University of Victoria

Department Discussion Paper DDP0707

ISSN 1914-2838

Department of Economics

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

Alan Mehlenbacher

Department of Economics, University of Victoria

Victoria, B.C., Canada V8W 2Y2

Abstract

This study uses a multiagent system to investigate how sealed-bid auction results vary across 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.

JEL classification: C15, C72, D83.

Keywords: agent-based computational economics; multi-dimensional value signals; sealed-bid auctions; impulse balance learning

Author Contact:
Alan Mehlenbacher, Dept. of Economics, University of Victoria, P.O. Box 1700, STN CSC, Victoria, B.C., Canada, V8W 2Y2; e-mail: amehlen@uvic.ca; Voice: (250) 721-8537; FAX: (250) 721-6214

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 2, I describe the auction model. Section 3 provides details of the learning model and its properties of convergence and sensitivity. Section 4 compares the results of learning model with results from lab experiments in other studies. Section 5 demonstrates the nonlinear variation of revenue and efficiency with the common value percent. Section 6 shows the results of requiring the auction winners to reveal the actual common value to the auction losers. Section 7 presents conclusions.

2 Auction Model

The multiagent system platform is described in Mehlenbacher (2007). In this section, I describe how the system implements values and the value signals for bidders (2.1), the levels of information feedback (2.2), and the number of periods and bidders (2.3).

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"

[1] The notation is summarized in Table 1.

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 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 (θ_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 $[\underline{S}_\mathrm{P}, \overline{S}_\mathrm{P}]$ of a distribution of the private value signals $v_{P,t}^i$ and a support $[\underline{S}_\mathrm{C}, \overline{S}_\mathrm{C}]$ for the common value v_C. In most experimental studies and the simulations in this paper $[\underline{S}_\mathrm{C}, \overline{S}_\mathrm{C}] = [\underline{S}_\mathrm{P}, \overline{S}_\mathrm{P}]$. 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 $\hat{v}^i_{C,t}$ 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 $[\underline{S}_C, \bar{S}_C]$. There is uncertainty among the bidders about what this common value is, and a larger ε 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

---

[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.

strategies for the asymmetric distributions (beta(4,2) and beta(2,4)) shift right and left respectively.

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 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

[3] The notation is defined in Table 1 and the information feedback is summarized in Table 2.

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 3.2 and 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 ε neighbourhood of the center of the support, but this has minimal effect on the results.

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 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

---

[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.

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 6 and 7).

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 Mehlenbacher (2007). 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 3.1 describes Selten's impulse balance learning method. The next two sections describe the adjustment rules for the downward impulses for winners (Section 3.2) and the upward impulses for losers (Section 3.3) . Section 3.4 presents results from using impulse balance learning and an improved learning method: impulse learning with loss aversion (ILA). Section 3.5 presents convergence and sensitivity analyses for the ILA method, and Section 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 $\varepsilon = 5$.

3.1 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 auctions with common values. Bids are adjusted using downward $a^i_{-,t}$ or upward $a^i_{+,t}$

---

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

[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 $m^i_t = b^i_t - b^{(2)}_t$.

adjustments or "impulses" that the agent calculates using profit π_t^i, foregone profit $\pi_{F,t}^i$, and money on the table m_t^i. For profitable winners $a_{-,t}^i$ is money on the table m_t^i, and for unprofitable winners it is the loss π_t^i. For losers, $a_{+,t}^i$ 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 λ_t^i as the ratio of its expected value of the upward impulse to the expected value of the downward impulse: $\lambda_t^i = \frac{E_t\left[a_+^i\right]}{E_t\left[a_-^i\right]}$. To determine its adjusted bid, the agent weights the impulses by a learning rate ϕ and the downward impulse weight λ_t^i. The bid for period *t+1* is then a revision of the previous bid $b_{t+1}^i = b_t^i + \phi(a_{+,t}^i - \lambda_t^i a_{-,t}^i)$. 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, $\gamma_t^i = \frac{b_t^i}{\hat{v}_t^i}$, so that we can discuss the value multiplier γ_t^i that can be compared with theoretical and experimental results.

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 π_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 π_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 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]

3.4 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 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 λ_t^i, which varies with the bidder value, could be expected to deal with this problem but it is not producing satisfactory results. Figure 2 shows that the values of λ_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 3.1 When λ_t^i is removed from the model, the results improve slightly as shown in Figure 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 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 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 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 1 shows that the results from impulse balance learning vary significantly with the initial values 0.95 ± 0.02, 0.85 ± 0.02, and 0.75 ± 0.02. However, Figure 4 shows that the ILA method is quite insensitive to the initial values. In Figure 1, the profit curve for an initial value 0.95 ± 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 ± 0.02 and then to 0.75 ± 0.02, the values for mid-value and high-value bidders increase considerably so that the curve becomes much flatter. In Figure 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 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 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.

3.5 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 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 ϕ.

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 ϕ answers this question. Figure 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 ϕ=0.1 to 0.0024 for ϕ=1.0. These results are quite insensitive to initial value, and I use $\phi = 0.5$ and initial values of 0.85 ± 0.02 in the simulations as arbitrary choices.

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 ε= 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 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., $\mathcal{E} = \dfrac{v_{P,t}^{winner} - \min\left\{v_{P,t}^{i}\right\}}{\max\left\{v_{P,t}^{i}\right\} - \min\left\{v_{P,t}^{i}\right\}}$. $\mathcal{E} = 1$ 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.

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 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 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 ε). Figure 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 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 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 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 11, and the results for highest-value-signal winning percent for pure common-value auctions are shown in Table 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 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 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.

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 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 ε=8, 12, 18, and 27 respectively. The results in Figure 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 ε for four bidders, but decreases with ε for seven bidders. The results shown in Figure 13 are consistent with their results for four bidders (profit increases with ε), but not for seven bidders (no change in profit with ε). Figure 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 14, and the highest-value-signal winning percent for pure common-value auctions is shown in Table 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 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 4 shows that the agents produce similar results for levels of uncertainty above $\varepsilon = 8$.

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.

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 9 for first-price auctions and Figure 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.

5.2 Revenue

See Figures 15 and 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 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.

5.3 Efficiency

See Figure 11 for first-price auctions and Figure 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 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.

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 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 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 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.

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 3.4). 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 3.5).

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), 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 5.1). The nonlinear revenue curves tend to decrease more rapidly at low common value percents (Result 5 in Section 5.2). The very nonlinear efficiency curves tend to stay high and then decrease rapidly for common value percents (Result 7 in Section 5.3). In addition, revenue in most cases decrease with increasing information feedback to the bidders (Result 6 in Section 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 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.

