



AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search
<http://ageconsearch.umn.edu>
aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

Assessing the performance of auctions for the allocation of conservation contracts: Theoretical and computational approaches

A. Hailu*, S. Schilizzi*, S. Thoyer**

* School of Agricultural and Resource Economics, University of Western Australia, 35 Stirling Highway, Crawley (Perth), WA 6009, Australia Atakelty.Hailu@uwa.edu.au

** Ecole Nationale Supérieure d'Agronomie de Montpellier, Laboratoire Montpelliérain d'Economie Théorique et Appliquée, 2 Place Viala, 34060 Montpellier cedex, France

Selected Paper prepared for presentation at the American Agricultural Economics Association Annual Meeting, Providence, Rhode Island, July 24-27, 2005

Copyright 2005 by A. Hailu, S. Schilizzi and S. Thoyer. All rights reserved. Readers may make verbatim copies of this document for non-commercial purposes by any means, provided that this copyright notice appears on all such copies.

Assessing the performance of auctions for the allocation of conservation contracts: Theoretical and computational approaches

Abstract

There is a growing interest in using auctions for purchasing public goods from private agents. Auctions are being trialed in Australia and elsewhere to allocate conservation contracts. The expectation is that competitive bidding will reduce information rents and increase cost-effectiveness. This paper examines how auctions would perform under different assumptions regarding the rationality of bidders. A theoretical model requires bidders to be rational and use Nash equilibrium strategies, while an agent-based model assumes boundedly rational bidders learning from experience. The study illustrates the synergies between economic theory and agent-based modelling. Our findings provide a cautionary message regarding the performance of conservation auctions.

1. Introduction

Interest in auctions as instruments for purchasing conservation services from landholders has recently increased throughout Australia, especially after the BushTender¹ biodiversity trial auctions in Victoria. In the BushTender trials, farmers were asked to bid competitively for agri-environmental contracts: they voluntarily submitted proposals for conservation activities and nominated a corresponding compensation payment in sealed-bid process. Conservation activities are translated into biodiversity indexes by appointed scientists. Submissions are then ranked on the basis of biodiversity index to bid ratios and contracts awarded until a pre-determined budget is exhausted. Given the budget and the bids, the selection procedure maximizes biodiversity conservation per dollar spent.

Currently, several additional auction trials are underway as part of the federal government's market-based instrument (MBI) pilot program. This interest is mainly driven by a commonly held belief that auctions allow the government to minimize information rents and thus are better than alternative approaches like fixed price schemes. For example, in a recently published paper, Stoneham et al (2003)

argue that the amount of biodiversity benefits acquired through the first round of BushTender auctions would have cost about seven times as much if a fixed price scheme had been used instead. However, whether these beliefs or predictions are warranted requires a more careful investigation.

This study examines the validity of the expectation regarding the superiority of auctions over fixed price schemes using both theoretical and computational models. Conservation auctions have limited budgets or are budget-constrained. Theoretical analysis can be used to derive equilibrium strategies for target-constrained² auctions under the assumption of perfectly rational agents. An agent-based model (ABM) is used to assess outcomes for both target-constrained and budget-constrained auctions in contexts where bidders are boundedly rational but can learn from experience. Because there are no clear theoretical predictions for a budget-constrained procurement auction, only the agent based model is then used to examine the performance of budget-constrained auctions.

The paper is structured as follows. In the second section, theoretical results are derived for a target constrained procurement auction. The agent-based model as well as the learning theory employed in that model are presented in the third section. The fourth section presents the main ABM results. First, results for a target-constrained auction obtained with the agent-based model are presented and compared to those derived from the theoretical model. Results simulated for budget-constrained auctions are then compared to target-constrained auctions having the same budget. The value of doing this stems from the fact that, although theory has primarily studied the target-constrained auction, in practice governments use budget-constrained auctions. The pattern of bidding behaviours both learnt and optimal consistently point to a need for a more careful assessment of the potential benefits of conservation auctions. Section five offers some conclusions and policy messages.

