Optimal Auctions Revisited

Dov Monderer and Moshe Tennenholtz
Faculty of Industrial Engineering and Management
Technion–Israel Institute of Technology
Haifa 32000
Israel
e-mail: dov@ie.technion.ac.il,moshet@ie.technion.ac.il

Abstract

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. We present an upper bound on the revenue obtained by a seller in any auction with a fixed number of participants, and we show that this bound may be a least upper bound in some setups. We further show that the revenue obtained by standard auctions (e.g., English auctions) approaches the theoretical bound, when the number of participants is large. 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 if the participants are sufficiently risk-seeking.

A preliminary report on this work appears in the proceedings of AAAI-98.

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 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 [8] and game theory [5]. In particular, the design of protocols for strategic interactions is the subject of the field termed *mechanism design* [5, 16, 8]. 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 (see e.g. [2, 3, 18, 4, 7, 20]). In addition, the idea of using microeconomic tools for problem solving has also been the subject of market-oriented approaches to distributed planning [22]. Many basic principles and ideas for the design of multi-agent AI systems grew up from the mechanism design literature. Much of the research in mechanism design has been devoted to the study of auctions [10, 12, 15, 17, 23]. There are two reasons for that. One is the popularity of auctions as selling mechanisms, and the other is the understanding that many of the features studied in the auctions literature shed light on other economic mechanisms [5]. Evidently, auctions have turned to be most popular mechanisms 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 (i.e., the winner and the prices 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., [15, 12]). 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 trade 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 [19]. 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., [16, 17, 9]) 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 using 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?

We show:

1. The expected highest valuation of an object from the point of view of the participants is an upper bound on the expected gain a seller can obtain in any auction with risk-averse (and in particular with risk-neutral) agents. This result is not obvious given that the bids of agents in equilibria (in some 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. There are cases (information structures) in which the above-mentioned upper bounded is a least upper bound.

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

Another basic characteristic of electronic commerce is the relatively high risk for the buyers, and in particular the doubtful commitment power of the seller. 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 (in relatively standard auctions) change when we remove the risk-aversion assumption? We show:

4. Without the risk-aversion assumption, there exists an auction mechanism (a third-price auction) in which the seller can 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 an upper bound on the revenue obtained by the seller in any auction for risk-averse (or risk-neutral) agents. We also show that there are cases where the above-mentioned upper bounded is a least upper bound. In Section 4 we also show that this upper bound can be almost matched in English auctions when the number of agents is large enough. In Section 5 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* (or an *auction mechanism*) for n potential participants with the names $\{1, 2, \ldots, n\}$ is characterized by 4 parameters,

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

M, g, c, d, where M is a 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 $g_i(m) = c_i(m) = d_i(m) = 0$. An auction procedure is *anonymous* if its rules do not depend on the names of the players. More precisely, for every one-to-one function $\pi : \{1, 2, \cdots, n\} \to \{1, 2, \cdots, n\}$,

$$f_i(m_{\pi(1)}, m_{\pi(2)}, \cdots, m_{\pi(n)}) = f_{\pi(i)}(m_1, m_2, \cdots, m_n),$$

where f stands for g,c, or d. In this paper we deal only with anonymous auction procedures.

Classical auction theory associates a (Bayesian) game with each auction procedure and analyzes the behavior of the agents under the equilibrium assumption. In order to associate a game with an auction procedure we have to define two additional parameters, the information structure and the utility functions of the players in the game. In this paper the information structure is given by a vector $\bar{F} = (F_i)_{i=1}^n$ of distribution functions on R_+. Let $(\hat{v}_i)_{i=1}^n$ be mutually independent non-negative random variables such that the distribution of $\hat{v}_i$ is F_i, that is, $F_i(v) = Prob(\hat{v}_i \leq v)$. If $\hat{v}_i = v_i$, v_i is interpreted as the maximal willingness to pay of Agent i and it is called the type of i.[2] The unique probability measure induced by F_i on R_+ is denoted by P_i. Let P denote the product probability measure of $(P_i)_{i=1}^n$ on R_+^n and let P_{-i} denote the product probability measure defined by $(P_j)_{j \in N \setminus \{i\}}$ on $R^{N \setminus \{i\}}$, where $N = \{1, 2, \cdots, n\}$. The utility structure is given by a vector $\bar{u} = (u_i)_{i=1}^n$ of utility functions for money. Agent i has the utility function $u_i : R \to R$, normalized with $u_i(0) = 0$. It is assumed that u_i is increasing and 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