8 Acknowledgements

I am indebted to David Scoones, Linda Welling, Don Ferguson, and Tony Marley for their valuable suggestions. I am also grateful to the Social Sciences and Humanities Research Council of Canada for scholarship support during the course of this research.

9 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.

Mehlenbacher, A., 2007. Multiagent system platform for auction experiments. University of Victoria, Economics Department Discussion Paper, DDP0706.

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.

10 Tables

Table 1. Sealed-Bid Model Notation Summary	
Symbol	Description
b_t^i	Bid price of bidder i in auction t.
$b_t^{(1)},\ldots,b_t^{(n)}$	Ordered bid prices in a sealed-bid auction where $b_t^{(1)}$ is the highest bid.
ε	Radius of the support for the common value signal. .
γ_t^i	Value multiplier: $b_t^i = \gamma_t^i \hat{v}_t^i$.
λ_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.
ϕ	Learning rate of bidder i at period t.
p_t^i	Payment made by bidder i, given that it wins.
π_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.
θ_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 2. Information Levels (incremental)		
Level	Description	Feedback
I1	Number of bidders	n
	Value signal support	$[\underline{S}_{\mathrm{C}}, \overline{S}_{\mathrm{C}}]$
	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. Summary of Comparison with Lab Experiments (Section 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 4. Highest-Value-Signal Winning Percent for Pure Common Value Information Level I3

Payment	Uncertainty ε	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%

11 Figures

Figure 1. Impulse Balance Learning: Profit by Common Value Signal
Common Value, First Price, I3

Initial value multiplier 0.95 ± 0.02	Initial value multiplier 0.85 ± 0.02	Initial value multiplier 0.75 ± 0.02
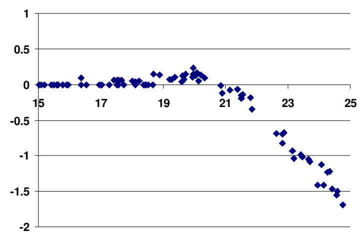	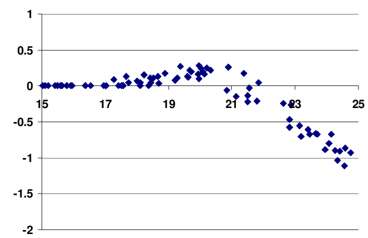	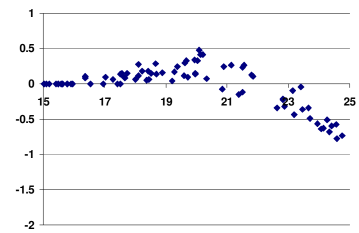

Figure 2. Impulse Balance Learning: λ_t^i by Common Value Signal
Common Value, First Price, I3

Initial value multiplier 0.95 ± 0.02	Initial value multiplier 0.85 ± 0.02	Initial value multiplier 0.75 ± 0.02
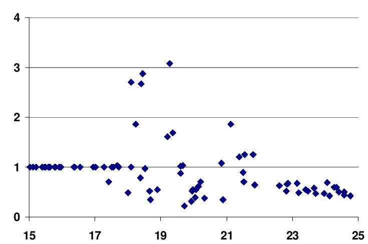	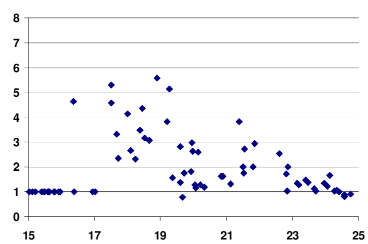	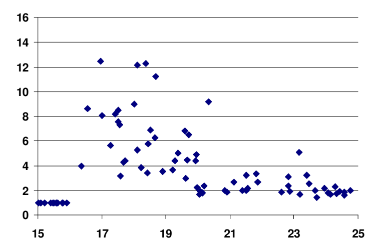

Figure 3. Learning Alternatives: Profit by Common Value Signal
Common Value, First Price, I3, Initial value multiplier 0.95 ± 0.02

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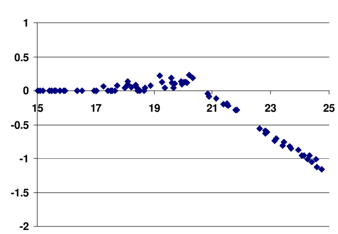	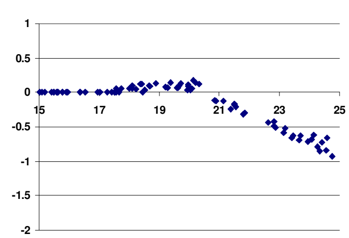	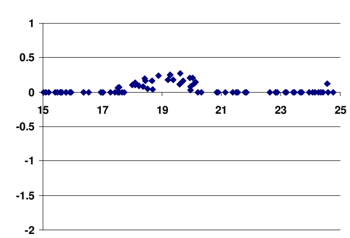

Figure 4. ILA learning: Profit by Common Value Signal
Common Value, First Price, I3

Initial value multiplier 0.95 ± 0.02	Initial value multiplier 0.85 ± 0.02	Initial value multiplier 0.75 ± 0.02

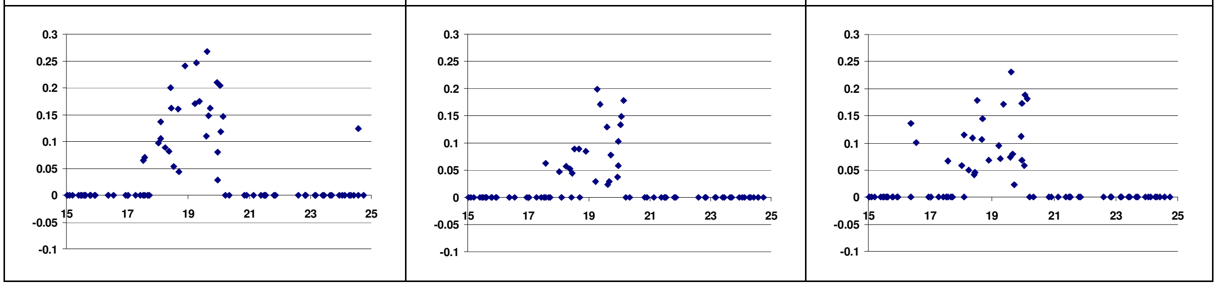

Figure 5. Bidders with Varying Value Signals
Profit by Common Value Signal
Initial value multiplier 0.95 ± 0.02
Common Value, First Price, I3

