Optimal Auctions Revisited

Dov Monderer and Moshe Tennenholtz
Faculty of Industrial Engineering and Management
Technion–Israel Institute of Technology
Haifa 32000
Israel

Abstract

The Internet offers new challenges to the fields of economics and artificial intelligence. This paper addresses several basic problems inspired by the adaptation of economic mechanisms, and auctions in particular, to the Internet. Computational environments such as the Internet offer a high degree of flexibility in auctions' rules. This makes the study of optimal auctions especially interesting in such environments. Although the problem of optimal auctions has received a lot of attention in economics, only partial solutions are supplied in the existing literature. We present least upper bounds (l.u.b) R_n on the revenue obtained by a seller in any auction with n participants. Our bounds imply that if the number of participants is large then the revenue obtained by standard auctions (e.g., English auctions) approach the theoretical bound. Our results heavily rely on the risk-aversion assumption made in the economics literature. We further show that without this assumption, the seller's revenue (for a fixed number of participants) may significantly exceed the upper bound.

1. Introduction

The Internet exhibits forms of interactions which are not captured by current studies and theories in Economics. The highly distributed nature of the Internet, the relatively easy access to economic trades carried out in remote locations, and the ability of defining various types of Internet trades by individual users, lead to new kind of settings for which new theories should be developed and evaluated.

The Internet is a distributed environment where self-interested parties may interact. The strategic interaction among agents is a major topic of study in Microeconomics (Kreps 1990) and Game Theory (Fudenberg & Tirole 1991). In particular, the design of protocols for strategic interactions is the subject of the field termed *mechanism design* (Fudenberg & Tirole 1991; Myerson 1991). Research on strategic aspects of multi-agent activity in Artificial Intelligence has grown rapidly in the recent years. Work in AI has mostly concentrated on the design of protocols for agents' interaction. Hence, work in AI shares much in common with work on mechanism design in Economics (Bond & Gasser 1988; Demazeau & Muller 1990; Rosenschein & Zlotkin 1994; Durfee 1992). Many basic principles and ideas grew up from the mechanism design literature. Much of the research in mechanism design has been devoted to the study of auctions (McAfee & McMillan 1987; Wolfstetter 1996). There are two reasons for that. One is the popularity of auctions as an economic trade, and the other is the understanding that many of the features studied in the auctions literature shed light on other economic mechanisms (Fudenberg & Tirole 1991). Evidentially, auctions have turned to be a most popular strategic tool in electronic commerce as well, and the number of different auctions carried out in the Internet is huge.

In an auction a good is sold to a single buyer taken from a set of potential buyers (or more generally, several goods are sold to a set of agents), according to some bidding rules and rules for determining the auction's outcome (e.g., the exact price to be paid). There are several auctions that have been found to be representative ones, and which are most widely used in Economic trades. Basic theories regarding these auctions have been developed and few basic results on the seller's revenue in different types of auctions have been obtained (e.g., (Myerson 1982; Milgrom & Weber 1982)). Some preliminaries regarding basic auction theory will be presented in the next section. Given the popularity of auctions in the Internet, and their fundamental role in Economic theories, the adaptation of auction theory to computational settings becomes a task of considerable importance. Naturally, when one considers economic trades in computational settings, he may wish to consider the effects of computational bounds on agents' behavior in auctions. This is the approach taken for example in (Sandholm 1996). However, as we show, the computational setting also suggests the need for a careful study of basic issues in economics that have (somewhat surprisingly) been neglected.

The Internet is a computational setting, with flexible software which is used by relatively sophisticated users. This makes the study of optimal auctions (e.g., (Myerson 1991; Riley & Samuelson 1981; Maskin & Riley 1984)) highly relevant to this setting. In an optimal auction the seller chooses an auction mechanism that maximizes her revenue in equilibrium. A major claim against optimal (non-standard) auctions is that they are hard to handle, but given the computational environment this task becomes much more reasonable and doable. All that is needed in order to define and manage a new kind of auction is a simple software; typical Internet users will have little problem in participating in such auctions. Indeed, the reader may easily find in the Internet various novel modifications of existing auctions. This leads to the following central questions. What is an upper bound on the revenue one can obtain by an optimal auction, can it be obtained in certain cases, and can it be obtained by classical auctions?

