

Applying Evolutionary Search to a Parametric Family of Auction Mechanisms

Andrew Bye

HP Labs,
Filton Road, Stoke Gifford,
Bristol, BS34 8QZ
andrew.bye@hp.com

Abstract. In this paper we describe an evolution-based method for evaluating auction mechanisms, and apply it to a space of mechanisms including the standard first- and second-price sealed bid auctions. We replicate results known already in the Auction Theory literature regarding the suitability of different mechanisms for different bidder environments, and extend the literature by establishing the superiority of novel mechanisms over standard mechanisms, for commonly occurring scenarios. Thus this paper simultaneously extends Auction Theory, and provides a systematic method for further such extensions.

1 Introduction

Auctions are an important class of mechanisms for resolving multi-agent allocation problems of various types. There exists a substantial body of work (see, e.g. Klemperer (1999) for a review) regarding the theory underlying auctions, most of which focuses on the problem of how to design them so as to achieve some desired outcome for the auctioneer. In situations where the auctioneer plays the role of seller, this outcome is often revenue maximization, and many results of a qualitative nature are known regarding the suitability of different mechanisms under different assumptions on the economic scenario under consideration.

In parallel with this work, researchers have begun investigating how to design autonomous agents capable of participating in auctions (Preist, Bye & Bartolini (2001), Anthony et al. (2001), Boutilier et al. (1999b), Bye et al. (2002), Cliff & Bruten (1998), Walsh et al. (2002), etc.). Often such study is motivated by the possibility that suitable autonomous agents will be superior to humans in making (possibly quite complex) economic decisions, and indeed (Das et al. 2001) report human experiments to substantiate this possibility. When the agents acting in markets are non-human, the space of potential market designs increases markedly, since mechanisms that might seem “non-sensical” or difficult to interpret for humans can be considered.

Recently, a few papers have begun to address the confluence of these ideas, taking inspiration from the Auction Theory work on mechanism design, extending it into new design spaces that might have been infeasible before, and adding

a degree of automation to the design process: for example, Cliff (2002) describes an application of Genetic Algorithms (Holland 1975) to the choice of a continuous parameter Q_s governing the probability that a seller will be chosen to shout in a given round of a stylized continuous double auction.

This paper continues the direction of such work, examining a space of auction mechanisms that includes the standard first- and second-price auctions, using GAs applied to a multi-agent system to evolve good players for each mechanism under consideration. We find that under several classes of non-pathological conditions (e.g. bidders are risk-averse, and are unaware of how many players they will face in a given auction), there exist exotic sealed bid mechanisms which are expected to return significantly higher revenue to the auctioneer than either the first- or second-price sealed bid mechanisms. See Section 4 for more details.

The paper is laid out as follows: in the next two sections we discuss the methods used, introducing relevant Game Theory concepts as needed. In Section 4 we describe the results of our experiments. In Section 5 we describe in more detail the relationship between these results and others in the literature, and we conclude in Section 6.

2 Auction Theory

2.1 Terms and Notation

In this paper we study sealed-bid auctions, in which a good is put up for sale, and each potential buyer submits a bid to the auctioneer; the auctioneer chooses a winner, and allocates payments to each agent. In most variants of this type of auction, the good is awarded to the buyer who submits the highest bid, and only the winner pays. In a **first-price** auction, the winner’s payment is equal to her bid; in a **second-price** auction, the winner’s payment is equal to the second highest bid.

In order to analyse how bidders might be expected to behave in such auctions, we need to specify how they are motivated, i.e. what is the good worth to each agent. We use a model in which agents are only interested in their own awards and payments, and there is some intrinsic monetary “value” v associated to the good. All outcomes can therefore be represented by a single number, the monetary gain the agent makes. This is $v - x$ for a win with payment x , and $-x$ for a non-win with the same payment. The risk preferences of agents are differentiated by use of a von Neumann – Morgenstern utility function u , so that an agent strictly prefers a selection of possible outcomes x_i with corresponding probabilities p_i , over a second selection of possible outcomes y_j with corresponding probabilities q_j , if and only if

$$\sum_i p_i u(x_i) > \sum_j q_j u(x_j). \quad (1)$$

In this representation, assuming twice-differentiability of u , an agent for which $u''(x) = 0$ is known as **risk-neutral**; if $u''(x) < 0$, the agent is **risk-averse**, and if $u''(x) > 0$, the agent is known as **risk-seeking** (see Fig. 1).