2. Nash Equilibrium bids for a target constrained auction

Vickrey (1962) formulated a Nash equilibrium bidding model in the case of multiple unit sealed-bid discriminative price auctions (when agents bid only for one unit) and demonstrated that the Revenue Equivalence Theorem holds for risk-neutral bidders with individual values for the auctioned objects drawn

from a uniform distribution. Harris and Raviv (1981) generalized the Vickrey model for bidders' valuations drawn from general distribution functions and when all bidders have identical concave utility functions. All subsequent extensions (Milgrom and Weber, 1982; Cox et al, 1984) have focused on "selling" auctions. In the literature, optimal bid formulas have been explicitly given for direct or selling auctions (e.g. Cox et al. 1984) but not for procurement or reverse auctions. We customized the Vickrey-Harris-Raviv approach to model the Nash equilibrium risk neutral bid functions in a procurement multiple unit auction, relevant for government conservation schemes. We first do the calculation for a single unit auction then extend it to a multiple unit action.

Nash equilibrium bidding strategy for a single-unit procurement auction in a discriminative sealed bid setting

Let n risk-neutral bidders compete to sell one unit of a good to the auctioneer: let v_i be the monetary value of this good to bidder i . Assume that each v_i is drawn (with replacement) from a distribution with density $f(\cdot)$ and probability distribution function $F(\cdot)$ whose support is the interval $[0, v^{\sup}]$. Suppose bidder 1 with value v bids b , and all $(n-1)$ rival bidders k bid according to the strictly monotonous increasing equilibrium strategy $B(v_k)$.

The expected gain of bidder 1 is:

$$E(v, b) = (b - v) \Pr[B(v_k) > b] \quad \forall k \neq 1, \text{ or}$$

$$E(v, b) = (b - v) [1 - F(B^{-1}(b))]^{n-1} \quad (1)$$

Maximizing (1) with respect to b yields the following first-order conditions:

$$[1 - F(B^{-1}(b))]^{n-1} - (b - v)(n-1)[1 - F(B^{-1}(b))]^{n-2} \frac{f(B^{-1}(b))}{B'(v)} = 0 \quad (2)$$

At equilibrium, $b=B(v)$

$$[1-F(v)]^{n-1} - (B(v)-v)(n-1)f(v) \frac{(1-F(v))^{n-2}}{B'(v)} = 0$$

$$B'(v) = (n-1)(B(v)-v) \frac{f}{1-F}(v) \quad (3)$$

For a uniform distribution between 0 and 1, $F(v) = v$, $f(v) = 1$ and (3) simplifies to:

$$B'(v) = (n-1) \frac{B(v)-v}{1-v} \quad (4)$$

Therefore, the optimal bidding strategy is given by:

$$B(v) = \frac{n-1}{n}v + \frac{1}{n} \quad (5)$$

The optimal bidding strategy is one of overbidding ($b > v$). This overbidding declines when the number of bidders (n) increases.

As a comparison, the optimal bidding strategy for a “selling” (or direct) sealed bid auction is:

$$B(v) = \frac{n-1}{n}v \quad (\text{Wolfstetter, 1996})$$

Contrary to what one might have expected, the optimal bid formulae for direct and procurement auctions are not symmetrical.

Generalization to multiple unit procurement auctions

Consider a multiple unit reverse auction with n bidders and m units demanded by the auctioneer, each bidder wanting to sell *at most* one unit. Each bidder submits a bid for a single unit with the understanding that each of the m^{th} lowest bidders will sell a unit of the good at a price equal to his own bid (discriminative sealed bid auction).

The probability that a bid b by bidder 1 will win is the probability $G(B^{-1}(b))$ that at least $(n-m)$ of the values drawn by the rivals are greater than $B^{-1}(b)$. This probability is given by $1 - H(F)$ where $H(F)$ is the distribution function of the m^{th} order statistics for a $(n-1)$ sample from distribution F :

$$G(B^{-1}(b)) = \frac{(n-1)!}{(m-1)!(n-1-m)!} \int_{B^{-1}(b)}^{v^{\text{sup}}} F(v)^{m-1} (1-F(v))^{n-1-m} f(v) dv \quad (6)$$

The expected gain of bidder 1 of value v and bid b is:

$$E(v, b) = (b-v) \cdot G(B^{-1}(b)) \quad (7)$$

The first-order conditions for the maximization problem in (7) are:

$$G(B^{-1}(b)) + (b-v) \frac{G'(B^{-1}(b))}{B'(v)} = 0$$

At equilibrium, $b = B(v)$

$$B'(v) \cdot G(v) + B(v) \cdot G'(v) = v \cdot G'(v)$$

$$B(v) = - \frac{\int_v^{v^{\text{sup}}} u G'(u) du}{G(v)}$$

Therefore, the optimal bid is:

$$B(v) = \frac{\int_v^{v^{\text{sup}}} u G'(u) du}{\int_v^{v^{\text{sup}}} G'(u) du} \quad (9)$$