Another basic feature of electronic sales is the lack of commitment power of the seller, and the relatively high-risk for the buyers. This suggests that one may wish to relax the assumption that the buyers are risk-averse (or risk-neutral) agents, that is (always) taken in economics. How does the expected revenue of the seller (from relatively standard auctions) change when we remove the risk-aversion assumption?

We show:

1. The expected highest valuation of an object from the point of view of the participants is a least upper bound on the gain a seller can obtain in any auction for risk-averse (and in particular for risk-neutral) agents. This result is not obvious given that the bids of agents in equilibria (in relatively standard auctions) may be higher than their actual valuation for the good. This result solves a basic problem in the theory of optimal auctions.

2. If the number of participants is large, then the seller's revenue in an English auction approaches the upper bound.

3. We suggest that the modeling of agents as risk-seeking may be appropriate for the Internet setting. We show that in this case a third-price auction enables the seller to obtain a revenue that is higher than the expected highest valuation for the good.

In Section 2 we introduce some preliminaries. In particular, we describe the classical model of auctions. In Section 3 we present a least upper bound on the revenue obtained by the seller in any auction for risk-averse (or risk-neutral) agents. We also show that this upper bound can be almost matched in English auctions when the number of agents is large enough. In Section 4 we discuss the potential need for modeling agents as having a risk-seeking attitude, and show that in this case a third-price auction can lead to a seller's revenue that is higher than the previously mentioned upper bound.

2. Preliminaries

Consider a seller who wishes to sell a particular good, where there are n agents denoted by $1, 2, \ldots, n$ who wish to buy this good. An auction is a procedure in which participants submit messages (typically monetary bids) for the good.[1] The auction's rules specify the type of messages, and as a function of the messages submitted by the participants they determine the winner and the payments to be made by the participants (to the auction organizer). Formally, an auction procedure for n potential participants, $N = \{1, 2, \ldots, n\}$ is characterized by 4 parameters, M, g, c, d, where M is the set of messages, $g = (g_1, \ldots, g_n)$ with $g_i : M^N \to [0, 1]$ for all i and $\sum_{i=1}^{n} g_i(m) \leq 1$ for all m, and $c = (c_1, \ldots, c_n); d = (d_1, \ldots, d_n)$ with $c_i, d_j : M^n \to R$ for all i, j. Participant i submits a message $m_i \in M$. Let $m = (m_1, m_2, \ldots, m_n)$ be a vector of messages, then the organizer conducts a lottery to determine the winner, in which the probability that i is the winner equals $g_i(m)$. The winner, say j, pays $c_j(m)$ and every other participant i pays $d_i(m)$. It is assumed that M contains the null message e, which is interpreted as non-participation. It is further assumed that if $m_i = e$, then $g_i(m) = c_i(m) = d_i(m) = 0$. Classical auction theory associates a (Bayesian) game with each auction procedure and analyses the behavior of the agents under the equilibrium assumption. To do this we have to define the information structure of the game. Let $(\hat{v}_i)_{i=1}^{n}$ be mutually independent non-negative random variables. When $\hat{v}_i = v_i$, v_i is interpreted as the maximal willingness to pay of Agent i. v_i is called the type of i.[2] The distribution of $\hat{v}_i$ is denoted by F_i. That is, $F_i(v) = Prob(\hat{v}_i \leq v)$. We take F_i to be the uni-

[1] These messages can also refer to complete strategies. Hence, the analysis does not refer only to one-shot interactions.

