Department of Finance, 2007, "Debt management strategy 2007-2008," 18 pages.

Hortacsu, A., Sareen, S., 2005, "Order flow and the formation of dealer bids: information flows and strategic behavior in the Government of Canada securities auctions," NBER Working Paper, 11116, 51 pages.

Lu, D., Yang, J., 2003, "Auction participation and market uncertainty: evidence from Canadian treasury auctions," Conference paper presented at the Canadian Economics Association 37th Annual Meeting, 21 pages.

Malvey, P. F., Archibald, C. M., 1998, "Uniform-price auctions: update of the Treasury experience," U.S. Department of the Treasury, 19 pages.

Merrick, J. J., Naik, N.Y. , Yadav, P.K., 2005, "Strategic trading behavior and price distortion in a manipulated market: anatomy of a squeeze," *Journal of Financial Economics*, 77: 171-218.

Nyborg, K.G., Strebulaev, I.A., 2004, "Multiple unit auctions and short squeezes," *Journal of Financial Studies*, 17(2), 545-480.

Nyborg, K.G., Sundaresan, S., 1996, "Discriminatory versus Uniform treasury auctions: evidence from when-issued transactions," *Journal of Financial Economics*, 42, 63-104.

Sareen, S., 2006, "Commodity bundling in government securities auctions," Duke University, Working Paper.

Swierzbinski, J., Borgers, T., 2004, "The design of Treasury bond auctions: some case studies," In: Janssen, M. C. W. (Ed), *Auctioning Public Assets: Analysis and Alternatives*, Cambridge University Press, 230-256.

U.S. Department of Treasury: http://www.treasurydirect.gov/instit/research/research.htm

Viswanathan, S., J.D. Wang, 2004, "Inter-dealer trading in financial markets," *Journal of Business*, 77(4): 987-1040.

Wang, J. D., Zender, J. F., 2002, "Auctioning divisible goods," *Economic Theory*, 19: 673-705.

## 5.9 Tables

| Table 5.1. Treasury Auction Models | | | | | | |
|---|---|---|---|---|---|---|
| | Wang & Zender (2002) | Daripa (2001) | Viswanathan & Wang (2004) | Chatterjea & Jarrow (1998) | Nyborg & Strebulaev (2004) | This Model |
| Number of when-issued sellers | – | – | – | 2 | 1 | n |
| Auction | + | + | + | + | + | + |
| Secondary | – | – | + | + | + | + |
| Multi-Unit | – | – | – | – | + | + |
| Primary dealer constraints | – | – | – | – | – | + |
| Number of auction winners | n | n | n | 2 | 1 | n |
| Payment Rules | 2 | 2 | 2 | 2 | 2 | 5 |

Table 5.2. Treasury Model Notation Summary

| | |
|---|---|
| $\alpha_{i,t}^{SO}$ | Share for dealer i of the overall allocation at the stopout price. |
| $b_{i,t}^{j}$ | Price bid in auction for security j |
| $\delta_{i,t}^{S}$, $\delta_{i,t}^{D}$ | Quantity not cleared in the secondary market: $\delta_{i,t}^{S} = (\Delta_{i,t}^{S} - \tau_{i,t}^{S})$, $\delta_{i,t}^{D} = (x_{i,t}^{D} - \tau_{i,t}^{D})$ |
| $\varepsilon_{D}$, $\varepsilon_{C}$ | Difference between common values for security D and C from security S. |
| $p_{t}^{cut}$, $p_{t}^{avg}$ | Auction cutoff price and average price. |
| $CO_t$ | Percentage of issue quantity that is customer orders. |
| $\Delta_{i,t}^{S}$, $\Delta_{i,t}^{D}$ | Quantity not allocated in the auction, i.e., difference between the bid quantity and the allocation. |
| i | Dealer i. |
| n | Total number of dealers |
| $n_P$, $n_R$ | Number of primary and regular dealers respectively. |
| $p_{i,t}^{j}$ | Payment in auction for security j |
| $P_j$ | Market price for security j. |
| $\pi_{i,t}^{j}$ | Profit at end of period t across the markets for security j. |
| $\pi_{i,t}^{F,j}$ | Foregone profit at end of period t across the markets for security j. |
| $\phi$ | Bid price learning rate. |
| $\chi$ | Bid quantity learning rate. |
| $q_{i,t}^{S}$ | Quantity sold by short sellers in the when-issued market. |
| $q_{i,t}^{j}$ | Quantity bid in auction for security j. |
| $Q_t$, $Q_t^j$ | Total auction issue quantity, and total quantity traded for security j. |
| $\overline{q}_i$, $\underline{q}_i$ | Maximum and minimum total auction bid quantity allowed for dealer i. |
| $q_{i,t}^{CO}$ | Customer order quantity for dealer i. |
| t | Time period t $\in$ {0, 1, …, T}, where T is very large. |
| $\tau_{i,t}^{j}$, $\tau_{i,t}^{S}$ | Quantities that clear the secondary market, |
| $\overline{\theta}_i$, $\underline{\theta}_i$ | Maximum and minimum allowed fraction of issue for which a dealer can bid. |
| $v^j$ | Values for security $j \in \{S, D, C\}$, where $v_j = P_j$ |
| $x_{i,t}^{j}$ | Quantity allocated in auction for security j. |

Table 5.3. Bids from a Bank of Canada Auction
(Source: Lu and Yang, 2003)

| Bidder | Yield (%) | Amount ($ million) | Awarded ($ million) | Bidder | Yield (%) | Amount ($ million) | Awarded ($ million) |
|---|---|---|---|---|---|---|---|
| FKA | 5.148 | 25 | 25 | LJB | 5.205 | 100 | |
| FPB | 5.165 | 50 | 50 | RWE | 5.208 | 25 | |
| KVI | 5.168 | 100 | 100 | AXJ | 5.21 | 25 | |
| LZE | 5.17 | 50 | 50 | YPQ | 5.21 | 200 | |
| AXG | 5.18 | 31.5 | 31.5 | RWE | 5.213 | 100 | |
| AXJ | 5.18 | 50 | 50 | YPQ | 5.22 | 200 | |
| KVI | 5.18 | 10 | 10 | AXG | 5.23 | 100 | |
| KVI | 5.184 | 525 | 525 | TOA | 5.236 | 250 | |
| FPB | 5.188 | 75 | 75 | LZE | 5.248 | 100 | |
| AXJ | 5.19 | 50 | 50 | RWE | 5.248 | 200 | |
| LJB | 5.19 | 150 | 150 | FKA | 5.25 | 400 | |
| LJB | 5.19 | 300 | 300 | FPB | 5.25 | 200 | |
| LZE | 5.19 | 100 | 100 | JQS | 5.25 | 200 | |
| TOA | 5.198 | 250 | 250 | RFB | 5.3 | 100 | |
| FPB | 5.199 | 75 | 75 | | | | |
| AXG | 5.2 | 100 | 91.042 | | | | |
| AXJ | 5.2 | 25 | 22.76 | | | | |
| LJB | 5.2 | 20 | 18.208 | | | | |
| LJB | 5.2 | 100 | 91.042 | | | | |
| LJB | 5.2 | 250 | 227.604 | | | | |
| YPQ | 5.2 | 225 | 204.844 | | | | |

Table 5.4. Adjustment rules for S, D, C:
Full Allocation with Full Clearing (D)

| $\pi_{i,t}^{j} \geq 0$ | $\pi_{i,t}^{j} < 0$ |
|---|---|
| $\uparrow$ q(1) | $\downarrow$ b( $\pi$ ), $\downarrow$ q(q) |
| j1 | j2 |

Table 5.5. Adjustment rules for S: Partial/Nil Allocation

| Full Clearing | | | Partial/Nil Clearing | |
|---|---|---|---|---|
| $\pi_{i,t}^{F,S} \geq 0$ | | $\pi_{i,t}^{F,S} < 0$ | $\pi_{i,t}^{F,S} \geq 0$, | $\pi_{i,t}^{F,S} < 0$, |
| $\pi_{i,t}^{S} \geq 0$ | $\pi_{i,t}^{S} < 0$ | | | |
| $\uparrow$ b( $\pi^F$ ) | $\uparrow$ b( $\pi^F$ ), $\downarrow$ q(1) | $\downarrow$ b( $\pi^F$ ), $\downarrow$ q(Δ) | $\uparrow$ b( $\pi^F$ ), $\downarrow$ q(δ) | $\downarrow$ b( $\pi^F$ ), $\downarrow$ q(δ,Δ) |
| S3 | S4 | S5 | S6 | S7 |

Table 5.6. Adjustment rules for D: Full Allocation with Partial/Nil Clearing

| $\pi_{i,t}^{D} \geq 0$ | $\pi_{i,t}^{D} < 0$ |
|---|---|
| $\downarrow$q(δ) | $\downarrow$b( $\pi$ ), $\downarrow$q(δ) |
| D3 | D4 |

| Table 5.7. Adjustment rules for D: Partial Allocation, Some Clearing | | | | | |
|---|---|---|---|---|---|
| Full Clearing | | | Partial Clearing | | |
| $\pi_{i,t}^{F,D} \geq 0$ | | $\pi_{i,t}^{F,D} < 0$ | $\pi_{i,t}^{F,D} \geq 0$ | | $\pi_{i,t}^{F,D} < 0$ |
| $\pi_{i,t}^{D} \geq 0$ | $\pi_{i,t}^{D} < 0$ | | $\pi_{i,t}^{D} \geq 0$ | $\pi_{i,t}^{D} < 0$ | |
| ↑b($\pi^F$) | ↓b($\pi$), ↓q(x) | ↓b($\pi^F$), ↓q(Δ) | ↑b($\pi^F$),↑↓q(1) | ↓b($\pi$), ↓q(x) | ↓b($\pi^F$), ↓q(Δ) |
| D5 | D6 | D7 | D8 | D9 | D10 |

| Table 5.8. Adjustment rules for D: Partial Allocation, Nil Clearing | | |
|---|---|---|
| $\pi_{i,t}^{F,D} \geq 0$ | | $\pi_{i,t}^{F,D} < 0$ |
| $\pi_{i,t}^{D} \geq 0$ | $\pi_{i,t}^{D} < 0$ | |
| ↑b($\pi^F$), ↑q(1) | ↓b($\pi$), ↓q(x) | ↓b($\pi^F$), ↓q(Δ) |
| D11 | D12 | D13 |

