

1

Are Auctions More Efficient Than Fixed Price Schemes When Bidders Learn?

by

Atakelty Hailu †

Steven Schilizzi †

Abstract:

Auctions are increasingly being considered as a mechanism for allocating conservation contracts to private landowners. This interest is based on the widely held belief that competitive bidding helps minimize information rents. This study constructs an agent-based model to evaluate the long-term performance of conservation auctions under settings where bidders are allowed to learn from previous outcomes. The results clearly indicate that the efficiency benefits of one-shot auctions are quickly eroded under dynamic settings. Furthermore, the auction mechanism is not found to be superior to fixed payment schemes except when the latter involve the use of high prices.

Keywords:

AUCTIONS; GOVERNMENT PROCUREMENT; EFFICIENCY OF PUBLIC SPENDING; BIODIVERSITY CONSERVATION; LEARNING; AGENT-BASED MODELLING

† School of Agricultural & Resource Economics, The University of Western Australia, 35 Stirling Highway, Crawley WA 6009. Email: Atakelty.Hailu@uwa.edu.au and Steven.Schilizzi@uwa.edu.au

We gratefully acknowledge several useful comments from an anonymous referee and Sophie Thoyer.

Australian Journal of Management, Vol. 29, No. 2 December 2004, © The Australian Graduate School of Management

1. Introduction

Biodiversity conservation on private land is an example of the supply by private agents of a public good, for which there is no market to reward a landowner's efforts (Stoneham & Chaudri 2000). Reduced land degradation, deforestation, and water quality of lakes and rivers are other such environmental public goods (or services) for which there are no readily available markets.

The traditional policy approach to these conservation issues has been regulation or negotiated contracts with fixed payment schemes. The public authority would typically forbid, constrain, or impose certain activities or outcomes, and try to enforce the regulations with fines and other legal sanctions. On-farm bush clearing in Australia is subject to such a policy. Where regulatory approaches are not used, the policy approach to conservation by private agents is through negotiated contracts, with fixed payment schemes. This has been the dominant approach in the European Union.

Both the regulatory and fixed-price approaches have their limits. Regulations can be costly to enforce and lead to economically inefficient resource allocation. The fixed payment approach suffers from the information asymmetries that exist between landowners and government. Landowners usually know their supply costs far better than government does. This allows the former to extract information rents. Recognition of such limitations has in recent years increased governments' interest in economic or market-based instruments such as auctions (Murtough, Aretino & Matsyck 2002).

The idea of introducing competitive bidding among landowners has motivated interest in conservation auctions (Latacz-Lohmann & van der Hamsvoort 1997, 1998). The risk of the bid being rejected was thought to mitigate the temptation to overbid above one's true costs. Latacz-Lohmann and van der Hamsvoort show that auctioning such contracts provides the government and society at large with two important benefits: the revelation of otherwise private information on the costs of conservation, and increased efficiency of public spending, due to the reduced extraction of information rents by landowners. Even more importantly, the auction mechanism is shown to be economically more efficient than the traditional fixed-price scheme.¹ Cason, Gangadharan and Duke (2003) go on to show that efficiency is further increased if some information is withheld from landowners in the auctioning process, namely the ecological scoring of their environmental services.

However, these expectations are primarily based on outcomes relating to one-shot auctions. In practice, conservation auctions would be repeated auctions where bidders have the opportunity to learn from previous outcomes. Whether a repeated auction would continue to be efficient or not is an open question. Given the complexity of the repeated auction, the theory provides insufficient guidance for its design and for predicting its outcomes.

A number of issues need to be clarified in the context of auctioning conservation contracts. Will the auction mechanism perform better over time than a (possibly negotiated) fixed-price scheme? Do results previously obtained for a single auction carry forward to repeated auctions? Does increased competition among bidders really reduce private information rents?

1. This result holds subject to the higher transaction costs of an auction mechanism.

Regarding the behaviour of landowners, the key question is what they can learn over time from previous auctions that might defeat the purpose of the auction. Landowners are assumed to maximise profits, and are therefore caught between the temptation to bid higher and the risk of not being selected by the auctioneer. It can be safely assumed that they use available information to their best advantage. As Klemperer (2002) reminds us, the opportunity for learning and strategizing can easily invalidate the predictions of one-shot auction theory.

This study conducts an agent-based computational experiment to assess the performance of repeated auctions under circumstances where the bidders learn from previous experience. Some of the key questions that this study attempts to answer are:

- (a) Does the competitive bidding process succeed in keeping information rents to a minimum in spite of bidder learning? In other words, is competitive bidding efficient when learning is present?
- (b) How efficient is the repeated auction mechanism relative to the much simpler fixed-price scheme?²

Our results indicate that the expected efficiency advantages of the auction might not survive as bidders learn to extract information rents over time, making the outcomes of the auction less attractive relative to simpler mechanisms like fixed-price schemes.

The remainder of the paper is organized as follows. Section 2 reviews the theoretical and experimental study of auctions, identifying a role for agent-based modelling. Section 3 describes the agent-based model of auctioning conservation contracts and specifies the learning algorithm employed. The results of the study are presented and interpreted in section 4. Section 5 summarizes and concludes the paper.

2. The Study of Auctions: A Role for ABM

To date, a complete study of auctions in their natural settings cannot be said to exist. Rather, we have a theory of highly stylised and simplified auction settings. The theory comes only at the cost of a number of simplifications, making its predictions of little practical importance or relevance to practitioners (Aschenfelter 1989; Rothkopf & Harstad 1994). As Box 1 shows, there are many parameters needed to characterise an auction, so that there are numerous types of different auctions. In the auctioning of conservation contracts studied here, some are key features while other choices are optional. Key parameters appear with two ticks and describe essential characteristics of conservation contracts. For instance, because of their long-term nature, conservation auctions have repetitiveness as an essential feature.

2. We do not consider here the weight of an auction's higher transaction costs.

Box 1
Characterisation of Conservation Auctions

CRITERIA	AUCTION CATEGORY DESCRIPTORS			
Repetition	One-shot		Repeated ✓✓	
Timing	Simultaneous ✓		Sequential	
Items	Single		Multiple ✓✓	
	Indivisible	Divisible	Identical	Different ✓✓
Auction type	One-sided ✓✓		Two-sided	
	Open		Sealed-bid ✓	
	1 st price Dutch	2 nd price English	1 st price Standard ✓	2 nd price Vickrey ✓
Reserve price	With reserve price		Without reserve price ✓	
CRITERIA	BIDDER DESCRIPTORS			
Value type	Private value ✓✓		Common value	
	IPV ✓ – Affiliated		Pure – Almost	
Bidder numbers	Fixed ✓		Variable	
	Exogenous ✓		Endogenous	
	Known		Uncertain	
Bidder info	Symmetric		Asymmetric ✓✓	
	Risk neutral		Risk averse	
Bidder strategy	<i>for sequential & repeated auctions:</i>			
	No learning		Learning ✓✓	
Bidder strategy	No collusion		Collusion	
	No gaming		Gaming of auction	
CRITERIA	AUCTION INFORMATION DISCLOSURE BY AUCTIONEER			
Announce	Winning bid or bids		or not ✓	
	Reserve price		or not ✓	
	Best estimate price forecasts		or not ✓	

Note: ✓✓ = Key parameter for conservation auctions;
 ✓ = Optional parameter for conservation auctions; and
 No tick = Parameter could take either value.