[2] We use the independent-private-value model of information. There exist more complicated information models in which an agent does not know his own type, and the agents' types are correlated (see Milgrom and Weber (1982)). We believe that the independence assumption is the right one in the Internet auctions setup where there are many anonymous participants.

form distribution on some interval $[a_i, b_i]^3$. The probability measure induced by F_i on R_+ is denoted by P_i. Let P denote the product probability measure of $(P_i)_{i\in N}$ on R_+^N and let P_{-i} denote the product probability measure defined by $(P_j)_{j\in N\setminus\{i\}}$ on $R^{N\setminus\{i\}}$. Each agent i has a utility function for money $u_i : R \to R$, normalized with $u_i(0) = 0$. It is assumed that Agent i is an expected utility maximizer. It is further assumed that if Agent i with the type v_i receives the item and pays x_i, his utility is $u_i(v_i - x_i)$. Agent i is risk-averse if u_i is a concave function. A risk-averse agent weakly prefers a certain amount to a lottery whose expected payoff equals this amount. Agent i is risk-neutral if u_i is linear. Such an agent is indifferent between a lottery and its expected payoff.[4] A strategy for agent i is a function $b_i : R_+ \to M$, where $b_i(v_i)$ is the message submitted by i when his type is v_i. Let $b = (b_1, b_2, \ldots, b_n)$ be an $n-tuple$ of strategies. b is in equilibrium if for every agent i and for every v_i, the expected utility of agent i given that his type is v_i and given that each agent j, $j \neq i$ uses b_j is maximized over $m_i \in M$ at $m_i = b_i(v_i)$. Before expressing the above verbal description with the appropriate formula we remark that this definition makes sense only if certain technical conditions are imposed on all functions under discussion. For simplicity we do not explicitly present these conditions, which are quite common in the economic literature.[5] For a vector of strategies $b = (b_i)_{i\in N}$ we denote $b_{-i} = (b_j)_{j\neq i}$, for $v \in R_+^N$ we denote $b(v) = (b_1(v_1), b_2(v_2), \ldots, b_n(v_n))$ and for agent i we denote $b_{-i}(v_{-i}) = (b_j(v_j))_{j\neq i}$. Thus, b is in equilibrium if for every agent i and for every type v_i,

$$\max_{m_i\in M_i} E_{P_{-i}}(u_i(v_i - c_i(m_i, b_{-i}))g_i(m_i, b_{-i})+$$

$$u_i(-d_i(m_i, b_{-i}))(1 - g_i(m_i, b_{-i})))$$

[3] In the economic literature other distributions are discussed as well. We find uniform distributions quite natural for the Internet setting. Nevertheless, some of our results hold for any distribution. In particular, our proof about the upper bound on the expected revenue can be extended to the case of arbitrary (twice continuously differentable) distributions.

[4] Notice that a risk-neutral agent is a specific instance of a risk-averse agent. A risk-neutral agent has (up to an increasing linear transformation) a specific utility function $(u_i(x) = x)$, while the set of concave utility functions that represent risk-averse agents is huge. In fact, the standard assumption in Economics is that agents/buyers are risk-averse, and the assumption of risk-neutral agents is taken only as an approximation to risk-averse agents, for mathematical convenience.

[5] These conditions refer to the structure of the utility and distribution functions. In particular, it is assumed that both are twice continuously differentable.

is attained at $m_i = b_i(v_i)$, where $E_{P_{-i}}$ denotes the expected value operator with respect to P_{-i}.

Under the equilibrium assumption (i.e, the assumption that economic agents use equilibrium strategies), if the auction game $G = G(n, A, \bar{u}, \bar{F})$, where $\bar{u} = (u_i)_{i=1}^n$ and $\bar{F} = (F_i)_{i=1}^n$, has a unique equilibrium profile $b = (b_1, b_2, \ldots, b_n)$, then the seller expected revenue is denoted by R_G. That is,

