

Bidding Strategies for Realistic Multi-Unit Sealed-Bid Auctions

Ioannis A. Vetsikas and Nicholas R. Jennings

School of Electronics and Computer Science

University of Southampton

Southampton SO17 1BJ, UK

{iv,nrj}@ecs.soton.ac.uk

Abstract

When autonomous agents decide on their bidding strategies in real world auctions, they have a number of concerns that go beyond the models that are normally analyzed in traditional auction theory. Oftentimes, the agents have budget constraints and the auctions have a reserve price, both of which restrict the bids the agents can place. In addition, their attitude need not be risk-neutral and they may have uncertainty about the value of the goods they are buying. Some of these issues have been examined individually for single-unit sealed-bid auctions. In this paper, we extend this analysis to the multi-unit case, and also analyze the multi-unit sealed-bid auctions in which a combination of these issues are present, for unit-demand bidders. This analysis constitutes the main contribution of this paper. We then demonstrate the usefulness in practice of this analysis; we show in simulations that taking into account all these issues allows the bidders to maximize their utility. Furthermore, using this analysis allows a seller to improve her revenue, i.e. by selecting the optimal reserve price.

Introduction

Auctions have become commonplace; they are used to trade all kinds of commodity, from flowers and food to industrial commodities and keyword targeted advertisement slots, from bonds and securities to spectrum rights and gold bullion. Once the preserve of governments and large companies, the advent of online auctions has opened up auctions to millions of private individuals and small commercial ventures. Given this, it is desirable to develop autonomous agents that will let the masses participate effectively in such settings, even though they do not possess professional expertise in this area. To achieve this, however, these agents should account for the features of real-world auctions that expert bidders take into consideration when determining their bidding strategies.

While game theory is widely used in multi-agent systems as a way to model and predict the interactions between rational agents in auctions, the models that are canonically analyzed are rather limited. As discussed below, some work has been done towards extending these models to incorporate features that are important in real auctions, but this work

invariably looks at each feature separately; additionally the cases examined are almost all instances of single-unit auctions. While this is useful for economists and perhaps expert bidders, who can integrate the lessons learned using human intuition and imagination, an automated agent cannot do this. It is therefore necessary to analyze the strategic behavior in multi-unit (m^{th} and $(m+1)^{th}$ price) auction models that incorporate all the relevant features. To this end, we have looked at a number of auctions, ranging in scope from the eBay auctions (held mainly between individuals) to B2B auctions (used by businesses to procure materials and commodities), with various different rules, ranging from the traditional English auction to the position auction used by Google Adwords. Despite their differences, a number of common features are present. We list the most important of these below and highlight what is already known about each of them.

First, *budget constraints* are very important, whenever businesses and individuals place bids, because they limit the upper range of these bids. Here, we will assume that the available budget constitutes an absolute spending limit. Now, this case has been examined for single-unit auctions (Che and Gale 1998), but not for multi-unit ones;¹ it has also been proven that the revenue generated by a 1st price auction is always higher than that of the equivalent 2nd price one.

Second, bidders may adopt different *attitudes towards risk*. Essentially, this indicates whether bidders are conservative or not, and their willingness to take risk in order to gain additional profit. Normally bidders are assumed to be “risk-neutral”, meaning their utility equals their profit. However, they can also be “risk-averse”, “risk-seeking”, or even have a more complicated risk attitude; for example, the bidding behavior of bidders on eBay suggests a complicated attitude towards risk. The equilibrium strategy of a risk-averse agent participating in a 1st price auction has been analyzed before, e.g. in (Liu, Goodwin, and Koenig 2003).

Third, setting a *reserve price* (i.e. a minimum transaction price) in the auction is a common way for the seller to increase her profit; auctions taking place in traditional auction

¹The “dual” problem of designing a truthful mechanism for budget constrained bidders has been examined in (Borgs et al. 2005).

houses typically have a reserve price. This case has been examined for single-unit sealed-bid auctions in (Myerson 1981; Riley and Samuelson 1981).² In (Maskin and Riley 1984), this problem is examined for risk-averse bidders.