[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/or the agents' types are correlated [12]. We believe that the independence assumption is the right one in the Internet auctions setup where there are many anonymous participants.

$u_i(v_i - x_i)$. 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.[3] Agent i is *risk-seeking* if u_i is convex. A *strategy* for agent i is a measurable function $b_i : R_+ \to M$, where $b_i(v_i)$ is the message submitted by i when his type is v_i. We identify two strategies of i that coincide almost surely with respect to P_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 almost every v_i (w.r.t P_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 formulas 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[4]. 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 almost every type v_i (w.r.t P_i),

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

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

is attained at $m_i = b_i(v_i)$, where $E_{P_{-i}}$ denotes the expected value operator with respect to P_{-i}. Some pathological auction games do not possess equilibrium strategies. In this paper we do not deal with auction games that

[3]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.

[4]For example, no problem arise if we assume a discrete model in which the message space M is finite and each $\hat{v}_i$ takes only finitely many values. In the economic literature it is common to deal with a continuous setup and to require smoothness conditions. In particular, it is assumed that the utility functions and the distribution functions are twice continuously differentiable.

are over complicated in the sense that they do not possess an equilibrium. To summarize, an *auction game* is a 4-tuple $G = (n, A, \bar{F}, \bar{u})$, where $n \geq 2$, A is an anonymous auction procedure for n potential participants, $\bar{F}$ is an information structure, and $\bar{u}$ is a utility structure. G is called *information-symmetric* if there exists a distribution function F such that $F_i = F$ for all i. G is called *utility-symmetric* if $u_i = u_j$ for all i, j.

Under the equilibrium assumption (i.e, the assumption that economic agents use equilibrium strategies), if the auction game $G = G(n, A, \bar{u}, \bar{F})$ has a unique equilibrium profile $b = (b_1, b_2, \ldots, b_n)$, then the seller expected revenue in equilibrium 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. As we said, we do not consider auction games that are over complicated in the sense that they do not have an equilibrium. Consider a fixed number of participants n, a fixed information structure $\bar{F}$, and a fixed vector of utility functions $\bar{u}$ with $u_i = u$ for all i and $u(x) = x$ for all x. Myerson [15] solved the optimality problem $\max_A R_G$ and in addition he showed that if the distribution functions are the uniform distributions on bounded intervals, an optimal auction procedure is a second-price auction with an appropriate reservation price determined by the information structure. Maskin and Riley [9] 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.

2.1 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 $d_i(x) = 0$.[5] 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, $(b_1, b_2, \ldots, b_n) = (s, s \ldots, s)$, where s is the unique solution of the differential equation defined on $[0, 1]$ by (see [17]):

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

In particular, if the agents are risk-neutral and $F(x) = x$ is the uniform distribution on $[0, 1]$, then $s(v) = \frac{n-1}{n}v$ 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-highest bid in 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 highest 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 [23, 6, 14] and [13]). That is, $b_i(v_i) > v_i$.

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

English Auctions (Ascending 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. The precise formulation of the dynamic game associated with the auction procedure is not an easy task. Milgrom and Weber [12] provided one such formulation for a modification of the Japanese variant of the English auction[6]. It is interesting to note that English auctions (at least in the variant (and setup) discussed in [12]) 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 (Descending 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 in particular, the expected revenue of the seller in both auctions coincide.

[6]In the Japanese variant, the price is raised continuously, and a bidder who wishes to be active at the current price presses a button. When he releases the button, he has withdrawn from the auction. We do not know the source of the term "Japanese" used in [12].