$$R_G = E_P\left(\sum_{i=1}^{n}(g_i(b)c_i(b) + (1 - g_i(b))d_i(b)\right).$$

When the auction game has more than one equilibrium profile, we denote by R_G the revenue of the seller in the worst case, that is the greatest lower bound of the revenues obtained in some equilibrium. We do not consider auction procedures A that are over complicated such that the associated game does not have an equilibrium. Consider a fixed information structure $\bar{F}$, a fixed vector of utility functions $\bar{u}$ with $u_i = u$ for all i and $u(x) = x$ for all x, and a fixed number of participants n. Myerson (Myerson 1982) has solved the optimality problem $\max_A R_G$ and proved that an optimal auction procedure is a second-price auction with an appropriate reservation price. Maskin and Riley (Maskin & Riley 1984) have investigated the optimality problem with a fixed symmetric information structure and a fixed symmetric utility structure in which $u_i = u$ for all i, and u is a concave function. Their results indicate, that if an optimal solution exists, it should involve a very sophisticated auction procedure. In the next section we will present a least upper bound $R_n(\bar{F}) = \sup_{A,\bar{U}} R_G$ to the seller's revenue, where $\bar{u}$ ranges over all utility structures of risk averse agents, and A ranges over all auction rules. A seller, whose revenue is close to this upper bound does not have to worry about optimality. Obviously $R_n(\bar{F})$ is an upper bound for every fixed utility structure, but it is not a necessarily a least upper bound for a fixed utility structure.

Classical Auction Mechanisms

In this section we discuss some classical auction mechanisms and show how they can be described in the framework mentioned above.

First Price Auctions One of the most popular auction mechanisms is the first-price auction. In such an auction, each participant submits a bid in a sealed envelop. The agent with the highest bid wins the object and pays his bid, all other participants pay nothing. Ties are broken with some lottery mechanism. In a first-price auction, $M = R_+ \cup \{e\}$, and for $x \in R_+^N$, $g_i(x) = 0$, if $x_i < w(x) = \max\{x_j : j \in N\}$, and

$g_i(x) = \frac{1}{k(x)}$ if $x_i = w(x)$ and $k(x)$ denotes the number of agents j for which $x_j = w(x)$. Also, $c_i(x) = x_i$ and $d_i(x) = 0$.[6] In the next section we will use the standard equilibrium analysis of first-price auctions that we now present: We assume that all agents are symmetric ($u_i = u$) and $F_i = F$ for all i, and that F is supported in the interval $[0, 1]$ in the sense that $F(0) = 0$, $F(1) = 1$, and $F'(x) > 0$ for all $x \in [0, 1]$. It can be shown that if u is twice continuously differentiable, $u'(x) > 0$ for all x, and $\frac{u}{u'}$ is increasing, then there exists a unique equilibrium, $(b_1, b_2, \dots, b_n) = (b, b \dots, b)$, where b is the unique solution of the differential equation (see (Riley & Samuelson 1981)):

$$b'(v) = (n-1)\frac{u(v - b(v))F'(v)}{u'(v - b(v))F(v)}; \;\; b(0) = 0.$$

In particular, if the agents are risk-neutral and $F(x) = x$ is the uniform distribution on $[0, 1]$, then $b(v) = \frac{n-1}{n}v$ for all $v \in [0, 1]$.

k-price auctions In a k-price auction each participant submits a bid in a sealed envelop, the winner is the one with the highest bid and he pays the k-order statistics of the vector of bids. For example, in a second-price auction if the three highest bids were 10, 9 and 8, then the winner is the one whose bid was 10, and he pays 9. However, if the first three bids were 10,10, 9, then the winner is selected with a probability 0.5 from the two agents with the bid of 10, and he pays 10. It can be easily verified (and it is well-known) that in a second-price auction, the strategies $b_i(v_i) = v_i$ are in equilibrium for every information and utility structures. k-price auctions, $k \geq 3$ have the interesting feature that in equilibrium the agents overbid (see (Wolfstetter 1996)). That is, $b_i(v_i) > v_i$.