For a uniform distribution between 0 and 1, we have:

$$G(v) = \frac{(n-1)!}{(m-1)!(n-1-m)!} \int_v^1 u^{m-1} (1-u)^{n-m-1} f(u) du \quad (10)$$

$$b(v) = \frac{\int_v^1 u^m (1-u)^{n-m-1} du}{\int_v^1 u^{m-1} (1-u)^{n-m-1} du} \quad (11)$$

An important feature of the optimal bid as determined in (11) is that the level of overbidding is high for low value bidders. Overbidding decreases as the value increases, with the bids from high value bidders asymptotically approaching their respective values (see Figure 1(a), solid line). For example, in an auction involving 100 bidders with values uniformly distributed between 0 and 1 competing for the sale of 30 units, the level of overbidding is 300% for a value of 0.075 and only 11.70% for a value of 0.305.

3. Agent-based and learning models

Auction theory has been used to identify optimal auction design, but the results are usually only valid under very restrictive assumptions concerning both the environment of the auction and the rationality of players. As Arifovic and Ledyard (2002) emphasize, the three main roadblocks on auction and mechanism design theory are that: “(1) much of the theory is about one-shot games while many applications involve repeated play against the same (or very similar) opponents; (2) there is no generally accepted view as to the right model of individual and group behaviour, (3) we have not yet incorporated computational limitations, of either the mechanisms or the agents, into the theory”. In fact, experimental results (Erev and Roth, 1998, Camerer, 2003: chapter 6) tend to demonstrate that the way people play is better captured by learning models rather than by the Nash-Equilibrium predictions. So, in practice, what we would observe is people learning over time, not people landing on the Nash equilibrium at the outset of the game.

Initial theoretical work has sought to demonstrate the convergence of learning games with game theoretic predictions (Jordan, 1995; Milgrom and Roberts, 1991; Roth and Erev, 1995). However, the availability of repeated game equilibria solutions is limited and more complicated settings may become analytically intractable. The need to use alternative methods to generate the outcomes of the learning processes has led to the rapid development of computational methods using artificial agents. A number of

testbed experiments have been conducted to compare learning algorithms with theoretical results in order to validate the use of computational agent-based techniques. Arifovic and Ledyard (2002) analyse the Groves-Ledyard mechanism for the provision of public goods in a repeated play. They develop a computational testbed with a behaviour based on experimentation and replication and compare the outcomes with the predicted Nash equilibrium. Brenner (2000) studies the dynamics of market prices under the assumption of boundedly rational agents who learn by imitation, satisficing and experience (the Variation-Imitation-Decision model). He compares the learning process to the game-theoretic predictions under the assumption of rational agents and demonstrates that the simulated price dynamics converge to the equilibrium price in the long run. Hon-Snir and Monderer (1998) analyse a repeated first price auction (in which the types of the players are pre-determined) and demonstrate that if every player is using a belief-based learning scheme, then after a large number of rounds, the players' bids are in equilibrium with the bids of the one shot auction in which the types are commonly known. This paper develops a computational model based on artificial learning agents and compares its long run outcomes to the theoretical or Nash equilibrium predictions of target constrained auctions. The ABM model is then used to study the outcomes of a budget constrained auction.

3.1 The learning model

Different learning models have been developed over the last several decades. A typology of learning models presented by Camerer (2003) shows the relationship between these learning algorithms. The models differ in terms of their information requirements or assumptions. The reinforcement-learning algorithm is chosen for this study as it is particularly suitable for modelling bidding behaviour without requiring that players be knowledgeable about forgone payoffs associated with strategies that they did not select.

The reinforcement-learning algorithm was developed by Roth and Erev (1995) based on the reinforcement principle that is widely accepted in the psychology literature. Erev and Roth (1998) extend

and use this learning algorithm to model behaviour from 12 experimental studies³ of repeated games with unique nontrivial mixed equilibrium in mixed strategies. They find that the predictions of the reinforcement learning model for the choices of experimental subjects are at least as good as equilibrium predictions. Their algorithm or modified versions of it have been used in several agent-based studies of electricity auction markets (Nicolaisen *et al* 2001; Bunn and Oliveira, 2001).

