Comput Econ (2010) 35:189–209
DOI 10.1007/s10614-010-9199-x

What Format for Multi-Unit Multiple-Bid Auctions?

Agent-Based Simulation of Auction Performance and Nonlinear Bidding Behaviour

Atakelty Hailu · Sophie Thoyer

Accepted: 14 January 2010 / Published online: 4 February 2010

Abstract This paper uses computational experiments where bidders learn over nonlinear bidding strategies to compare outcomes for alternative pricing format for multi-unit multiple-bid auctions. Multi-unit multiple-bid auctions, in which bidders are allowed to submit multiple price-quantity bids, are promising mechanisms for the allocation of a range of resources. The main advantage of such auctions is to avoid the lumpy bid problem which arises when bidders can only compete on the basis of one bid. However, there is great uncertainty about the best auction formats when multi-unit auctions are used. The theory can only supply the expected structural properties of equilibrium strategies and the multiplicity of potential equilibria makes comparisons across auction formats difficult. Empirical studies and experiments have improved our knowledge of multi-unit auctions but they remain scarce and most experiments are restricted to two bidders and two units. Moreover, they demonstrate that bidders have limited rationality and learn through experience. This paper constructs an agent-based computational model of bidders to compare the performance of alternative procurement auction formats under circumstances where bidders submit continuous bid supply functions and learn over time to adjust their bids in order to improve their net incomes. The setting is for independent private values. We show that bidding behaviour displays more interesting patterns than is depicted in the theoretical literature and that bidding

This research was conducted when S. Thoyer was working at UWA (SARE) as an invited academic.

A. Hailu (✉)
School of Agricultural and Resource Economics, University of Western Australia,
35 Stiriling Highway, M089, Crawley, Perth, WA 6009, Australia
e-mail: ahailu@are.uwa.edu.au

S. Thoyer
ENSAM–LAMETA, 2 place Viala, 34060, Montpellier cedex 1, France
e-mail: thoyer@ensam.inra.fr

patterns depend on the interplay between heterogeneity in the bidder population and the degree of rationing in the auction. Results indicate that the three auction formats have similar performance for most levels of competition but that their performances differ when competition is weak. This ranking is dependent on whether the population of bidders is homogenous or heterogeneous.

Keywords Agent-based modelling · Agent-based computational economics · Computational experiments · Auction design · Multi-unit auctions · Multiple-bid auctions · Procurement auctions · Vickrey auctions · Uniform auctions · Discriminatory auctions

JEL Classification C900 · D440

1 Introduction

Multiple-bid multi-unit auctions are increasingly being used to allocate resource. Well known application areas include electricity markets (Wolfram 1998) and the allocation of Treasury bills or foreign exchange (Tenorio 1999). These auctions enable the auctioneer to sell (buy) several units of the same good through a tender process where bidders are allowed to submit bids in the form of demand schedules (supply schedules). These auctions are more flexible than single-bid auctions that limit bids to single quantity-price pair bids and thus help avoid the 'lumpy bid' problem inherent in single-bid auctions (Tenorio 1993). In the literature, the term 'multi-unit auction' is simply used to refer to the multi-unit multiple-bid auction. For brevity, we will use this shorter name in this paper.

There is great uncertainty about the performance of alternative auction formats in the case of multi-unit multiple-bid auctions. As a result, the choice between discriminatory (or pay-as-bid) and uniform price formats continues to be controversial. This is the case for Treasury bill auctions organized in Europe and in the US, where policymakers have switched from discriminatory to uniform payment formats in the hope of improving the allocative efficiency of the auctions and increasing budgetary revenues (Binmore and Swierzbinski 2000).

Economic theory does not provide much guidance on the relative efficiency of alternative formats in a multiple-bid setting. The uncertainty is even greater when the bidder population is heterogeneous or when bidder marginal values (costs) are not constant (Ausubel and Cramton 2002). Although the theory provides some insights into the possible structural properties of bidder strategies, the multiplicity of equilibria makes theoretical comparisons of different formats difficult. Empirical studies are relatively scarce (Wolfram 1998; Tenorio 1993). Experiments are restricted to very simplified settings in which two bidders compete for two units and have a flat demand. Moreover, they demonstrate that bidders have limited calculation capacities and learn from experience through repeated play instead of landing on the equilibrium strategies at the outset of the game. The case for experimental and computational approaches to further our understanding of multiple-bid auction design is therefore strong (Binmore and Swierzbinski 2000, p. 407).

Agent-based computational economics (ACE) provides a useful and inexpensive research tool for examining the performance of auctions under different contexts and for comparing the relative performance of different auction designs. It is increasingly used to complement theoretical and experimental studies in economics (Tesfatsion 2002).[1] This paper constructs an agent-based model to examine the performance of three alternative formats for multi-unit auctions: discriminatory, uniform and a generalized Vickrey pricing. The simulated auction market is cast as a procurement auction where a government agent buys services from a population of bidders with private independent values reflecting different production capacities and different cost structures. Bidders submit supply schedules indicating the amount of services they would provide at different prices. Auctions are repeated and bidders use genetic algorithm learning to update their individual bid functions with the objective of increasing their net incomes.

The performance of each auction format is evaluated for different levels of competition,[2] with the demand level from the purchasing agency ranging in magnitude from 12.5 to 75% of the aggregate capacity of the bidders. The comparative analysis is also undertaken for different levels of heterogeneity in the size and cost structures of individual bidders. The paper is organized as follows. We first review the various auction formats and the structural properties of bidding strategies implied by theoretical analysis. The generalized Vickrey is the only payment format for which equilibrium bidding strategies can be theoretically calculated: Ausubel (2005) has demonstrated that truthful bidding is a weakly dominant strategy. Discriminatory and uniform formats lead to a coordination problem. It is therefore analytically intractable to identify potential equilibria and to measure theoretically the relative efficiency of these auctions. In the third section, we develop an agent-based model (ABM) of boundedly rational bidders revising their bid choices using a genetic learning algorithm. In the fourth section, we present the results from the computational experiments and compare bidding behaviours, budgetary outlays and efficiency of allocation for the three formats. The simulation results provide some confirmation of analytical predictions. But they also indicate that bidding behaviours display more interesting patterns depending on the interplay between the nature of heterogeneity in the bidder population, the intensity of competition, and the type of the auction. We demonstrate that the different pricing formats lead to similar budgetary efficiency except when the level of competition is low. At low competition levels, the ranking of the formats depends on the nature of the bidder population. In the fifth section, we summarize the paper and draw some general recommendations.

2 Multiple-Bid Auctions

In this paper, we concentrate on the case of simultaneous procurement auctions, for multiple identical units, with independent private values. We also assume that the

[1] Tesfatsion's web site at http://www.econ.iastate.edu/tesfatsi/ace.htm is an excellent source of information on ACE research in economics.