Impulse Balance	ILA

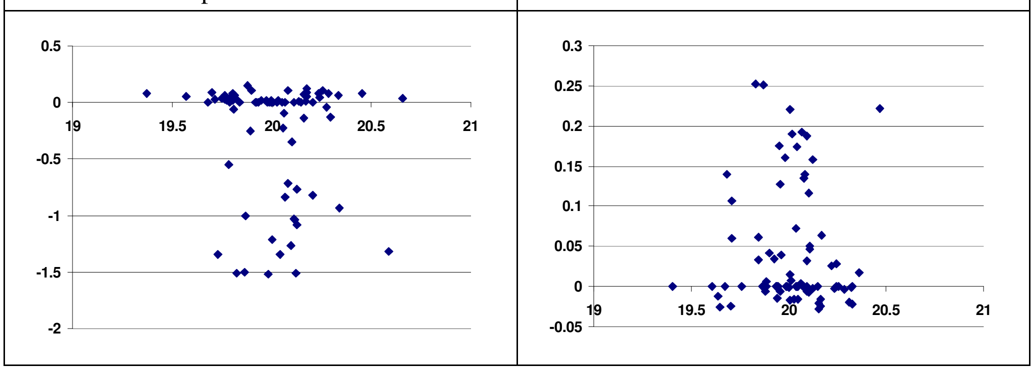

Figure 6. Convergence of Value Multiplier
Value Multiplier by Period
Common Value, First Price, I3 Information

Impulse Balance	ILA
Initial Value multiplier: 0.95 ± 0.02	Initial Value multiplier: 0.95 ± 0.02

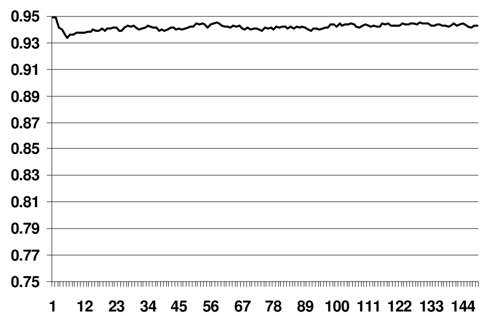

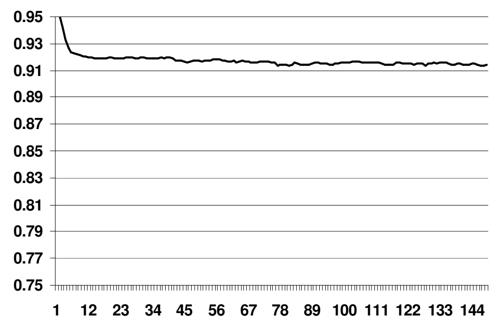

Impulse Balance	ILA
Initial Value multiplier: 0.85 ± 0.02	Initial Value multiplier: 0.85 ± 0.02

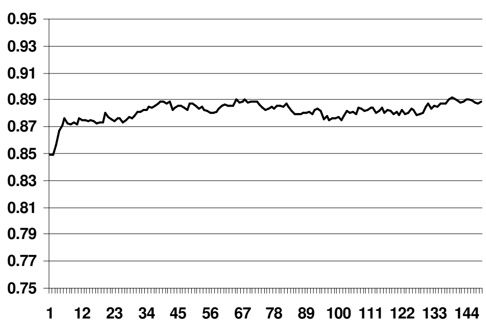

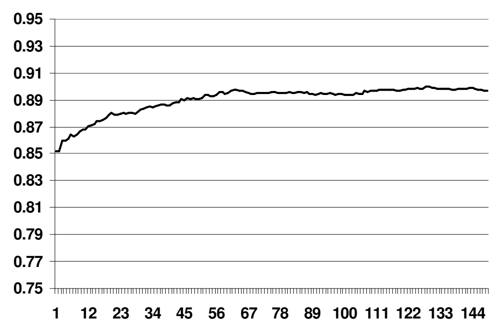

Impulse Balance	ILA
Initial Value multiplier: 0.75 ± 0.02	Initial Value multiplier: 0.75 ± 0.02

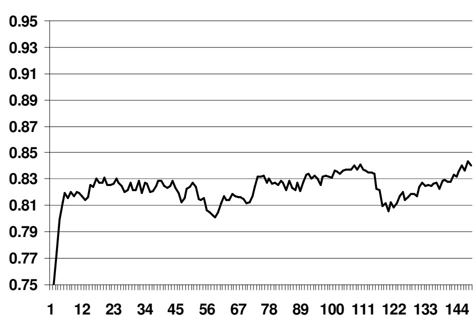

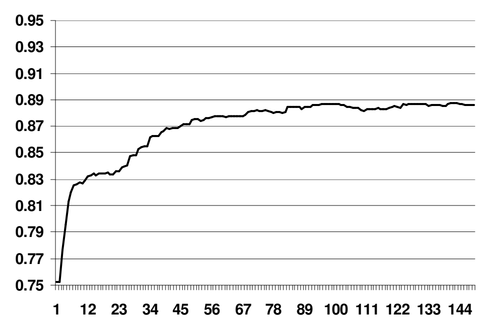

Figure 7. Sensitivity to Learning Rate
Value Multiplier by Period
Common Value, First Price, I3 Information
(I1: Solid, I2: Dot, I3: DashDot)