Roth and Erev's algorithm is based on the following four principles rooted in the psychology of learning (Erev and Roth 1998): *the law of effect*, *the power law of practice*, *experimentation* and *recency*. The *law of effect* asserts that the tendency to choose an action is strengthened (*reinforced*) or weakened depending upon whether the action produces favourable results or not. This principle implies that choice behaviour is probabilistic. The *power law of practice* refers to the fact that learning curves tend to be initially steep. *Experimentation* (or *generalization*) implies that strategies that are similar to previously chosen successful ones will be employed more often. Experimentation prevents players from quickly being locked into particular choices. *Recency* (or *forgetting*) requires that recent experience have more impact on behaviour than past experience.

The main features of the algorithm can be described using the following three equations. If the propensity of player n to choose strategy j at time t is denoted by $q_{nj}(t)$, the propensity updating function can be written as (Erev and Roth 1998, p. 863):

$$q_{nj}(t+1) = (1 - \phi)q_{nj}(t) + E_k(j, R) \quad (12)$$

where: ϕ is the recency parameter, R is the reinforcement (or profit) from previous choice of strategy k , while $E_k(j, R)$ is the following three step generalization function⁴:

$$\begin{aligned} E_k(j, R) &= R(1 - \varepsilon) && \text{if } j = k \\ &= R \cdot \frac{\varepsilon}{2} && \text{if } j = k \pm 1 \\ &= 0 && \text{otherwise} \end{aligned} \quad (13)$$

where ε is an experimentation parameter.

The probability that player n uses his k^{th} strategy depends on the propensities, as given by:

$$p_{nk}(t) = \frac{q_{nk}(t)}{\sum_j q_j(t)} \quad (14)$$

Therefore, this learning algorithm has three parameters, namely, the recency parameter, the experimentation parameter as well as a scale parameter (s) that determines the initial propensities⁵. The parameters values that provided the best data for the 12 games studied in Erev and Roth (1998) were used in this study. These values are 0.1, 0.2 and 9, respectively. Our experiments with alternative values for these parameters show that the convergence between theoretical and learnt bids is insensitive to the values of these parameters.

3.2 Structure of agent-based model

The agent-based model used in this study has two types of agents representing the actual players in a real auction, namely landholders (farmers) and the government. Each landholder has an environmental quality score and an opportunity cost associated with putting the land being offered under conservation. The government agent has a conservation target described in terms of environmental units to be protected.

Each auction round involves the following three major steps or activities. First, landholders choose bids. Landholder choices are the result of the reinforcement learning process described in section 3.1: the learning that a bidder undertakes focuses on the mark-up that he can apply to his true opportunity cost of providing the conservation service. Each bidder has 100 strategies, relating to mark-ups ranging from 0 to 1.0 in step of length 0.01. Thus the bid choice of a bidder depends on his opportunity costs, his previous bid choices as well as the profits or rewards obtained for those choices. Second, the government agent ranks bids based on environmental quality score to bid ratios⁶ and awards conservation contracts to bidders starting from the one offering the highest environmental quality per dollar until the conservation target is met. In the case of budget-constrained auctions, contracts would be offered based on the ranking until the budget is exhausted. Third, landholder agents update their contract status based on the message

from the government agent. This leads to updated values of propensities for choice that determine future bid choices.

In the agent-based model, the number of bidders and the initial distribution of their opportunity costs (uniform distribution between 0 and 1) remains the same. However, the opportunity costs of individual bidders are allowed to vary within a given range. This was done to partly mimic the fact that in the theoretical model a bidder knows about the distribution of opportunity costs of his rivals but not the particular set of bidders he is playing against. This differs from the model in Hailu and Schilizzi (forthcoming) where the opportunity costs of the players were assumed to remain the same. Results were generated by averaging 100 replications each using a different random seed.

4. Auction performance

We first compare and contrast simulated budget-constrained results to theoretical predictions for the target-constrained auctions. Target-constrained outcomes from the agent-based model are then compared to equivalent budget-constrained auction results. The impact of bidder learning on auction performance is measured using different criteria:

- *returns to budget*: the biodiversity quality score gain for a given budgetary expenditure
- *information rents*: the difference between payments made to farmers and their true opportunity costs (known only to themselves)
- *auction efficiency*: the opportunity cost of a given amount of biodiversity gain. The auction is said to be efficient when the ratio of social opportunity cost to biodiversity gain is minimum.