[2] The level of competition here is measured as the ratio of demand by the government agency over aggregate supply by bidders. It reflects the degree of demand rationing but it does not include the impact of changes in the number of bidders.

number of units that the auctioneer wishes to buy is fixed (as opposed to a budget-constrained auctioneer or an auctioneer with a downward-sloping demand curve). Since most of the literature on multi-object multiple-bid auctions describes selling auctions, here we describe and summarize briefly what would be the equivalent predictions for a procurement auction in the case of a continuous rather than discrete bid specification. The continuous bid specification can be interpreted as one relating to the purchase of perfectly divisible units and was initially developed by Wilson (1979) who referred to it as "auctions of shares".

There are two main options for conducting multi-unit auctions—open-cry and sealed-bid. We will focus on the latter. In sealed bid auctions, each bidder is asked to submit multiple bids indicating the price he is willing to accept for different quantities sold. In effect, these multiple bids are equivalent to an inverse supply function. To describe the allocation procedures in the different auction formats, we need to define the concept of residual demand. Let's define the supply schedule of bidder i as $Q^i = S^i(b)$ with b the per-unit bids. We can then define the residual demand facing bidder i, $D^{-i}(b)$, as the total demand Q_d by the government agency less the sum of the amounts offered by all other bidders j for each level of bid price.

$$D^{-i}(b) = \max\left\{0,\, Q_d - \sum_{j\neq i} S^j(b)\right\}$$

In all auction formats, the allocation problem is solved by awarding each bidder the quantity Q^{*i} at which his supply schedule intersects his residual demand.

$$Q^{*i} = D^{-i}(b*) = S^i(b*)$$

However, the three formats differ in the calculation of the payments for the winners (see Fig. A1 in Appendix):

- In a discriminatory auction, each bidder is paid an amount equal to the sum of his actual winning bids (or the area under his supply schedule up to Q^{*i}).
- In a uniform-price auction, all units sold earn the clearing price equating aggregate supply to demand. Therefore, infra-marginal units receive payments that are higher than the corresponding bids.

In a generalized Vickrey auction, each successful bidder is paid the entire area under the residual demand up to Q^{*i}. This form of payment is the generalization of the second price Vickrey payment from the single unit auction case. Each winner is paid the amount corresponding to what the auctioneer would have had to pay to other bidders, had he not participated in the auction.[3] Since his payment does not depend in his bids, it can be predicted intuitively that his dominant strategy is to bid truthfully, i.e. to bid

[3] A simple example can be given in the discrete case. Let's assume an auction with two bidders, each with a supply of four units. Their bid schedules are (1, 3, 6, 7) and (2, 4, 5, 9) for bidders 1 and 2 respectively. Demand by auctioneer is 4 units. Bidder 1 will sell two units and will get paid R1 = 9 + 5 = 14 and bidder 2 will also sell two units for a total payment R2 = 7 + 6 = 13.

Table 1 Structural properties of equilibrium strategies for various formats of procurement multi-unit auctions

Sealed-bid format	Structural property of equilibrium strategies and efficiency[a]
Discriminatory	Scope for "flat supply"[b] and and for "supply inflation"[c] ; Inefficient allocation
Uniform-price	"Supply inflation[b]"; Coordination at a high price equilibrium; Inefficient allocation
Generalized Vickrey	Truthful bidding is a weakly dominant strategy; Efficient allocation

[a] Efficiency here refers to the social opportunity cost of the allocation of resources. An efficient reallocation is obtained when goods are bought (sold) from (to) bidders with the lowest marginal production costs (highest marginal utility)
[b] The bid curve is almost flat, above the true cost curve
[c] Equivalent to demand reduction (or bid shading) in a selling auction: bidding is sincere on the first unit then there is differential increasing overbidding

his true opportunity costs. This has been demonstrated for the single unit case by Vickrey and more recently by Ausubel for the multi-unit case (Ausubel 2005). However, it is possible for bidders to coordinate bidding for clearing prices that are above those implied by cost curves even under this pricing scheme as our results demonstrate.

2.1 Equilibrium Strategies, Efficiency, and Revenue

For multiple-bid auctions, no closed form expressions of the bidding strategies are available in the general case and most studies have therefore focused on the structural properties of the equilibrium strategies (summarized in Table 1). Wilson (1979), Back and Zender (1993), Engelbrecht-Wiggan and Kahn (1998), Tenorio (1999) and Ausubel and Cramton (2002), have analyzed the outcomes of different multi-unit auction formats and shown that the revenue equivalence theorem does not extend to the case of multiple-bid auctions.

Using simplified settings, the studies demonstrate[4] the issue of bid shading (or demand reduction) associated with a uniform-price multiple-bid selling auction. Their findings agree: although it is a dominant strategy to bid truthfully for the first unit (or, in the continuous case, when quantity tends to zero), it is efficient for the bidders to shade their bids for additional quantities. Moreover, the amount of bid shading increases with quantities offered. "The reason for this differential shading is that the incentive to win units at any price below marginal value is offset by the incentive to reduce the price paid on infra-marginal units that are won anyway" (Ausubel and Cramton 2002, p. 23). The latter becomes increasingly important when quantities increase, which explains the increasing bid shading. The consequence is that bids no longer correlate with opportunity costs, leading to efficiency losses. All demonstrations can be carried over with no restriction to the procurement case. (See last column of Table 1). This phenomenon has been verified empirically by Wolfram (1998) who

[4] All demonstrations are made for a model where bidders' values are private and independently distributed and ex-ante symmetric (the distribution of information is uniform across bidders in the pre-auction situation).

conducted an econometric analysis of "supply inflation" in the electricity procurement auction in England and Wales.

It is also demonstrated, in a setting with two bidders and two units, that there is an incentive, in a discriminatory format, to submit flatter supply curves than in a uniform price auction. If bidders are risk neutral, submitting entirely flat supply curves is a possible equilibrium (Back and Zender 1993), although drastic demand reduction is also a possible outcome, especially when the difference in the marginal values of the two units is high (Engelbrecht-Wiggan and Kahn 1998; Krishna 2002). Since there are different classes of equilibrium strategies, it is difficult to analyze how bidders coordinate or even compare the efficiency of the two formats. Tenorio (1993) compares the discriminatory and uniform price auction formats in the Zambian foreign exchange market, which successively implemented the two. He demonstrates that the uniform price auction yields higher average revenue to the auctioneer. His case study includes some form of affiliated values, however.

Only in the generalized Vickrey payment is truthful bidding a weakly dominant strategy, resulting in efficient allocation. On the other hand, for uniform and discriminatory formats, only increased competition can lead to the reduction of strategic behaviour and to more truthful bidding (Ausubel and Cramton 2002; Swinkels 1999). The generalized Vickrey is rarely employed in practice because the payment rule is not easily understood by bidders. Therefore, it is crucial that more results be produced to help decision-makers to make a choice between discriminatory and uniform payments. In particular, two issues are of interest for them: to assess how these formats compare for different levels of competition and for different types of heterogeneity in the bidder population. There are two sources of heterogeneity which are worth exploring. The first one is associated with heterogeneity in the supply capacity of bidders and the second is related to heterogeneity in supply costs. The theoretical literature does not provide answers on the impact of these types of heterogeneity. There is a need, therefore, to turn to experiments and simulations in order to further our understanding.