2.1 *The Theoretical and Analytical Study of Auctions*

Since Vickrey's seminal contribution in 1961, the game theoretic approach has become the predominant tool for the study of auctions in the economics literature.³ Vickrey's work laid the foundations for the theoretical analysis of auctions, which, over the four decades since, has developed into a large body of literature. A number of reviews of this literature have been written, providing an overview of its achievements and of its shortcomings. Major reviews include Cassady's book (1967) and survey papers by Engelbrecht-Wiggans (1980), McAfee and McMillan (1987), Milgrom (1985, 1989), Wilson (1992), Rothkopf and Harstad (1994), and Klemperer (1999, 2002).

Given our focus on the practical problems of auction design and implementation, and on the prediction of outcomes, Rothkopf and Harstad's (1994) and Klemperer's (2002) critical reviews provide a good vantage point from which to view the achievements and shortcomings of this literature. Most of the theory on auctions has focused on the one-shot auction of a single indivisible good, when real settings most often involve repeated auctions with multiple units (Klemperer 1999, 2002). Theory has also targeted issues like independence of bidder values, revenue equivalence among auction types, bidder risk aversion and informational symmetry, to the detriment of issues closer to real world settings, like bidder learning, variable number of bidders, collusive behaviour and opportunistic 'auction gaming' (Bower & Bunn 2001), all of which can in practice have drastic consequences (Klemperer 2002). It has been clearly demonstrated that relaxing assumptions leads to striking changes in model predictions (Rothkopf & Harstad 1994); in particular, auction outcomes are very sensitive to their informational structure. This observation is particularly important for the new and emerging field of study, auctions for biodiversity conservation on private land (Stoneham, Crowe, Platt, Chaudri, Soligo & Strappazzon 2000).

Underlying the gap between theory and reality lie standard assumptions about human rationality and capacity to process information. The game theory on which the analysis of auctions is predicated rests on the assumption of perfect rationality and on infinite and costless information processing capacity⁴. The literature seemed, until recently, to have forgotten the warnings made by Nash (1951) in his seminal paper. Nash was well aware of the poor descriptive and normative content of game theoretical models. Instead, he suggested they could be used as intellectual (mathematical) tools to build more appropriate models of human behaviour. Recently, the literature on limited and bounded rationality has relaxed these rationality assumptions, introducing human learning to the research agenda on auctions.

In the case of conservation auctions, the gap between theory and reality becomes pronounced due to the dynamic nature of repeated auctions. In other words, these auctions are repeated games (Hausch 1986, 1993; Engelbrecht-Wiggans 1994; Bikhchandani 1988). However, the literature on repeated games is not very useful for studying repeated auctions (Rothkopf & Harstad 1994). First,

3. Before 1961, the strategic interactions between bidders were not directly considered and each bidder was viewed as playing a game against nature; that is, without considering the other players' strategies. Friedman (1956) exemplifies this early approach.

4. As exemplified by the rationalisability (or infinite regress) assumption necessary for the identification of many game equilibria.

the number of action choices (or strategies) available to bidders is not finite. Second, the context of the auction (or the game) changes as participants gain information from past auctions. Third, the number and the identity of bidders do not necessarily stay the same over time. Moreover, the full rationality assumptions underlying game theory have two important consequences: they cannot easily be used in repeated auctions to yield tractable models, and they exclude the possibility of learning (Jeitschko 1998). Examples of this can be found in Black and de Meza (1992); Gale and Hausch (1994); Engelbrecht-Wiggans 1994; and Bernhardt and Scoones 1994. Those studies who did consider learning did so in the context of common value auctions (Engelbrecht-Wiggans & Weber 1983; Hausch 1986; Bikhchandani 1988), but then the learning was about the value of the auctioned item, not about competitors' types and valuations or bidder behaviour (Jeitschko 1998).

A major obstacle to improved problem representation under standard approaches is the resulting loss of analytical tractability. Relaxing simplifying assumptions quickly leads to intractable models or confusing results. To a large extent, these limitations can be attributed to the tools available to economists. More recently economists have turned to experimental economics and agent-based modelling to incorporate increased realism and complexity in modelling auctions.

2.2 *Experimental and Computational Studies of Auctions*

2.2.1 Experimental Economics Experiments using human-subjects are relatively recent and have allowed economists to test theoretical predictions.⁵ The experimental study of auctions consists of running artificial auctions in laboratory settings, often (but not always) with university students. These auctions are in a sense 'real', involving real people, real goods and real money, while at the same time allowing for the control of most or all factors affecting the outcomes.

Experiments do, however, suffer from the somewhat artificial setting they thus set up. Although real goods and real money are involved, the stakes are usually small and the participants typically inexperienced by comparison with auctions occurring in natural settings. This leads to the problem of 'parallelism' (Friedman & Sunder 1994, pp. 15–16). Parallelism does not affect the validity of the experiments themselves relative to the theoretical proposition that needs to be tested, but rather the extrapolation of experimental results to the real world. Thus care must be taken when interpreting experimental results obtained in the context of the assumptions usually made in the theoretical literature.

Furthermore, in parallel to the analytical literature, the experimental literature on auctions, as reviewed by Kagel (1995), has focused on simpler auction types. Issues explored include: the effects of varying the number of bidders in the auction; the impact of uncertainty of the number of bidders; the effect of bidder values 'affiliation'; the effects of reporting price information back to bidders; bidders' attitude to risk; and the potential for and consequences of collusion. Little work has been done on repeated auctions and on the impact of bidder learning and 'gaming' of the auction.

Laboratory experiments have allowed the testing of some major predictions made by auction theory, by showing, for example, that with more competitors,

5. One of the first to do so, since the early 1960's, was Vernon L. Smith (see Smith 1990).

bidders bid higher (more aggressively) in first price auctions, lower in third price auctions, and do not change their bids in second price (Vickrey) auctions. As for the impact of concealing information about the number of rivals, experiments have shown that doing so raises the average market price of the item being auctioned. This is predicted by theory, provided that bidders are assumed to have constant or decreasing absolute risk aversion, a variable that is unobservable.

This last result is illustrative of the problems that can arise when comparing theoretical predictions with observations using human subjects in the study of auctions. In order to interpret experimental findings, one may need to rely on unobservable variables such as agents' degree of risk aversion. This reliance leads to circular reasoning, since knowing whether agents exhibit risk aversion or not has to be inferred from their bidding behaviour. Resorting to computational experiments, where the attributes and behaviour rules of agents are perfectly known or explicitly modelled, can break this logical circle.