Learning Rate 0.1

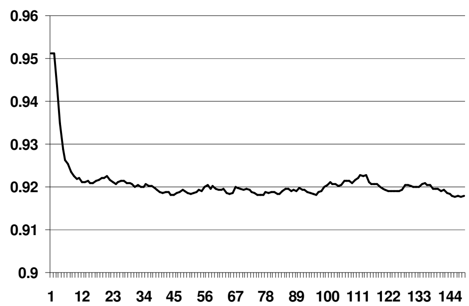

Learning Rate 0.3

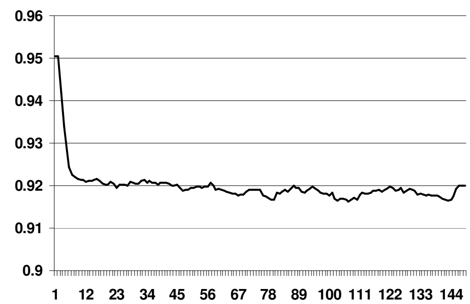

Learning Rate 0.5

Learning Rate 0.7

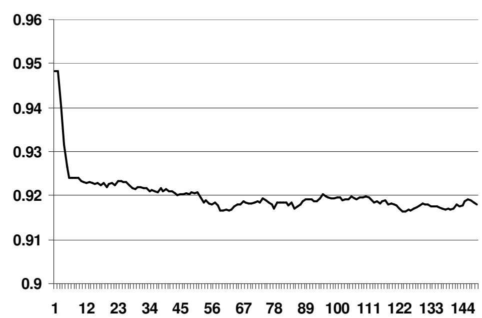

Learning Rate 0.9

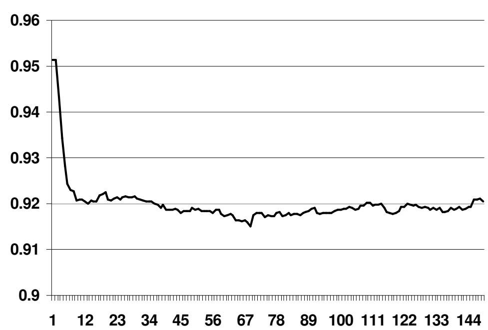

Learning Rate 1.0

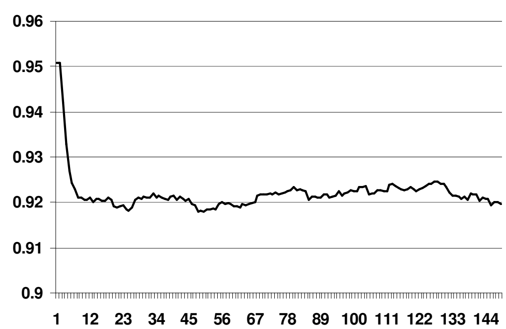

Figure 8. Value Multiplier: First-price auctions
Value Multiplier by Common Value Percent
(I1: Solid, I2: Dot, I3: DashDot)

Four Bidders	Seven Bidders
$\varepsilon = 8$	$\varepsilon = 8$

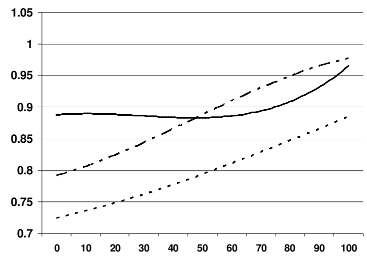

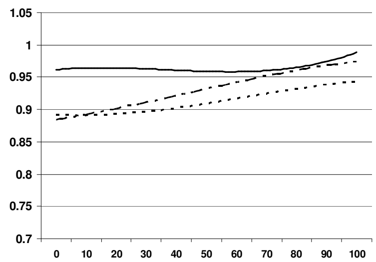

$\varepsilon = 12$	$\varepsilon = 12$

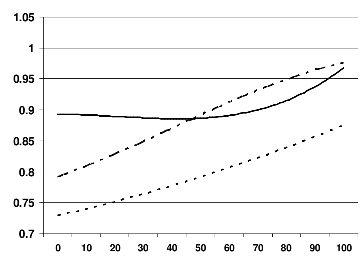

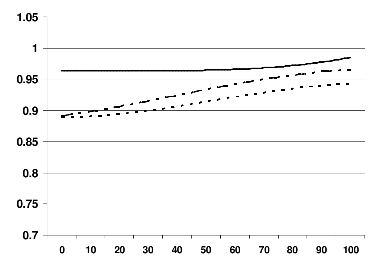

$\varepsilon = 18$	$\varepsilon = 18$

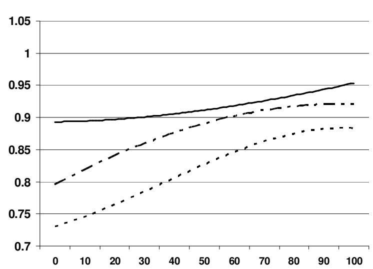

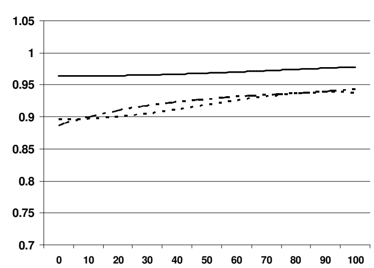

$\varepsilon = 27$	$\varepsilon = 27$

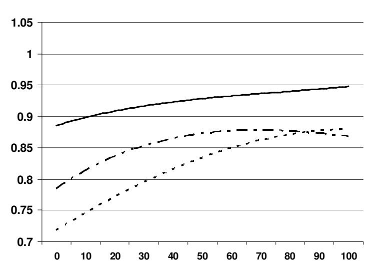

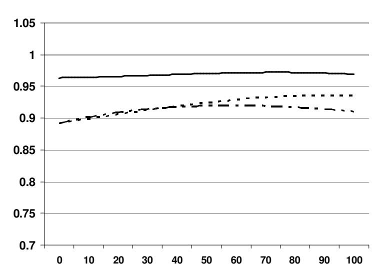

Figure 9. Profit: First-price auctions Profit by Common Value Percent (I1: Solid, I2: Dot, I3: DashDot)	
Four Bidders	Seven Bidders
$\varepsilon = 8$	$\varepsilon = 8$
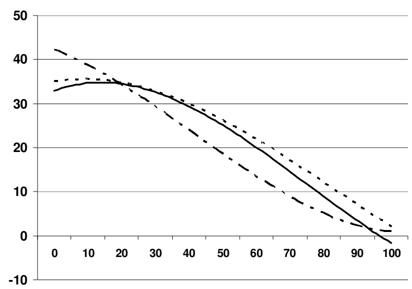	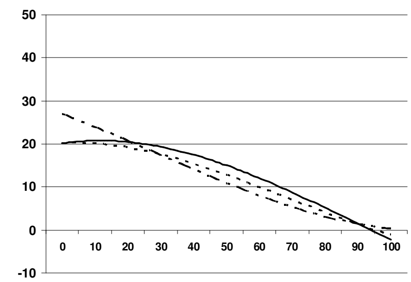
$\varepsilon = 12$	$\varepsilon = 12$
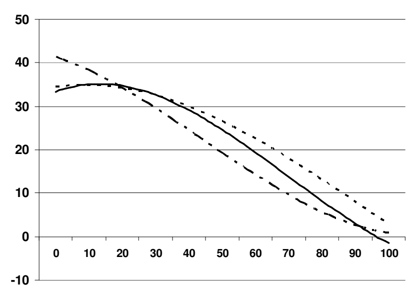	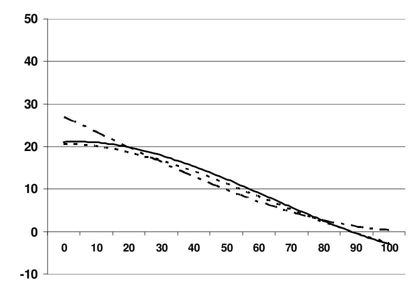
$\varepsilon = 18$	$\varepsilon = 18$
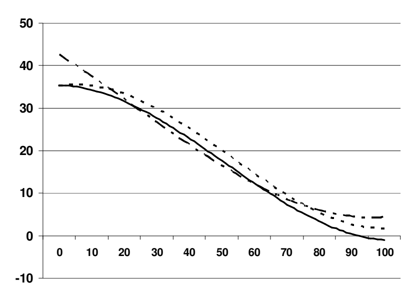	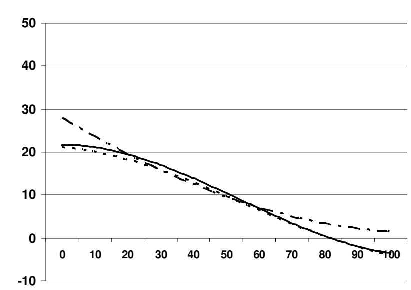
$\varepsilon = 27$	$\varepsilon = 27$
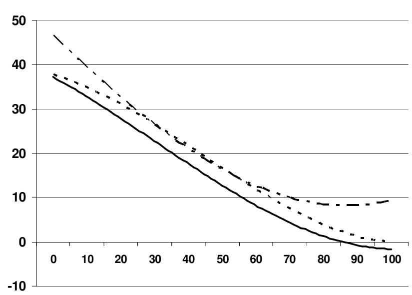	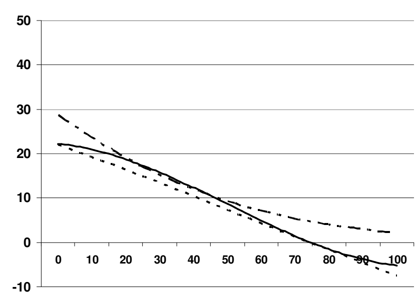