2.2 Experiments

One way to make up for the unavailability of predictive analytical results is to turn to experimental methods. Experiments can be designed to confirm theoretical results but their function can also be to carry a problem beyond the analytical capabilities of theoretical analysis or to fill a void when theory is incomplete or not available.

Given both the weakness of the theory on multiple-bid auctions and the increasing use of such auctions in economic life, a growing number of researchers have tried to investigate bidding behaviour with laboratory experiments. Alemsgeest et al. (1998) demonstrate that, in the two unit case, an ascending clock auction (i.e. discriminatory in the sealed bid setting) generates less revenue than the uniform sealed-bid auction, due to strategic bid shading. Kagel and Levin (2001) also compare uniform-price sealed bid and open ascending auctions, with a real player with flat demand for two units playing against a robot with unit demand. Their findings also highlight the issue of demand reduction and show that in the open format, bids converge towards equilibrium values more rapidly than in a sealed bid format, as if the "clock could enhance learning". They confirm that the Ausubel format leads to more truthful bidding.

Engelmann and Grimm (2003) compare bidding behaviour under five auctions formats and for a case where two bidders with flat demand curves compete to buy two units. Their experiments demonstrate that demand reduction is more acute in uniform open than in uniform sealed bid auctions, and that the Ausubel format eliminates bid shading. They also find that, in clear contrast to theoretical prediction, bidders in discriminatory auctions place substantially different bids on the first and second unit, even when their valuations for the two units are close. They suspect that it might be due to myopic zero profit aversion on the part of the bidders but do not prove it.

Most experiments are conducted under simple settings. Human experiments can also be extremely costly and complicated to run when exploring issues such as competition or heterogeneity amongst bidders. One way to deal with these difficulties is to employ computational experiments (Duffy 2006).

3 The Modelling of Bidding Strategies with Artificial Learning Agents

Unlike conventional or deductive approaches, the starting point in agent-based computational economics (ACE) is the specification of agent attributes and behaviours rather than equations or equilibrium conditions describing the system under study. Therefore, ACE is suited to the study of systems where modelling outcomes can be gainfully enriched through the explicit incorporation of phenomena like agent heterogeneity and through the relaxation of other restrictive assumptions that are normally imposed in theoretical analysis for tractability purposes (Epstein and Axtell 1996; Tesfatsion 2002). Studies applying ACE to the study of auctions include Andreoni and Miller (1995), Nicolaisen et al. (2001), Bower and Bunn (2001), Bunn and Oliveira (2001), Hailu and Schilizzi (2004), and Hailu and Thoyer (2006, 2007). The model presented in this paper differs from previous models because it tackles the issue of multiple-bid auctions: competing bidders submit continuous bid supply functions in the auction and employ genetic algorithms to update their bidding strategies.

3.1 Structure of Agent Based Model

Our auction model has a population of agents selling goods in a sealed-bid auction to a single buyer, the government agent. The government agent has a fixed target or demand level. Each seller is characterized by a (true) non-decreasing supply function and a given supply capacity indicating the maximum amount of good it has for sale. The government agent does not know the true supply functions of the different bidders and makes selection based on submitted or declared supply bid functions. Over time, sellers learn to choose, in a repeated process, the supply bid functions that maximize their expected net incomes.

Each auction round involves two stages. In the first stage, the government collects bid functions from the sellers, calculates the residual demand for each bidder and determines the equilibrium quantities bought from each of them at the intersection of their bid supply and their residual demand. In the second stage, payments to individual bidders are determined according to the auction format in use. Sellers use the results of the auction to compute their net incomes and to update the probabilities with which they choose their bid strategies for the next round. The strategy choice probabilities of a bidder therefore depend on his opportunity costs as well as on the history of choices he has made and rewards obtained for those choices.

For the sake of simplicity, it is assumed that the true supply function of a seller i is linear and can be written as:

$$P_i = a_i^0 + b_i^0 Q_i$$

with $0 \leq Q_i \leq ms_i$, where ms_i is the capacity of bidder i, a_i^0 is his entry price (equivalent to his fixed costs) and, b_i^0 is the supply slope.

3.2 Seller Choice Strategies and Learning Algorithm

We allow bid curves to be non-linear.[5] To accommodate different types of nonlinearities, we use a general Box–Cox functional form for these bid curves. The learnt bid curve is assumed to be represented as follows:

$$\beta_i(Q_i) = a_i + \frac{b_i(Q_i^{c_i} - 1)}{c_i}$$

with $\beta_i(Q_i)$being the strategic bid of player i. If c_i is equal to unity, this form reduces to a linear bid curve with entry price of $a_i - \frac{b_i}{c_i}$ and a slope of b_i. For example, truthful bidding would be represented by a bid curve with the following features: $c_i = 1$, $b_i = b_i^0$ and $a_i = a_i^0 + b_i^0$. Bid curves with different types of curvature and entry prices can be generated through the bidder's selection of values for these three parameters. The bid curves could also combine steep with flat portions, making Box–Cox formulation highly flexible. A set of curves with different curvature and the associated parameter values are shown in Fig. 1 below as an example of the flexibility of this formulation for learnt bid curves.

Thus, there are three dimensions to the seller's choice strategy: a_i, b_i, and c_i. The learning algorithm described below will allow bidders to progressively explore different combinations of these learnt parameters to retain the best values based on the performance of past bids.

A constraint is imposed on the choice of strategies so that the chosen bid function does not have any section falling below the true cost function (bidding below true costs

[5] We are thankful to the reviewer who suggested nonlinear bidding curve as well as an alternative learning algorithm (genetic algorithms). The original version of the paper limited learnt bids to linear curves and used reinforcement learning. As we discuss below, the genetic algorithm allows for a better refinement in the learning space and is computationally less demanding.

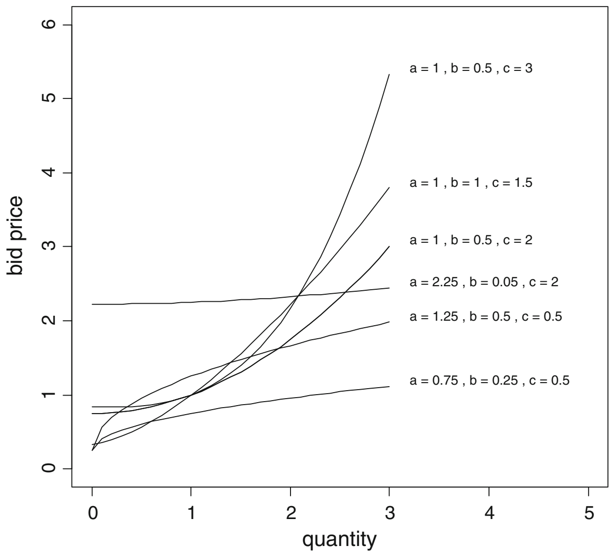

Fig. 1 Nonlinearity and choice of Box–Cox parameter values

is a dominated strategy and is therefore not included in the bidder's choice set). We also impose that bidders won't use extreme overbidding strategies by restricting their bids to less than ten times the marginal cost of supply of the most expensive unit by the most expensive supplier. This wide range was allowed to avoid artificially limiting the learnt bids to a narrow region.