The possibility of collusion amongst bidders has been little studied in the theoretical literature, partly because auction theory is primarily based on non-cooperative game theory. Introducing cooperative behaviour in a non-cooperative context introduces complexities that can make the problem analytically intractable, unless some drastic simplifications are made. This poses a problem to experimentalists if they need a clear theoretical benchmark for interpreting laboratory outcomes. Few instances of collusion, and in particular, of long lived and successful collusions, have been observed in laboratory settings (Kagel 1995). This is in stark contrast to what happens in the real world, where collusion is a constant worry for auctioning agents (Cassady 1967; Rothkopf & Harstad 1994; Klemperer 2002). Where conspiracies among bidders were allowed, if not facilitated, in the lab, interpreting the results proved to be difficult and subject to alternative explanations (Isaac & Walker 1985; Kagel 1995, pp. 652–4). These studies showed however the existence and durability of collusion should be linked, at least in common value auctions, to the announcement of a reserve price. Unfortunately for practitioners, the theory predicts that reserve prices in this case will be announced, when in practice they nearly never are. Thus again, theory and empirical observations cannot be put together because of the existence of some unobserved variable linking the two. Computational experiments can be useful here to bridge the gap.

2.2.2 Agent-Based Modelling Computational economics has at least two functions with respect to theory. Its complementary function is to carry a problem beyond the analytical capabilities of theoretical analysis. The supplementary function is to fill in a void when theory is not available, or when it provides insufficient guidance. This is the case with repeated auctions, where bidders interact over time and have opportunities for learning and strategizing.

The 1990s have witnessed the rapid growth of interest in agent-based computational economics (ACE). ACE is the application to economic problems of agent-based modelling (ABM) which is the study of artificial societies of interacting autonomous agents that directly emulate the behaviours of individuals, institutions and environmental components that make up the system being studied (Epstein & Axtell 1996; Parunak, Savit, & Riolo 1998; Tesfatsion 2002). Unlike conventional or equation-based approaches, the starting point in ACE is the

specification of agent attributes and behaviours rather than equations relating system level variables to describe the dynamics of the system. Therefore, ABM is better suited to the study of systems where modelling outcomes can be gainfully enriched through the explicit incorporation of phenomena like agent heterogeneity, local interactions and networking, inductive learning, as well as through the relaxation of restrictive assumptions that are normally imposed under conventional analysis for tractability purposes (Tsfatsion 2002).

The last few years have seen a dramatic increase in the application of ACE to fields as diverse as the evolution of behavioural norms (Axlerod 1997; Axtell, Epstein & Young 2000); financial markets and trade networks (Kirman 1997, 1999; Tsfatsion 1997); labour markets (Pingle & Tsfatsion 2001); electricity auction markets (Nicolaisen, Petrov, & Tsfatsion 2001; Bunn & Oliveira 2001; Bower & Bunn 2001), and structural change in agriculture (Balman, Happe, Kellermann & Kleingarn 2001; Berger 2001).

While pioneering agent-based modelling work in economics began at least with Thomas Schelling's elegant studies from the 1960's and 1970's (Schelling 1978), interest in ABM has grown rapidly as a result of increasing computer power and the emergence of object-oriented programming (OOP) languages such as SmallTalk, C++, Objective-C, and the increasingly popular Java. OOP languages provide a natural means for representing agents as software objects encapsulating attributes and behaviours (rules). Interest in ABM has also been aided by the growing interest in complex systems analysis following work at the Santa Fe Institute and IIASA Laxenburg; the contributions of experimental economics (Kagel & Roth 1995); and, most importantly perhaps, the recent realisation that computational and experimental methods can usefully complement theoretical and analytical approaches as is the case in other sciences (Pingle & Tsfatsion 2001).

2.2.3 Economic Experiments and Agent-Based Models as Complementary Tools Experiments and agent-based models can be compared on a number of criteria (Vriend 2002). In terms of the approach to inference, the experimental and ABM methods are 'bottom-up' rather than 'top-down' approaches to science, allowing for the emergence of aggregate or macro-level phenomena from the interaction of individual agents. For this reason, both methods can be used to analyse problems that are complex and difficult to tackle analytically. The two approaches differ on a number of criteria as indicated in table 1. Some of the major differences between the two relate to:

- Model construction or experimental setup (specification of behavioural rules);
- Degree of control the researcher has over the experiment; and
- Temporal length of the problems addressed in the experiment.

The temporal length of the analysis is especially relevant for the study of conservation auctions. The two approaches can be employed in a complementary manner. For example, experiments can be useful in generating data needed to characterize the behaviour of computational agents while agent-based models can be used to thoroughly explore the longer term outcomes, something which would be too tedious and too expensive with experiments.

Table 1
Comparison of Economic Experiments and Agent-Based Models

Criteria	Agent-based Models	Economic Experiments
Approach to inference, or micro-macro relationship	Individual based or ‘bottom up’ approach allowing for the emergence of aggregate or macro phenomena from the interaction of <i>software agents</i> representing economic entities	Individual based or ‘bottom up’ approach allowing for the emergence of aggregate or macro phenomena from the interaction of <i>human participants</i>
Specification of behavioural rules	Allows for the specification of behavioural rules underlying choices. Rules of game and behaviours of agents can be evolved	Need to clarify rules of the game and incentives but not rules for making choices. Rules and incentives usually held constant for a given session
Informational problems	Adequacy of agents’ behavioural rules (Are they correctly specified?)	Observability of variables from theoretical models (e.g. risk aversion, beliefs)
Degree of control	Investigator has complete control over attributes and behaviours of computational agents	Incomplete, as participating individuals might have different perceptions than those intended by the researcher
Explanation of agents’ choices	Known, or can be traced back	Need to be inferred from subject’s choices
Temporal length of analysis	Flexible. Long-term analysis causes no difficulties	Short, due to cost and time considerations, except (sometimes) over the Internet
Representativeness / realism	Subject to the accuracy of investigator’s specification of details and scope of analysis	Problems of ‘Parallelism’, especially magnitude of stakes (much smaller in the lab than in real world) and subjects’ own previous experience
Data	Need to calibrate model and create initial population representing subject population	Generated by participants’ choices and strategies; can provide the foundation for calibrating computational agents
Cost	Inexpensive, except for analyst’s time	Higher costs of conducting experiments as well as any incentives provided to participants

3. The Agent-Based Auction Model

3.1 Overview

The agent-based model used in this study is simple. Two types of agents representing the actual players in a real auction are included in the model. These are:

- (a) Farmer agents bidding for environmental conservation contracts. Each farmer has an environmental quality value and an opportunity cost associated with putting the land being offered under conservation.