Figure 10. Profit with Many Bidders
Profit by Common Value Percent
First-price Auctions, $\varepsilon = 27$
(I1: Solid, I2: Dot, I3: DashDot)

12 Bidders

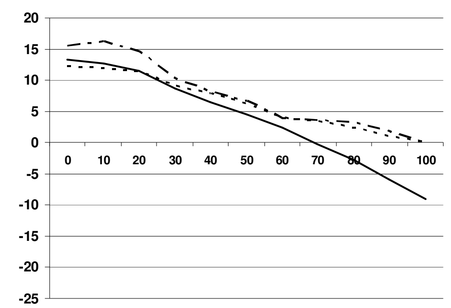

25 Bidders

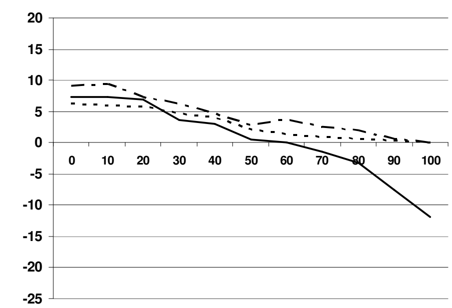

50 Bidders

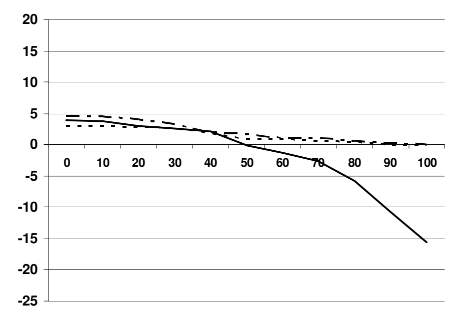

100 Bidders

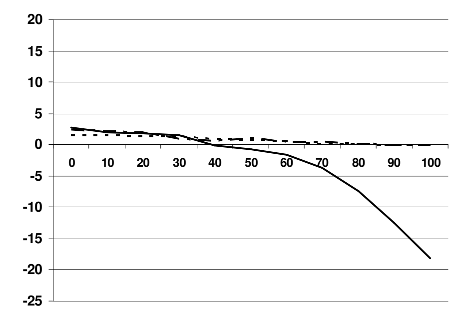

200 Bidders

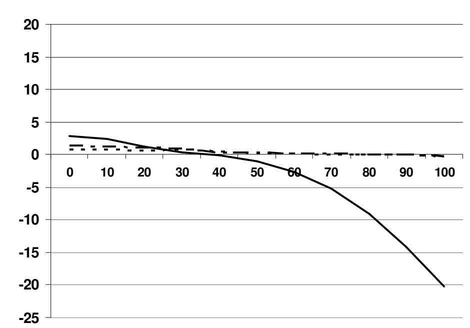

300 Bidders

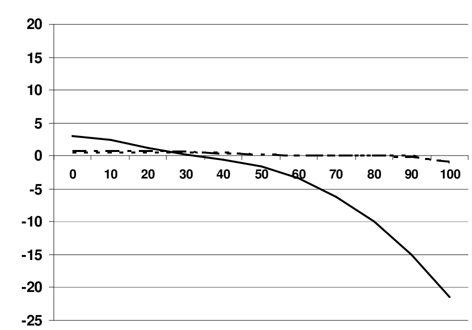

Figure 11. Efficiency: First-price auctions Efficiency by Common Value Percent (I1: Solid, I2: Dot, I3: DashDot)	
Four Bidders	Seven Bidders
$\varepsilon = 8$	$\varepsilon = 8$
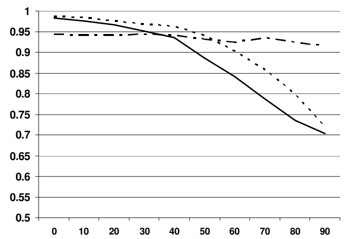	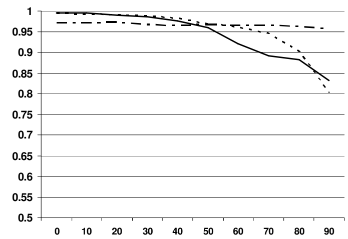
$\varepsilon = 12$	$\varepsilon = 12$
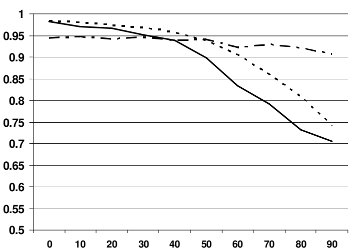	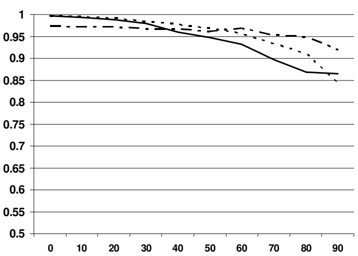
$\varepsilon = 18$	$\varepsilon = 18$
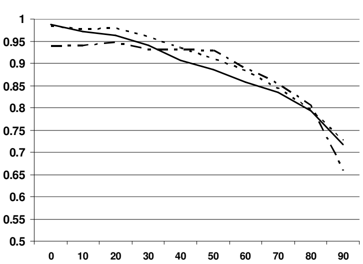	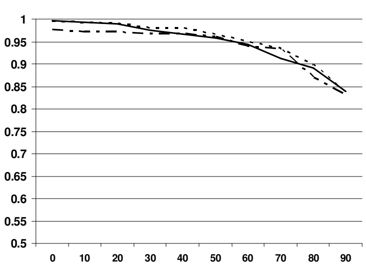
$\varepsilon = 27$	$\varepsilon = 27$
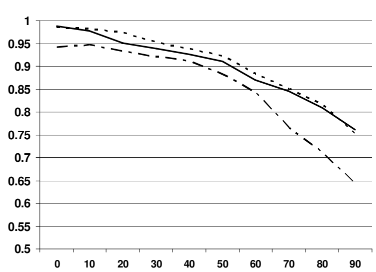	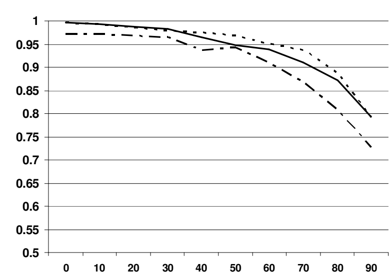