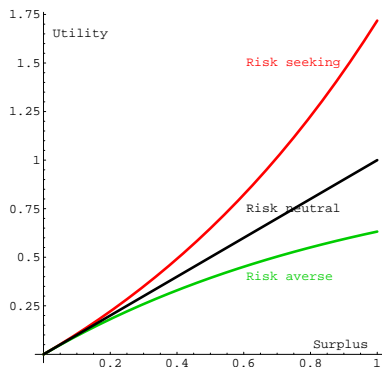


Fig. 1. Typical utility functions for a risk- seeking, neutral and averse agent.

The value of a good to an agent can be independent of the value of the good to other agents, or it can be derived from information about how other agents value the good. In the former case we say that the agent has a **private** value for the good, and in the latter case that there is some **common** value component. We treat both these cases by postulating that each bidder receives a “signal”, and that the value of the good to the agent is some specified function of *all* the agents’ signals. Since a bidder only necessarily knows her own signal, her decision problem may in general involve guess work about the worth of the good.

2.2 Revenue for the Auctioneer

In this paper we study mechanisms from the point of view of the amount of money they are expected to make the seller. Perhaps the most important result in this area is the **Revenue Equivalence Theorem** (see e.g. Krishna 2002), which states that if (concerning the environment)

- all agents are risk-neutral,
- all bidders’ signals are picked from a common, known distribution,

and if (concerning the mechanism),

- in equilibrium, the good always goes to the bidder with the highest signal,
- any bidder whose signal is the lowest possible expects to make nothing,

then the expected revenue to the seller is the same, independent of the mechanism.

This rather surprising result means that, subject to these hypotheses, it *doesn’t matter* what type of auction a seller runs, he should expect to make the same amount of money whatever the mechanism. But of course there *are* many different auction mechanisms in use, of extremely variable type, because at least one of the hypotheses on which the Revenue Equivalence Theorem rests is

often violated. It is known, for example, that most people are not risk-neutral¹, and in the case when the bidders are risk-averse, it makes more sense for a seller to run a first-price auction.

A method commonly used to establish an ordering on auctions for different types of buyer, is to treat the problem as a non-cooperative game, and solve for the game’s Nash equilibrium. The difficulty with doing this is that the equations used to define such an equilibrium might well be intractable. In this paper we pursue an alternative method for determining an ordering: we *simulate* a population of buyers, and play the game many times with random selections from this population. The resulting averaged returns for the seller are estimates of the true expected return. As such, the results they give are not as satisfactory as those derived from Game Theory.

3 Methods

The basic methodology pursued was to instantiate a group of agents according to various environmental and agent-preference parameters, and let them compete in a specified auction according to specified strategies, logging the utility extracted by each agent as a result. Since many of the environmental parameters require some degree of randomization in the agent instantiation (e.g. randomized private value for the good), this procedure was repeated a large number of times, so as to generate an estimate of the expected utility to an agent of using a given strategy in a given context.

3.1 Context Parameterization: Environment, Preferences and Mechanism

We chose to investigate a space of mechanisms very similar to the first- and second-price sealed bid auctions specified earlier.

Definition 1. Let $w = (w_1, \dots, w_n)$ be a vector of n real numbers. A w -price auction is a sealed bid auction in which the highest bidder wins the good, and pays

$$\frac{\sum_{j=1}^N w_j bid_j}{\sum_{j=1}^N w_j} \quad (2)$$

where N is the minimum of n and the number of bidders, and bid_1, bid_2, \dots are the bids, ordered highest to lowest.

In this paper we examine a one-dimensional sub-space of w -price auctions, namely those of type $w = (1 - w_2, w_2)$. In this parameterization, $w_2 = 0$ is a standard first price auction, $w_2 = 1$ a standard second-price auction, and all other values of w_2 correspond to non-standard auction types that have not previously been studied.