- (b) A government agent which selects winning farmers and awards contracts based on the criteria applying under the particular auction format being used. The government agent has a fixed budget.

Each auction round incorporates the following three major steps or activities.

- Step 1: Farmers construct their bids. The bids farmers make depend on their respective opportunity costs, their previous bid prices as well as their success or failure in the previous auction. For example, if a farmer agent was successful in the previous bid, then it tends to bid the same or a higher price. In the very first period, farmers have no prior experience and start by bidding their true opportunity costs. The details of the learning algorithm used to select the mark-ups are described in more detail below.
- Step 2: The government agent ranks the bids submitted by farmers based on the auction criteria, selects winners accordingly and informs each bidder whether it has been successful or not.
- Step 3: Farmer agents update their contract status based on the message from the government agent.

Bidders are assumed to submit their true opportunity costs as bids (i.e. bid honestly) in the first round. This assumption was made because it is reasonable to assume that, before they acquire information about the auction environment, bidders are likely to select bids that are based on their respective opportunity cost rather than bids that are random. Experimental evidence also lends support to the assumption that bidders' initial bids are close to their opportunity costs. Of course, the final outcomes or learnt bids will not depend on how bidders start bidding. However, the distinction is important in terms of the efficiency of auctions over time. For example, bids submitted in the first round of the BushTender auction (Stoneham, Chaudhri, Ha, & Strappazzon 2003) follow a pattern that would be substantially altered over time as a consequence of learning.

3.2 Auction Implementation

A discriminatory sealed-bid auction format targeting program objectives rather than enrolment is used by the government agent. Under this format each winning bidder is paid only its bid amount, and thus different bidders may be receiving different payments for the same service. The selection process under this auction type involves two steps:

- Step 1: Ranking all bids received from farmers based on the ratio of environmental quality per bid dollar. The environmental quality or benefit value attached to each bid is an ecological score or index value assigned by scientists to the farmer's tendered activities; it is not a dollar measure of the benefits.
- Step 2: Awarding conservation contracts to bidders starting from the one offering the highest environmental quality per dollar until the budget is fully allocated. The government agent has no reserve price. These features capture the details of the conservation auctions (BushTender and others) currently in operation in Australia; none of them have reserve prices.

The discriminatory auction format is expected to minimize the extraction of information rents from bidders enabling the government to award more contracts than is possible under a fixed payment scheme in which each bidder would be paid the same price per environmental quality (i.e. the bid to environmental quality ratio of the marginal or last ranked winner). The discriminatory auction format has been used in the US CRP program and in the BushTender trials in Victoria.

The details of the experimental setup used to generate the results reported in this paper are as follows. The population of bidders is 100. For simplicity, the bidder population was set to be heterogeneous only in opportunity costs, while the environmental benefit score of the conservation activities that each bidder offers to undertake was set equal to 1.0. These opportunity cost, benefit, and budget values should be considered as a normalization of the actual values, which does not affect the results. The opportunity cost values were randomly drawn from a uniform distribution between 0.5 and 1.5. The government budget is fixed at \$30 (e.g. million). The level of the budget was chosen to be roughly equal to 30% of total opportunity cost. Thirty successive discriminatory auctions are simulated. One hundred runs or replications of these 30 successive auctions (using different random seeds) were used to generate the average results discussed below.⁶

For comparison purposes, fixed-price schemes were also simulated with identical parameter specifications as for the auction. Two fixed prices, namely, 0.90 and 1.00, were used. These fixed prices relate to 90% and 100% of the average opportunity cost of the population of bidding agents. Results obtained under these schemes allow us to see how the results of the auction compare to a take-it-or-leave-it price scheme based on estimates of average opportunity cost levels.

3.3 Learning Algorithm

The repeated auction constitutes an environment where the bidders can acquire and utilize information from previous auction outcomes. A variety of learning theories or algorithms have been used to model the behaviour of economic agents. But reinforcement learning and direction learning theories are more relevant in this context because of the information input they require. These two algorithms do not require that the bidder know the forgone payoffs for strategies that he did not select in previous moves. Roth and Erev's reinforcement algorithm (Roth & Erev 1995; Erev & Roth 1998) is based on the reinforcement principle that is widely accepted in the psychology literature. Under this learning algorithm, an agent's tendency to implement an action is strengthened (*reinforced*) or weakened depending upon whether the action produces favourable results or not. The algorithm also allows for exploration or *experimentation* with new alternatives. The Roth-Erev algorithm or modified versions of it have been used in several agent-based studies of electricity markets (e.g. Nicolaisen, Petrov & Tesfatsion 2001; Bunn & Oliveira 2001; Hailu & Schilizzi 2003).

The Roth-Erev algorithm does not enforce a direction on the adjustment made in choices between periods. Learning direction theory (Selten & Stoecker 1986; Selten, Abbink & Cox 2001), on the other hand, asserts that ex-post rationality is the strongest influence in adaptive behaviour. According to this theory, more

6. Reference to results from one particular run relate to the run with a random seed of 1972.

frequently than randomly expected, parameter changes, if they occur, are oriented towards additional payoffs that might have been gained by other actions. This theory is qualitative and does not specify how bids should be adjusted between rounds.

In this study, we enforce a direction on the bid adjustments based on ex-post rationality. The following very simple learning rule for probabilistically changing agent's prices is used:

1. If an agent wins a contract in the previous auction, it will maintain the same bid or increment it by 10%.
2. If an agent loses in the previous auction, it will maintain the same bid or lower it by 10%.

The agent is as likely to change the bid as he is to maintain it for the next round. In addition, agents do not bid below opportunity cost and any bids above 300% percent of individual opportunity cost are not considered. The upper limit is based on the ratio of the lowest to highest opportunity costs and is never reached by agents in the simulations.

4. Analysts of Results

Results from the agent-based model are presented and analysed below. We shall first present a brief description of how we would expect bidder agents to behave in the auction environment. This will be followed by a discussion of the bidding patterns observed in the agent-based model. The performance of the auction is then presented and compared to that of a fixed-price scheme in terms of efficiency and equity outcomes.

4.1 Expected Bidding Behaviour

It is easier to think through the optimal or learnt bidding patterns if we start by distinguishing among infra-marginal winners, the marginal winner, the marginal loser and infra-marginal losers. The marginal winner is the bidder ranked last by the auctioneer. The marginal loser is the bidder ranked highest among the losers. If there are n winners, the marginal winner is the one ranked n^{th} and the marginal loser is the one ranked $(n+1)^{\text{th}}$. Since the optimal or learnt bids should be non-decreasing in opportunity costs, the marginal winner should be the one with the highest opportunity cost among the winners. As the last winner, the marginal winner's bid must be somewhere between his opportunity cost and the opportunity cost of the marginal loser. That is, the marginal winner would lose if he submitted a bid that is equal to or greater than the opportunity cost of the marginal loser and the latter bid his opportunity cost. Therefore, the best strategy for the marginal bidder is to bid just under the opportunity cost of the marginal loser.