Fourth, there may be *uncertainty in the bidders' valuation* of the offered commodity. For example, when businesses bid in the Google Adwords keyword auction, they can't precisely know the additional revenue that advertising in this way will bring them, and therefore they can't evaluate the actual economic value of the ad. Nevertheless, it can be assumed that the agent has some idea about his own value and this can be represented by a probability distribution. In the literature, this problem has been mostly looked at from the point of view of having a cost for introspection, which allows the agent to determine his valuation more precisely (Larson and Sandholm 2001; Thompson and Leyton-Brown 2007).³ However, in many practical settings, introspection is simply not possible, because of the lack of further relevant data, or excessive costs that cannot be justified by the increased accuracy.⁴

The last important feature is considering a bidder's desire to *purchase multiple items*, with a different valuation for each. In this case, it is known that bidders should shade their bids, compared to the case when only one item is desired, even to the point of bidding for less items than desired, in order to gain more profit (strategic demand reduction) (Weber 1997). To date, however, an optimal strategy is not known for this feature; it is open problem. This is the reason why we make the usual assumption that each agent wishes to buy only one unit (*unit-demand bidders*), like e.g. in (Vetsikas and Jennings 2007).

Given this background, in this paper we make the following contributions:

- We derive novel equilibria for the multi-unit m^{th} price sealed-bid auction case, when these features are looked at separately.
- We combine, for the first time, two of the features that we look at. Specifically, we derive the dominant strategy for the case of uncertainty in the valuation that bidders have, when the bidders are not risk-neutral, in the setting of the $(m + 1)^{th}$ price auction.⁵

²In our previous short paper (Vetsikas and Jennings 2008), we presented some initial stages of this work. More specifically the equilibrium strategies that exist in the cases when there is a reserve price and that agents are not risk-neutral (examined separately) are presented; this is why these cases are only briefly discussed in this paper.

³There are few notable exceptions where no introspection is assumed, such as (Parkes 1999), which examines how to select the auction that yields the highest revenue and/or efficiency, when bidders have uncertainty in their valuations.

⁴In most of the above settings, the dominant strategy in the case of a 2^{nd} price auction is some variation of truth-telling (Krishna 2002). This result is generalized trivially to the multi-unit variant (the $(m + 1)^{th}$ price auction).

⁵We do not analyze the equilibrium, in the presence of valuation uncertainty, in the m^{th} price auction because our model needs to be enriched in order to be able to analyze this case; we leave this for future work.

- After doing this, we combine all the features. We first derive the equilibrium strategies for both the m^{th} and $(m + 1)^{th}$ price auction, in the presence of budget constraints, reserve prices and any bidder risk attitude. Then we also include uncertainty of bidders' valuation for the case of the $(m + 1)^{th}$ price auction and derive the dominant strategy.
- Finally, we demonstrate the usefulness in practice of this analysis; we show, using simulations, that taking into account all these features allows the bidders to improve their utility, as opposed to using a strategy which doesn't account for all of them. Furthermore, a seller can maximize her revenue by selecting the optimal reserve price.⁶

The Multi-Unit Auction Setting

In this section we formally describe the auction setting to be analyzed and define the objective function that the agents wish to maximize. We also give the notation that we use.

In particular, we will compute and analyze the symmetric Bayes-Nash equilibria⁷ for sealed-bid auctions where $m \geq 1$ identical items are being sold; these equilibria are defined by a strategy, which maps the agents' valuations v_i to bids b_i . The two most common settings in this context are the m^{th} and $(m + 1)^{th}$ price auctions, in which the top m bidders win one item each at a price equal to the m^{th} and $(m + 1)^{th}$ highest bid respectively. We assume that there is a reserve price $r \geq 0$ in our setting; this means that bidders, who wish to participate in the auction, must place bids $b_i \geq r$.

We assume that N indistinguishable bidders (where $N \geq m$) participate in the auction and they have a private valuation (utility) v_i for acquiring any one of the traded items; these valuations are assumed to be i.i.d. from a distribution with cumulative distribution function (cdf) $F(v)$, which is the same for all bidders. In the case that there is uncertainty about the valuation v_i , the agent knows that it is drawn from distribution $G_i(v_i)$, but not the precise value. As the valuations v_i are independent, we can assume that any uncertainty that a bidder has about his own valuation is independent of the uncertainty he has about other agents' valuations.

We also assume that each bidder has a certain budget c_i , which is known only to himself and which limits the maximum bid that he can place in the auction. The available budgets of the agents are i.i.d. drawn from a known distribution with cdf $H(c)$.