¹ Most people tend to act in a risk-averse manner in their daily lives.

The space of agent preferences and environmental variables which we explored was motivated by examining exceptions to the Revenue Equivalence Theorem; we allowed variable group size, variable risk preference, and correlated (non-independent) bidders' signals. In addition, we allowed the degree of commonality in values to be altered.

1. The number of agents in each trial was either an arbitrary fixed number, or was chosen with uniform probability from a set of consecutive integers bigger than 2. In most experiments the fixed number was chosen to be 6, and the range $\{2, 3, 4, 5, 6\}$.
2. The signals (t_1, \dots, t_n) of a group (a_1, \dots, a_n) of bidders were chosen to be a weighted sum of a shared random signal and a sequence of independent random signals, with each such signal coming from a uniform distribution on $[0, 1]^2$. Thus independent variables S, X_1, \dots, X_n were generated, and the signal t_i for agent a_i was chosen to be $cS + (1 - c)X_i$, where $c \in [0, 1]$ parameterizes the degree of correlation between agents' signals.
3. The calculation of the utility extracted from winning the good depends on two properties of the agents involved: their risk preferences, and the degree of commonality in value. For risk, we chose to use Constant Absolute Risk utility functions:

$$u_\alpha(x) = \begin{cases} \frac{1}{\alpha}(e^{\alpha x} - 1) & \text{if } \alpha \neq 0, \\ x & \text{if } \alpha = 0, \end{cases} \quad (4)$$

α is zero for risk-neutral agents, negative for risk-averse agents and positive for risk-seeking agents. Fig. 1 plots these functions for $\alpha = -1, 0, 1$.

To model common value, we assumed that the monetary value to agent a_i of winning the good was given by $d \cdot (\sum_j v_j)/n + (1 - d)v_i$, where d is a parameter controlling the degree of common value, with $d = 0$ representing purely private values, and $d = 1$ purely common values. Thus the utility reward to agent a_i of winning the good at bid b , conditional on signals t_j for all players, was

$$u_\alpha \left(d \left(\sum_j t_j \right) / n + (1 - d)v_i - b \right). \quad (5)$$

3.2 Strategy Optimization

The above variables w_2, c, d, α specify the context in which the agents have to act, but not how they should act in that context. The most challenging piece of analysis in Mechanism Design is always figuring out how a bidder is likely

² We also performed experiments with the family of distributions

$$B_{m,k}(x) = m \cdot x^{m-k}(1-x)^k (m - 1k - 1) \quad (3)$$

(for $k = 0, \dots, m$). The results generated with these distributions were not qualitatively different from those generated for the uniform distribution.

to behave. The standard Game Theory approach is to enumerate all strategies that an agent might pursue, and determine a Nash equilibrium: a strategy *from which deviation is not rational*, i.e. which is expected-utility maximizing given that the other agents' behaviour is fixed.

As mentioned before, this process is often impossible, either because of the intractability of the strategy space, or because the equations which need to be solved to determine a deviation-proof strategy are too complex.

In this paper we take an empirical approach to finding good strategies, whereby each agent in a population of bidders is equipped with a bidding function which can be modified through evolution to adapt to the necessities of the game. As the agents play the game, successful strategies are bred preferentially, and thus the entire population improves. There is constant pressure to improve, because if an agent's deviation from the norm gives it a slightly higher expected utility, then it will be slightly more likely to breed than average, and so its genes will be preferentially reproduced into the next generation.

The main drawback to this approach is that it can neither be guaranteed that the population will evolve a good strategy within a reasonable period of time, nor that the solution on which the population eventually converges is a global rather than local optimum. Thus we gain formal simplicity at the cost of computation. We run the entire process of evolution many times independently, and reduce the effect of mutation as time goes by, so as to encourage convergence.