An infra-marginal bidder would win as long as his bid is not higher than that of the marginal winner. The net income of an infra-marginal bidder is maximized when he bids the highest he can without losing the auction, that is, when he aligns his bid with that of the marginal winner. Finally, we would expect the bidders with higher costs (the marginal and the infra-marginal losers) to submit bids equal to their opportunity costs or higher. The expected bidding behaviours are similar to

the equilibrium pricing strategies in a Bertrand oligopoly where the lower cost supplier prices the commodity at the marginal cost of the higher cost producer (see Elberfeld & Wolfstetter 1996).

4.2 *The Extraction of Rents: How Does it Happen?*

The most fundamental observation in the results presented below relates to the mechanism by which a simple learning process involving the use of past information and experimentation with new bid prices enables bidders to extract rents. Over time, bidders utilize the information they obtain from the auctions, either exploiting further their previously successful price choices or using the feedback from the previous auction to experiment with new bid mark-ups. This process of learning and adjustment leads the successful or infra-marginal⁷ bidders to mark up their bids up to that of the marginal bidder.

The following two diagrams help explain the outcomes of this learning process. The opportunity cost of involvement for the bidders is indicated on the x-axis. Contracting bidders start by bidding (and getting paid) different prices, reflecting their individual opportunity costs. The program payment rates are marked by asterisks in figures 1(a) and 1(b). In the first period the bid to environmental quality rates for the 42 winners fall in the range of 0.50 to 0.93 (see fig. 1a). But these differences in bid rates disappear over time. The spread in these prices is reduced as the infra-marginal bids catch up with the marginal winning bid eventually forming a narrow band of bid rates ranging from 0.90 to 0.99, as shown in figure 1b (for the 15th period).⁸ Just above this narrow band of 33 winning bids for that round is another band of ‘active’, but currently not selected, competitors. These two narrow bands represent the two components of the ‘active’ bidders some of who replace each other from one period to the next—generating a sort of ‘basket of crabs’ effect. The alignment of bid rates around those of the marginal winners can also be seen in figures 2 and 3. Figure 2 shows that the number of winning bidders that submit low bids diminishes over time. In period 5, a majority bid below 0.8; in period 15, a minority bid below 0.8; and in the final period (30), none bid below 0.8. Figure 3 illustrates the effect of this bid alignment process on the supply curve of environmental benefits. Initially (e.g. periods 1 and 5), the use of the auction is likely to generate efficiency gains since it would allow the government to exploit the heterogeneity in the opportunity costs and bids in the discriminatory auction. With bidder learning, however, the heterogeneity in bids at the lower end disappears and the supply curve that the government faces is a kinked supply curve with a flat section. In other words, the procuring agency loses most of the benefits from the heterogeneity in opportunity costs.


The observations regarding the dynamics of rent extraction made in this study are similar to those in Hailu and Schilizzi (2003) who investigate the performance of auctions but for an environmental quality that grows (along a logistic curve) in response to conservation. In that study, benefits are also defined as the level rather than increments in environmental quality. The growth in the resource reinforces the adjustments resulting from learning and leads to much higher rent extraction rates

7. An infra-marginal bidder is one who is ranked higher by the auctioneer than the marginal or lowest ranked winner.

8. It should be noted, however, that the convergence occurs earlier than period 15.

(Hailu & Schilizzi 2003). However, the alignment of bid rates for a growing environmental quality occurs on a falling marginal cost of benefit provision in the problem investigated in that work.

Figure 1
Distributions of Bid and Government Rates for Periods 1 and 15



Note:
 * Gov. payment to env. quality ratio
 ◇ Bid to env. quality ratio

Figure 2
Bid-Benefit Ratio Versus Percentage of Times That a Bidder Was Selected

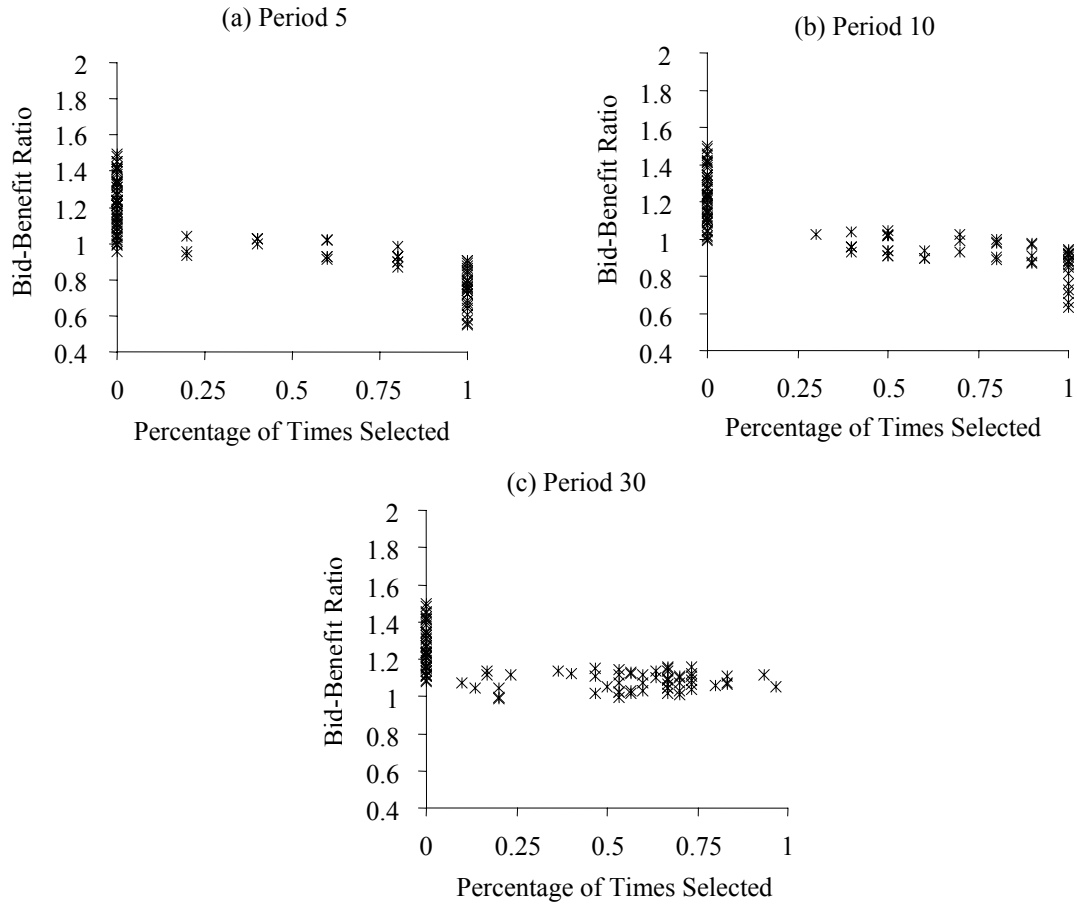
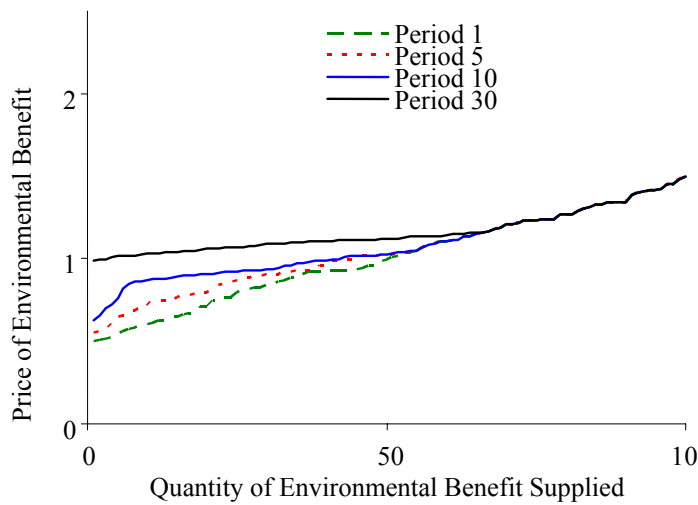