4.1 Convergence of learnt bids to the Nash equilibrium optimal bids

The learning process described in our model leads to a bidding behaviour which is very close to what would be the optimal bid strategy in a one shot game under the hypothesis of perfect rationality. The learnt bids curve converges towards the optimal bids curve, at least for those bidders who have had ample opportunity to learn. These are bidders that have been selected most of the time, or the bidders with the

lowest opportunity costs (Figure 1a). As the frequency of selection falls, the distance between learnt bids and optimal bids increases.

Figures 1(a) and (b) about here

However, Figure 1(b) shows that convergence is not perfect. Instead, learnt bids remain slightly above optimal bids. Part of the reason lies in a statistical effect, whereby there is no bidder who gets selected 100% of the time; however infrequently, it is rejected by the auction when its bid is too high. Since Figures 1 (a) and (b) represent averages over 100 repetitions, they include runs where individual bidders were not selected because of high bids. However, the main reason is that, whereas it is not possible for a bidder to bid below his optimal bid and not get selected, it is quite possible for him to bid above his optimal bid and still get selected. This can be verified by examining individual runs (rather than averages over 100 runs). It is this dissymmetry in the probability of being selected, conditional on one's bid relative to one's optimal bid, which explains the slight discrepancy observed in Figure 1(b).

In conclusion, we demonstrate that the learnt bids converge to the optimal bids. Even with a simple adjusting behaviour based only on his own past experience, an agent with bounded rationality can finally adopt a bidding behaviour similar to what his decisions would have been in a perfect rationality setting. However, we also need to demonstrate that the learning is robust and that this result is not contingent on initial conditions. One of them is the issue of the bidding strategy in the first round, when the agent has not had an opportunity to learn from his past mistakes or successes.

4.2 Are initial bidding strategies important for final outcomes?

There are no clear theoretical grounds for answering the question of what should be the initial bidding strategy of players. Rather, it appears to be an empirical issue. Do bidders start by bidding honestly (i.e submitting their true opportunity costs) or do they start by bidding above their opportunity costs? More importantly, do initial bidding styles affect the final outcomes of the auction?

It turns out that starting bids do not matter in terms of final outcomes. Learnt bids still converge to optimal bids, and do so at approximately the same rate. In addition, as shown in the first three columns of

Table 1, auction performance remains nearly exactly the same, whether in terms of total program outlay (budgetary costs), total forgone profits (measured by the sum of opportunity costs of all winners⁷), or net income transfers rates (calculated as the share of payments to landholders above and beyond their true opportunity costs⁸).

Table 1 about here

In conclusion, our learning model does converge to the Nash Equilibrium and it is robust to different initial bidding strategies. In particular, it does not seem to matter whether bidders start by bidding honestly or not. We are thereafter confident in using this model to explore questions related to auction performance for which there is no simple theoretical guideline for calculating the optimal bid.

4.3 Budget constraint versus target constraint

As highlighted in the introduction, the literature has predominantly studied the target-constrained (TC) auction, where a predetermined number of units are to be sold or, in our case, purchased. From theoretical work done on budget-constrained (BC) auctions, it appears there is no dominant solution to the optimal bid problem (Müller and Weikard, 2002), unless exogenous elements are introduced such as expected bid caps or perceived reserve prices (Latacz-Lohmann and Van der Hamsvoort, 1997). And yet, in practice, BC auctions are the rule in policy applications, in particular in government procurement auctions. For political and budget-control requirements, a budget constraint is the usual choice and it is, therefore, of interest to compare the performance of these two auction designs.

The BC auction was run with a budget constraint equal to the total program outlay implied by the TC auction (with a target of 30 units as indicated earlier). This budget was \$9.46. The population of bidders remained the same between the two experiments. Table 2 summarizes the results.

Table 2 about here

The BC auction and the TC auction lead to similar results (see Table 2). The budget of \$9.46 implied by the target-constrained auction for the purchase of 30 units results in a purchase of about 29.8 units. These 29.8 units have a total opportunity cost of 4.75 whereas the 30 units purchased under the target auction have a total opportunity cost of 4.80. Both lead to similar informational rents. Differences in the results from the two types of auctions are related to the fact that the budget might not be fully utilized in the budget constrained auction because partial contracts are not allowed.⁹

These results tentatively suggest that we may be able to use our theoretical understanding of target based auctions to make predictions about budget constrained auctions. Since most procurement auctions are budget constrained, this has practical value, while also extending the scope of auction theory beyond its current analytical capabilities.