| Table 5.9. Adjustment rules for D: Nil Allocation | |
|---|---|
| $\pi_{i,t}^{F,D} \geq 0$ | $\pi_{i,t}^{F,D} < 0$ |
| ↑b($\pi^F$) | ↓b($\pi^F$), ↓q(q) |
| D14 | D15 |

| Table 5.10 Adjustment rules for Security C: Partial/Nil Allocation | | | |
|---|---|---|---|
| Partial Allocation | | Nil Allocation | |
| $\pi_{i,t}^{F,C} \geq 0$ | $\pi_{i,t}^{F,C} < 0$ | $\pi_{i,t}^{F,C} \geq 0$ | $\pi_{i,t}^{F,C} < 0$ |
| $\uparrow \mathrm{b}(\pi^F)$ | $\downarrow \mathrm{b}(\pi^F), \downarrow \mathrm{q}\{\Delta\}$ | $\uparrow \mathrm{b}(\pi^F)$ | $\downarrow \mathrm{q}\{\Delta\}$ |
| C3 | C4 | C5 | C6 |

| Table 5.11. List of adjustments: Security S | | |
|---|---|---|
| Rule | Price Adjustment | Quantity Adjustment |
| S1 | $a_{i,t}^{S}=0$ | $z_{i,t}^{S}=1$ |
| S2 | $a_{i,t}^{S}=\dfrac{\pi_{i,t}^{S}}{P_S q_{i,t}^{S}}$ | $z_{i,t}^{S}=-q_{i,t}^{S}$ |
| S3 | $a_{i,t}^{S}=\dfrac{\pi_{i,t}^{F,S}}{P_S q_{i,t}^{S}}$ | $z_{i,t}^{S}=0$ |
| S4 | $a_{i,t}^{S}=\dfrac{\pi_{i,t}^{F,S}}{P_S q_{i,t}^{S}}$ | $z_{i,t}^{S}=-1$ |
| S5 | $a_{i,t}^{S}=\dfrac{\pi_{i,t}^{F,S}}{P_S q_{i,t}^{S}}$ | $z_{i,t}^{S}=-\Delta_{i,t}^{S}$ |
| S6 | $a_{i,t}^{S}=\dfrac{\pi_{i,t}^{F,S}}{P_S q_{i,t}^{S}}$ | $z_{i,t}^{S}=-\delta_{i,t}^{S}$ |
| S7 | $a_{i,t}^{S}=\dfrac{\pi_{i,t}^{F,S}}{P_S q_{i,t}^{S}}$ | $z_{i,t}^{S}=-(\delta_{i,t}^{S}+\Delta_{i,t}^{S}\mathbf{1}_{(p_{i,t}^{S}>P_D)})$ |

| Table 5.12. List of adjustments: Security D | | | | | |
|---|---|---|---|---|---|
| Rule | Price Adjustment | Quantity Adjustment | Rule | Price Adjustment | Quantity Adjustment |
| D1 | $a_{i,t}^{D}=0$ | $z_{i,t}^{D}=1$. | D9 | $a_{i,t}^{D}=\frac{\pi_{i,t}^{D}}{P_D q_{i,t}^{D}}$ | $z_{i,t}^{D}=-q_{i,t}^{D}$. |
| D2 | $a_{i,t}^{D}=\frac{\pi_{i,t}^{D}}{P_D q_{i,t}^{D}}$ | $z_{i,t}^{S}=-q_{i,t}^{D}$. | D10 | $a_{i,t}^{D}=\frac{\pi_{i,t}^{F,D}}{P_D q_{i,t}^{D}}$ | $z_{i,t}^{D}=-\Delta_{i,t}^{D}$ |
| D3 | $a_{i,t}^{D}=0$ | $z_{i,t}^{D}=-\delta_{i,t}^{D}$. | D11 | $a_{i,t}^{D}=\frac{\pi_{i,t}^{F,D}}{P_D q_{i,t}^{D}}$ | $z_{i,t}^{D}=1$. |
| D4 | $a_{i,t}^{D}=\frac{\pi_{i,t}^{D}}{P_D q_{i,t}^{D}}$ | $z_{i,t}^{D}=-\delta_{i,t}^{D}$. | D12 | $a_{i,t}^{D}=\frac{\pi_{i,t}^{D}}{P_D q_{i,t}^{D}}$ | $z_{i,t}^{D}=-x_{i,t}^{D}$. |
| D5 | $a_{i,t}^{D}=\frac{\pi_{i,t}^{F,D}}{P_D q_{i,t}^{D}}$ | $z_{i,t}^{D}=0$. | D13 | $a_{i,t}^{D}=\frac{\pi_{i,t}^{F,D}}{P_D q_{i,t}^{D}}$ | $z_{i,t}^{D}=-\Delta_{i,t}^{D}$ |
| D6 | $a_{i,t}^{D}=\frac{\pi_{i,t}^{D}}{P_D q_{i,t}^{D}}$ | $z_{i,t}^{D}=-x_{i,t}^{D}$. | D14 | $a_{i,t}^{D}=\frac{\pi_{i,t}^{F,D}}{P_D q_{i,t}^{D}}$ | $z_{i,t}^{D}=0$ |
| D7 | $a_{i,t}^{D}=\frac{\pi_{i,t}^{F,D}}{P_D q_{i,t}^{D}}$ | $z_{i,t}^{D}=-\Delta_{i,t}^{D}$. | D15 | $a_{i,t}^{D}=\frac{\pi_{i,t}^{F,D}}{P_D q_{i,t}^{D}}$ | $z_{i,t}^{D}=-q_{i,t}^{D}$ |
| D8 | $a_{i,t}^{D}=\frac{\pi_{i,t}^{F,D}}{P_D q_{i,t}^{D}}$ | $z_{i,t}^{D}=\mathbf{1}_{(p_{i,t}^{D}<P_C)}-\mathbf{1}_{(p_{i,t}^{D}>P_C)}$. | | | |

| Table 5.13. List of adjustment:s Security C | | |
|---|---|---|
| Rule | Price Adjustment | Quantity Adjustment |
| C1 | $a_{i,t}^{C} = 0$ | $z_{i,t}^{C} = 1$. |
| C2 | $a_{i,t}^{C} = \frac{\pi_{i,t}^{C}}{P_C q_{i,t}^{C}}$ | $z_{i,t}^{C} = -q_{i,t}^{C}$. |
| C3 | $a_{i,t}^{C} = \frac{\pi_{i,t}^{F,C}}{P_C q_{i,t}^{C}}$ | $z_{i,t}^{C} = 0$. |
| C4 | $a_{i,t}^{C} = \frac{\pi_{i,t}^{F,C}}{P_C q_{i,t}^{C}}$ | $z_{i,t}^{C} = -\Delta_{i,t}^{C}$. |
| C5 | $a_{i,t}^{C} = \frac{\pi_{i,t}^{F,C}}{P_C q_{i,t}^{C}}$ | $z_{i,t}^{C} = 0$. |
| C6 | $a_{i,t}^{C} = 0$ | $z_{i,t}^{C} = -\Delta_{i,t}^{C}$. |

| Table 5.14. Revenue Differences from Discriminatory (n=30) | | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | Mean Differences | | | | t-statistic | | | |
| $\varepsilon_C$ $\varepsilon_D$ | Average | Spanish | Uniform | Vickrey | Average | Spanish | Uniform | Vickrey |
| .02.02 | 0.159 | -0.317 | -0.338 | 0.002 | 185 | -753 | -1168 | 0.01 |
| .02.04 | 0.134 | -0.318 | -0.335 | -0.106 | 14 | -961 | -1472 | -0.47 |
| .02.06 | 0.112 | -0.320 | -0.334 | -62.971 | 8 | -973 | -1232 | -1.02 |
| .02.08 | 0.090 | -0.315 | -0.332 | -60.703 | 5 | -853 | -1052 | -0.98 |
| .02.10 | 0.058 | -0.316 | -0.332 | 1.495 | 2 | -1154 | -1536 | 5.14 |
| .02.12 | 0.107 | -0.318 | -0.333 | 1.772 | 4 | -1054 | -1390 | 5.63 |
| .02.14 | 0.033 | -0.318 | -0.333 | 1.746 | 1 | -769 | -1056 | 13.70 |
| .02.16 | -0.008 | -0.319 | -0.333 | 1.864 | 0 | -1140 | -1513 | 228.41 |
| .02.18 | -0.019 | -0.320 | -0.333 | 2.124 | 0 | -917 | -1472 | 231.08 |
| .02.20 | -0.153 | -0.321 | -0.333 | 2.359 | -4 | -874 | -1005 | 231.60 |
| .04.02 | 0.292 | -0.642 | 0.272 | 0.285 | 237 | -596 | 57 | 97.22 |
| .04.04 | 0.317 | -0.646 | -0.678 | -0.312 | 158 | -923 | -1211 | -0.94 |
| .04.06 | 0.296 | -0.644 | -0.673 | -0.391 | 18 | -813 | -1015 | -1.57 |
| .04.08 | 0.271 | -0.634 | -0.667 | -36.950 | 15 | -798 | -1044 | -0.98 |
| .04.10 | 0.265 | -0.633 | -0.665 | -0.264 | 14 | -886 | -1155 | -0.79 |
| .04.12 | 0.254 | -0.634 | -0.666 | -0.377 | 10 | -1018 | -1378 | -1.36 |
| .04.14 | 0.167 | -0.637 | -0.666 | 0.510 | 4 | -948 | -1213 | 2.47 |
| .04.16 | 0.192 | -0.639 | -0.667 | 1.467 | 5 | -1054 | -1432 | 9.43 |
| .04.18 | 0.280 | -0.639 | -0.667 | 1.838 | 9 | -1190 | -1565 | 14.74 |
| .04.20 | 0.207 | -0.641 | -0.668 | 2.260 | 6 | -972 | -1242 | 207.30 |
| .06.02 | 0.404 | -0.947 | 0.389 | 0.396 | 229 | -443 | 78 | 114.18 |
| .06.04 | 0.464 | -0.979 | 0.372 | 0.395 | 261 | -976 | 5 | 7.97 |
| .06.06 | 0.478 | -0.970 | -1.014 | -1.014 | 329 | -1040 | -1386 | -1385.80 |
| .06.08 | 0.477 | -0.957 | -1.004 | -0.321 | 78 | -910 | -1109 | -0.75 |
| .06.10 | 0.446 | -0.954 | -1.001 | 0.194 | 26 | -853 | -1082 | 0.46 |
| .06.12 | 0.496 | -0.951 | -1.000 | -0.634 | 53 | -1052 | -1349 | -2.44 |
| .06.14 | 0.421 | -0.950 | -0.998 | -62.951 | 15 | -937 | -1261 | -1.02 |
| .06.16 | 0.413 | -0.953 | -0.999 | -10.548 | 12 | -1146 | -1470 | -1.13 |
| .06.18 | 0.373 | -0.954 | -0.999 | -1.170 | 11 | -719 | -792 | -6.24 |
| .06.20 | 0.320 | -0.958 | -1.001 | -0.335 | 8 | -1055 | -1323 | -1.26 |