English auctions The English auction is the most popular open auction. In such an auction, there is an initial reservation price (which can be zero) and at each time every player can increase the bid publicly. The auction is over if for a certain (fixed in advanced) time period no one increases the bid. The agent with the last bid wins the object and pays the last bid. Such an auction, like all other open auctions, is analyzed in the framework described above by defining the message space M to be the set of strategies (protocols) in the dynamic game induced by the mechanism. Though formally, participants do not actually have to submit a protocol, it is obvious that had they have to do it, they will submit their true protocol, because this protocol is optimal with respect to their beliefs concerning the behavior of the other participants. It is interesting to note that English auctions are equivalent to second-price auctions in the sense that for every vector of types, both auctions yield the same winner and the same payoffs to the organizer. In particular, the seller's revenue in both auctions is identical.[7]

Dutch Auctions In a Dutch auction, the auctioneer initially calls for a very high price, and then continuously lowers the price until some bidder stops the auction and claims the good for that price. Dutch auctions are strategically equivalent to first-price auctions, and their analysis (and in particular, the expected revenue of the seller) coincide.

3. An Upper Bound on Optimal Auctions

In this section we supply an upper bound on the revenue obtained by the seller in any auction for risk-averse agents. We will be able to show that no matter how sophisticated an auction is, although agents may have bids that are higher than their actual evaluation of a good [8] the seller's revenue can not be higher than the expected highest valuation of the auction's participants for this good.

Theorem 1 *Let $n \geq 2$. Let $\bar{F} = (F_i)_{i=1}^n$ be a fixed information structure. Then for every utility structure of risk-averse agents $\bar{u} = (u_i)_{i=1}^n$, for every auction procedure $A = (M, g, c, d)$, and for every equilibrium strategies $b = (b_i)_{i=1}^n$, the expected revenue of the seller from the auction is bounded by the expected highest type. That is,*

$$R_G \leq E_P(\max(\hat{v}_1, \hat{v}_2, \dots, \hat{v}_n)),$$

where $G = G(n, A, \bar{u} = (u_i)_{i=1}^n, \bar{F} = (F_i)_{i=1}^n)$ is the auction game.

Moreover, the above-mentioned upper bound is a least upper bound.

Proof: Consider agent i, who participates in the auction, and denote the expected revenue from agent i by $\bar{R}_i$. If agent i is of type v_i, then his expected utility $B_i(v_i)$ when bidding $b(v_i)$, while the bidding strategies of the other agents are $b_j, j \neq i$, is given by:

$$B_i(v_i) = E_{P_{-i}}(u_i(v_i - c_i(b_i(v_i), b_{-i}))g_i(b_i(v_i), b_{-i}) +$$

$$u_i(-d_i(b_i(v_i), b_{-i}))(1 - g_i(b_i(v_i), b_{-i}))).$$

[6] If the seller announces a participation fee of $c > 0$, then $c_i(x) = x_i + c$ and $d_i(x) = c$.

[7] This equivalence principle does not hold if we do not assume the independence of types (see (Milgrom & Weber 1982)).

[8] See the discussion of k-price auctions in the previous section.

As i can always choose not to participate we conclude that $B_i(v_i) \geq 0$. Because u_i is concave (i.e., $u_i(\alpha x + (1-\alpha y) \geq \alpha u_i(x) + (1-\alpha)u_i(y))$, we have that

$$u_i(E_{P_{-i}}((v_i - c_i(b_i(v_i), b_{-i}))g_i(b_i(v_i), b_{-i})+$$

$$(-d_i(b_i(v_i), b_{-i})(1 - g_i(b_i(v_i), b_{-i}))) \geq 0.$$

Because u_i is increasing, $u_i(0) = 0$, and u is concave, Jensen inequality (i.e., $u_i(E(\cdot)) \geq E(u_i(\cdot))$) implies that:

$$E_{P_{-i}}((v_i - c_i(b_i(v_i), b_{-i}))g_i(b_i(v_i), b_{-i})+$$

$$(-d_i(b_i(v_i), b_{-i})(1 - g_i(b_i(v_i), b_{-i})))) \geq 0.$$

Let $R_i(v_i)$ be the expected revenue from Agent i given that his type is v_i. The last inequality yields

$$\bar{R_i} = E_{P_i}(R_i) \leq E_P(\hat{v}_i g_i(b)).$$

Therefore

$$R_G = \sum_{i=1}^{n} \bar{R_i} \leq E_P(\sum_{i=1}^{n} \hat{v}_i g_i(b)) \leq E_P(\max(\hat{v}_i)_{i=1}^{n}).$$

We prove that our bound in an l.u.b in the case where each F_i is the uniform distribution over $[0,1]$. The proof of the general case is similar. Consider a first-price auction, and assume $u_i = u_\alpha$ $(0 < \alpha < 1)$ where $u_\alpha(x) = x^\alpha$ for every $x \geq 0$.[9] We will show that when α approaches zero the expected gain in this auction approaches the above-mentioned upper bound.