Figure 12. Value Multiplier: Second-price auctions
Value Multiplier by Common Value Percent
(I1: Solid, I2: Dot, I3: DashDot)

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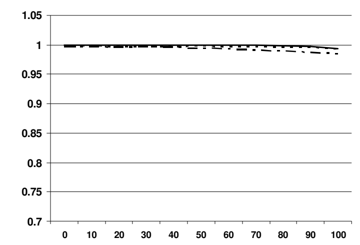	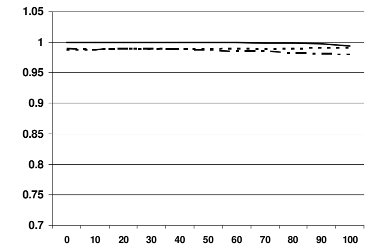
$\varepsilon=12$	$\varepsilon=12$
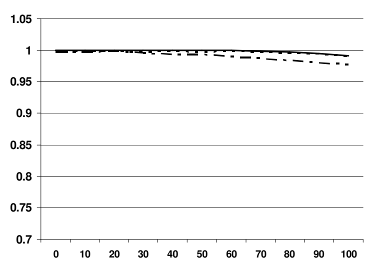	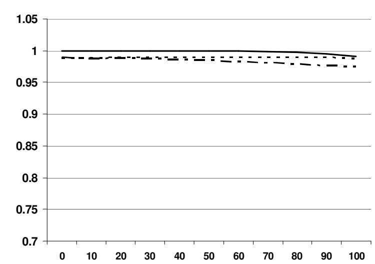
$\varepsilon=18$	$\varepsilon=18$
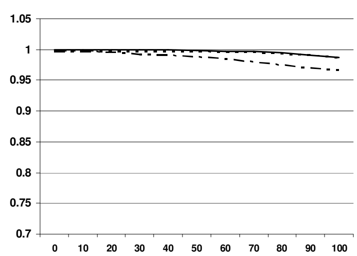	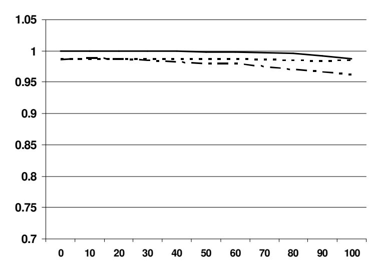
$\varepsilon=27$	$\varepsilon=27$
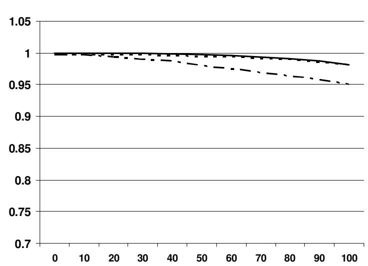	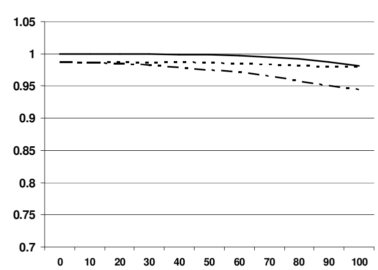

Figure 13. Profit: Second-price auctions
Profit by Common Value Percent
(I1: Solid, I2: Dot, I3: DashDot)

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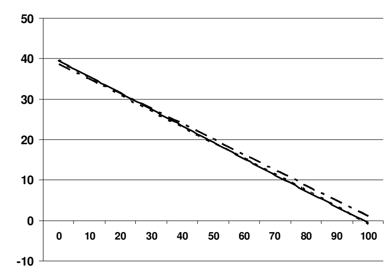	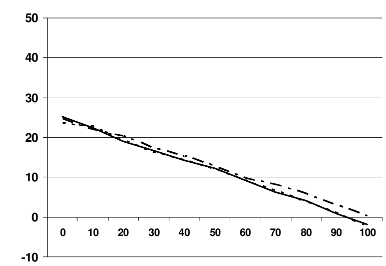
$\varepsilon=12$	$\varepsilon=12$
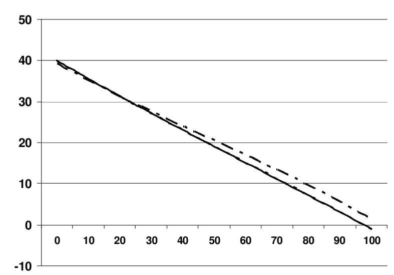	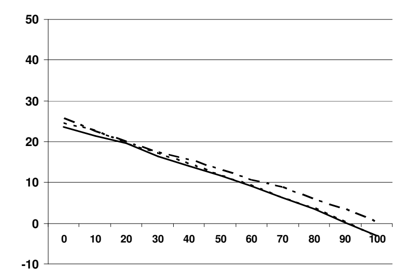
$\varepsilon=18$	$\varepsilon=18$
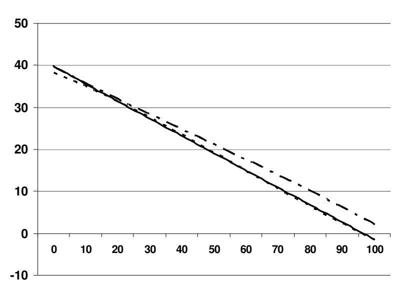	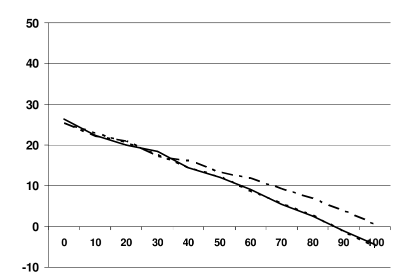
$\varepsilon=27$	$\varepsilon=27$
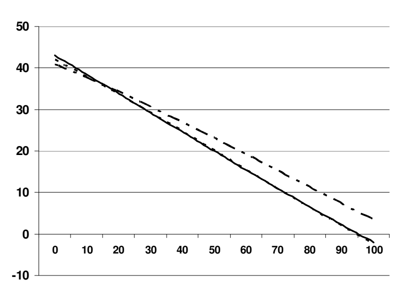	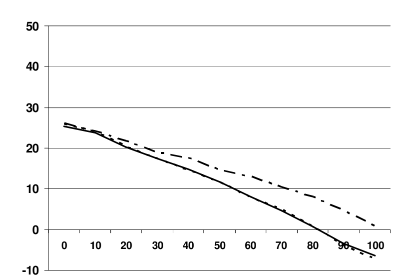