| Table 5.14. Differences from Discriminatory (continued) | | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | Mean Differences | | | | t-statistic | | | |
| $\varepsilon_C$ $\varepsilon_D$ | Average | Spanish | Uniform | Vickrey | Average | Spanish | Uniform | Vickrey |
| .08.02 | 0.496 | -1.237 | 0.475 | 0.467 | 212 | -310 | 76 | 77.78 |
| .08.04 | 0.603 | -1.301 | 0.573 | 0.584 | 218 | -810 | 61 | 92.32 |
| .08.06 | 0.622 | -1.303 | -1.364 | 0.616 | 223 | -899 | -1154 | 108.86 |
| .08.08 | 0.650 | -1.285 | -1.346 | -0.860 | 248 | -790 | -974 | -3.55 |
| .08.10 | 0.646 | -1.277 | -1.340 | -6.630 | 85 | -958 | -1227 | -1.04 |
| .08.12 | 0.600 | -1.274 | -1.336 | -0.954 | 25 | -881 | -1112 | -6.45 |
| .08.14 | 0.615 | -1.272 | -1.337 | -17.685 | 27 | -877 | -1209 | -1.00 |
| .08.16 | 0.612 | -1.270 | -1.334 | -21.354 | 34 | -1054 | -1269 | -1.04 |
| .08.18 | 0.496 | -1.271 | -1.334 | -1.251 | 13 | -978 | -1239 | -15.52 |
| .08.20 | 0.560 | -1.273 | -1.336 | -1.076 | 15 | -826 | -1052 | -4.78 |
| .10.02 | 0.577 | -1.511 | 0.536 | 0.539 | 208 | -271 | 68 | 68.75 |
| .10.04 | 0.722 | -1.626 | 0.706 | 0.704 | 235 | -602 | 89 | 84.10 |
| .10.06 | 0.771 | -1.634 | 0.701 | 0.739 | 190 | -785 | 48 | 26.29 |
| .10.08 | 0.800 | -1.613 | -1.689 | -1.308 | 180 | -830 | -1003 | -7.43 |
| .10.10 | 0.807 | -1.603 | -1.681 | -24.770 | 141 | -839 | -1053 | -1.12 |
| .10.12 | 0.799 | -1.599 | -1.677 | -0.728 | 109 | -1034 | -1211 | -1.68 |
| .10.14 | 0.755 | -1.592 | -1.673 | -1.419 | 37 | -801 | -977 | -11.62 |
| .10.16 | 0.735 | -1.591 | -1.674 | -1.354 | 24 | -1009 | -1287 | -8.22 |
| .10.18 | 0.767 | -1.591 | -1.671 | -1.616 | 28 | -957 | -1242 | -9.35 |
| .10.20 | 0.621 | -1.591 | -1.669 | -1.700 | 15 | -1059 | -1329 | -25.68 |
| .12.02 | 0.650 | -1.778 | 0.616 | 0.603 | 233 | -225 | 80 | 71.44 |
| .12.04 | 0.841 | -1.950 | 0.798 | 0.826 | 238 | -624 | 75 | 115.57 |
| .12.06 | 0.914 | -1.953 | 0.876 | 0.889 | 250 | -708 | 62 | 80.75 |
| .12.08 | 0.938 | -1.954 | -1.458 | 0.811 | 160 | -958 | -7 | 8.19 |
| .12.10 | 0.958 | -1.931 | -2.022 | -1.471 | 140 | -980 | -1251 | -4.43 |
| .12.12 | 0.961 | -1.923 | -2.016 | -22.797 | 102 | -724 | -834 | -1.22 |
| .12.14 | 0.952 | -1.911 | -2.009 | -20.183 | 102 | -860 | -1064 | -1.05 |
| .12.16 | 0.975 | -1.907 | -2.006 | -1.194 | 77 | -874 | -1111 | -3.69 |
| .12.18 | 0.945 | -1.908 | -2.007 | -1.539 | 55 | -923 | -1219 | -5.00 |
| .12.20 | 0.928 | -1.911 | -2.009 | -14.366 | 40 | -819 | -997 | -1.12 |

| Table 5.14. Differences from Discriminatory (continued) | | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | Mean Differences | | | | t-statistic | | | |
| $\varepsilon_C$ $\varepsilon_D$ | Average | Spanish | Uniform | Vickrey | Average | Spanish | Uniform | Vickrey |
| .14.02 | 0.709 | -1.990 | 0.643 | 0.683 | 228 | -198 | 50 | 77.93 |
| .14.04 | 0.942 | -2.258 | 0.932 | 0.932 | 257 | -557 | 94 | 121.37 |
| .14.06 | 1.037 | -2.285 | 0.995 | 1.024 | 212 | -659 | 99 | 134.85 |
| .14.08 | 1.078 | -2.281 | 1.004 | 0.787 | 184 | -897 | 62 | 3.82 |
| .14.10 | 1.110 | -2.266 | -2.369 | -2.366 | 183 | -994 | -1283 | -609.23 |
| .14.12 | 1.107 | -2.257 | -2.364 | -1.469 | 144 | -879 | -1173 | -3.81 |
| .14.14 | 1.107 | -2.241 | -2.352 | -97.697 | 103 | -1014 | -1225 | -1.48 |
| .14.16 | 1.120 | -2.235 | -2.349 | -1.143 | 85 | -778 | -917 | -3.32 |
| .14.18 | 1.119 | -2.223 | -2.340 | -1.301 | 70 | -893 | -1130 | -3.80 |
| .14.20 | 1.120 | -2.229 | -2.346 | -1.731 | 81 | -783 | -945 | -4.72 |
| .16.02 | 0.770 | -2.174 | 0.675 | 0.712 | 173 | -169 | 40 | 51.77 |
| .16.04 | 1.049 | -2.554 | 0.978 | 1.009 | 228 | -474 | 60 | 75.29 |
| .16.06 | 1.163 | -2.616 | 1.115 | 1.145 | 231 | -708 | 97 | 108.46 |
| .16.08 | 1.215 | -2.614 | 1.148 | 1.190 | 169 | -717 | 78 | 74.59 |
| .16.10 | 1.237 | -2.605 | -1.973 | -1.361 | 111 | -778 | -7 | -3.98 |
| .16.12 | 1.265 | -2.581 | -2.704 | -2.179 | 110 | -784 | -903 | -6.89 |
| .16.14 | 1.260 | -2.570 | -2.696 | -1.708 | 94 | -922 | -1171 | -4.87 |
| .16.16 | 1.261 | -2.562 | -2.690 | -129.145 | 109 | -1010 | -1518 | -1.47 |
| .16.18 | 1.281 | -2.556 | -2.686 | -63.681 | 107 | -884 | -1092 | -1.04 |
| .16.20 | 1.313 | -2.546 | -2.679 | -64.562 | 101 | -839 | -1110 | -1.03 |
| .18.02 | 0.829 | -2.379 | 0.736 | 0.749 | 206 | -143 | 43 | 52.11 |
| .18.04 | 1.145 | -2.840 | 1.110 | 1.087 | 250 | -424 | 77 | 83.62 |
| .18.06 | 1.274 | -2.934 | 1.238 | 1.256 | 243 | -642 | 78 | 110.75 |
| .18.08 | 1.358 | -2.933 | 1.306 | 1.306 | 151 | -700 | 79 | 104.01 |
| .18.10 | 1.372 | -2.951 | 1.331 | 1.028 | 149 | -879 | 47 | 4.91 |
| .18.12 | 1.397 | -2.915 | -3.048 | -14.645 | 91 | -936 | -1221 | -1.22 |
| .18.14 | 1.431 | -2.903 | -3.040 | -1.515 | 142 | -930 | -1140 | -2.82 |
| .18.16 | 1.403 | -2.884 | -3.026 | -1.774 | 122 | -798 | -1013 | -4.85 |
| .18.18 | 1.434 | -2.876 | -3.023 | -63.933 | 90 | -820 | -1154 | -1.03 |
| .18.20 | 1.456 | -2.877 | -3.024 | -1.763 | 101 | -813 | -1064 | -3.92 |
| .20.02 | 0.880 | -2.628 | 0.765 | 0.761 | 168 | -138 | 35 | 38.12 |
| .20.04 | 1.229 | -3.149 | 1.183 | 1.174 | 206 | -405 | 88 | 71.53 |
| .20.06 | 1.383 | -3.257 | 1.323 | 1.385 | 234 | -808 | 87 | 125.05 |
| .20.08 | 1.469 | -3.261 | 1.455 | 1.489 | 170 | -701 | 86 | 149.13 |
| .20.10 | 1.518 | -3.290 | 1.497 | 1.506 | 126 | -909 | 65 | 80.41 |
| .20.12 | 1.571 | -3.261 | -1.774 | -2.218 | 118 | -699 | -4 | -5.28 |
| .20.14 | 1.560 | -3.226 | -3.381 | -2.659 | 95 | -912 | -1228 | -8.65 |
| .20.16 | 1.567 | -3.214 | -3.377 | -63.976 | 113 | -733 | -910 | -1.04 |
| .20.18 | 1.564 | -3.206 | -3.369 | -1.049 | 106 | -970 | -1392 | -1.84 |
| .20.20 | 1.588 | -3.193 | -3.363 | -9.670 | 91 | -728 | -944 | -1.27 |