[7]This equivalence principle does not hold if we do not assume the independence of types (see [12]).

3 An Upper Bound on the Seller's Expected Revenue

In this section we supply an upper bound on the expected 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 mechanism is, although agents may have bids that are higher than their actual evaluation of a good [8] the seller's revenue cannot be higher than the expected highest valuation for this good of the auction's participants. We further show that there exists an information structure for which our bound is a least upper bound.

Theorem 1 *Let $G = (n, A, \bar{F}, \bar{u})$ be an auction game with risk-averse participants that possesses an equilibrium. Then for every equilibrium strategy profile $b = (b_1, b_2, \cdots, b_n)$, the expected revenue of the seller from the auction is bounded by the expected highest type. That is,*

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

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:

$$B_i(v_i) = E_{P_{-i}}\left(u_i(v_i - c_i(b_i(v_i), b_{-i}))g_i(b_i(v_i), b_{-i}) + u_i(-d_i(b_i(v_i), b_{-i}))(1 - g_i(b_i(v_i), b_{-i}))\right).$$

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

$$E_{P_{-i}}\left(u_i((v_i - c_i(b_i(v_i), b_{-i}))g_i(b_i(v_i), b_{-i}) + (-d_i(b_i(v_i), b_{-i})(1 - g_i(b_i(v_i), b_{-i})))\right) \geq 0.$$

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

$$u_i(E_{P_{-i}}\left((v_i - c_i(b_i(v_i), b_{-i}))g_i(b_i(v_i), b_{-i}) + (-d_i(b_i(v_i), b_{-i})(1 - g_i(b_i(v_i), b_{-i}))\right) \geq 0.$$

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

Because u_i is increasing and $u_i(0) = 0$,

$$E_{P_{-i}}\left((v_i - c_i(b_i(v_i), b_{-i}))g_i(b_i(v_i), b_{-i}) + (-d_i(b_i(v_i), b_{-i})(1 - g_i(b_i(v_i), b_{-i})))\right) \geq 0.$$

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

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

because $\sum_{i=1}^n g_i(b) \leq 1$. ∎

We now show that when each F_i is the uniform distribution on a given bounded interval, then the upper bound of Theorem 1 is a least upper bound (l.u.b.).

Proposition 1 *Let $n \geq 2$. Let A be the first-price auction mechanism for n participants. Let $0 \leq A < B < \infty$, and let $\bar{F} = (F_i)_{i=1}^n$ be the symmetric information structure in which F_i is the uniform distribution over $[A, B]$ for all i. Then*

$$\sup_{\bar{u}} R_{n,A,\bar{F},\bar{u}} = E_P(\max(\hat{v}_i)_{i=1}^n),$$

where $\bar{u}$ ranges over all utility structures of risk-averse agents for which the associated auction game possesses equilibrium.

Proof:

By Theorem 1,

$$\sup_{\bar{u}} R_{(n,A,\bar{F},\bar{u})} \leq E_P(\max(\hat{v}_i)_{i=1}^n).$$

We show that equality holds in the above inequality when there are two participants ($n = 2$), and each F_i, $i = 1, 2$, is the uniform distribution over $[0, 1]$. The proof of the general case is similar. For every $m > 2$ consider the auction game G_m in which the utility functions satisfy $u_i = u_m$, $i = 1, 2$, where

$$u_m(x) = 1 - e^{-mx} \quad for \quad every \quad x \in R.$$

We will show that when m approaches infinity the expected gain in the auction game G_m approaches $E_P(\max(\hat{v}_1, \hat{v}_2))$. From the equilibrium equation for first price auction given in Section 2 we can deduce that the auction game G_m possesses a unique equilibrium (b_1, b_2) such that $b_i = s_m$ for all i and

$$s'_m(v) = \frac{1 - e^{-m(v - s_m(v))}}{mve^{-m(v - s_m v)}}.$$

Therefore,

$$s'_m(v) = \frac{e^{mv} - e^{ms_m(v)}}{mve^{ms_m(v)}},$$

or

$$(ve^{ms_m(v)})' = e^{mv}.$$

As $s_m(0) = 0$,

$$s_m(v) = \frac{1}{m} \ln \left(\frac{e^{mv} - 1}{mv} \right), \quad 0 < v \leq 1.$$