From the equilibrium equation is Section 2 we can deduce that in equilibrium $b_i = b$ for all i and

$$b'(v) = \frac{n-1}{\alpha} \cdot \frac{(v - b(v))}{v}.$$

This implies that

$$b'(v) \cdot v + \frac{n-1}{\alpha} \cdot b(v) = \frac{n-1}{\alpha} \cdot v.$$

By multiplying both sides of the equality by $v^{\frac{n-1}{\alpha}-1}$ we get that:

$$b'(v) \cdot v^{\frac{n-1}{\alpha}} + \frac{n-1}{\alpha} \cdot b(v) \cdot v^{\frac{n-1}{\alpha}-1} = \frac{n-1}{\alpha} \cdot v^{\frac{n-1}{\alpha}}.$$

[9] Notice that u_α is not defined for $x < 0$, and that it can not be extended to a concave function on the whole real line. Therefore, the complete proof makes use of appropriate modifications of the u_α's. For simplicity, we omit these modifications in the proof presented here.

The above is equivalent to:

$$(b(v) \cdot v^{\frac{n-1}{\alpha}})' = \frac{n-1}{\alpha} \cdot v^{\frac{n-1}{\alpha}}.$$

By taking the integral of both sides we get:

$$(b(v) \cdot v^{\frac{n-1}{\alpha}}) = \frac{n-1}{\alpha} \cdot \frac{v^{\frac{n-1}{\alpha}} + 1}{\frac{(n-1)}{\alpha} + 1}.$$

Hence, we have:

$$b(v) = v \cdot \frac{n-1}{n-1+\alpha}.$$

Therefore, the expected revenue of the seller is

$$E(b(v_{max})) = \frac{n-1}{n-1+\alpha} E(v_{max}).$$

where $v_{max} = \max(\hat{v}_i)_{i=1}^{n}$. Hence, the seller's revenue approaches our bound when α approaches 0. ▮

Optimality of English Auctions

The above theorem has shown a least upper bound on the expected revenue of the seller in any auction for risk-averse agents. A natural question is whether this least upper bound can be obtained by standard auctions. In the following theorem we use the fact that the revenue of the seller in English auctions does not depend on the utility functions. We can show:

Theorem 2 *Let $(F_i)_{i=1}^{\infty}$ be a sequence of distribution functions supported in a fixed bounded interval. Let A_n be an English auction for n participants with $\bar{F} = (F_1, \ldots, F_n)$. Let $R(A_n)$ denote the expected revenue for the seller in the auction A_n. Then, R_n approaches $E_P(\max(\hat{v}_1, \hat{v}_2, \ldots, \hat{v}_n))$, when n approaches infinity (in the sense that the limit of the ratio approaches 1, when n approaches ∞).*[10]

Hence, the optimal expected revenue can be almost obtained in English auctions, when the number of agents is large enough (which one may expect in Internet auctions). Our results have also an interpretation which is somewhat complementary to the revenue equivalence principle mentioned before. The famous revenue equivalence theorem implies that we can not improve upon the revenue obtained in classical auctions when the agents are risk-neutral. Our results tell us that, at least in what seem to be a setup which is appropriate for many Internet auctions, this is also true when we have arbitrary risk-averse agents.

[10] Proof will appear in the full paper.

4. Risk-Seeking Agents

The number of Internet auctions grow relatively fast, and various modifications for the classical auctions are considered. However, as we have shown, under the classical assumptions made in Economics one can not go far in optimizing the seller's revenue. In this part of the paper we wish to consider a different kind of setting that may explain some of the phenomena that occur in the Internet setting, and which explore the potential of obtaining a revenue that is higher than the highest participants' valuation.