## 5.10 Figures

Figure 5.1. Market Model Concepts

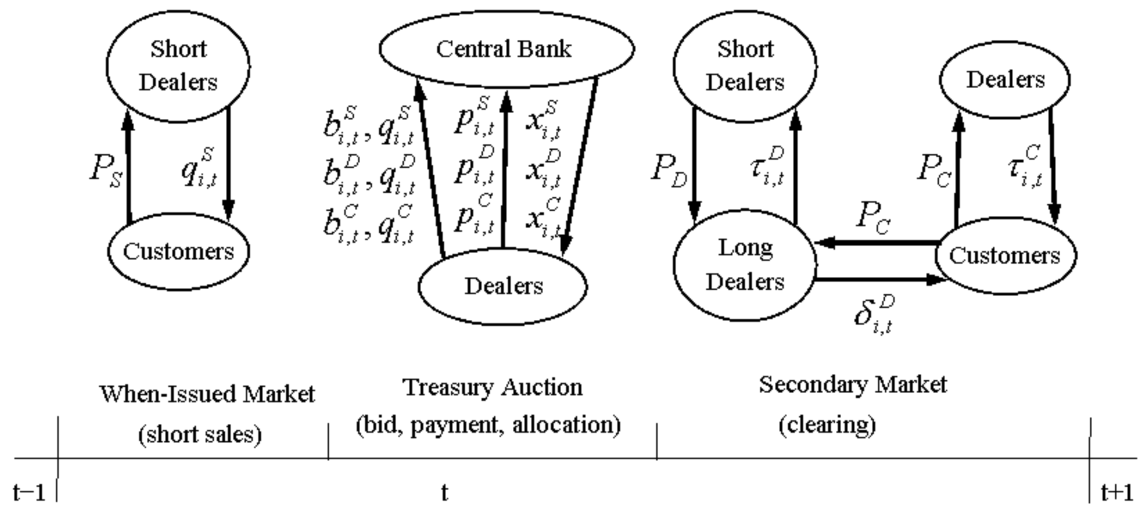


Figure 5.2. Learning Model Concepts

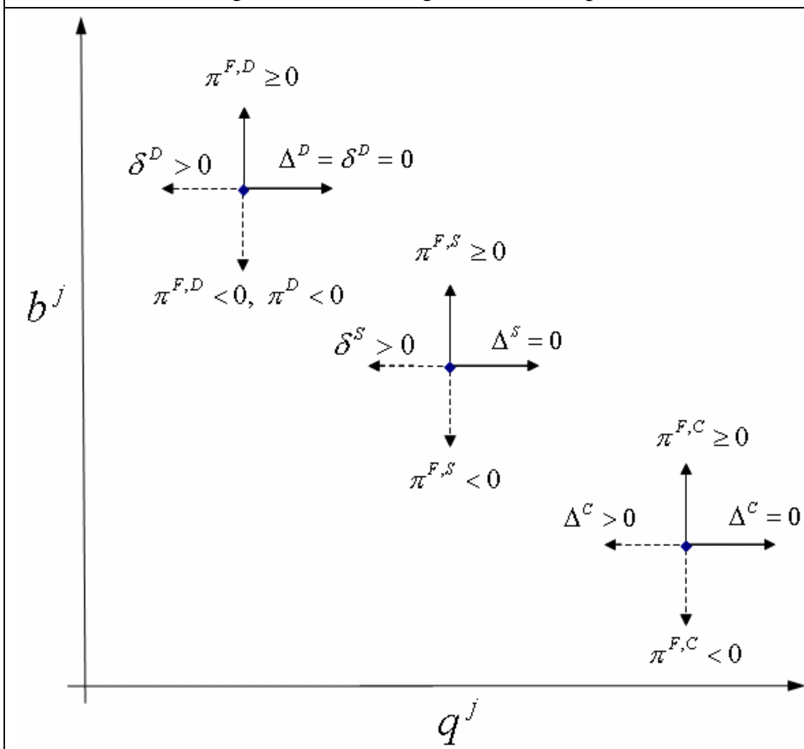

Figure 5.3. Sensitivity to Initial price and Learning Rates

$\varepsilon_D = \varepsilon_C = 0.20$

(Light area: revenue variation less than 0.1%;
Dark areas: revenue varies from light area by more than 0.1%.)

| Payment | Initial Price and Price Learning Rate | Initial Price and Quantity Learning Rate | Price Learning Rate and Quantity Learning Rate |
|---|---|---|---|
| Discrim | 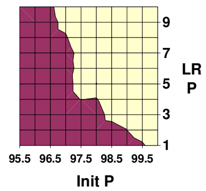 | 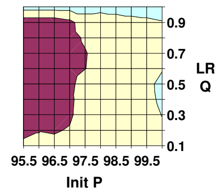 | 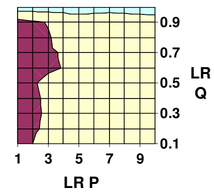 |
| Uniform | 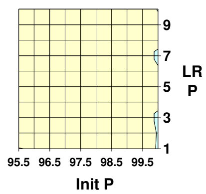 | 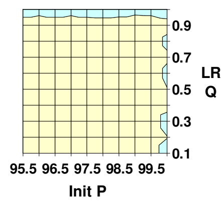 | 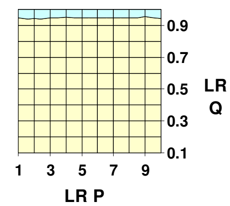 |
| Average | 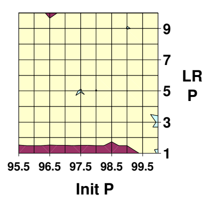 | 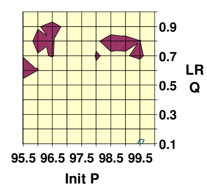 | 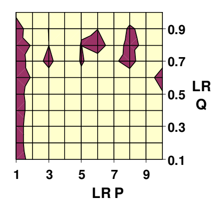 |
| Spanish | 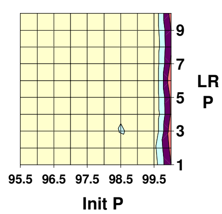 | 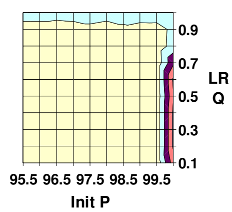 | 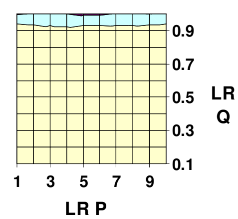 |
| Vickrey | 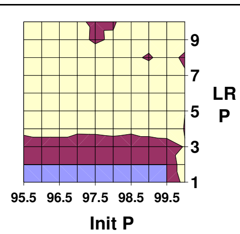 | 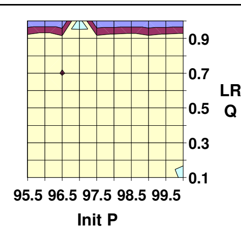 | 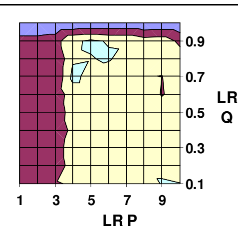 |

Figure 5.4. Convergence

$\varepsilon_D = \varepsilon_C = 0.10$

(S: black; D: medium gray; C: light gray)

| *Discriminatory* | *Uniform* |
|---|---|
| Bid Price | Bid Price |
| 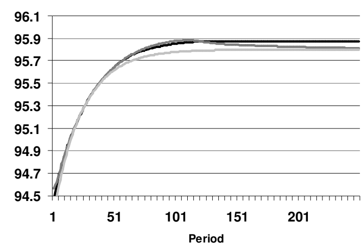 | 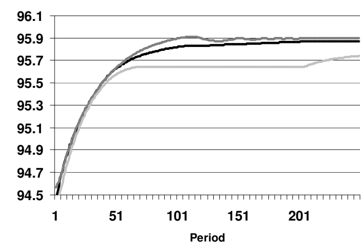 |
| Bid Quantity | Bid Quantity |
| 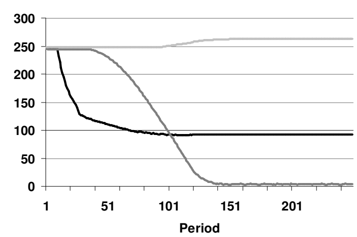 | 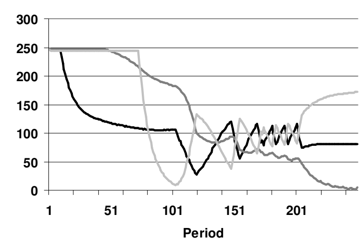 |
| Allocation | Allocation |
| 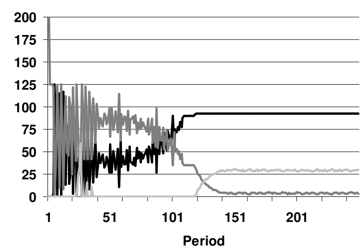 | 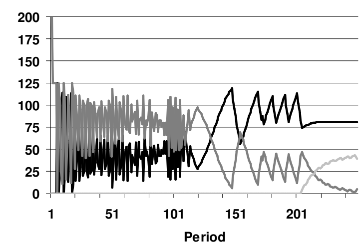 |
| Clearing | Clearing |
| 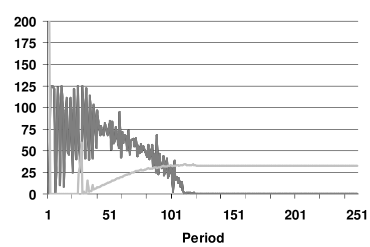 | 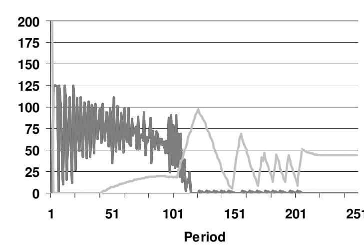 |

Figure 5.4. Convergence (continued)