The convergence of boundedly rational bidders to the Nash equilibrium, the insensitivity of auction outcomes to initial bidding strategies, and the similarity of target and budget constrained auctions are results that have been obtained using a specific learning algorithm – reinforcement learning – and a specific format for the agent-based model: a constant number of 100 bidders with a uniform distribution of opportunity costs, each varying within a given range. However, other simulations we have done with a different learning algorithm (based on learning direction theory) and other ABM model configurations, suggest our results should not depend on these specifications, but that they are rather general. All predict that learning leads to bidding patterns where low opportunity cost agents align their bids with the more expensive or marginal winners.

5. Conclusion

This study has investigated the performance of a special category of auctions used in the area of resource conservation, namely multi-unit procurement auctions. They represent the case where a government agency wishes to buy public goods or services from private agents, namely environmental services from landholders. Although auctions have been studied extensively, the auction contexts explored are generally

kept simple for analytical tractability. This paper has proposed an alternative approach, agent-based modelling, to overcome analytical difficulties and complement theoretical approaches. As a prerequisite, however, it is necessary to check the consistency of results from agent-based models with those predicted from economic theory whenever these are known. The goal of this exercise is to observe the extent to which theoretical predictions are reproduced in agent-based models. Once a consistency is established for simple cases, the limitations imposed by analytical tractability can be overcome through the use of agent-based modelling for the investigation of richer and more complex auction settings.

For the case of multi-unit auctions under standard settings, optimal bid predictions from Nash equilibrium analysis were successfully reproduced by the agent-based model populated with boundedly rational agents that learn according to the reinforcement algorithm. These results are not sensitive to the manner in which bidders start bidding (honestly or randomly) in the initial round. The agent-based model was then used to show that budget-constrained and target-constrained auctions lead to similar outcomes. This equivalence is interesting, as most theoretical analysis is based on target constrained auctions whereas real conservation auctions are constrained by available budgets.

The key result of this paper is that learnt bids are characterized by the alignment of infra-marginal bids with those of the marginal bid. This leads to the extraction of rents similar to those that would occur under fixed price schemes where all contracting landholders are paid the same price. The policy implication of this is that current expectations regarding the performance of conservation auctions relative to fixed price schemes need to be reassessed. This cautionary message is further strengthened by the fact that auctions are likely to involve higher transaction costs. Whether natural resource management agencies should invest effort into appropriately estimating opportunity costs and then use those results to offer simpler but tailored take-it-or-leave-it type contracts is something that should be given a more careful consideration.

References

- Arifovic J. and Ledyard, J. 2002, 'Computer testbeds: the dynamics of Groves-Ledyard mechanisms'. Working Paper, Simon Fraser University and California Institute of Technology. Dec 31, 2002, 28 p.
- Brenner, T, 2000, 'The Dynamics of Prices – Comparing Behavioural Learning and Subgame Perfect Equilibrium, Working paper, Max Planck Institute for research into Economic systems , 33 pages
- Bunn D. and Oliveira F. 2001, 'Agent-based simulation – an application to the new electricity trading arrangements of England and Wales', *IEEE Transactions on Evolutionary Computation*, Vol 5 No 5, pp. 493-503
- Camerer, C.F. 2003, *Behavioral Game Theory. Experiments in Strategic Interaction*. (Chapter 6: Learning). Russell Sage Foundation, New York; Princeton University Press.
- Cox, J.C., Smith, V. L. and Walker, J.M. 1984, 'Theory and Behavior of Multiple Unit Discriminative Auctions', *The Journal of Finance*, Vol 39, No4, pp. 983-1010
- Erev, I. and Roth, A.E. 1998, 'Predicting how people play games with unique, mixed strategy equilibria', *American Economic Review*, 88, pp. 848-881.
- Hailu A. and Schilizzi S. (forthcoming), 'Are auctions more efficient than fixed price schemes when bidders learn?', *Australian Journal of Management*.
- Harris M. and Raviv A. 1981, 'Allocation mechanisms and the design of auctions', *Econometrica* vol. 49, pp. 1477-99.
- Hon-Snir S. and Monderer D., 1998, 'A learning approach to auctions', *Journal of Economic Theory*, vol. 82, pp. 65-88.
- Jordan J.S. 1995, 'Bayesian Learning in Repeated Games', *Games and Economic Behavior*, vol. 9, no 1, pp. 8-20.
- Latacz-Lohmann U. and van der Hamsvoort C. 1997, 'Auctioning conservation contracts, a theoretical analysis and an application', *American Journal of Agricultural Economics*, vol. 79, pp. 407-418.
- Milgrom P. R. and Weber R. J., 1982, 'A theory of auctions and competitive bidding', *Econometrica* vol. 50, pp. 1089-1122.
- Milgrom and Roberts, 1991, 'Adaptive and sophisticated learning in normal form games', *Games and Economic Behaviour*, vol. 3, pp. 82-100.
- Müller, K. and H-P. Weikard, 2002, 'Auction mechanisms for soil and habitat protection programmes', in: K. Hagedorn (ed.), *Environmental Co-operation and Institutional Change*. Cheltenham: Edward Elgar, pp. 200-211.
- Nicolaisen, J., V. Petrov, and L. Tesfatsion, 2001, 'Market power and efficiency in a computational electricity market with discriminatory double-auction pricing', *IEEE transactions on Evolutionary Computation*, vol. 5 no. 5, pp. 504-523.