Let $\hat{v}_{max} = \max\{\hat{v}_1, \hat{v}_2\}$, then the expected revenue of the seller is $R_m = E(s_m(\hat{v}_{max}))$. As the distribution of $\hat{v}_{max}$ is $G(v) = v^2$, we have that

$$R_m = \int_0^1 s_m(v) 2v dv.$$

Using standard calculus techniques, it can be shown that $s_m(v) \leq v$ and that $lim_{m \to \infty} s_m(v) = v$ for all $v \in [0, 1]$. Therefore, by the Lebesque's convergence theorem,

$$\lim_{m \to \infty} R_m = \int_0^1 2v^2 dv = E(\hat{v}_{max}).$$

Hence, $\sup_{\bar{u}} R_{n,A,\bar{F},\bar{u}} = E(\hat{v}_{max})$. ▮

4 Almost Optimal Auctions

Consider a seller who wishes to conduct an auction in a setup with many potential buyers (e.g., the Internet). The seller knows the number of potential participants[9] as well as the symmetric information structure. That is he knows F such that $F_i = F$ for all i. However, he cannot possibly know the utility structure. In order to maximize his revenue, the seller may behave as a Bayesian agent, that is he can assign a probability to each possible vector of utility functions and compute his expected revenue for each possible auction procedure. Even if the seller can assign a "reasonable" probability distribution over utility structures, the Bayesian approach cannot be considered as a practical solution because of the very high level of computational complexity. Another possible solution is to assume that all agents are risk-neutral in their relation to money in the relevant range and to use Myerson's result mentioned in Section 2. That is, to conduct a second-price auction with an appropriate reservation price. In this section we show that when F is supported in a bounded interval, this approach makes sense, not only because it is the best the seller can do in the risk-neutral case, but rather because the second-price auction procedure is almost optimal for every utility structure of risk-averse agents if the number of participants is sufficiently large.

Assume that the distribution function F satisfies $F(A) = 0$, $F(B) = 1$, F is continuously differentiable in $[A, B]$, and $F'(x) > 0$ for every $x \in (A, B)$. Let $(\hat{v}_i)_{i=1}^{\infty}$ be a sequence of i.i.d. random variables, where each $\hat{v}_i$ is distributed according to F. We denote by v_{max}^n the first-order statistics (that is, the highest type) of $(\hat{v}_i)_{i=1}^{n}$, and we denote by v_{sec}^n, the second-order statistics (that is, the second-highest type). Let $R_{max}^n = E(v_{max}^n)$ denote the expected value of the first-order statistics amongst $(\hat{v}_i)_{i=1}^{n}$, and let $R_{sec}^n = E(v_{sec}^n)$ denote the expected value of the second-order statistics. It can be shown that for every utility structure $(u_i)_{i=1}^{n}$, the second-price auction game with the symmetric information structure given by F has a unique equilibrium in which $b_i(v_i) = v_i$ for every $v_i \in [A, B]$. Therefore the expected revenue of the seller in this game equals R_{sec}^n. Thus, the expected revenue of the seller in a second-price auction game does not depend on the utility structure.

Theorem 2 *Let $0 \leq A < B < \infty$. Let F be a distribution function on R_+ that satisfies $F(A) = 0$, $F(B) = 1$, and $F'(x) > 0$ for every $A < x < B$. Let $(\hat{v}_i)_{i=1}^{\infty}$ be a sequence of independent and identically F-distributed random*

[9] Our results can be extended to the case of random number of participants. See Section 7.

variables. For every $n \geq 2$, let A_n be the second-price auction mechanism for n participants, let $\bar{F}(n) = (F_1, \ldots, F_n)$ be the information structure for n participants in which $F_i = F$ for every $1 \leq i \leq n$, and let $\bar{u}(n) = (u_{i,n})_{i=1}^{n}$ be a sequence of utility functions.