Figure 14. Efficiency: Second-price auctions
Efficiency by Common Value Percent
(I1: Solid, I2: Dot, I3: DashDot)

Four Bidders

$\varepsilon = 8$

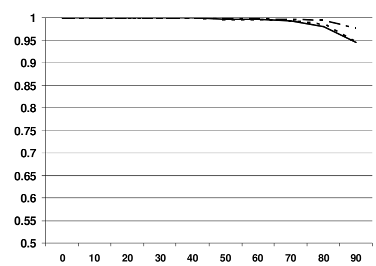

$\varepsilon = 12$

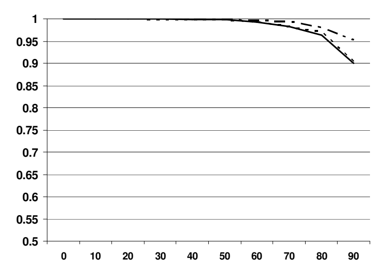

$\varepsilon = 18$

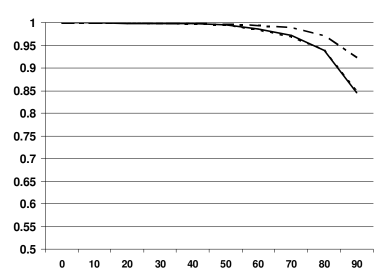

$\varepsilon = 27$

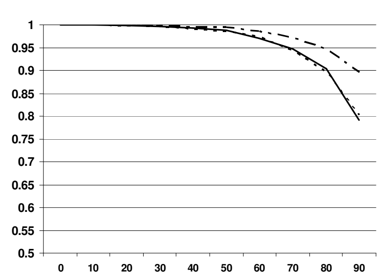

Seven Bidders

$\varepsilon = 8$

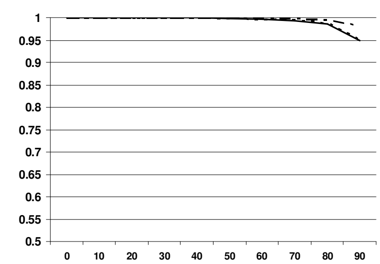

$\varepsilon = 12$

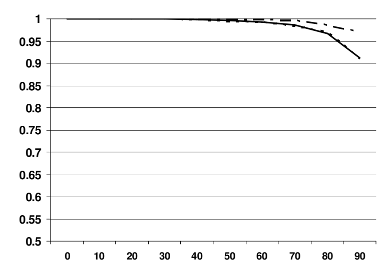

$\varepsilon = 18$

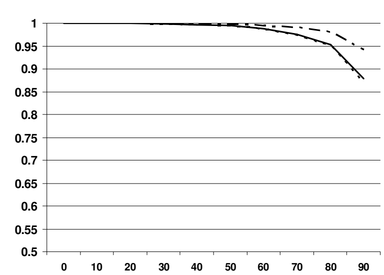

$\varepsilon = 27$

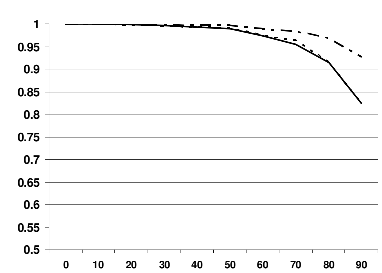

Figure 15. Revenue: First-Price Auctions
Revenue by Common Value Percent
(I1: Solid, I2: Dot, I3: DashDot)

<table>
<tr><th>Four Bidders</th><th>Seven Bidders</th></tr>
<tr><td>$\varepsilon = 8$</td><td>$\varepsilon = 8$</td></tr>
<tr><td>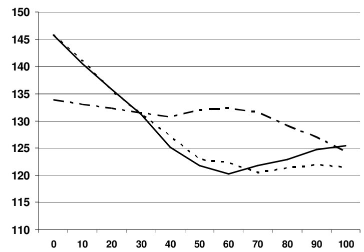
</td><td>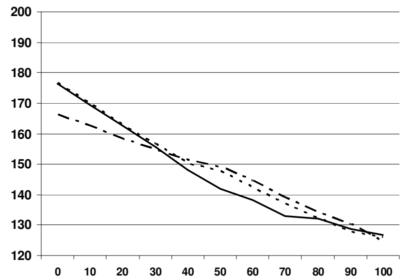
</td></tr>
<tr><td>$\varepsilon = 12$</td><td>$\varepsilon = 12$</td></tr>
<tr><td>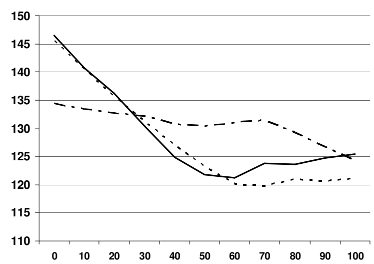
</td><td>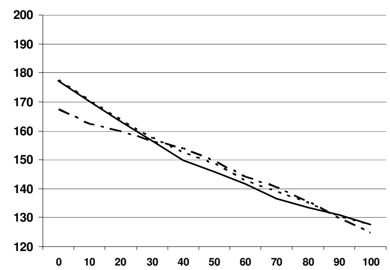
</td></tr>
<tr><td>$\varepsilon = 18$</td><td>$\varepsilon = 18$</td></tr>
<tr><td>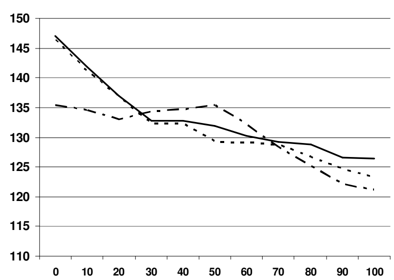
</td><td>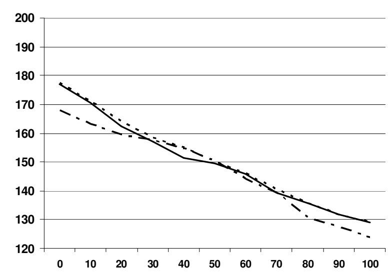
</td></tr>
<tr><td>$\varepsilon = 27$</td><td>$\varepsilon = 27$</td></tr>
<tr><td>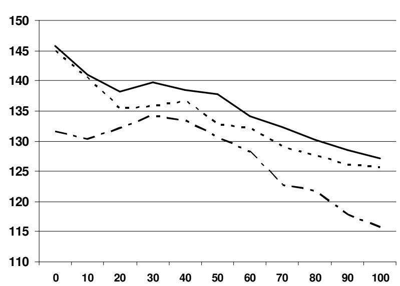
</td><td>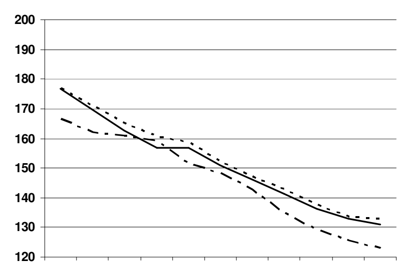
</td></tr>
</table>