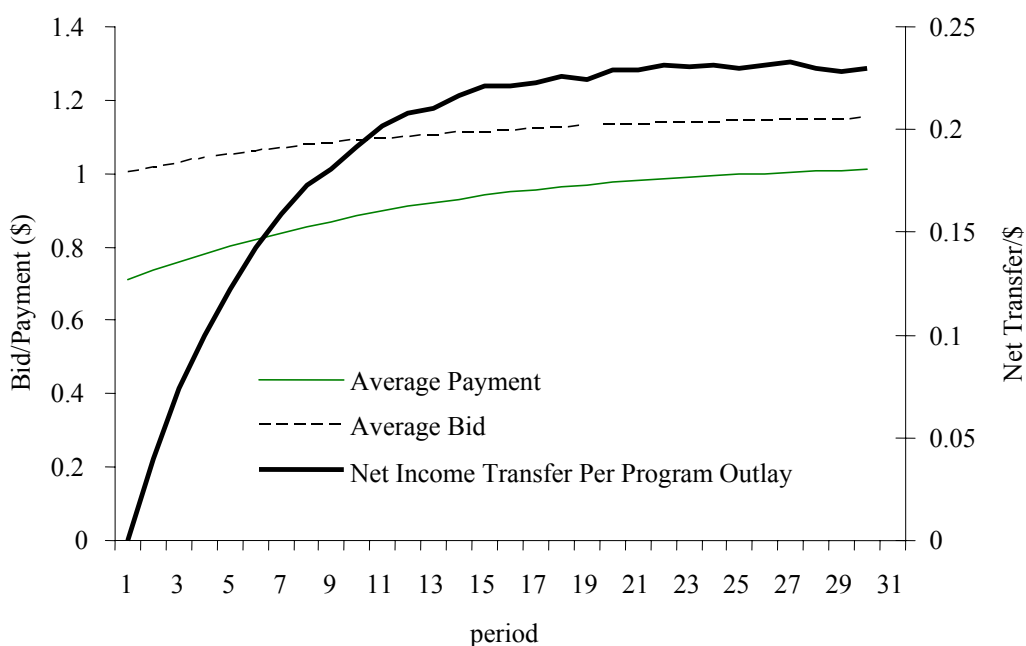
Figure 3
Environmental Benefit Supply Curves Over Time



4.3 Auction Performance

4.3.1 *Efficiency of Auction* Total environmental benefits as well as benefits per program outlay deteriorate over time with the auction. This trend is accompanied by two trends with equity and efficiency implications. First, there is a crowding out of program participants as the number of bidders hired by the auctioneer falls down. After bids are fully adjusted, the number of winning bidders is reduced by 15% compared to what it was in the beginning of the program. By the 30th round, the number of winning bidders is about 70% of its initial level. Second, the proportion of payments over and above true opportunity costs increase over time. The proportion of net income transfers in program payments rises to about 20% within the first 10 rounds and stabilizes at about 22% by the 15th round; in other words, *more than one-fifth of the payments constitute information rents*. See figure 4. This shows that the efficiency performance of the auction is lower than one would expect based on the literature on single shot auctions.

Figure 4
Some Performance Indicators for the Discriminatory Auction



4.3.2 *Program Participation and Distribution of Payments* About 44% of the bidders are never selected. Generally, these tend to be the high cost bidders. Most of the remaining 56 bidders manage to procure three or more contracts from the 30 auction rounds, with the average number of contracts acquired being 17. Only two of these bidders acquire more than 25 contracts while 8 acquire less than 10. All of these 56 bidders have acquired contracts in the last 15 rounds. These numbers clearly indicate the intensity of the competition among the ‘active’ bidders, reminding us of a ‘basket of crabs’ where each individual manages to get on top but is quickly pushed back under.

This ‘active’ set of bidders includes bidders with the lowest opportunity costs of conservation. The total opportunity cost of the first 41 lowest opportunity cost

bidders is within the reach of the government budget of \$30. In other words, under perfect information, this group would have been hired in its entirety to undertake the conservation works. However, after the bid adjustment on the part of bidders has occurred, the auction mechanism manages to hire the services of only 28 bidders (i.e. 69% compared to the perfect information setting without rent extraction). The auction transforms the group of 'active' bidders into a subgroup that continuously extracts rents from the auction mechanism.

4.3.3 Fixed-Price Scheme Versus Discriminatory Auction A widely held belief is that an auction outperforms a fixed-price scheme in terms of economic efficiency. This may be true of single shot auctions, but, as will now appear clearly, is unlikely to be true of repeated auctions with learning. To test this belief, we used the agent-based model to compare the performances of the auction and fixed-price mechanisms. To the extent that bidders initially bid close to their opportunity costs and learn over time, the results highlight elements of the auction dynamics that would not be fully appreciated in the context of single shot auctions. These are summarized in table 1.

The performance of the auction relative to the fixed-price scheme depends on the level of the fixed-price employed in the latter. With a fixed-price set at 90% of the average opportunity cost and looking at the results from fully adjusted bids in round 30, the auction is found to be inferior in terms of efficiency, participation, and the provision of environmental benefits. The auction provides lower rates of participation and environmental benefits per dollar of program outlay (by a factor of 0.89). The proportion of net income transfer in program payments is higher under the auction than under the fixed-price scheme. In other words, the informational rents extracted are higher, by 5%. The outcomes of the auction are more similar to those of the fixed payment scheme with higher fixed-prices (e.g. column 5 of table 1) and, therefore, higher built-in net income transfers.