Then

$$\lim_{n\to\infty} \frac{R_{G_n}}{R^n_{max}} = 1,$$

where G_n is the second-price auction game $(n, A_n, \bar{F}(n), \bar{u}(n))$

Proof: As $R_{G_n} = R^n_{sec}$, we have to show that

$$\lim_{n\to\infty} \frac{R^n_{sec}}{R^n_{max}} = 1.$$

Without loss of generality we can assume that $[A, B] = [0, 1]$. We will prove that both, $\lim_{n\to\infty} R^n_{max} = 1$ and $\lim_{n\to\infty} R^n_{sec} = 1$. By the Lebesque's dominance convergence theorem it suffices to show that $lim_{n\to\infty} v^n_{max} = 1$ almost surely and that $\lim_{n\to\infty} v^n_{sec} = 1$ almost surely. It can be easily deduced from Theorem 1 (Page 251) in [21], that a sequence of random variables $(X_n)_{n=1}^{\infty}$ converges to a random variable X if for every $0 < \epsilon < 1$,

$$\sum_{n=1}^{\infty} Prob(|X - X_n| \geq \epsilon) < \infty.$$

It can be easily verified that the distribution function of v^n_{max} is $F(x)^n$, and that the distribution function of v^n_{sec} is $nF(x)^{n-1} - (n-1)F(x)^n$ (see e.g., [1]). Therefore, for $0 < \epsilon < 1$

$$\sum_{n}^{\infty} Prob(1 - v^n_{max} \geq \epsilon) = \sum_{n=1}^{\infty} F(1-\epsilon)^n < \infty,$$

and

$$\sum_{n}^{\infty} Prob(1 - v^n_{sec} \geq \epsilon) = \sum_{n=1}^{\infty} (nF(1-\epsilon)^{n-1} - (n-1)F(1-\epsilon)^n) < \infty,$$

because $F(1-\epsilon) < 1$. ∎

As English auctions are equivalent to second-price auctions, Theorem 2 (combined with Theorem 1) shows that in the information-symmetric independent-private-value model, when agents are risk-averse, the optimal expected revenue can be almost obtained in an English auction, when the number of agents is large enough (which one may expect in Internet auctions). Moreover, this result holds uniformly for all utility structures, in the sense that the number of participants that is needed to guarantee ϵ-optimality does not depend on the utility functions of the participants.

Our results have also an interpretation which is somewhat complementary to the famous revenue equivalence principle [15]. 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 this implication is almost true in a setup with many participants (e.g., Internet auctions) and arbitrary utility structures of risk-averse agents (which is the most general case discussed in the literature).

5 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 [17, 9]. 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 [13] 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 [23, 6, 14].

We will discuss the equilibrium set of a third-price auction with three participants ($n = 3$) in which $F_i = F$ for all $1 \leq i \leq 3$, where F is the uniform distribution on $[0, 1]$. We assume that each participant uses the utility function $u_i = w_\alpha$, $\alpha > 1$, where $w_\alpha(x) = x$ when $x \leq 0$, and $w_\alpha(x) = \alpha x$ for $x > 0$. By a slight modification in the proof of Theorem AT in [14] [10] it can be shown that a continuous function s_α defined on $[0, 1]$ constitutes a symmetric equilibrium strategy in the auction game if and only if $s_\alpha(0) = 0$, s_α is increasing, and

$$\int_{t=0}^{v} w_\alpha(v - s_\alpha(t))dt = 0$$

for every $v \in [0, 1]$. Solving this integral equation yields the unique solution:

$$s_\alpha(v) = (1 + \sqrt{\alpha})v, \quad v \in [0, 1].$$

Because of the simple structure of the utility functions, it can be shown that $(b_1, b_2, b_3) = (s_\alpha, s_\alpha, s_\alpha)$ is also the unique continuous equilibrium in this