$\varepsilon_D = \varepsilon_C = 0.10$

(S: black; D: medium gray; C: light gray)

| *Average* | *Spanish* | *Vickrey* |
|---|---|---|
| Bid Price | Bid Price | Bid Price |
| 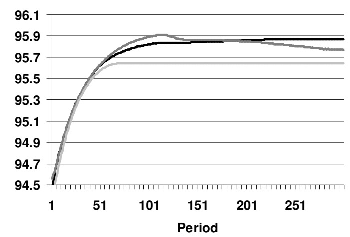 | 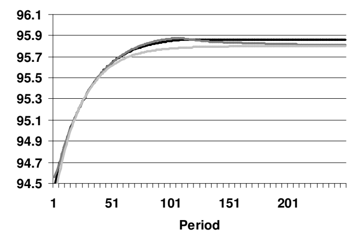 | 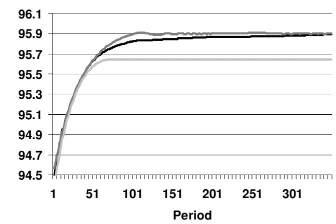 |
| Bid Quantity | Bid Quantity | Bid Quantity |
| 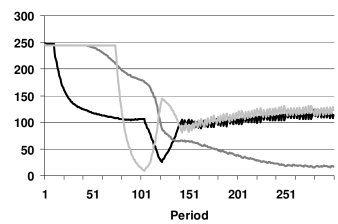 | 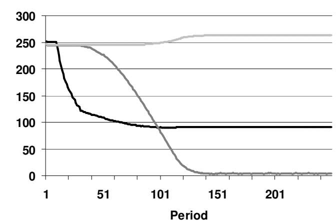 | 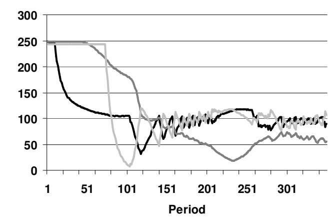 |
| Allocation | Allocation | Allocation |
| 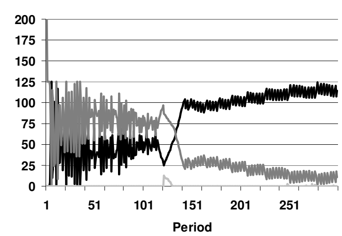 | 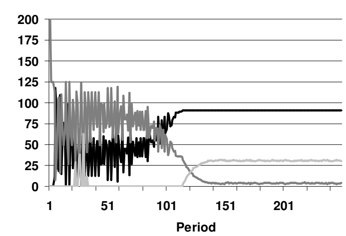 | 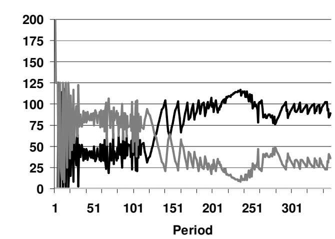 |
| Clearing | Clearing | Clearing |
| 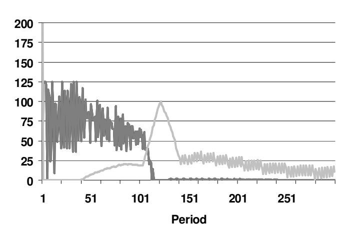 | 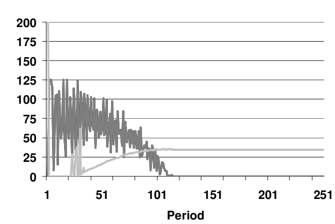 | 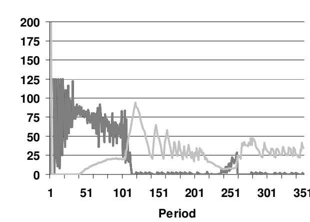 |

Figure 5.5. Endogenous Agent Variation: Bid Price

$$\varepsilon_D = \varepsilon_C = 0.10$$

(S: black; D: medium gray; C: light gray)

(Primary Dealers: Agents 1-10; Regular Dealers: Agents 11-20)

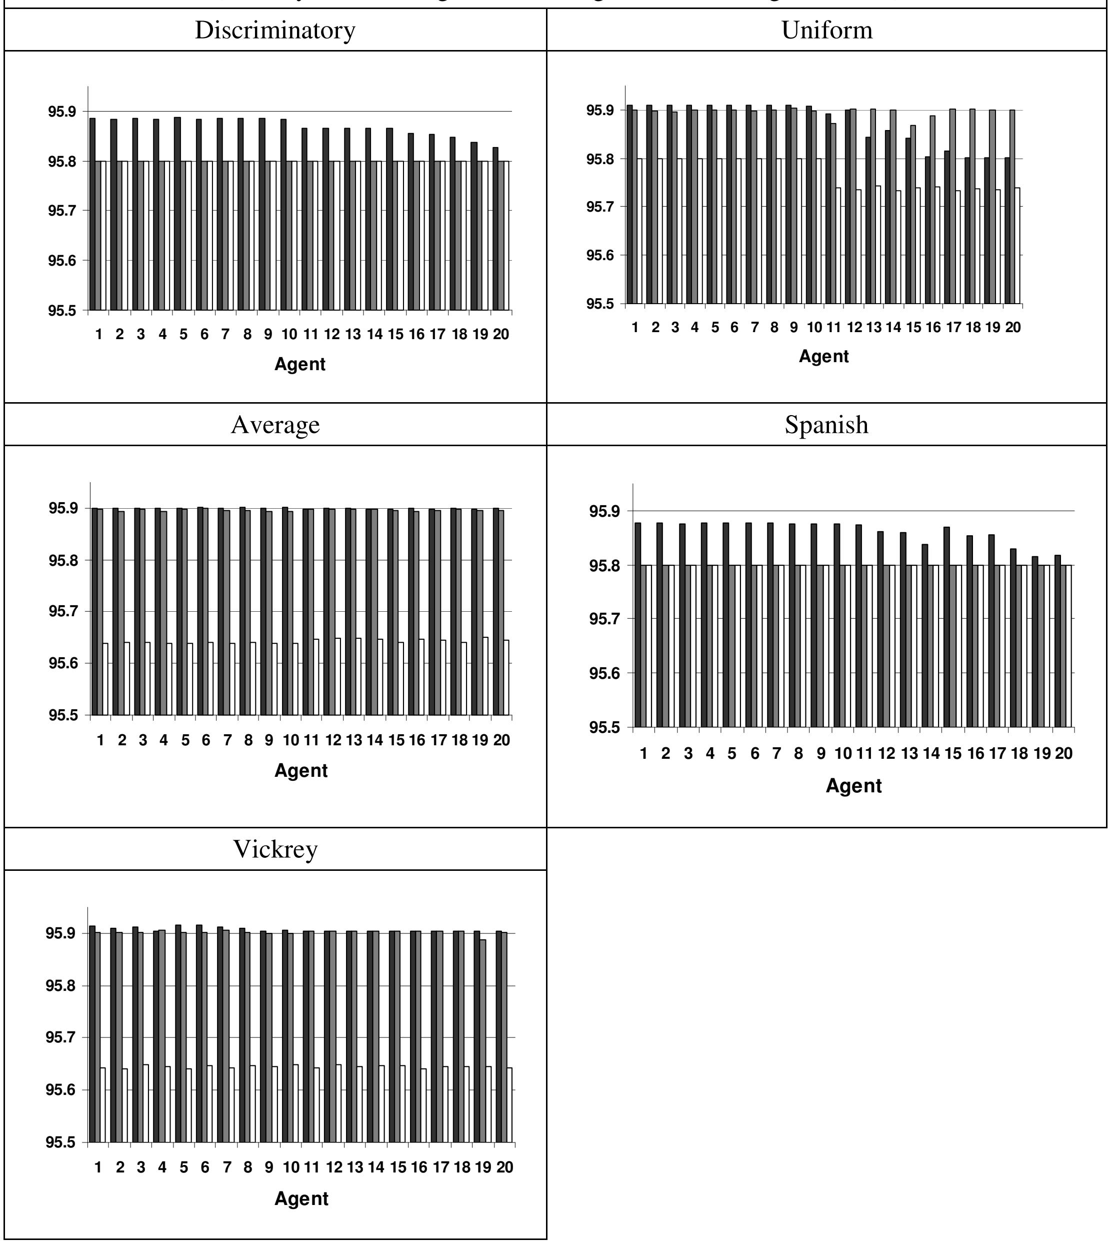

Figure 5.6. Endogenous Agent Variation: Quantities

$$\varepsilon_D = \varepsilon_C = 0.10$$

(S: black; D: medium gray; C: light gray)

(Primary Dealers: Agents 1-10; Regular Dealers: Agents 11-20)