- Roth, A.E. and I. Erev, 1995, 'Learning in extensive form games: experimental data and simple dynamic models in the intermediate term', *Games and Economic Behavior*, vol. 8, pp. 164-212.
- Stoneham G., V. Chaudhri, A. Ha, L. Strappazzon 2003, 'Auctions for conservation contracts: an empirical examination of Victoria's BushTender trial', *Australian Journal of Agricultural and Resource Economics*, vol. 47, no. 4, pp. 477-500.
- Vickrey W., 1962, 'Auctions and bidding games', in *Recent Advances in Game Theory*, Princeton University Press, pp. 15-27.
- Wolfstetter E., 1996, 'Auctions: an introduction', *Journal of Economic Surveys*, vol. 10, no. 4, pp.. 367-420.

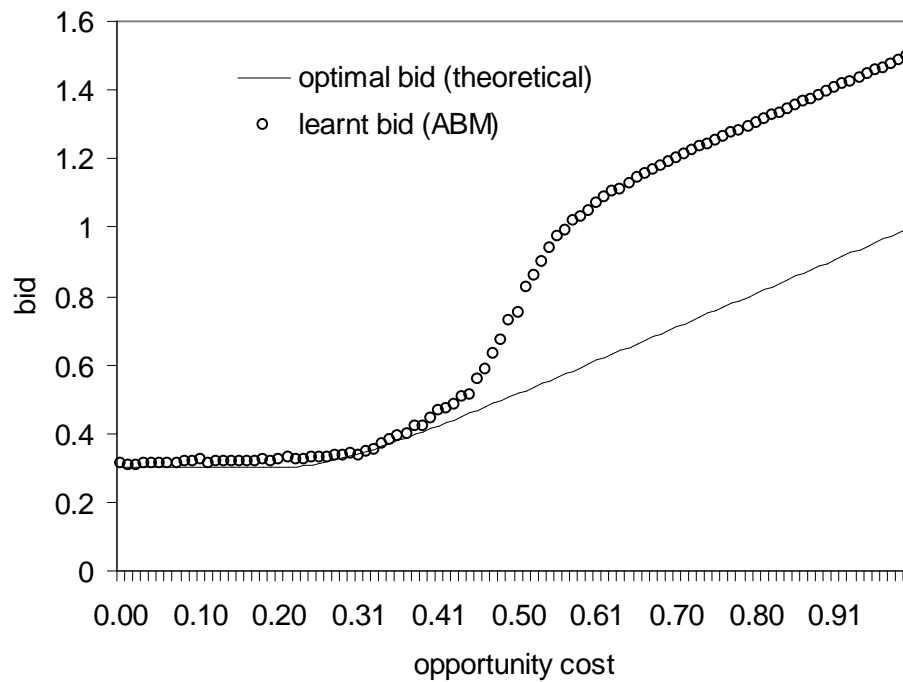


Figure 1(a). Learnt versus optimal bids as a function of opportunity cost (100 by 1 type bidders in a procurement auction for 30 units): All bidders