3.3 The Learning Algorithm

Different learning models have been developed over the last several decades. A typology of learning models presented in Camerer (2003) shows the relationship between these learning algorithms and how certain variants are special cases of others. The models differ in terms of their information requirements. Learning algorithms that do not require that bidders have knowledge of forgone payoffs associated with strategies that they did not select are the most useful for the auctions studied here. This is because bidders learn from individual experience based on the results from their previously utilized strategies. Reinforcement-learning algorithm (Roth and Erev 1995; Erev and Roth 1998) is a candidate and has been used in several auction studies (e.g. Nicolaisen et al. 2001; Bunn and Oliveira 2001).

Genetic algorithm (GA) learning has similar minimal information requirements. Compared to reinforcement learning, the genetic algorithm (GA) has two advantages. First, reinforcement learning is an algorithm for selecting desirable strategies from a fixed set of alternatives. Tested strategies that generate positive rewards have a higher probability of being selected. This strategy to learning, while useful, can introduce bias in that strategies are played because they were found to be rewarding in previous

Table 2 Description of the four populations of bidders

	Bidders	Capacity	Entry price (a^0b^0/c^0)	Supply slope b^0
Population 1	8 identical bidders	0.5	0.5	0
Population 2	4 small	0.25	0.5	0
	4 large	0.75	0.5	0
Population 3	8 identical bidders	0.5	0.5	0.5
Population 4	2 small-low cost	0.25	0.5	0.25
	2 small-high cost	0.25	0.5	0.75
	2 large-low cost	0.75	0.5	0.25
	2 large-high cost	0.75	0.5	0.75

rounds and not because they are the best possible.[6] The genetic learning algorithm does not suffer from such bias. In our simulations, we find a much greater variability in auction outcomes across replications when reinforcement learning used. Second, GA does not require the specification of a discrete set of strategies (or parameter tuples) over which bidders learn. With GA, one only needs to specify the range for the parameter values but learnt optimal parameter values are considered as continuous numbers and can take any value in the range. As a result, the simulation results are not 'contaminated' by the artificiality of the discreteness in the parameter value sets. Of course, one can employ fine steps to reduce this problem in the reinforcement learning; however, that increases greatly the number of feasible strategies slowing the simulation dramatically.

A population of 100 chromosomes is used for each bidder in the GA. Both cross-over and mutation operations are used to evolve the population. The algorithm is elitist, with the fittest 5% retained and used in the cross-over operator to generate the remaining members in the next generation of chromosomes. The mutation operator is the non-uniform mutation operator described in Michalewicz (1996, pp. 103–104), allowing mutation rates to go down as the simulation progresses.

4 Simulation Results and Discussion

4.1 Simulation Settings

Bidding under the Vickrey, discriminatory and uniform auctions was simulated in our computation experiments for different levels of competition and heterogeneity in the size and cost structures of the bidder population. The level of competition was varied by changing demand while keeping aggregate supply capacity constant at 4.0. Six levels of demand, ranging from 0.5 to 3.0, were used. These correspond to 12.5, 25, 37.5, 50, 62.5 and 75% of available aggregate capacity. Auction performance was simulated for the following four bidder populations (Table 2):

[6] We are thankful to the reviewer for attracting our attention to this shortcoming.

Population 1: a homogeneous population of bidders with similar capacity and costs, each with a flat supply cost curve ($a^0 = 0.5, b^0 = 0, c^0 = 1$, ms = 0.50)

Population 2: a population where 50% of the bidders are small capacity bidders (ms = 0.50) and 50% are large capacity bidders (ms = 0.75), each with a with flat supply cost curves ($a^0 = 0.5, b^0 = 0, c^0 = 1$)

Population 3: a homogeneous population with rising marginal cost curves ($a^0 = 1.0, b^0 = 0.5, c^0 = 1$, ms = 0.50), and

Population 4: a highly heterogeneous population consisting of 4 small (each with a capacity of 0.25) and 4 large bidders (each with a capacity of capacity 0.75) all with rising marginal cost curves ($b^0 = 1$). Within each group, half of the bidders have a supply slope of 0.25 and the other half have a supply slope of 0.75.

4.2 Convergence of Simulated Strategies

Simulations are run over a large number of rounds until bidders have had ample time to learn and adjust their bids. The genetic algorithm learning is allowed to run until auction outcomes have converged, with convergence defined as the absence of change in outlay for a period of at least ten rounds. In the majority of the cases, the median number of rounds required for the results of the genetic algorithm learning to converge was 104 rounds.

Since all simulations are undertaken with 100 replications using different random seeds, the strategies and the performance of the auctions are evaluated based on the average values obtained from these 100 replications. First, the bidding strategies observed are presented and compared to available theoretical predictions. Then the performance results of the three auction formats are compared in terms of budgetary outlays per unit and in terms of the social efficiency of allocation as measured by the cost of production per unit.

4.3 Results for a Homogenous Population of Bidders with Flat Supply

This population has the simplest structure. The bidders are homogeneous in capacity and the level of marginal cost is constant. Since the level of marginal cost is constant and identical for all bidders, the competition at any price level involves the entire aggregate capacity.

The simulated bidding behaviour for this population indicates that bidding is almost truthful with regards to entry price (i.e., entry price is very close to 0.5) in all auction formats when rationing is tight (i.e. the ratio of demand to aggregate capacity is less than 50% or that demand is less than 2.0). Bidder entry prices are progressively inflated as the demand level rises with the degree of overbidding on the entry pricing being higher with the discriminatory auction, relative to the other auctions.

Beyond the entry price, all auction formats lead to overbidding, with bid prices lying clearly above the costs. However, the nature of the overbidding and its relationship to demand levels is dependent on auction format. Bid curves become flatter with discriminatory pricing, with the lower portions of these curves becoming almost horizontal, as the level of demand increases. Completely flat curve bidding is the strategy used at the highest demand level under the discriminatory auction. This is consistent with the literature (Engelbrecht-Wiggan and Kahn 1998), which indicates that, for discriminatory auctions, the expected structural properties of equilibrium strategies are higher entry price and flat bidding on the subsequent units (*high flat bidding* henceforth). A flattening of the supply curve improves bidder revenue as the prices received for infra-marginal units are brought closer to that of the marginal unit. So under this auction format, there is an incentive for bidders to organize their bids at the auction clearing price. However, when competition is tight, these flat supply curves are susceptible to price undercutting by rivals and the bidder can easily be priced out with small changes in others' bids. Under such circumstances, bidders have an incentive to ensure winning by bidding with truthful entry prices while at the same time earning positive net income by inflating prices on subsequent units. This is what is referred to as *supply inflation* in the literature. Thus, supply inflation under discriminatory pricing is not precluded by theory (Krishna 2002), and it is exactly what is observed in our simulations for the case of this homogenous populations for entry prices when demand is very low but also at the end portions of flatter bid curves (for higher quantity levels) when demand is high.