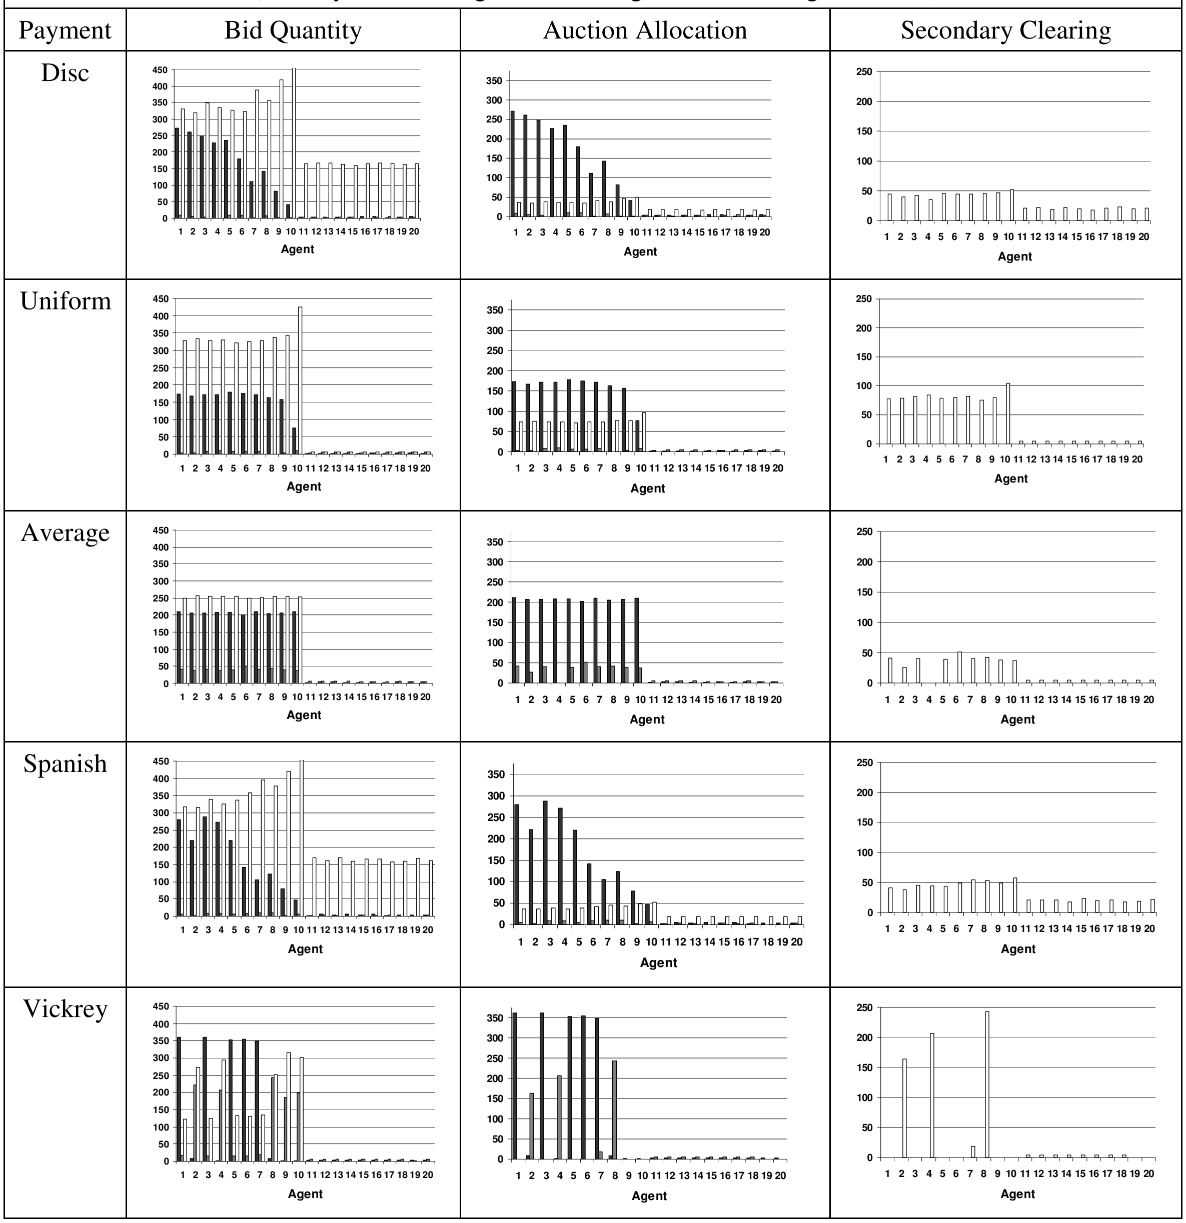

Figure 5.7. Revenue
(000,000)

**Discriminatory Payment**

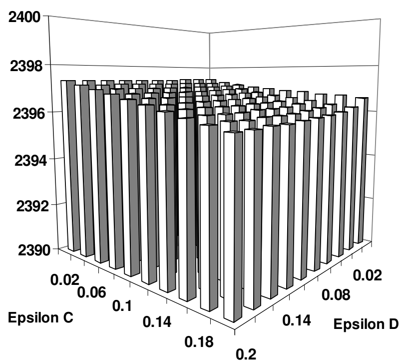


**Uniform Payment**

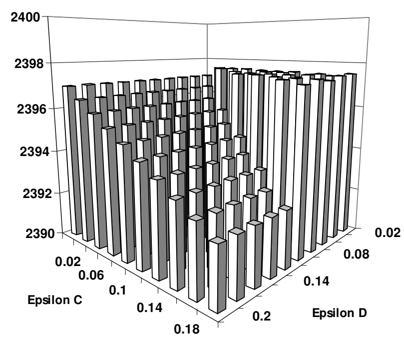


**Average Payment**

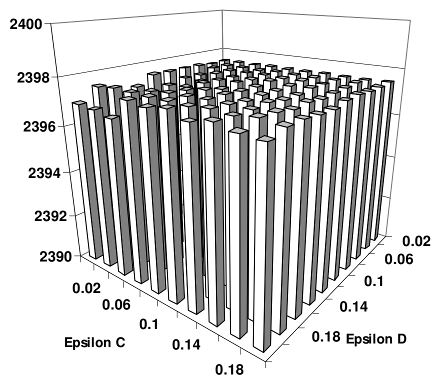


**Spanish Payment**

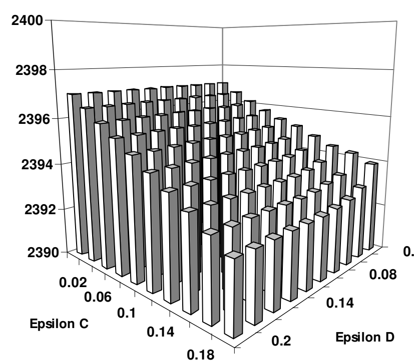


**Vickrey Payment**

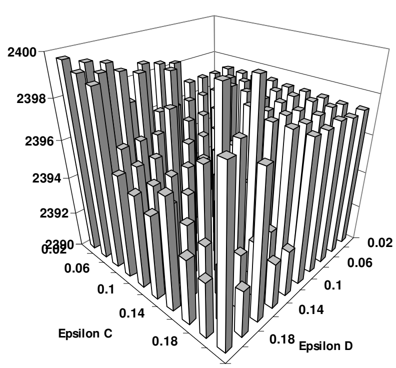

Figure 5.8. Revenue Differences from Discriminatory
(millions)

Uniform Payment

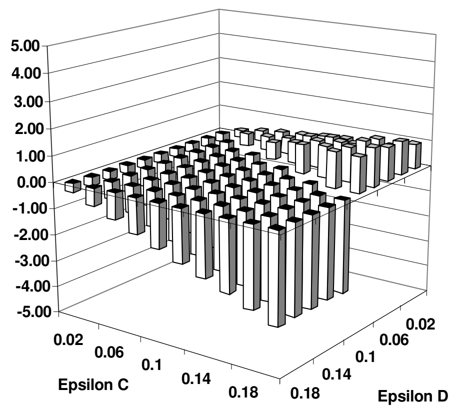


Average Payment

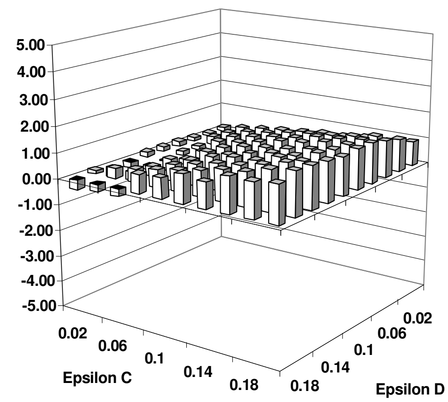


Spanish Payment

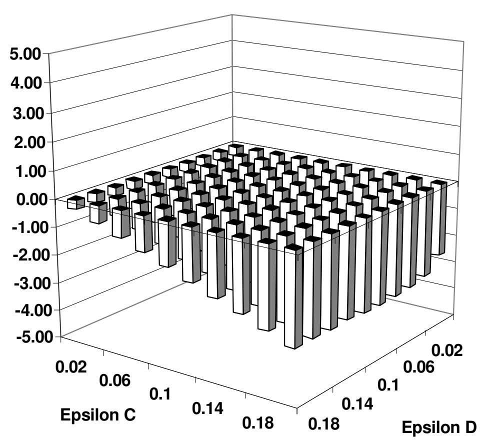


Vickrey Payment

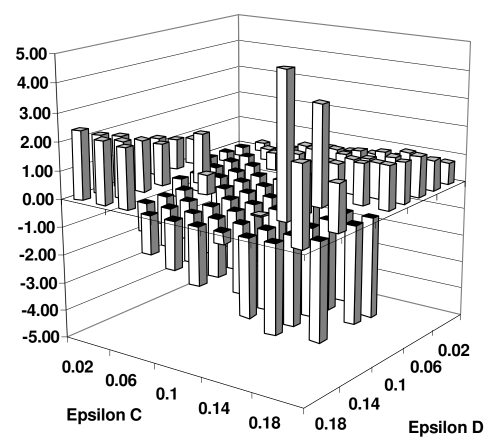

Figure 5.9. When-Issued Trading (Security S)