According to utility theory, every rational agent has a strictly monotonically increasing utility function $u(\cdot)$ that maps profit into utility; the alternative with the highest expected utility is the preferred outcome. This function determines the agent's risk attitude. Some functions $u(x)$

⁶The use of equilibrium analysis in order to design autonomous, intelligent agents has been done before, e.g. in (Liu, Goodwin, and Koenig 2003) and (Vetsikas and Jennings 2007).

⁷The Bayes-Nash equilibrium is the standard solution used in game theory to analyze auctions. The equilibria being symmetric means that all agents use the same bidding strategy. This is a common assumption made in game theory, in order to restrict the space of strategies that we examine. It is likely that in addition to the symmetric equilibria we compute there are also asymmetric ones.

used widely in economics are: $u(x) = x^\alpha, \alpha \in (0, 1)$ (CRRA), $u(x) = 1 - \exp(-\alpha x), \alpha > 0$ (CARA), and $u(x) = -\gamma x, \gamma \in [0, 1]$, all of which indicate risk-averse bidders.

We also use the following additional notation in the proofs: $Z(x)$ is the probability distribution of any opponent's bid b_j . Thus $Z(x) = \text{Prob}[b_j \leq x]$, and $B^{(k)}$ is the k^{th} order statistic of these bids of the opponents. Since there are $(N - 1)$ opponents for each agent, the distribution $\Phi_k(x)$ of $B^{(k)}$ can be computed as (Rice 1995):

$$\Phi_k(x) = \sum_{i=0}^{k-1} C(N-1, i) (Z(x))^{N-1-i} (1-Z(x))^i \quad (1)$$

where the notation $C(n, k)$ is the total number of possible combinations of k items chosen from n . As shown in (Vetsikas and Jennings 2007), $\forall N \geq m$ the following holds:

$$\Phi'_m(x) = (N-m)(\Phi_m(x) - \Phi_{m-1}(x)) \frac{Z'(x)}{Z(x)} \quad (2)$$

Equilibria in the Presence of Reserve Prices and Varying Risk Attitudes

In this section we present the equilibria that exist in an m^{th} price auction, in the cases when there is a reserve price $r \geq 0$ and that agents are not risk-neutral. We examine each case separately in the following theorems. More details on these two cases are given in (Vetsikas and Jennings 2008).

Theorem 1 *In the case of an m^{th} price sealed-bid auction, with reserve price $r \geq 0$, with N participating risk-neutral bidders, in which each bidder i is interested in purchasing one unit of the good for sale with inherent utility (valuation) for that item equal to v_i , where v_i are i.i.d. drawn from $F(v)$, the following bidding strategy constitutes a symmetric Bayes-Nash equilibrium:*

$$g(v_i) = v_i - (F(v_i))^{-(N-m)} \int_r^{v_i} (F(z))^{N-m} dz \quad (3)$$

Theorem 2 *In the case of an m^{th} price sealed-bid auction with N participating bidders, in which each bidder i is interested in purchasing one unit of the good for sale with inherent utility (valuation) for that item equal to v_i , where v_i are i.i.d. drawn from $F(v)$, and a risk attitude which is described by utility function $u(\cdot)$, the bidding strategy $g(v)$, which constitutes a symmetric Bayes-Nash equilibrium, is the solution of the differential equation:*

$$g'(v_i) = \frac{u(v_i - g(v_i)) - u(0)}{u'(v_i - g(v_i))} (N-m) \frac{F'(v_i)}{F(v_i)} \quad (4)$$

with boundary condition $g(0) = 0$.

Equilibria in the Case of Budget Constraints

In this section we examine the equilibria for the case when agents have budget constraints only. We assume that the agents are risk-neutral and that the reserve price is $r = 0$.

Theorem 3 *In the case of an m^{th} price sealed-bid auction with N participating risk-neutral bidders, in which each bidder i is interested in purchasing one unit of the good for sale with inherent utility (valuation) for that item equal to*

v_i , and has a budget constraint c_i , where v_i and c_i are i.i.d. drawn from $F(v)$ and $H(c)$ respectively, the following bidding strategy constitutes a symmetric Bayes-Nash equilibrium:

$$b_i = \min\{g(v_i), c_i\} \quad (5)$$

where $g(v)$ is the solution of the differential equation:

$$g'(v_i) = \frac{(1 - H(g(v_i)))F'(v_i)}{\frac{1 - (1 - H(g(v_i)))(1 - F(v_i))}{(N-m)(v_i - g(v_i))} - (1 - F(v_i))H'(g(v_i))} \quad (6)$$

with boundary condition $g(0) = 0$.