The link between genomes and bidding function was as follows: A gene consisted of a sequence (g_1, g_2, \dots, g_k) of real-valued "control points" assigned to evenly spaced input signals $(0, 1/k, \dots, 1)$. The bid output for an input signal t was generated by interpolating between the control points:

$$bid(t) = \begin{cases} g_1 & \text{if } t < 0 \\ g_l + (k \cdot t - l)(g_{l+1} - g_l) & \text{if } t \in [l/k, (l+1)/k) \\ g_{k+1} & \text{if } t \geq 1 \end{cases} \quad (6)$$

This representation was chosen over others (e.g. power series representation, GP etc.) because it combines useful features of the domain, while placing very few restrictions on the space of all such functions. Specifically,

1. Stability: These functions are stable under small random mutations: changing the data c_i does not make a huge difference to the output values generated by the bid function.
2. Locality: A change in a value g_l has no effect on the function for signals above $(l+1)/n$ or below $(l-1)/n$, so each g value, or sequence of g values, represents a partial solution for a certain input range.

In addition, the functions generated are guaranteed to be continuous.

These data g_l were then mutated and recombined according to a standard Genetic Algorithm, where the fitness of a given genome was determined relative to other genomes by participation in a sequence of randomized games. Specifically, the evaluation of a population of genomes was according to the following algorithm:

```

For each of a large number of iterations {
  while (not all agents have played in this round) {
    select some as-yet-unplayed agents to play a game
    generate random signals for the agents
    get bids for each agent, according to their genome
    select a winner and determine payments
    accumulate the corresponding utility rewards
  }
}

```

An agent's fitness was equal to its accumulated utility from all the games. This process is modular with respect to the contextual parameters specified in Section 3.1.

The Genetic algorithm was simply

```

generate a population of  $N$  random genomes
for each generation {
  assess the fitness of every individual by
    playing a large number of games as above
  rank the genomes by this fitness measure
  repeat  $N$  times {
    select two genomes from the old population,
      favouring highly ranked genomes
    select a random point in the genome, and
      combine the first half of one with the
      second half of the other
    with pre-selected probability, mutate each
      of the genes by an amount picked from a
      pre-selected distribution3
    place this new genome into the new population
  }
}

```

Thus we did not necessarily preserve the fittest individual from each generation. Notice that the fitness function is stochastic, so genomes can gain unfair advantage from being lucky (in the selection of their signals, for example). Notice also that the fitness is measured relative to other agents. This means that the most successful agent strategy is not necessarily that which gives greatest expected return, since it may (for example) be incentive compatible in such an environment to deliberately disadvantage oneself if in doing so one's opponents are even more disadvantaged. An example of this is that agents with very low signals are (in most environments) incented to bid higher than their valuation: they are very unlikely to win the good (and hence have negative surplus), whereas they are much more likely to decrease the winner's surplus, and hence increase their own relative fitness.

Thus this process finds good players at the repeated competitive game, not at the one-shot game. It is hoped (and we shall demonstrate, in some cases) that

³ The mutation probability was constant (at 0.8), and the maximal size of the mutation μ was reduced as time went on, being given, in generation G , by $\mu = \mu_0(G_0/(G_0 + G))^\lambda$, for fixed $\mu_0 = 0.05$, $G_0 = 20$, and $\lambda = 1.5$

this effect is very small, so that conclusions about the one shot game can still be made; it is worth noting that in real auctions with real players (humans or corporations), exogenous effects such as these are commonplace.

4 Results

Shown below (Fig. 2) is a graph showing the genes of the best individual in a population of 360, plotted against generation number, for an example evolutionary run for risk-neutral agents with independent private values, competing in a second price auction. In this context, the optimal bidding strategy is to bid one's signal, so given that $k = 5$ (i.e. the genome consists of 5 control values), the expected-utility-maximizing genome is $(0, 0.25, 0.5, 0.75, 1.0)$. As can be seen, the best individual is initially far from perfect, and varies greatly over the first few generations, since mutation is (relatively) high, and the population very diverse. As time wears on, however, the population discovers a bidding strategy that is close to optimal: the final population's best individual's bid is always within 5% of its optimal value.




Fig. 2. Genome for best individual in an example evolutionary run, plotted as a function of generation number. The population is initialized to random rational strategies (i.e. agents in the initial population cannot lose money, initially).