In sum, the auction is not necessarily superior to fixed payment mechanisms with a built-in check on the amount of money that can be handed out to individual bidders. The auction limits the total, but not the individual amount of payments. As a result, bidders collectively determine the cut-off price, leading to the division of the budget into a smaller number of bigger chunks. This process of concentration reflects the collective marking-up of individual bids that occurs among successful bidders. Relative to the fixed-price scheme, the auction mechanism has obvious implications in terms of distributional equity.

5. Conclusions

This study has attempted to evaluate the performance of conservation auctions under dynamic settings using agent-based modelling. The economic agents represented in the model are farmers or landholders bidding to offer their environmental services and a government agent with a fixed budget awarding conservation contracts through auctions. The auction mechanism evaluated is a discriminatory sealed-bid auction in which a government agent ranks submitted bids according to their respective benefit-costs ratios, where the benefits are represented by environmental quality scores. Bidders are allowed to learn from previous auction results and experiment with alternative bidding strategies. The

performance of the auction mechanism is also compared to the performance of a fixed payment scheme under two alternative price scenarios.

Table 2
Performance of Environmental Conservation Discriminatory Auction Relative to Fixed-Price Schemes

	Relative to fixed-price payment set at 90% of mean opportunity cost		Relative to fixed-price payment set at 100% of mean opportunity cost	
	All period average	Period 30	All period average	Period 30
Proportion of winners among participants	0.99	0.89	1.09	0.97
Net income transfer per program outlay	0.86	1.05	0.76	0.92
Environmental benefit value per program outlay	0.99	0.89	1.10	0.99
Average payment to winners	1.02	1.12	0.92	1.01

The simulation results provide insights into the long-term performance of auctions that lead us to question commonly held assumptions. First, there is a marked difference between the short-term and long-term efficiency of conservation auctions. This would be true only if the agents' initial bids are closely related to their individual opportunity costs rather than being random bids. Repeated auctions involve the payment of substantial rents to winning bidders as the latter learn to mark up their bids matching the bid corresponding to the marginal winner. The auction mechanism also leads to a much lower program participation rate than would be possible without rent extraction. The total number of participating bidders is split into two main groups as the auction rounds proceed. The first group is mainly one of high-cost bidders that are quickly rendered 'inactive' and do not succeed in acquiring conservation contracts. The second group constitutes the 'active' bidders that have learnt to price their services within a narrow margin of the marginal or most expensive winner.

Second, the fixed-price scheme with reserve prices set at 90% of average opportunity cost outperforms the auction mechanism in terms of participation and efficiency. The auction mechanism compares more favourably only to fixed-price schemes with higher payment levels. This is because the higher reserve prices involve built-in net income transfers that are similar to those achieved by bidders who learn to 'game' the auction over time.

The main conclusion from this study is that with bidder learning the efficiency benefits of single-shot auctions do not necessarily extend to repeated auctions. Even as simple a learning process as directional learning can ensure that bidder prices adjust to extract almost all information rents despite competitive

bidding conditions.⁹ This is even an understatement, as in reality bidder learning is likely to involve interaction or communication among landholders rather than individual learning from private experience only, as was simulated in this study.

Finally, there are several interesting issues closely related to the subject of this study that have not been investigated here. First, the potential gains from auctions as opposed to other schemes might depend on the distribution of opportunity costs, number of bidders and the size of the government budget. Second, this study allowed for individual learning where each bidder makes adjustments using only his own previous bidding result. Learning in reality has social elements and involves communication among bidders. Such communication is likely to increase the speed of learning and might have significant implications for the performance of the auction mechanism. Thus, there is a need for further exploration of the impact of the structure of the bidder population and different auction parameters including the spatial scope, information revelation and frequency of auctions. The authors are currently investigating some of these issues.

(Date of receipt of final transcript: September, 2004.
Accepted by Robert Marks, Area Editor.)

References

- Aschenfelter, O. 1989, 'How auctions work for wine and art', *Journal of Economic Perspectives*, vol. 3, pp. 23–6.
- Axelrod, R. 1997, *The Complexity of Cooperation. Agent-Based Models of Competition and Collaboration*, Princeton Studies in Complexity, Princeton University Press, N.J.
- Axtell, R, Epstein, J. & Young, H.P. 2000, 'The emergence of classes in a multi-agent bargaining model', Working Paper No. 9, Center on Social and Economic Dynamics, Brookings Institution, USA.
- Balman, A., Happe, K., Kellermann, K. & Kleingarn, A. 2001, 'Adjustment costs of agri-environmental policy switchings. A multi-agent approach', paper presented to the American Agricultural Economics Association 2001 Annual Meeting, Aug. 5–8th, Chicago.
- Berger, T. 2001, 'Agent-based spatial models applied to agriculture: A simulation tool for technology diffusion, resource use changes and policy analysis' *Agricultural Economics*, vol. 25, pp. 245–60.
- Bernhardt, D. & Scoones, D. 1994, 'A note on sequential auctions', *American Economic Review*, vol. 84, pp. 653–57.
- Bikhchandani, S. 1988, 'Reputation in repeated second-price auctions', *Journal of Economic Theory*, vol. 46, pp. 97–119.
- Black, J. & de Meza, D. 1992, 'Systematic price differences between successive auctions are no anomaly', *Journal of Economics and Management Strategy*, vol 1, pp. 607–28.
- 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*, vol. 25, pp. 561–92.

9. The outcomes in this model are not specific to the learning algorithm used. We have checked this with reinforcement learning. The results remain the same but reinforcement learning leads to slower convergence than Selten's directional learning algorithm.