PROOF.(brief) Consider any bidder i . We assume that each opponent j ($j \neq i$) bids $b_j = \min\{g(v_j), c_j\}$. The distribution $Z(x)$ of the opponent's bid b_j is equal to:

$$Z(x) = 1 - (1 - F(g^{-1}(x)))(1 - H(x))$$

The k^{th} order statistic of these bids $B^{(k)}$ is drawn from distribution $\Phi_k(x)$, described by equation 1. Using this fact, we compute the expected utility of bidder i , who places bid b_i :

$$Eu_i(b_i) = (v_i - b_i)(\Phi_m(b_i) - \Phi_{m-1}(b_i)) + \int_0^{b_i} (v_i - \omega) \frac{d\Phi_{m-1}(\omega)}{d\omega} d\omega$$

The bid b_i which maximizes this utility is found by setting: $\frac{dEu_i(b_i)}{db_i} = 0 \Leftrightarrow -(\Phi_m(b_i) - \Phi_{m-1}(b_i)) + (v_i - b_i)\Phi'_m(b_i) = 0$. Since this bid maximizes the agent's utility, it must be equal to $b_i = g(v_i)$. Using this fact and equation 2 we eventually get differential equation 6. The bidder bids $b_i = 0$, when his valuation $v_i = 0$, thus the boundary condition is $g(0) = 0$.

Now it is time to consider how the available budget c_i changes the bid b_i . If $g(v_i) \leq c_i$, then the agent bids $b_i = g(v_i)$, as this maximizes the expected profit $Eu_i(b_i)$. When $g(v_i) > c_i$, it is $\frac{dEu_i}{db_i} > 0, \forall b_i < g(v_i)$, and the bid must be $b_i \leq c_i < g(v_i)$, therefore the bid b_i that maximizes Eu_i is $b_i = c_i$. Thus it is $b_i = \min\{g(v_i), c_i\}$. ■

In order to discuss features of the equilibrium strategy, we examine the special case when both the valuations v_i and the budget constraints c_i are drawn from a uniform distribution on $[0, 1]$ (i.e. $F = H = U[0, 1]$). This distribution is the canonical one used in auction theory for this purpose:

Corollary 1 *In the case that $F(v)$ and $H(c)$ are uniform distributions $U[0, 1]$, the equilibrium strategy is: $b_i = \min\{g(v_i), c_i\}$, where $g(v)$ is the solution of the d.e.:*

$$g'(v_i) = \frac{1 - g(v_i)}{\frac{v_i + g(v_i) - v_i g(v_i)}{(N-m)(v_i - g(v_i))} - (1 - v_i)} \quad (7)$$

In figure 1, we present the solution of this equation for different numbers of bidders, N , and items being sold, m . Because of the presence of the budget constraints, it is easier for bidder i to win the auction now; this happens, because an opponent j , with valuation v_j higher than bidder i , might not be able to outbid him, due to having a low budget constraint c_j . It is therefore expected that strategy $g(v)$ would suggest bidding less than the equilibrium strategy $g_\infty(v)$, when there is infinite budget (i.e. no constraints); this strategy is: $g_\infty(v) = \frac{N-m}{N-m+1}$ (see (Krishna 2002)).

It is also interesting to note that function $g(\cdot)$ deviates from $g_\infty(\cdot)$ most in the case when $N - m = 4$. This is due to two