The relationship between bidding curve slopes and demand exhibits a similar pattern under the Vickrey and uniform pricing formats. For low demand levels, bidding strategies are comparable to the ones observed in the discriminatory format (supply inflation), whereas for higher demand levels, bid curves are steeper (and include no flat portions) compared to those in the discriminatory auction. That is, both the Vickrey and the uniform auctions involve *supply inflation* strategies. In these auctions, supply inflation allows the marginal bidder (the one setting the clearing price) and all other winners to make greater profits. It is the behaviour expected by theory for uniform auctions. However, the propensity to bid with shallower curves that is evident at higher demand levels can be explained by the fact that each bidder has a lower probability of being the price setter as demand increases. Infra-marginal bidders have no incentives to inflate their bids as it does not impact the price they get. However, a flatter curve helps prevent the clearing price from slipping down greatly if competitors lower their bid curves. Therefore, the shallow bid curves combine features of both the supply inflation and high flat bidding strategies.

4.4 Results for a Population of Bidders with Flat Supply but Heterogeneous Sizes

The introduction of size heterogeneity has impact on bidding behaviour. We observe differences between the strategies adopted by small and high capacity bidders, especially at higher demand levels. Large bidders participate more in price setting compared to their smaller competitors. Heterogeneity has another important effect besides this. Larger bidders tend to use bid curves that have substantial shallow or completely flat bid sections in the case of discriminatory auctions. In the Vickrey and uniform auctions, large bidders adopt strategies that combine supply inflation and high flat bidding by using bid curves that have shallower sections (lower quantity levels) but rise rapidly. These reflect the predominant role of large suppliers in price setting.

4.5 Results for Homogeneous Population of Bidders with Upward Sloping Supply

When the marginal cost of production is positively sloped, bidding strategies are very similar across all formats. First, all bidding strategies involve overbidding across all demand levels. For example, bid entry prices are more than twice true entry costs under all pricing formats. Second, bidders adopt higher bid entry prices with flatter bid curves as the level of demand increases. However, supply curves include completely flat portions (for lower portions of the bid supply curve) only in the case of the discriminatory auction. The latter behaviour conforms to the theoretical predictions as explained above and is explained by reduced risks of being completely undercut by competitors when the demand level increases together with the incentives for overbidding on all quantities under discriminatory pricing. In the other two formats, the absence of relatively truthful entry prices of the type observed for populations 1 and 2 is an indication of the fact that bidders find aggressive pricing more rewarding when cost curves are increasing and, therefore, the degree of competition among units is not as intense as it is when cost curves are flat (i.e., marginal costs are constant).

4.6 Results for Population of Bidders with Heterogeneous Sizes and Supply Slopes

There are four types of bidders in this population, with bidders distinguished by size (0.25 and 0.75) and slope of supply cost curve (0.25 or 0.75). All have the same entry price of 0.5. As in the case of the homogeneous population with rising costs, entry bid prices are highly inflated for all auction formats and all bidders, regardless of size or cost structure. However, bidding patterns differ among bidders and also between auction formats even for the same type of bidder. In particular, bidding patterns for discriminatory auction are different from those for the other two auctions. The difference between bid curve patterns in the Vickrey and uniform auctions are minor.

As with the results for the other populations, truthful bidding is almost non-existent with a discriminatory auction. The most predominant overbidding is the theoretically expected "high flat" supply bid which is the strategy adopted by winning bidders. When demand is low (below 1.5), it is the bidders with flatter cost curves (both small and large) that win the auctions and are influential in setting the clearing price. As demand increases to 1.5 and beyond, bidders with steeper cost curve become successful and switch from "supply inflation" strategies to "high flat bidding", with large bidders who have steep cost curves adopting bid curves that have mainly flat portions with steeper upper portions. This behaviour is consistent with what has already been observed in more homogeneous population settings. It confirms that supply inflation is a strategy adopted by bidders who are likely to be totally priced out by rivals.

The role in pricing setting of bidders with cost advantages (flatter cost curves) declines as demand level increases (beyond 2.0) for both the Vickrey and uniform auctions. At lower demand levels these bidders are the price setters. Price setting in the intermediate demand level of 1.5 involves all bidders, but with bidders that have cost advantages successful in selling all their capacities. For higher demand levels, price setting is primarily done by higher cost suppliers. Low cost bidders tend to bid with curves that are shallower and lie under the price line. Therefore, the pattern with which the target demand is sourced among the bidders is different between discriminatory and the other auctions. In Vickrey and uniform auctions, the capacity of low cost suppliers tends to be exhausted before other bidders become successful as demand rises.

4.7 Bidding Strategies: A Summary

Examining the pattern of bidding behaviour across the different populations, the following general observations can be made.

(1) Overbidding is the norm under the discriminatory auction. The theoretically predicted high flat bidding (overbid on entry price and low or zero slope) becomes predominant when competition is weak. For a bidder, competition is weak either because demand is higher or because it has cost advantages relative to marginal units in the auction. As a result, high flat bidding is observed when demand is high regardless of the population, although the bidding curve might include a steep end segment. Supply inflation (truthful bidding on entry price but higher supply slopes) is the predominant overbidding strategy under the discriminatory auction when demand is low. This behaviour allows the bidder to minimize the risk of being completely undercut by competitors. Therefore, this bidding strategy is rational when the level of competition is intense because demand is low and/or marginal cost is constant pitting every unit for sale against every other. When marginal costs are rising, bidding under discriminatory auction involves higher entry prices with increasingly flatter curves.

(2) Under both the Vickrey and uniform pricing formats, the typical strategy when marginal costs are constant is relatively truthful entry prices with rising bid curves ("supply inflation"). When marginal costs are rising, however, bid entry prices involve overbidding, with the degree of overbidding rising with the demand level. Unlike in the case of the discriminatory auction, some bidders (especially small ones) might have bid curves lying below the auction clearing prices, i.e. the small (and/or more competitive) bidders free ride on the price setting achieved by larger (and/or less competitive) ones. As a result, the number of bidders participating in price setting is lower under these auctions compared to under discriminatory pricing when the population of bidders is heterogeneous. Coupled with the incentives that bidders have for inflating prices under the uniform auction, this phenomenon can generate higher auction clearing prices relative to other formats when competition is weak.