We first verified the method by calculating revenue landscapes for situations where the ordering of first and second price auctions is qualitatively known. When players are risk-neutral, signals independent, and the number of players fixed, the Revenue Equivalence Theorem says that all our mechanisms will generate the same expected revenue. Fig. 3 plots expected seller revenue against w_2 (in black). These results were obtained by evolving a population of 360 agents

for 200 generations⁴ The average auctioneer revenue at generation 200 was measured and the entire process repeated 200 times in order to average out effects from a particular evolutionary run.

The two lines in grey, above and below the plotted curve of average revenue, are plus and minus one standard deviation relative to the average, and give an indication of the magnitude of experimental uncertainty.

As can be seen, there is no experimentally significant difference in revenue between any of the mechanisms in the risk-neutral independent private values case.




Fig. 3. Sample revenues for risk-neutral agents operating in groups of fixed size 6, versus coefficient w_2 of second-highest bid in payment. Left-hand side is first price auction, right-hand side is second-price auction.

As mentioned in Section 2.2, when we modify the above by having risk-averse buyers, the first price auction becomes preferred. Fig. 4 shows this effect.

As Milgrom et al. demonstrate, when values are correlated⁵, we expect that the second-price auction will give greater revenue (Milgrom & Weber 1982). Fig. 5 demonstrates this effect occurring.

Much more interesting than confirming known results, is investigating regions of the environment space where there are no clear cut results. For example, if buyers have partial common values, and are risk averse, then either first or second price could be optimal for the seller, depending on the magnitude of the two effects.

⁴ 200 was chosen by empirically finding a sufficiently high number for the evolution process to reach equilibrium in a selection of test cases.

⁵ In fact Milgrom & Weber (1982) discusses not correlation but *affiliation* between bidders' signals. It can be shown that the joint signal distributions we have chosen to use in this paper satisfy this stronger affiliation condition.




Fig. 4. Average revenues for risk-averse agents operating in groups of fixed size 6.




Fig. 5. Average revenues for risk-neutral agents in fixed-size groups with values that are 50% correlated.

Fig. 6 shows the situation when bidders are risk-averse, with parameter -10 , and have a common value coefficient of 0.5 . In this case, the first-price auction is clearly superior to the second-price auction. More surprisingly, a $(0.3, 0.7)$ -price auction is superior to both first- and second-price auctions.

Fig. 7 shows the same situation when the common value coefficient is 0.9 , and risk aversion is -15 . In this case, second-price is superior to first-price, and once again a w -price auction is superior to both. These auction forms, in which the winner pays a weighted average of his own and the second player's bid are not studied in the literature, but in this common scenario, can be revenue maximizing, depending on the nature of the agents playing the game. The optimality of non-standard auctions in this risk-averse, partial common-value setting persists if the number of agents is variable, as is seen in Fig. 8.




Fig. 6. Average revenues for risk-averse agents ($\alpha = -10$) in fixed-size groups with values that are 50% common.




Fig. 7. Average revenues for risk-averse agents ($\alpha = -15$) in fixed-size groups with values that are 90% common.

In both of these scenarios, the graph of utility versus w_2 is flat (see Fig. 9): the agents themselves are indifferent as to which auction they participate in. Thus selecting a revenue-maximizing value of w_2 need not antagonize bidders, however, as Bergman & Tennenholtz (2002) show, we should expect the real dynamics of auction choice on the part of bidders to be affected by more than




Fig. 8. Average revenues for risk-averse agents ($\alpha = -15$) with values that are 90% common, when group size is random in the range $[2, 6]$.

just expected revenue: the variance of payments is crucial also. Clearly more work is needed to understand population dynamics in this new environment.




Fig. 9. Graph of expected utility earned by each agent as a function of w_2 for the equilibrium strategy in the scenario described in Fig. 7

5 Related Work

Auction Theory is a mature field, with a substantial literature. We shall not attempt an exhaustive review here; interested readers are referred to Klemperer (1999) for an overview.