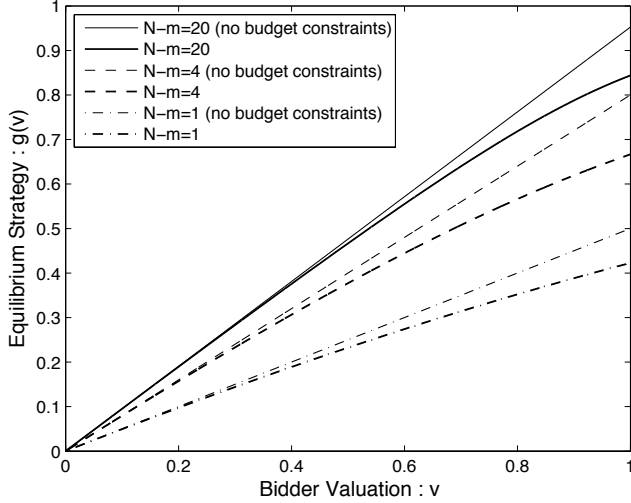


Figure 1: Equilibrium strategy $g(v)$ for the case when the valuations v and the budget constraints c are both drawn from the uniform distribution $U[0, 1]$; the equilibrium bid is $b = \min\{g(v), c\}$. The number of bidders, N , and the number of items being sold, m , take values $N - m = 1, 4, 20$. The equilibrium strategy $g_\infty(v) = \frac{N-m}{N-m+1}$, when there are no budget constraints is also presented.

facts, which would also affect to some degree the bidding strategy in the case of any other distributions of $F(v)$ and $H(c)$. First, when $N - m = 20$, there is a large number of opponents, and therefore a higher probability that one of them has a large budget. Second, when $N - m = 1$, the bids according to both $g()$ and $g_\infty()$ are usually smaller than the budget constraints c_i , which are drawn from $U[0, 1]$; as the bids are less constrained, the deviation is therefore smaller.⁸

In the case of an $(m+1)^{th}$ price auction, the agents submit truthful bids, if these are higher than the budget constraint c_i (Krishna 2002):

Theorem 4 *In an $(m+1)^{th}$ price auction, where the bidders have valuations v_i and budget constraints c_i , it is a dominant strategy to bid: $b_i = \min\{v_i, c_i\}$.*

Uncertainty in Agents' Valuations

In this section, we examine the case when the bidders do not precisely know their own valuations. Specifically, we assume that each bidder i knows that his own valuation v_i is drawn from some distribution $G_i(v_i)$. The mean of $G()$ is $\mu_{G_i} = EG_i(v_i)$. Therefore the bidder knows that his own valuation is centered around value μ_{G_i} , but he doesn't know it precisely. His uncertainty is thus represented by the distribution $G_i(v_i)$. We examine this case for the $(m+1)^{th}$ price auction. In the case that bidders are risk-neutral, it is a dominant strategy to bid as if they had value $v_i = \mu_G$ and

⁸We would like to point out that, as $N - m$ grows, the deviation does not approach zero, i.e. the budget constraints always lead to some reduction of the bids, even if it's a very small one.

the valuation uncertainty does not matter.⁹ This is the reason why we choose to examine this issue when the agents are not necessarily risk-neutral.

Theorem 5 *In an $(m+1)^{th}$ price auction, if a bidder knows only imprecisely his own valuation v_i , in that it is drawn from distribution $G_i(v_i)$, and his risk attitude is described by utility function $u_i()$, it is a dominant strategy to bid b_i , which is the solution of equation:*

$$\int_{-\infty}^{\infty} u(z - b_i) G_i'(z) dz = u(0) \quad (8)$$

PROOF. From the point of view of each bidder i , it knows that its opponents are going to be placing bids b_j . Let us assume that the highest of the opponents bids' b_{-i} is drawn from distribution $\Omega_i(x)$ (i.e. that $Prob[b_{-i} \leq x] = \Omega_i(x)$). This distribution depends on the opponents' bidding strategies, valuations and risk attitudes. We assume that this can be any function, provided that it is differentiable and increasing. The probability of winning by placing bid b_i is $\Omega_i(b_i)$, and in that case the utility of the bidder is $u(v_i - b_{-i})$, whereas his utility is $u(0)$, if it doesn't win, which happens with probability $(1 - \Omega_i(b_i))$. Therefore the expected utility of bidder i , when it has a known valuation $v_i = z$ and places bid b_i , is:

$$Eu_i(z, b_i) = \int_{-\infty}^{b_i} u(z - x) \Omega_i'(x) dx + (1 - \Omega_i(b_i)) u(0)$$

Using Bayes' rule, we calculate the expected utility $Eu_i(b_i)$ of bidder i , when his valuation is unknown:

$$Eu_i(b_i) = \int_{-\infty}^{b_i} \left(\int_{-\infty}^{\infty} u(z - x) G_i'(z) dz \right) \Omega_i'(x) dx + (1 - \Omega_i(b_i)) u(0)$$