(3) With nonlinear bidding, a hybrid strategy involving elements of both supply inflation and high flat bidding can be the best strategy. This involves the use of a shallow or flat curve at lower demand levels followed by a rapidly rising upper section. Such a bid curve pushes up auction clearing prices if the collective bidding of the population is such that the upper section is participating in price setting. The flatter lower section prevents prices from sliding down. In the

case of the discriminatory auction, the flatter section also improves revenue by raising payments for infra-marginal units.

4.8 Nash Equilibrium Property Tests of Learnt Strategies

The learnt bids were tested for best reply properties by checking, for one bidder at a time, if there is no other strategy that allows the bidder to increase his net income. The results on NE test pass counts indicate that bidders can coordinate bid curves using strategies that are not necessarily Nash equilibrium strategies. The percentage of learnt strategy choices that constitute a Nash equilibrium strategy set is higher in heterogeneous populations than among homogenous populations, reflecting the impact on strategies of size and cost differences in the former case. They are also highest among Vickrey auctions, followed by uniform auctions. For discriminatory auction, the pass rates are lower. This reflects that under the Vickrey and uniform auctions, a winning bidder's revenue is more likely to be determined by other bidder's bidding strategies. Further, NE test pass rates tend to be higher among bidders with cost disadvantages when demand levels are low, regardless of the auction format. This reflects that these bidders are priced out in the auction and there is very little that a unilateral move could alter in terms of bidder profits. NE pass rates are also high among bidders with cost advantages when demand is high under Vickrey (and to a lesser degree under uniform pricing), as these bidders tend to participate less in price setting and free ride on the prices set by other bidders; therefore, unilateral moves would not improve profits for these bidders under these circumstances because they are selling their full capacity at prices set by other bidders. The NE pass rates are summarized in Table 3.

Overall, more bid curve choices constitute NE strategies when the level of competition is tighter (demand is in the lower ranges). These figures confirm that the multiplicity of possible equilibria induce coordination failures. Bidders learn to coordinate their overbidding when demand is high. However, these coordinated bids do not constitute best reply strategies as individual bidders can improve their net incomes through unilateral deviation. With the Vickrey, for example, coordinating bidding choices with others so that the clearing price is high benefits all bidders. However, a bidder's net income might improve (but would never go down) if it reverts to a more truthful bidding strategy given the choices of its competitors. In the discriminatory auction, a bidder's revenue depends on his own bid, providing the bidder with the incentive to deviate if other bidders were to keep their bids fixed. Under the uniform, the bidder's revenue can depend on its own bidding strategy. Therefore, a bidder might have the same incentives to defect or 'free ride' on the price coordination choices of other bidders.

4.9 Auction Performance

The performance of an auction is measured using the following two criteria: budgetary outlay and the total production costs (allocative efficiency). The former measures the monetary transfers from the buyer to the bidders. The latter measures the

Table 3 Percentage of learnt bid strategies that pass Nash equilibrium tests

	Demand					
	0.5	1.0	1.5	2.0	2.5	3.0
Population 1						
Vickrey	77	64	49	40	43	61
Discriminatory	74	47	23	24	35	19
Uniform	66	61	49	31	34	35
Population 2						
Vickrey	80	62	50	44	52	67
Discriminatory	69	44	26	33	18	3
Uniform	69	59	43	40	38	31
Population 3						
Vickrey	65	59	52	42	31	59
Discriminatory	64	47	31	18	27	21
Uniform	61	53	51	44	32	32
Population 4						
Vickrey	75	69	71	68	64	70
Discriminatory	66	54	35	26	31	25
Uniform	70	64	68	55	52	36

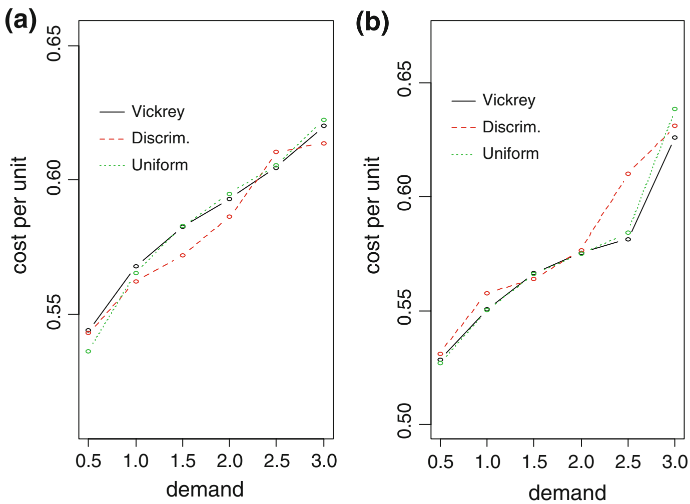

Fig. 2 Allocative efficiency, demand level and auction pricing format. **a** Bidder population 3, **b** bidder population 4

auction's social cost efficiency. From a social welfare perspective, the auction outcomes are more efficient if the product is purchased or sourced from lowest cost sources. This second criterion is relevant only for the last two populations (3 and 4) as any allocation is equally efficient when marginal costs are constant and identical

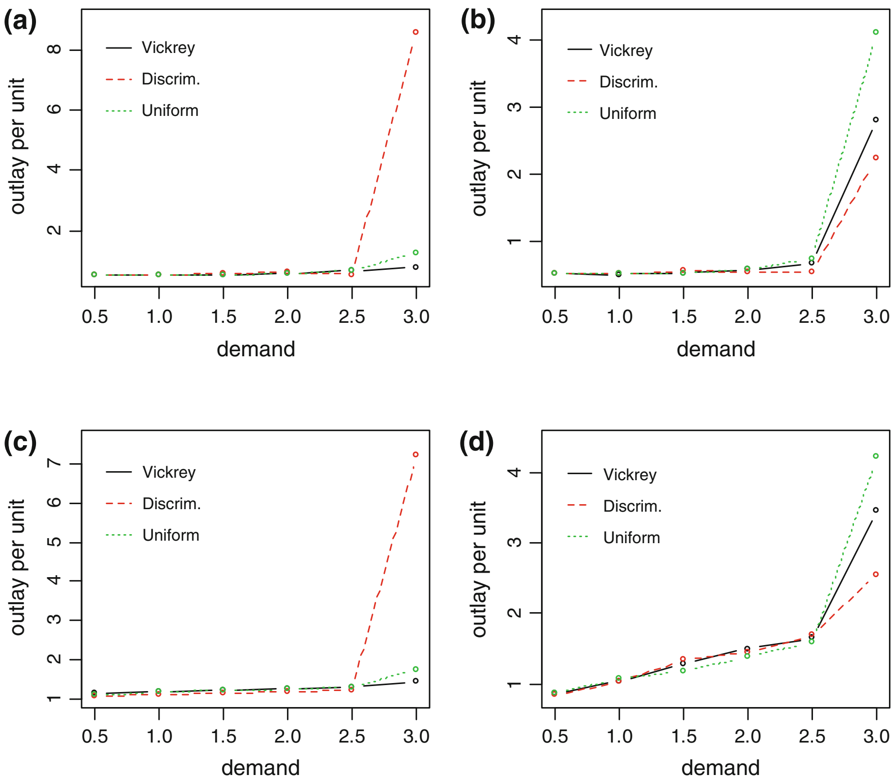

Fig. 3 Budgetary efficiency, demand level and auction pricing format. **a** Bidder population 1, **b** bidder population 2, **c** bidder population 3, **d** bidder population 4

for all bidders as in populations 1 and 2. The results from for allocative efficiency are plotted in Fig. 2 and indicate that the auction formats cannot be consistently ranked based on this criterion. In the case of population 3, for example, the discriminatory auction seems to generate lower social production costs than the other two formats for most demand levels. This advantage does not hold in the case of the most hetergoneous population (4), where the other two formats generate similar or lower costs.