Indeed, most of the Internet auctions are English auctions, which are known to be strategically equivalent to second-price auctions. There used to be some Dutch auctions (which are strategically equivalent to first-price auctions) but they seem to disappear. On the other hand, a central theorem of auction theory is that the expected revenue obtained by the seller in first-price auctions is greater than her expected revenue in second-price auctions, assuming the agents are risk-averse. So, how can one explain this situation? One way of explaining the related phenomena is by considering participants that are risk-seeking rather than risk-averse. Although this assumption does not appear in classical work in Economics, it may make sense in the Internet setting. Indeed, in the Internet setting agents attempt to buy items that they can not really see, and that are sold by an unknown seller, and with no real commitment of the seller. On the other hand, the buyers need to reveal their credit card information. Although we do not formally argue for this point, we believe that the modeling of agents as having a risk-seeking attitude rather than risk-aversion attitude should be treated carefully and seriously. In another paper (Monderer & Tennenholtz 1998) we discuss the risk-seeking assumption and show that if the agents are risk-seeking then second-price auctions lead to higher revenue to the seller than first-price auctions. We wish now to re-consider our upper bound on the seller's revenue in view of this different modeling perspective. Technically, the only difference from the previously mentioned setting, is that in the case of risk-seeking agents the utility function will be convex rather than concave.

It is well-known that a risk-neutral seller can sell lottery tickets with negative expected gain to a risk-seeking agent and obtain as a result very high gains. Our aim here is different. We wish to show that a value that is higher than the upper bound obtained, can be achieved in what can be considered as standard auctions. In particular, we consider the third-price auction, which is another (less studied) instance of k-price auctions.

We will be interested in the following utility function that captures risk-seeking attitude (and is complementary to the one discussed in (Sandholm 1996) with regard to risk-averse agents). We assume that $u(x) = x$ when $x \leq 0$ and that $u(x) = \alpha x$ for some constant $\alpha > 1$ where $x > 0$. We can show:

Theorem 3 *A third-price auction for risk-seeking agents can lead to an expected revenue which is higher than the expected highest participant's valuation for the good.*[11]

References

Bond, A. H., and Gasser, L. 1988. *Readings in Distributed Artificial Intelligence.* Ablex Publishing Corporation.

Demazeau, Y., and Muller, J. 1990. *Decentralized AI.* North Holland/Elsevier.

Durfee, E. 1992. What your computer really needs to know, you learned in kindergarten. In *10th National Conference on Artificial Intelligence*, 858–864.

Fudenberg, D., and Tirole, J. 1991. *Game Theory.* MIT Press.

Kreps, D. 1990. *A Course in Microeconomic Theory.* Princeton University Press.

Maskin, E., and Riley, J. 1984. Optimal auctions with risk-averse buyers. *Econometrica* 52(6):1473–1518.

McAfee, R., and McMillan, J. 1987. Auctions and bidding. *Journal of Economic Literature* 25:699–738.

Milgrom, P., and Weber, R. 1982. A thoery of auctions and competitive bidding. *Econometrica* 50:1089–1122.

Monderer, D., and Tennenholtz, M. 1998. Internet Auctions. Working paper.

Myerson, R. 1982. Optimal auction design. *Mathematics of Operations Research* 6:58–73.

Myerson, R. B. 1991. *Game Theory.* Harvard University Press.

Riley, J., and Samuelson, W. 1981. Optimal auctions. *American Economic Review* 71:381–392.

Rosenschein, J. S., and Zlotkin, G. 1994. *Rules of Encounter.* MIT Press.

Sandholm, T. 1996. Limitations of the vickrey auction in computational multiagent systems. In *2nd International Conference on Multi-Agent Systems.*

Wolfstetter, E. 1996. Auctions: An introduction. *Journal of Economic Surveys* 10(4):367–420.

[11] Proof will appear in the full paper.