Discriminatory Payment

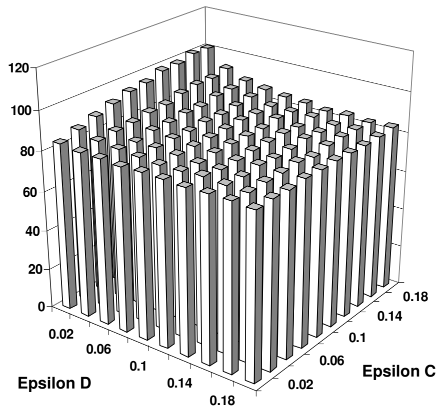


Uniform Payment

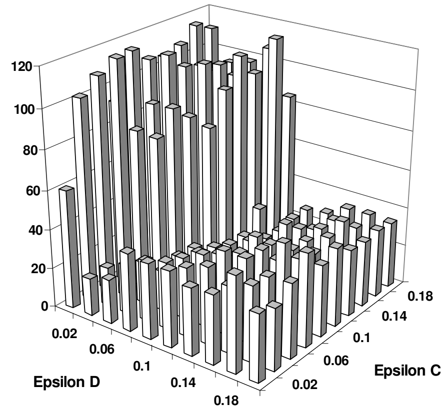


Average Payment

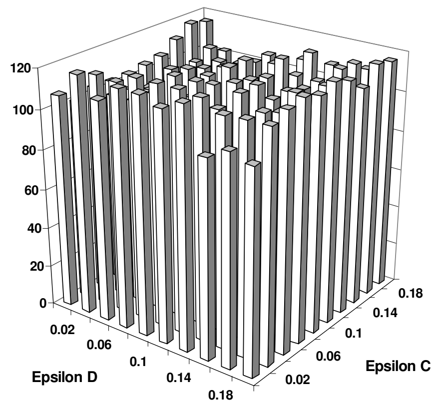


Spanish Payment

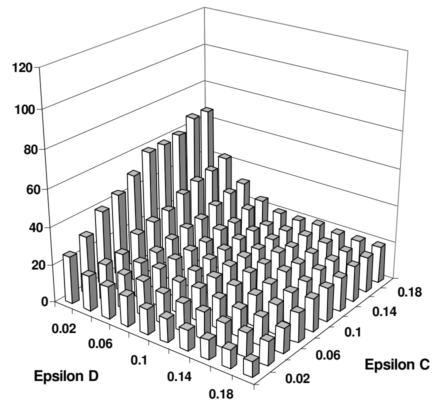


Vickrey Payment

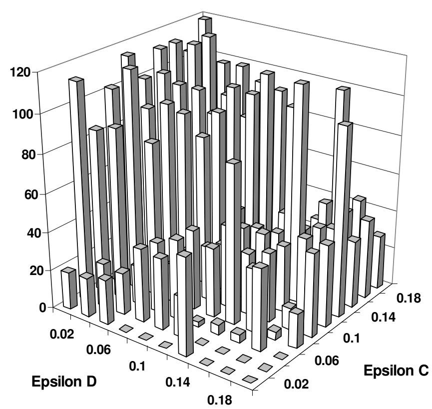

Figure 5.10. Secondary Market Trading (Security C)

Discriminatory Payment

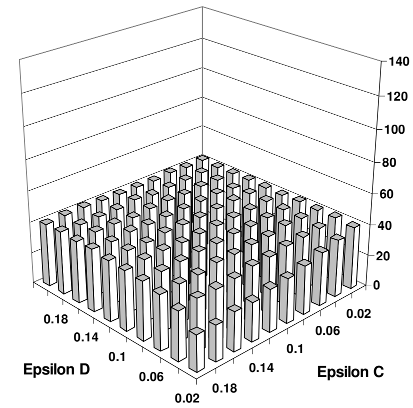


Uniform Payment

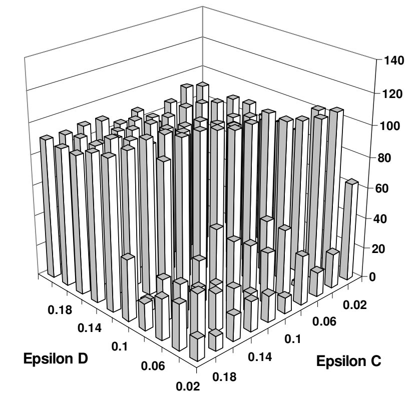


Average Payment

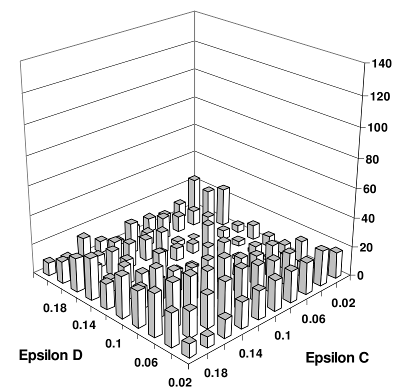


Spanish Payment

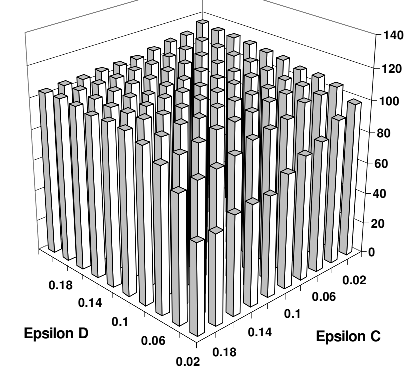


Vickrey Payment

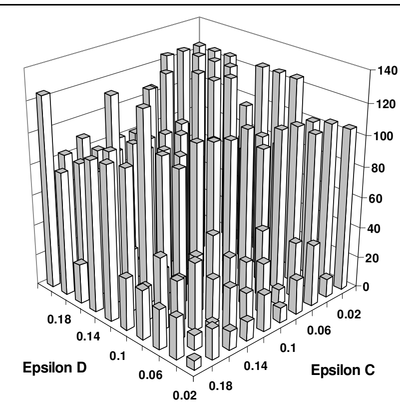

## 5.11 Appendix

**Profit for S1, S2:** For a full allocation and full clearing of security S, profit is

$\pi^S_{i,t} = (P_S - p^S_{i,t})q^S_{i,t}$.

The profit for security S is the difference between the revenue from the short sale ($P_S q^S_{i,t}$), the amount paid in the auction and the secondary market to acquire the shares $\left(p^S_{i,t}x^S_{i,t} + P_D\tau^S_{i,t}\right)$, and the amount of revenue refunded to the client if the unallocated shares are able to be acquired in the secondary market ($P_S\left(\Delta^S_{i,t} - \tau^S_{i,t}\right)$). Therefore, the full profit function is $\pi^S_{i,t} = P_S q^S_{i,t} - \left(p^S_{i,t}x^S_{i,t} + P_D\tau^S_{i,t}\right) - P_S\left(\Delta^S_{i,t} - \tau^S_{i,t}\right)$ = $P_S q^S_{i,t} - \left(p^S_{i,t}x^S_{i,t} + P_D\tau^S_{i,t}\right) - P_S\delta^S_{i,t}$. The main objective of a short selling agent is to achieve a full allocation in the auction ($x^S_{i,t} = q^S_{i,t}$) so that $\Delta^S_{i,t} = 0$ and $\tau^S_{i,t} = 0$. In this case the agent avoids paying a high squeeze price in the secondary market and avoids the risk of cancelling the unfilled customer orders. Full allocation results in the maximum possible profit for the current bid price: $\pi^S_{i,t} = P_S q^S_{i,t} - p^S_{i,t}x^S_{i,t} = (P_S - p^S_{i,t})q^S_{i,t}$. This level of profit is the benchmark that the agent uses to calculate foregone profits in the cases when security S is not fully allocated or fully cleared. Actual profit is negative when $p^S_{i,t} > P_S$.

**Foregone Profit for S3, S4, S5**: For a partial or nil allocation and full clearing of security S, foregone profit is $\pi^{F,S}_{i,t} = (P_D - p^S_{i,t})\Delta^S_{i,t}$.

If the agent is partially or nil allocated in the auction but clears the allocation shortfall $\Delta^S_{i,t}$ in the secondary market ($\tau^S_{i,t} = \Delta^S_{i,t} \rightarrow \delta^S_{i,t} = 0$), its profit function is $\pi^S_{i,t} = P_S q^S_{i,t} - \left(p^S_{i,t}x^S_{i,t} + P_D\tau^S_{i,t}\right) - P_S\delta^S_{i,t}$ = $P_S q^S_{i,t} - \left(p^S_{i,t}x^S_{i,t} + P_D\Delta^S_{i,t}\right)$. The foregone profit is the difference between this and the benchmark profit in the previous section so that

$\pi_{i,t}^{F,S} = (P_S - p_{i,t}^S)q_{i,t}^S - P_S q_{i,t}^S + \left( p_{i,t}^S x_{i,t}^S + P_D \Delta_{i,t}^S \right) = -p_{i,t}^S q_{i,t}^S + p_{i,t}^S x_{i,t}^S + P_D \Delta_{i,t}^S = (P_D - p_{i,t}^S)\Delta_{i,t}^S$.

The foregone profit is the extra amount paid in the secondary market to cover the auction shortfall.

**Foregone Profit for S6, S7**: For partial or nil clearing of a partial allocation of security S, the foregone profit is $\pi_{i,t}^{F,S} = (P_S - p_{i,t}^S)\Delta_{i,t}^S + (P_D - P_S)\tau_{i,t}^S$ =

$(P_D - p_{i,t}^S)\Delta_{i,t}^S - (P_D - P_S)\delta_{i,t}^S$.

If the agent is partially allocated in the auction ($0 < x_{i,t}^S < q_{i,t}^S$) and partially cleared in the secondary market ($0 < \tau_{i,t}^S < \Delta_{i,t}^S \to \delta_{i,t}^S > 0$), its profit function is

$\pi_{i,t}^S = P_S q_{i,t}^S - \left( p_{i,t}^S x_{i,t}^S + P_D \tau_{i,t}^S \right) - P_S \left( \Delta_{i,t}^S - \tau_{i,t}^S \right)$, the revenue from the short sales in the when-issued market, minus the amount paid in the auction, minus the amount paid in the secondary market, minus the amount refunded to the customers resulting from the amount not cleared in the secondary market. The foregone profit $\pi_{i,t}^{F,S}$ is the difference between this profit and the benchmark profit, so that $\pi_{i,t}^{F,S}$ =

$(P_S - p_{i,t}^S)q_{i,t}^S - P_S q_{i,t}^S + \left( p_{i,t}^S x_{i,t}^S + P_D \tau_{i,t}^S \right) + P_S \left( \Delta_{i,t}^S - \tau_{i,t}^S \right)$

= $P_S q_{i,t}^S - p_{i,t}^S q_{i,t}^S - P_S q_{i,t}^S + p_{i,t}^S x_{i,t}^S + P_D \tau_{i,t}^S + P_S \Delta_{i,t}^S - P_S \tau_{i,t}^S$ =

$-p_{i,t}^S q_{i,t}^S + p_{i,t}^S x_{i,t}^S + P_D \tau_{i,t}^S + P_S \Delta_{i,t}^S - P_S \tau_{i,t}^S = -p_{i,t}^S \Delta_{i,t}^S + P_D \tau_{i,t}^S + P_S \Delta_{i,t}^S - P_S \tau_{i,t}^S$ =

$(P_S - p_{i,t}^S)\Delta_{i,t}^S + (P_D - P_S)\tau_{i,t}^S$. Substituting $\Delta_{i,t}^S - \delta_{i,t}^S$ for $\tau_{i,t}^S$ yields

$(P_S - p_{i,t}^S)\Delta_{i,t}^S + (P_D - P_S)(\Delta_{i,t}^S - \delta_{i,t}^S) = -p_{i,t}^S \Delta_{i,t}^S + P_D(\Delta_{i,t}^S - \delta_{i,t}^S) + P_S \Delta_{i,t}^S - P_S(\Delta_{i,t}^S - \delta_{i,t}^S)$ =

$-p_{i,t}^S \Delta_{i,t}^S + P_D \Delta_{i,t}^S - P_D \delta_{i,t}^S + P_S \Delta_{i,t}^S - P_S \Delta_{i,t}^S + P_S \delta_{i,t}^S = (P_D - p_{i,t}^S)\Delta_{i,t}^S - P_D \delta_{i,t}^S + P_S \delta_{i,t}^S$ =

$(P_D - p_{i,t}^S)\Delta_{i,t}^S - (P_D - P_S)\delta_{i,t}^S$. Note that when $\delta_{i,t}^S = 0$, we have the result in the previous section.