The use of agents to investigate economic phenomena via simulation is a much newer field, known, broadly, as Agent-based Computational Economics (Tesfatsion & Judd forthcoming, 2006), a sub-field of which is concerned with designing agents that can operate in on-line auctions or negotiations. See, for example Anthony et al. (2001), Byde et al. (2002) with respect to 1-1 negotiation; Cliff & Bruten (1998), Preist & van Tol (1998), Gjerstad & Dickhaut (1998), Das et al. (2001), Tesauro & Bredin (2002) with respect to continuous double auctions; Preist, Bartolini & Philips (2001), Preist, Byde & Bartolini (2001), Byde (2001) with respect to sequences of English Auctions and Boutilier et al. (1999b), Boutilier et al. (1999a) with respect to sequences of sealed bid auctions. Walsh et al. (2002) gives an example of an approach to stretching the methods of Game Theory to deal with the sorts of large multi-stage games that economic agents participate in.

When it comes to investigating novel auction types automatically, or semi-automatically, the citations are much thinner on the ground. A general discussion of automated mechanism design appears in Conitzer & Sandholm (2002), which deals with issues of computational complexity. The work of Cliff (2001) is the first to explore evolutionary search methods over a parametric space of mechanisms. Cliff addresses the case of a continuous open-cry auction, using a Genetic Algorithm to adjust both the parameters of the bidding agents he uses, and the mechanism parameter, which in this case is the probability Q_s that in any given round, a seller will be chosen at random to make an offer. Besides being based on the continuous double auction, Cliff's work differs from ours in two significant ways. Firstly, the space of agent strategies explored is necessarily very restrictive⁶, whereas the strategy space our GA explores contains close approximations to all continuous bidding strategies. Secondly, although Cliff's choice of mechanism space was inspired by the experimental design used by Smith (1962), the Continuous Double Auction as it is used in such real-world institutions as the New York stock exchange is quite different – using order queues and bid improvement rules, for example; our cases $w_2 = 0, 1$ are faithful interpretations of first- and second-price sealed bid auctions, which are used in the world on a daily basis.

The work of Phelps et al. (2002) provides another approach to modifications of the continuous double auction, in which the modification is to the clearing rule, via use of Genetic Programming.

The space of mechanisms here is equivalent to a special case of a class of sealed-bid double auction mechanisms (“ k -price” auctions) first introduced by

⁶ This problem is difficult to address in the continuous open-cry auction, because such auctions have an intractably complicated strategy space: an agent will typically have many opportunities to act, for each of which the information space is the set of all previous actions by all agents.

Chatterjee & Samuelson (1983), and subsequently investigated by Satterthwaite & Williams (1989) among others under different assumptions on the independence, privacy, affiliation (e.g. Kadan 2004) or commonality of values.

6 Conclusions

In this paper we have described an application of evolutionary search to a family of auction mechanisms – an example of automated mechanism design. We have demonstrated that this technique can be used to explore a space of auction mechanisms, and by doing so in a specific setting that involves faithful versions of real-world mechanisms, have established the superiority of non-standard auction types in a variety of common environments.

The advantages of such a method for exploring auction design issues are clear: the agents discover good bidding strategies by evolution, without the need for complicated, possibly intractable, and certainly fragile mathematical analysis. In more complicated applications, the evolution process can implicitly take factors into consideration that might not have occurred to analysts. Additionally, the mechanism is tested for revenue generation against a small neighborhood of strategies, not just the Nash-equilibrium strategy. As a result, its sensitivity to agents' choice of strategy can be determined.