Figure 16. Revenue: Second-Price Auctions
Revenue by Common Value Percent
(I1: Solid, I2: Dot, I3: DashDot)

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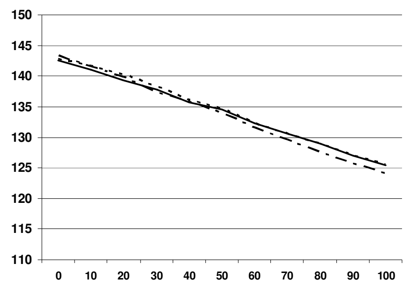	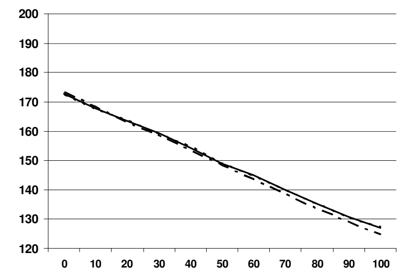
$\varepsilon = 12$	$\varepsilon = 12$
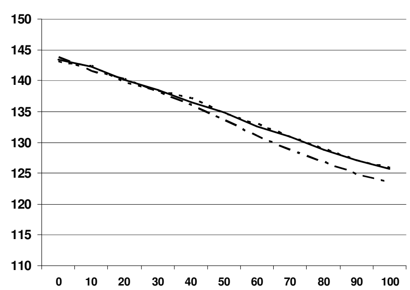	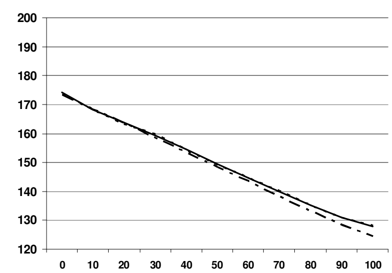
$\varepsilon = 18$	$\varepsilon = 18$
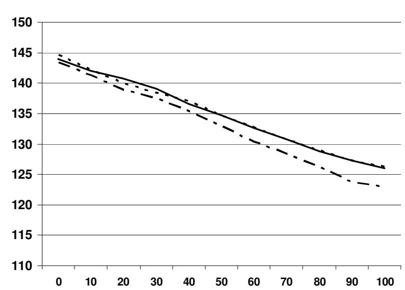	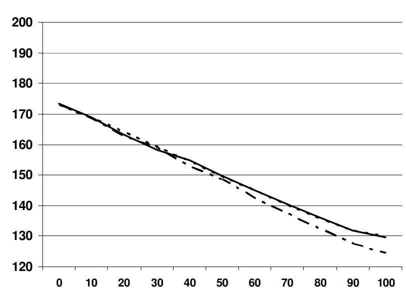
$\varepsilon = 27$	$\varepsilon = 27$
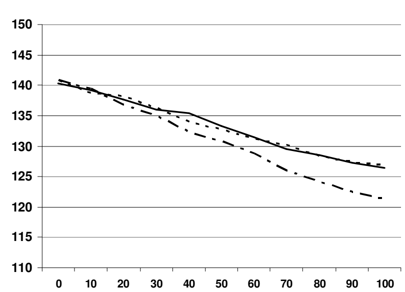	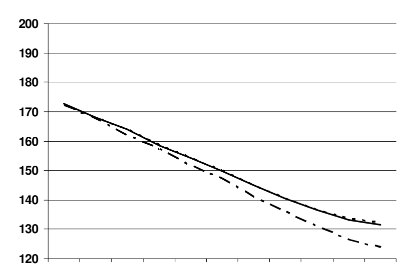

Figure 17. Revenue Effects of Revealed Common Value
Revenue Difference from No Information, by Common Value Percent
I3 Information
(Solid: $\varepsilon = 8$; Dot: $\varepsilon = 12$; Dash-Dot: $\varepsilon = 18$; Dash-Dot-Dot: $\varepsilon = 27$)

First-Price, Four Bidders

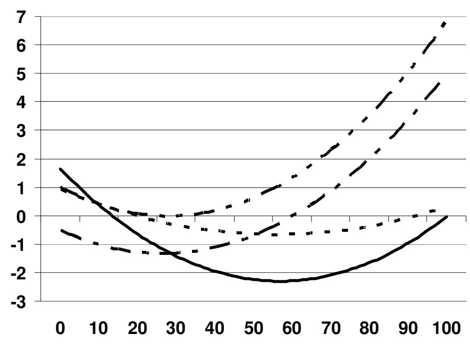

First-Price, Seven Bidders

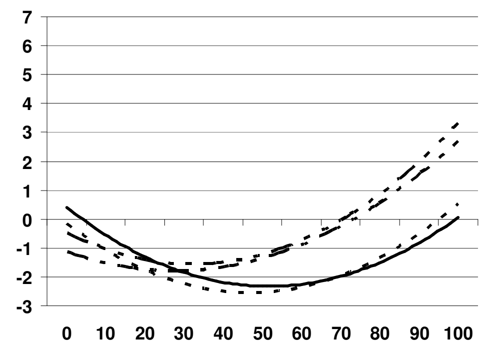

Second-Price, Four Bidders

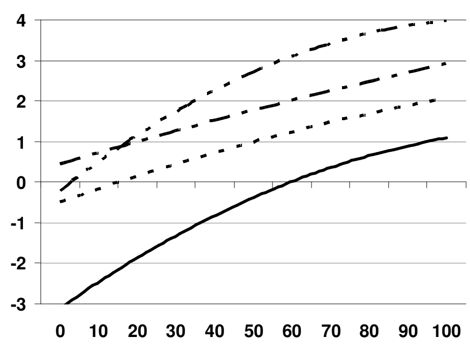

Second-Price, Seven Bidders