**Profit for D1, D2**: Profit for a full allocation and full clearing is $\pi_{i,t}^D = (P_D - p_{i,t}^D)q_{i,t}^D$.

The agent bidding for security type D has a full profit function $\pi_{i,t}^D = P_D\tau_{i,t}^D - p_{i,t}^D x_{i,t}^D$, earning revenue in the secondary market by selling $\tau_{i,t}^D$ securities at the price $P_D$ after buying $x_{i,t}^D$ securities at a price of $p_{i,t}^D$ in the auction. If the agent is fully allocated in the auction ($x_{i,t}^D = q_{i,t}^D$) and clears all of its allocation at price $P_D$ in the secondary market ($\tau_{i,t}^D = x_{i,t}^D$), the profit is $\pi_{i,t}^D = P_D q_{i,t}^D - p_{i,t}^D q_{i,t}^D = (P_D - p_{i,t}^D)q_{i,t}^D$. This level of profit is the benchmark that the agent uses to calculate foregone profits for security S. Actual profit is negative when $p_{i,t}^D > P_D$.

**Foregone profit for D3, D4:** Foregone profit with a full allocation of security D with partial clearing is $\pi_{i,t}^{F,D} = (P_D - P_C)\delta_{i,t}^D$ and with nil clearing is $(P_D - P_C)q_{i,t}^D$.

If an agent has a full allocation ($x_{i,t}^D = q_{i,t}^D$) that it partially clears ($\tau_{i,t}^D < q_{i,t}^D \to \delta_{i,t}^D > 0$) at price $P_D$, the profit function is $\pi_{i,t}^D = P_D\tau_{i,t}^D + P_C\delta_{i,t}^D - p_{i,t}^D x_{i,t}^D =$ $P_D\tau_{i,t}^D + P_C\delta_{i,t}^D - p_{i,t}^D q_{i,t}^D$. Since the foregone profit is the difference between this profit and the benchmark profit, $\pi_{i,t}^{F,D} = (P_D - p_{i,t}^D)q_{i,t}^D - P_D\tau_{i,t}^D - P_C\delta_{i,t}^D + p_{i,t}^D q_{i,t}^D = P_D(q_{i,t}^D - \tau_{i,t}^D) - P_C\delta_{i,t}^D =$ $(P_D - P_C)\delta_{i,t}^D$. For nil clearing $\delta_{i,t}^D = q_{i,t}^D$. Profit is negative when $P_D\tau_{i,t}^D + P_C\delta_{i,t}^D < p_{i,t}^D q_{i,t}^D$, but foregone profit is never negative since $P_D$ is always greater than $P_C$.

**Foregone Profit for D5, D6, D7:** Foregone profit for a partial allocation and full clearing (at price $P_D$) is $\pi_{i,t}^{F,D} = (P_D - p_{i,t}^D)\Delta_{i,t}^D$.

When there is full clearing of a partial allocation ( $\tau_{i,t}^{D} = x_{i,t}^{D}$), setting $\tau_{i,t}^{D} = x_{i,t}^{D}$ the profit is $\pi_{i,t}^{D} = (P_D - p_{i,t}^{D})x_{i,t}^{D}$ and foregone profit is $\pi_{i,t}^{F,D} = (P_D - p_{i,t}^{D})q_{i,t}^{D} - (P_D - p_{i,t}^{D})x_{i,t}^{D}$ = $(P_D - p_{i,t}^{D})\Delta_{i,t}^{D}$. Profit and foregone profit are negative when $p_{i,t}^{D} > P_D$.

**Foregone Profit for D8, D9, D10:** Foregone profit for a partial allocation and partial clearing is $\pi_{i,t}^{F,D} = P_D(q_{i,t}^{D} - \tau_{i,t}^{D}) - p_{i,t}^{D}\Delta_{i,t}^{D}$.

If an agent partially clears ( $\tau_{i,t}^{D} < x_{i,t}^{D} \rightarrow \delta_{i,t}^{D} > 0$) a partial allocation ( $0 < x_{i,t}^{D} < q_{i,t}^{D}$) at price $P_D$, the profit function is $\pi_{i,t}^{D} = P_D\tau_{i,t}^{D} + P_C\delta_{i,t}^{D} - p_{i,t}^{D}x_{i,t}^{D}$. The foregone profit is $\pi_{i,t}^{F,D} = (P_D - p_{i,t}^{D})q_{i,t}^{D} - P_D\tau_{i,t}^{D} - P_C\delta_{i,t}^{D} + p_{i,t}^{D}x_{i,t}^{D}$ = $P_D(q_{i,t}^{D} - \tau_{i,t}^{D}) - P_C\delta_{i,t}^{D} - p_{i,t}^{D}(q_{i,t}^{D} - x_{i,t}^{D})$ = $(P_D - P_C)\delta_{i,t}^{D} - p_{i,t}^{D}\Delta_{i,t}^{D}$. Profit is negative when $P_D\tau_{i,t}^{D} + P_C\delta_{i,t}^{D} < p_{i,t}^{D}x_{i,t}^{D}$, and foregone profit is negative when $(P_D - P_C)\delta_{i,t}^{D} < p_{i,t}^{D}\Delta_{i,t}^{D}$.

**Foregone Profit for D11, D12, D13:** Foregone profit for a partial allocation and nil clearing at price $P_D$ is $\pi_{i,t}^{F,D} = P_D q_{i,t}^{D} - P_C x_{i,t}^{D} - p_{i,t}^{D}\Delta_{i,t}^{D}$.

If an agent clears none ( $\tau_{i,t}^{D} = 0 \rightarrow \delta_{i,t}^{D} = x_{i,t}^{D}$) of a partial allocation ( $0 < x_{i,t}^{D} < q_{i,t}^{D}$) at price $P_D$, the entire allocation is sold at $P_C$, yielding a profit function $\pi_{i,t}^{D} = P_C x_{i,t}^{D} - p_{i,t}^{D}x_{i,t}^{D} = (P_C - p_{i,t}^{D})x_{i,t}^{D}$. The foregone profit is $\pi_{i,t}^{F,D} =$ $(P_D - p_{i,t}^{D})q_{i,t}^{D} - (P_C - p_{i,t}^{D})x_{i,t}^{D}$ = $P_D q_{i,t}^{D} - P_C x_{i,t}^{D} + p_{i,t}^{D}x_{i,t}^{D} - p_{i,t}^{D}q_{i,t}^{D}$ = $P_D q_{i,t}^{D} - P_C x_{i,t}^{D} - p_{i,t}^{D}\Delta_{i,t}^{D}$. Profit is negative when when $P_C < p_{i,t}^{D}$, and foregone profit is negative when a high payment leads to $P_D q_{i,t}^{D} - P_C x_{i,t}^{D} < p_{i,t}^{D}\Delta_{i,t}^{D}$.

**Foregone Profit for D14, D15:** Foregone profit for a nil allocation is $\pi_{i,t}^{F,D}$ = $(0.5P_D + 0.5P_C - p_{i,t}^{D})q_{i,t}^{D}$.

The profit with no allocation is zero. The foregone profit in this case depends upon the agent's expectation of clearing at price $P_D$ in the secondary market. Assuming equal probability of clearing and not clearing, the agents expected foregone profit is $\pi_{i,t}^{F,D} = (0.5P_D + 0.5P_C - p_{i,t}^D)q_{i,t}^D$, which is negative when $0.5P_D + 0.5P_C < p_{i,t}^D$.

**Profit for C1, C2***:* Profit for a full allocation and full clearing of security C is $\pi_{i,t}^C = (P_C - p_{i,t}^C)q_{i,t}^C$.

Security C is set up to be fully cleared in the secondary market, i.e., the allocation $x_{i,t}^C$ in the auction clears the secondary market at price $P_C$. Thus, an agent's only concern is to be fully allocated in the auction and to make a profit. The full profit function is $\pi_{i,t}^C = P_C x_{i,t}^C - p_{i,t}^C x_{i,t}^C$, and a fully allocated agent will have a profit of $\pi_{i,t}^C = (P_C - p_{i,t}^C)q_{i,t}^C$. Actual profit is negative when $p_{i,t}^C > P_C$.

**Foregone Profit for C3, C4:** Foregone profit for a partial allocation of security C is $\pi_{i,t}^{F,C} = (P_C - p_{i,t}^C)\Delta_{i,t}^C$.

In the case of partial allocation ($0 < x_{i,t}^C < q_{i,t}^C$) the profit is $\pi_{i,t}^C = (P_C - p_{i,t}^C)x_{i,t}^C$, and there is a foregone profit of $\pi_{i,t}^{F,C} = (P_C - p_{i,t}^C)(q_{i,t}^C - x_{i,t}^C)$ = $(P_C - p_{i,t}^C)\Delta_{i,t}^C$. Foregone profit is negative when $p_{i,t}^C > P_C$.

**Foregone Profit for C5, C6:** Foregone profit for a nil allocation of security C is $\pi_{i,t}^{F,C}$ $= (P_C - p_{i,t}^C)q_{i,t}^C$.

In the case of nil allocation, $\Delta_{i,t}^C = q_{i,t}^C$ so profit is zero and the foregone profit is the benchmark profit $(P_C - p_{i,t}^C)q_{i,t}^C$. Foregone profit is negative when $p_{i,t}^C > P_C$.