- 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*, vol. 5, no. 5, pp. 493–503.
- Cason, T., Gangadharan, L. & Duke, C. 2003, 'A laboratory study of auctions for reducing non-point source pollution', *Journal of Environmental Economics and Management*, vol. 46, pp. 446–71.
- Cassady, R. Jr. 1967, *Auctions and Auctioneering*, University of California Press, Berkeley, CA.
- Elberfeld, W. & Wolfstetter, E.G. 1996, 'A dynamic model of Bertrand competition with entry', December 1996, <http://ssrn.com/abstract=2174>
- Engelbrecht-Wiggans, R. 1980, 'Auctions and bidding models', *Management Science*, vol. 26, pp. 119–21.
- Engelbrecht-Wiggans, R. 1994, 'Sequential auctions of stochastically equivalent objects', *Economic Letters*, vol. 22, no. 1–2, pp. 87–90.
- Engelbrecht-Wiggans, R. & Weber, R.J. 1983, 'A sequential auction involving asymmetrically-informed bidders', *International Journal of Game Theory*, vol. 12, pp. 123–7.
- Epstein, J.M. & Axtell, R. 1996, *Growing Artificial Societies: Social Sciences From the Bottom Up*, Brookings Institution Press, Washington, D.C.
- Erev, I. & Roth, A.E. 1998, 'Predicting how people play games with unique, mixed-strategy equilibria', *American Economic Review*, vol. 88, pp. 848–81.
- Friedman, D. & Sunder, S. 1994, *Experimental Methods. A Primer for Economists*, Cambridge University Press.
- Friedman, L. 1956, 'A competitive bidding strategy', *Operations Research*, vol. 4, pp. 104–12.
- Gale, I.L. & Hausch, D.B. 1994, 'Bottom-fishing and declining prices in sequential auctions', *Games and Economic Behaviour*, vol 7, pp. 318–31.
- Hailu, A. & Schilizzi, S. 2003, 'Investigating the performance of market-based instruments for resource conservation: the contribution of agent-based modelling', paper presented at the 47th Annual Conference of the Australian Agricultural and Resource Society, Fremantle, Western Australia, 12–14 February.
- Hausch, D.B. 1986, 'Multi-object auctions: Sequential vs simultaneous sales', *Management Science*, vol. 32, pp. 1599–610.
- Hausch, D.B. 1993, 'A common value auction model with endogenous entry and information acquisition', *Economic Theory*, vol. 3, pp. 315–34.
- Isaac, R.M. & Walker, J.M. 1985, 'Information and conspiracy in sealed bid auctions', *Journal of Economic Behaviour and Organization*, vol. 6, pp. 139–59.
- Jeitschko, T.D. 1998, 'Learning in sequential auctions', *Southern Journal of Economics*, vol. 65, no. 1, pp. 98–112.
- Kagel, J.H. 1995, 'Auctions: A survey of experimental research', in *The Handbook of Experimental Economics*, (eds.) J.H. Kagel and A.E. Roth, Princeton University Press, N.J., pp. 501–85.
- Kirman, A.P. 1997, 'The economy as an interactive system', in *The Economy as an Evolving Complex System*, II, (eds.) W.B. Arthur, S.N. Durlauf, & D.A. Lane, Proceedings Volume XXVII, Santa Fe Institute Studies in the Sciences of Complexity, Addison-Wesley, Reading, MA, pp. 491–31.
- Kirman, A. 1999, 'Interactions and markets', in *Beyond the Representative Agent*, M. Gallegati & A. Kirman (eds), Edward Elgar, Cheltenham UK and Northampton MA, USA, pp. 1–44.
- Klemperer, P. 1999, 'Auction theory: A guide to the literature', *Journal of Economic Surveys*, vol. 13, no. 3, pp. 227–86.

- Klemperer, P. 2002, 'What really matters in auction design', *Journal of Economic Perspectives*, vol. 16, no. 1, pp 169–89.
- Latacz-Lohmann, U. & van der Hamsvoort, C. 1997, 'Auctioning conservation contracts: A theoretical analysis and an application', *American Journal of Agricultural Economics*, vol. 79, pp. 407–18.
- Latacz-Lohmann, U. & van der Hamsvoort, C. 1998, 'Auctions as a means of creating a market for public goods from agriculture', *Journal of Agricultural Economics*, vol. 49, no. 3, pp. 334–45.
- McAfee, R.P. & McMillan, J. 1987, 'Auctions and bidding', *Journal of Economic Literature*, vol. 25, June, pp. 699–738.
- Milgrom P. 1985, 'Auction theory', in *Advances in Economic Theory: Fifth World Congress*, B. Truman (ed.), Cambridge University Press.
- Milgrom P. 1989, 'Auctions and bidding: A primer', *Journal of Economic Perspectives*, vol. 3, no. 3, pp. 3–22.
- Murtough, G., Aretino, B. & Matsyck, A. 2002, 'Creating markets for ecosystems services', Productivity Commission Staff Research Paper, AusInfo, Canberra, Australia.
- Nash, J. 1951, 'Non-cooperative games', *Annals of Mathematics*, vol. 54, pp. 286–95.
- 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*, vol. 5, no. 5, pp. 504–23.
- Parunak, H.V.D., Savit, R. & Riolo, R.L. 1998, 'Agent-based modelling vs. equation-based modelling: A case study and users' guide,' in *Multi-Agent Systems and Agent-Based Simulation*, J. S. Sichman, R. Conte, and N. Gilbert (eds.), Springer, Singapore, pp. 11–25.
- Pingle, M. & Tesfatsion, L. 2001, 'Non-employment benefits and the evolution of worker-employer cooperation: Experiments with real and computational agents', Iowa State University Economic Report Series No. 55, June, pp. 52.
- Roth, A.E. & Erev, I. 1995, 'Learning in extensive form games: experimental data and simple dynamic models in the intermediate term', *Games & Economic Behavior*, vol. 8, pp. 164–212.
- Rothkopf, M.H. & Harstad, R.M. 1994, 'Modeling competitive bidding: A critical essay', *Management Science*, vol. 40, no. 3, pp. 364–84.
- Schelling, T.C. 1978, *Micromotives and Macrobehavior*, W.W. Norton & Co., New York, London.
- Selten, R. & Stoecker, R. 1986, 'End behaviour in sequence finite prisoners' dilemma supergames: A learning theory approach', *Journal of Economic Behavior and Organization*, vol. 7, pp. 47–70.
- Selten, R., Abbink, K. & Cox, R. 2001, 'Learning direction theory and the winners curse', Discussion Paper 10/2001, Department of Economics, University of Bonn, Germany.
- Smith, V.L. 1990, *Experimental Economics*, Schools of Thought in Economics series no. 7, Edward Elgar, Cheltenham UK and Northampton MA, USA.
- Stoneham, G., Crowe, M., Platt, S., Chaudri, V., Soligo, J. & Strappazon, L. 2000, 'Mechanisms for biodiversity conservation on private land', Department of Natural Resources and Environment, State of Victoria, Australia, July, p. 53.
- Stoneham, G. & Chaudri, V. 2000, 'Auction design for land use change in the Murray-Darling Basin', Department of Natural Resources and Environment, State of Victoria, Australia. November, p. 59.
- Stoneham, G., Chaudhri V., Ha A. & Strappazon, L. 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.

- Tesfatsion, L. 1997, 'A trade network game with endogenous partner selection', in *Computational Approaches To Economic Problems*, H. Amman, B. Rustem & A.B. Whinston (eds), Kluwer Acad. Press, Dordrecht, NL, pp. 249–69.
- Tesfatsion, L. 2002, 'Agent-based computational economics: Growing economies from the bottom up', *Artificial Life*, vol. 8, no. 1, pp. 55–82.
- Vickrey, W. 1961, 'Counterspeculation, auctions and competitive sealed tenders', *Journal of Finance*, vol. 16, pp. 8–37.
- Vriend, N.J. 2002, 'Was Hayek an ace?', *Southern Economic Journal*, vol. 68, no. 4, pp. 811–40.
- Wilson, R. 1992, 'Strategic analysis of auctions', in *Handbook of Game Theory*, R.J. Aumann & S. Hart, vol. 1, pp. 228–79.