[10] This theorem is proved under the assumption that the utility functions are twice continuously differentiable.

game. We do not know whether there exists non-continuous equilibrium in this game. Hence we have:

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.*

Proof: Assuming that the participants use continuous strategies, we have shown above that the expected revenue of the seller in the third-price auction described above is

$$R_\alpha = (1 + \sqrt{\alpha}) \cdot E(v_{[3]}),$$

where $E(v_{[3]})$ is the expected value of the third-order statistics (i.e., the expected value of the type of the agent with the third-highest evaluation). As $lim_{\alpha \to \infty} R_\alpha = \infty$, for sufficiently large α, $R_\alpha > E(\max(\hat{v}_1, \hat{v}_2, \hat{v}_3))$. ∎

Actually, in our case the expected first-order statistics is $\frac{3}{4}$, while the third-order statistics is $\frac{1}{4}$. Hence, if $\alpha > 4$ we get the desired result. Notice that if the agents are extremely risk-seeking (α is big), then the seller can extract a very high expected revenue in this game.

6 Remarks and Conclusion

Other Information Structures Theorem 1 can be extended to the general information structure case, where the players' types are correlated, and a player does not necessarily know his own type. We don't know under which conditions Theorem 2 remains valid in this case.

Random number of Participants It is not reasonable to model Internet auctions under the assumption that the number of participants is fixed and commonly known. However, since we use the independent-private-value model of information, all our results can be easily extended to the case in which the number of participants is randomly determined (see [11]).

Conclusion This paper has addressed several basic issues in the adaptation of economic mechanisms to computational and artificial environments, such as the Internet. There is much left to be done. Our upper bound can be refined in order to discuss particular classes of auctions and utility functions, and our result about risk-seeking agents can be further elaborated in order to find the optimal auction for risk-seeking agents. We believe that the introduction of the related issues, which have led to non-trivial and somewhat unpredictable results, clarify the importance and necessity of the related studies, and serve as essential steps in the adaptation of economic mechanisms to computational environments.

References

[1] B.C. Arnold, N. Balakrishman, and H.N. Nagaraja. *A First Course in Order Statistics*. John Wiley and Sons, 1992.

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

[3] Y. Demazeau and J.P. Muller. *Decentralized AI*. North Holland/Elsevier, 1990.

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

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

[6] J.H. Kagel and D. Levin. Independent private value auctions: bidder behavior in first-, second-, and third-price auctions with varying numbers of bidders. *Economic Journal*, 103:868–879, 1993.

[7] S. Kraus and J. Wilkenfeld. The Function of Time in Cooperative Negotiations. In *Proc. of AAAI-91*, pages 179–184, 1991.

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

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

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

[11] R.P. McAfee and J. McMillan. Auctions with a stochastic number of bidders. *Journal of Economic Theory*, 43:1–19, 1987.

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

[13] D. Monderer and M. Tennenholtz. Internet Auctions: Are They Gamblers' Attraction? Working paper, 1998.

[14] D. Monderer and M. Tennenholtz. K-Price Auctions. to appear in Games and Economic Behavior, 2000.

[15] R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6:58–73, 1981.

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

[17] J.G. Riley and W.F. Samuelson. Optimal auctions. *American Economic Review*, 71:381–392, 1981.

[18] Jeffrey S. Rosenschein and Gilad Zlotkin. *Rules of Encounter.* MIT Press, 1994.

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

[20] Tuomas W. Sandholm and V.R. Lesser. Equilibrium Analysis of the Possibilities of Unenforced Exchange in Multiagent Syustems. In *Proc. 14th International Joint Conference on Artificial Intelligence*, pages 694–701, 1995.

[21] A.N. Shiryayev. *Probability.* Springer-Verlag, 1984.

[22] Michael P. Wellman. A market-oriented programming environment and its application to distributed multicommodity flow problems. *Journal of Artificial Intelligence Research*, 1:1–23, 1993.

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