Judging by budgetary outlay, the three pricing formats generate very similar outcomes except in for the case of the highest demand level. See Fig. 3. The similarity in outlays at most demand levels confirms that a revenue equivalence effect is at work even under the multi-unit auction. The stark differences at the highest demand level, however, highlight the limitations of the theory. For populations with constant costs, the discriminatory auction generates much higher auction prices and higher budgetary outlays than the other two auctions. This is the case for both flat and rising supply cost curves. Further, the Vickrey is slightly better than the uniform auction at the highest demand level. Overall, the Vickrey would be the best choice under bidder populations 1 and 2.

When bidders are heterogeneous, the similarity in budgetary performance of the auctions still holds for most demand levels. However, the uniform auction generates the lowest level of budgetary efficiency at the highest demand level. The discriminatory auction performs best at the highest level of demand. However, the advantages of this auction over the other two formats is not as strong as the relative disadvantage that it has at the same demand level when bidder populations are homogeneous.

5 Conclusions

Economic theory does not provide an analytical description of the equilibrium bidding strategies under multi-unit uniform and discriminatory auctions. The choice of auction format continues to be a controversial issue. The objective of this paper is to contribute towards filling this knowledge gap by using computational experiments to simulate bidding behaviour and auction performance for three formats: uniform, discriminatory and generalized Vickrey auctions.

The paper started by discussing theoretical predictions for the three auction types and the knowledge gaps that exist. Findings from some studies using human experiments were also discussed. An agent-based model was then formulated to simulate bidding among a population of agents that use genetic algorithm learning to optimize their nonlinear bid curves based on individual auction experience. The bidders learn over a strategy space with three parameters: the intercept, the slope parameter and the Box–Cox nonlinearity parameter. The experiments are undertaken for seven different demand levels (ranging in magnitude from 12.5 to 75.0% of aggregate supplier capacity) and for four different types of bidder populations, with the most heterogeneous one consisting of four groups of bidders differentiated by size and marginal cost slopes. All bidders in all populations have the same true entry prices.

Our results indicate that bidding behaviour cannot be completely characterized by auction format. It also depends on the nature of the bidder population and the level of competition. In particular, bidding strategies are sensitive to the heterogeneity of size amongst bidders especially in the case of the uniform and Vickrey auctions. These auctions induce two types of strategies: truthful bidding and supply inflation (i.e. true entry price but increasing overbidding on the subsequent units). When the population of bidders is heterogeneous and demand levels higher, these auctions lead to patterns where smaller and/or more competitive are involved less in price setting. Supply inflation is observed mostly for high levels of competition and predominantly for large capacity bidders. It is the strategy adopted by the bidders who are likely to be the price setters. The supply inflation strategy can also include a flat section at lower quantity levels. The flat sections help the bidder prevent prices from sliding down if other bidders lower their bid curves to below the price line. On the contrary, when bidders are less likely to be price setters, they tend to bid with curves that are completely below the price line thus “free-riding” on the risks taken by their bigger or more expensive counterparts.

The discriminatory auction, on the other hand, never leads to truthful bidding: two types of overbidding behaviours are observed: supply inflation and high flat bidding (i.e. high leant entry price and flat supply bid). The high flat bidding expected by theory is found when levels of competition are low. However, supply inflation is a frequent strategy when competition levels are high. Our results provide evidence of such bidding behaviour among all types of bidders at high competition levels, and even at high demand levels for bidders with less competitive cost structures. This bidding behaviour has also been observed in human experimental studies (Engelmann and Grimm 2003). An intuitive explanation can be provided for this deviation from the high flat bidding predicted by theory. High flat bids have the capacity to improve bidder revenue as the prices received for all units sold are brought closer. However, this strategy increases the risk that the bidder is completely priced out by rivals. Therefore, when a bidder faces stiff competition as a result of its similarity with others or because of its less competitive cost structure, a strategy of supply inflation rather than high flat bidding allows it to avoid zero gain outcomes.

The picture provided by these simulations is more complex than the partial view that the theory provides in relation to the structural properties of equilibrium strategies under the three formats. It indicates that attention should be granted to the level of competition as well as the heterogeneity of the bidding population, not only in terms of cost structure but also in terms of size.

The analysis of the relative performance of auctions in terms of budget outlays also delivers a strong message. The discriminatory auction, which is commonly used in practice, can be the most expensive when bidders are homoegenous and competition is weak. Vickrey is the least expensive procurement auction in these cases. The uniform auction can be the least attractive in terms of budgetary outcomes when populations are heterogeneous and competition is very weak. For most low and intermediate demand levels, however, the three auction formats deliver very similar budgetary outcomes. Finally, the ranking of the auction formats on the basis of social cost efficiency (allocative efficiency) depends on the type of heterogeneity and the degree of competition. These results suggest that the choice of an auction format should be tailored to the bidder population and predicted competition level. Auction simulations could be useful for assessing the potential performance of the different formats as the theoretical predictions seem to provide little on the variety of bidding strategies and auction ranking outcomes when competition levels and bidder populations are varied.

Appendix

See Fig. A1.