References

- Anthony, P., W. Hall, V. Dang & N.R. Jennings (2001), Autonomous agents for participating in multiple on-line auctions, *in* 'Proc. IJCAI Workshop on E-Business and the Intelligent Web', pp. 54–64.
- Bergman, A. & M. Tennenholtz (2002), 'On the natural selection of market choice', *Autonomous Agents and Multi-Agent Systems* **5**(4), 387 – 396.
- Boutilier, C., M. Goldszmidt & B. Sabata (1999a), Continuous value function approximation for sequential bidding policies, *in* 'Proc. UAI '99'.
- Boutilier, C., M. Goldszmidt & B. Sabata (1999b), Sequential auctions for the allocation of resources with complementarities, *in* 'Proc. IJCAI '99'.
- Byde, A. (2001), A dynamic programming model for algorithm design in simultaneous auctions, *in* 'Proc. 2nd Int. Workshop on Electronic Commerce (WELCOM-01)', LNCS, Springer Verlag.
- Byde, A., C. Preist & N. Jennings (2002), Decision procedures for multiple simultaneous auctions, *in* 'Proc. 1st Int. Conf. Autonomous Agents and Multi-Agent Systems'.
- Chatterjee, K. & W. Samuelson (1983), 'Bargaining under incomplete information', *Operations Research* (31), 835–851.
- Cliff, D. & J. Bruten (1998), Less than human: Simple adaptive trading agents for CDA markets, *in* 'Proc. of the 1998 Symposium on Computation in Economics, Finance, and Engineering: Economic Systems'.
- Cliff, Dave (2001), Evolution of market mechanism through a continuous space of auction-types, Technical Report HPL-2001-326, HP Labs.
- Cliff, Dave (2002), Evolution of market mechanism through a continuous space of auction-types ii: Two-sided auction mechanisms evolve in response to market shocks, Technical Report HPL-2002-128, HP Labs.

- Conitzer, V. & T. Sandholm (2002), Complexity of mechanism design, in 'Proc. 18th Conference on Uncertainty in Artificial Intelligence (UAI)'.
- Das, R., J.E. Hanson, J.O. Kephart & G. Tesauro (2001), Agent-human interactions in the continuous double auction, in 'Proc. 17th Int. Joint Conf. on Artificial Intelligence'.
- Gjerstad, S. & J. Dickhaut (1998), 'Price formation in double auctions', *Games and Economic Behaviour* **22**(1), 1–29.
- Holland, J. (1975), *Adaption in Natural and Artificial Systems*, U. Michigan Press.
- Kadan, Ohad (2004), Equilibrium in the two player, k -double auction with affiliated private values, Technical Report 2004.12, Fondazione Eni Enrico Mattei. available at <http://ideas.repec.org/p/fem/femwpa/2004.12.html>.
- Klemperer, P. (1999), 'Auction theory: A guide to the literature', *Journal of Economic Surveys* **13**(3), 227–286.
- Krishna, Vijay (2002), *Auction Theory*, AP.
- Milgrom, M. P. & R. J. Weber (1982), 'A theory of auctions and competitive bidding', *Econometrica* (50), 1089 – 1122.
- Phelps, Steve, Peter Mc Burnley, Simon Parsons & Elizabeth Sklar (2002), Co-evolutionary auction mechanism design, in 'Lecture Notes in AI', Vol. 2531, Springer Verlag.
- Preist, C., A. Byde & C. Bartolini (2001), Economic dynamics of agents in multiple auctions, in 'Proc. 5th Int. Conf. on Autonomous Agents'.
- Preist, C., C. Bartolini & I. Philips (2001), Algorithm design for agents which participate in multiple simultaneous auctions, in F.Dignum & U.Cortes, eds, 'Agent Mediated Electronic Commerce III', Lecture Notes in AI, Springer Verlag.
- Preist, C. & M. van Tol (1998), Adaptive agents in a persistent shout double auction, in 'Proc. 1st Int. Conf. on the Internet, Computing and Economics', ACM Press.
- Satterthwaite, M. & S. R. Williams (1989), 'Bilateral trade with the sealed bid k -double auction: Existence and efficiency', *Journal of Economic Theory* (48), 107–133.
- Smith, V. L. (1962), 'Experimental study of competitive market behaviour', *J. of Political Economy* **70**, 111 – 137.
- Tesauro, G. & J. L. Bredin (2002), Strategic sequential bidding in auctions using dynamic programming, in 'Proc. AAMAS02', pp. 591 – 598.
- Tesfatsion, Leigh & Kenneth L. Judd (forthcoming, 2006), *Agent-Based Modeling*, Vol. 2 of *Handbooks in Computational Economics*, Elsevier.
- Walsh, W. E., R. Das, G. Tesauro & J. O. Kephart (2002), Analyzing complex strategic interactions in multi-agent games, in 'Proc. AAAI-02 Workshop on Game Theoretic and Decision Theoretic Agents', pp. 109–118.