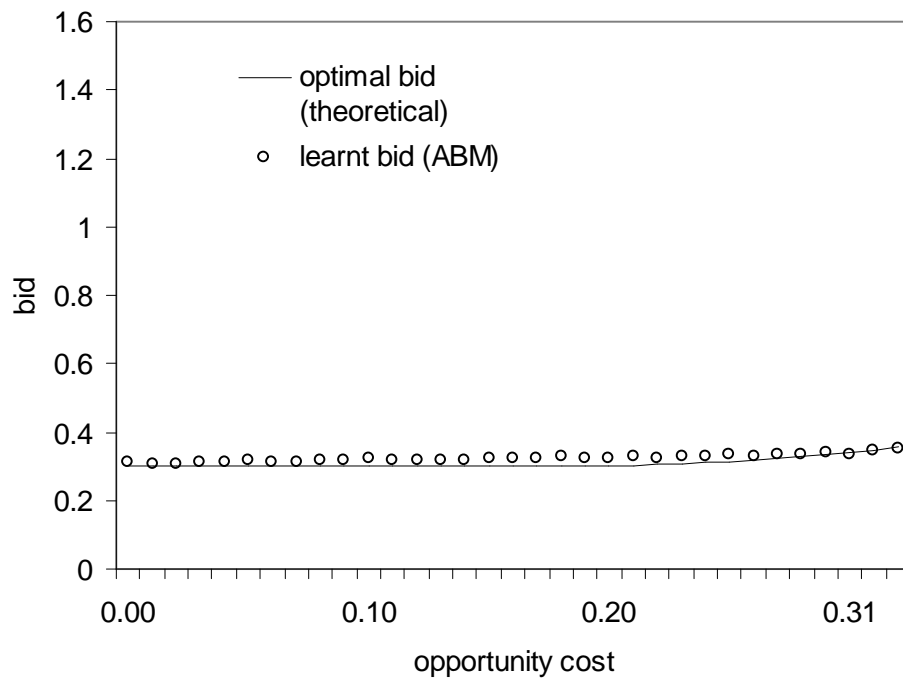


Figure 1(b). Learnt versus optimal bids as a function of opportunity cost (100 by 1 type bidders in a procurement auction for 30 units): winning bidders (>95% of the time)

Table 1: Results from a fixed target option (30 units) and 100 bidders: A case of partly random population with opportunity costs varying around a mean

Population of bidders		Total Program Outlay	Total Forgone Profits	Net Income Transfer Rate	MSD for All Bidders	MSD for All Winners
	Honest initial bids	9.56	4.80	0.50	0.16	4.8e-4
	Random initial bids	9.46	4.81	0.49	0.16	4.9e-4
	Outcome under perfect information system	4.50	4.50	0	0.01	0.03

Note: MSD stands for the mean squared deviation of the learnt bid from the optimal bid of the Nash equilibrium solution.

Table 2: Budget-constrained vs. target-constrained auctions

Type of auction	Total program outlay	Total opportunity costs	Net income transfer rate (info rents)	% of winners	Value for money*
Target constrained	9.46	4.81	49%	30%	3.13
Budget ¹⁰ constrained	9.46	4.75	48.5%	29.8%	3.20

Note (*): Benefit per \$ of program outlay.

Footnotes

¹ The BushTender trial is a pilot auction which was conducted as of 2001 to assess empirically the feasibility and efficiency of auction mechanisms for the allocation of biodiversity conservation contracts. See Stoneham et al (2003).

² Target constrained auctions refer to cases where the auctioneer is attempting to purchase or sell a fixed number of units.

³ Eleven of these games were conducted by different researchers in the period between 1960 and 1995.

⁴ For strategy sets without linear order, the generalization function should be specified as a two-step function. See Erev and Roth (1998, p. 863).

⁵ The initial propensity for any given strategy is set as the product of the scale parameter and an expected profit from bidding (arbitrarily set at 0.05). The results do not depend on this value.

⁶ The calculation of the environmental quality score in the Bush Tender pilot was exogenous and prior to the auction. Scientists collapsed a number of ecological dimensions into one single number.

⁷ It is a measure of the efficiency of the auction mechanism: the lower total forgone profits, the more efficient the mechanism at picking up the lowest opportunity cost contributors and the higher the global welfare

⁸ It is the equivalent of an informational rent paid by the auctioneer to the winners.

⁹ In a BC auction, the m winners are chosen so that the sum of the m payments is equal to or less than the budget but less than the sum of the $m+1$ payments. This discretization explains the discrepancy observed between the budget in TC and BC auctions.

¹⁰ This is from a fixed population case where the opportunity costs do not change over the different rounds.