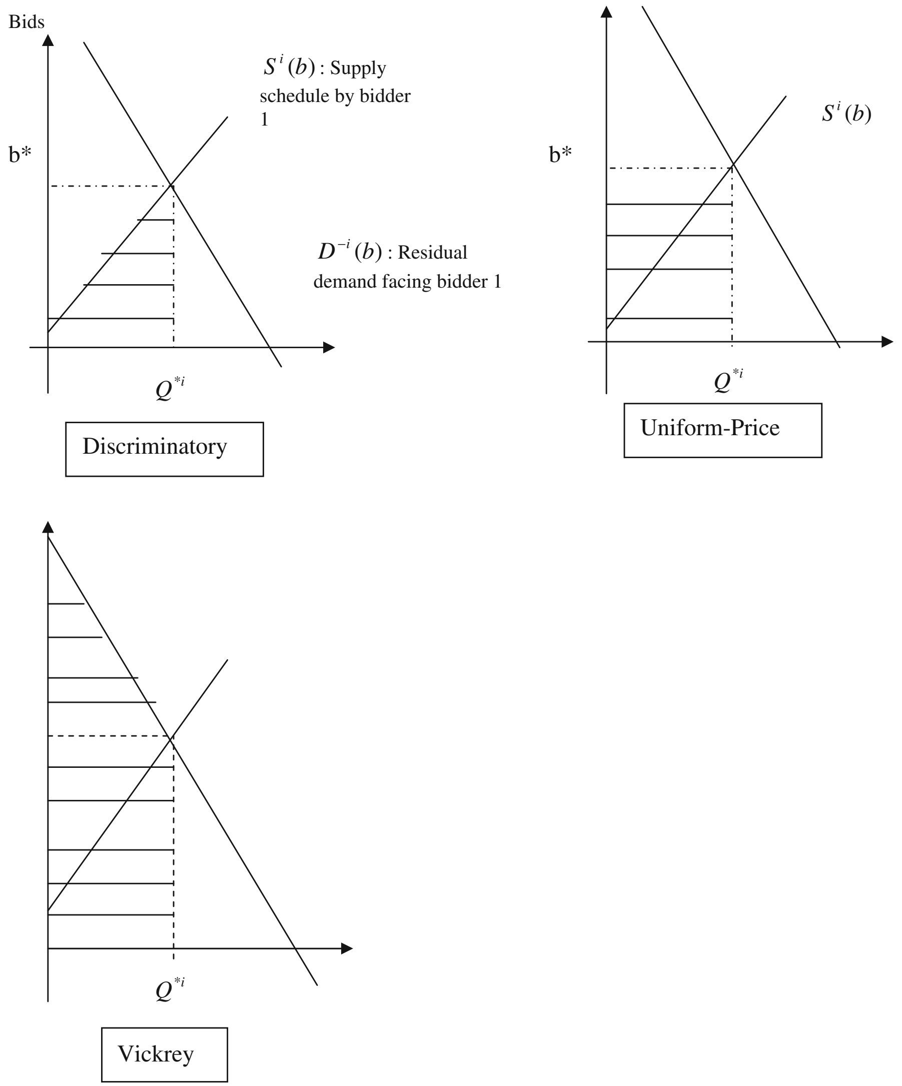

Fig. A1 Total payments under different auction formats

References

Alemsgeest, P., Noussair, C., & Olson, M. (1998). Experimental comparisons of auctions under single and multi-unit demand. *Economic Inquiry, 36*, 87–97.

Andreoni, J., & Miller, J. H. (1995). Auctions with artificial adaptive agents. *Games and Economic Behavior, 58*, 211–221.

Ausubel, L. M. (2005). An efficient ascending-bid auction for multiple objects. *American Economic Review, 94*, 1452–1475.

Ausubel, L. M., & Cramton, P. (2002) Demand reduction and inefficiency in multi-unit auctions. Working paper, University of Maryland.

Back, K., & Zender, J. F. (1993). Auctions of divisible goods: On the rationale for the treasury experiment. *Review of Financial Studies, 6*, 733–764.

Binmore, K., & Swierzbinski, J. (2000). Treasury auctions: Uniform or discriminatory?. *Review of Economic Design, 5*, 387–410.

Bower, J., & Bunn, D. W. (2001). Experimental analysis of the efficiency of uniform price versus discriminatory auctions in the England and Wales electricity market. *Journal of Economic Dynamics and Control, 25*, 561–592.

Bunn, D. W., & Oliveira, F. S. (2001). Agent-based simulation: An application to the new electricity trading arrangements of England and Wales. *IEEE Transactions on Evolutionary Computation, 5*, 493–503.

Camerer, C. F. (2003). *Behavioral game theory: Experiments in strategic interaction*. New York: Princeton University Press: Russell Sage Foundation.

Duffy, J. (2006). Agent-based models and human-subject experiments. In K. L. Judd & L. Tesfatsion (Eds.), *The handbook of computational economics* (Vol. 2). Elsevier: Amsterdam (*forthcoming*).

Engelbrecht-Wiggan, R., & Kahn, C. (1998). Multi-unit auctions with uniform prices. *Economic Theory, 12*, 227–239.

Engelmann, D., & Grimm, V. (2003). Bidding behaviour in multi-unit auctions—an experimental investigation and some theoretical insights. Working paper, CSERGE.

Epstein, J. M., & Axtell, R. (1996). *Growing artificial societies: Social sciences from the bottom up*. Washington DC: Brookings Institution Press.

Erev, I., & Roth, A. E. (1998). Predicting how people play games with unique, mixed strategy equilibria. *American Economic Review, 88*, 848–881.

Hailu, A., & Schilizzi, S. (2004). Are auctions more efficient than fixed price schemes when bidders learn?. *Australian Journal of Management, 29*, 147–168.

Hailu, A., & Thoyer, S. (2006). Multi-unit auction format design. *Journal of Economic Interaction and Co-ordination, 1*, 129–146.

Hailu, A., & Thoyer, S. (2007, September). Designing multi-unit multiple bid auctions: An agent-based computational model of uniform, discriminatory and generalized Vickrey auctions. *Economic Record, 83*(S1), S57–S72.

Kagel, J, & Levin, D. (2001). Behavior in multi-unit demand auctions: Experiments with uniform price and dynamic Vickrey auctions. *Econometrica, 69*(2), 413–454.

Krishna, V. (2002). *Auction theory*. Academic Press, Elsevier: London, 303 pp.

Michalewicz, Z. (1996). *Genetic algorithms + data structures = evolution programs*. New York: Springer-Verlag.

Nicolaisen, J., Petrov, V., & Tesfatsion, L. (2001). Market power and efficiency in a computational electricity market with discriminatory double-auction pricing. *IEEE Transactions on Evolutionary Computation, 5*, 504–523.

Roth, A. E., & Erev, I. (1995). Learning in extensive form games: Experimental data and simple dynamic models in the intermediate term. *Games and Economic Behavior, 8*, 164–212.

Swinkels, J. (1999). Asymptotic efficiency for discriminatory private value auctions. *Review of Economic Studies, 66*, 509–528.

Tenorio, R. (1993). Revenue equivalence and bidding behaviour in a multiunit auction market: An empirical analysis. *The Review of Economics and Statistics, 75*, 302–314.

Tenorio, R. (1999). Multiple unit auctions with strategic price-quantity decisions. *Economic Theory, 13*, 247–260.

Tesfatsion, L. (2002). Agent-based computational economics: Growing economies from the bottom up. *Artificial Life, 8*, 55–82.

Wilson, R. (1979). Auctions of shares. *The Quarterly Journal of Economics, 93*, 675–689.

Wolfram, C. (1998). Strategic bidding in a multiunit auction: An empirical analysis of bids to supply electricity in England and Wales. *Rand Journal of Economics, 29*, 703–725.