The bid b_i which maximizes the expected revenue is found by setting: $\frac{dEu_i(b_i)}{db_i} = 0$, which leads to equation 8. The solution of this equation is the bid b_i which maximizes the expected utility of bidder i . ■

From equation 8, it follows that risk-averse agents bid less than the mean μ_{G_i} of their valuation distribution $G_i()$, while risk-seeking agents do the opposite. Furthermore, risk-averse bidders will bid even less as the variance of distribution $G_i()$ increases, i.e. the more uncertain they get about their valuation, and the opposite happens to risk-seeking bidders. We can prove that, for any symmetric distribution $G_i()$, these statements always hold, in the following propositions:

Proposition 1 *In an $(m+1)^{th}$ price auction, if a bidder knows only imprecisely his own valuations v_i , in that it is drawn from symmetric distribution $G_i(v_i)$, and his risk attitude is described by utility function $u_i()$, it is a dominant strategy to bid:*

- $b_i < \mu_{G_i}$, if the bidder is risk-averse (i.e. $u_i()$ is concave),
- $b_i = \mu_{G_i}$, if the bidder is risk-neutral (i.e. $u_i()$ is a linear function), and
- $b_i > \mu_{G_i}$, if the bidder is risk-seeking (i.e. $u_i()$ is convex).

⁹The fact that risk-neutral agents bid the mean value μ_G that they have for their valuation has been observed in some of the related work, e.g. in (Larson and Sandholm 2001; Thompson and Leyton-Brown 2007; Parkes 1999).

PROOF. Equation 8 becomes $b_i - \mu_{G_i} = 0$, for linear function $u_i(x)$, which is the case of risk-neutral bidders. For risk-averse bidders, using the fact that $u(\cdot)$ is concave and that $G(\cdot)$ is symmetric, it can be shown that:

$$\int_{-\infty}^{\infty} u(z - \mu_{G_i})G'_i(z)dz < u(0)$$

This means that if we set $b_i = \mu_{G_i}$, the left hand side of equation 8 is less than the right hand side, and cannot be the solution; the solution must increase the left hand side of the equation, so it must be smaller: $b_i < \mu_{G_i}$. For risk-seeking bidders, a similar argument, as for the case of risk-averse bidders, yields: $b_i > \mu_{G_i}$. ■

Proposition 2 Assume two bidders with valuations v_i and v_j , which are drawn from the symmetric distributions $G_i(v_i)$ and $G_j(v_j)$; these distributions belong to the same class of distributions (i.e. they are both uniform) and they have the same mean μ , but they have different variance: $\sigma_i < \sigma_j$. The bidders have the same risk attitude, which is described by utility function $u(\cdot)$: if $u(\cdot)$ is concave (i.e. risk-averse bidders), then the bidders' bids are: $b_i > b_j$, and if $u(\cdot)$ is convex (i.e. risk-seeking bidders), then the bidders' bids are: $b_i < b_j$.

PROOF. We give a brief proof for concave function $u(\cdot)$; for convex $u(\cdot)$ the proof is similar. Let

$$F_i(x) = \int_{-\infty}^{\infty} u(z - x)G'_i(z)dz \Leftrightarrow$$

$$F_i(x) = \int_0^{\frac{1}{2}} \left(u(G_i^{-1}(p) - x) + u(G_i^{-1}(1-p) - x) \right) dp$$

and define $F_j(x)$ in the same way. It is then: $F_i(b_i) = u(0) = F_j(b_j)$. Since $G_i(\cdot)$ and $G_j(\cdot)$ are symmetric, belong to the same class of distributions, have the same mean μ , and furthermore $\sigma_i < \sigma_j$, the following formulas hold:

$$\forall p \in [0, \frac{1}{2}] : G_i^{-1}(p) - G_j^{-1}(p) = G_j^{-1}(1-p) - G_i^{-1}(1-p)$$

$$G_j^{-1}(p) < G_i^{-1}(p) < G_i^{-1}(1-p) < G_j^{-1}(1-p)$$

Subtracting b_i from all terms in these equations, and using the fact that $u(\cdot)$ is concave, we get:

$$u(G_i^{-1}(p) - b_i) - u(G_j^{-1}(p) - b_i) >$$

$$u(G_j^{-1}(1-p) - b_i) - u(G_i^{-1}(1-p) - b_i)$$

From this equation it follows that: $F_j(b_i) < F_i(b_i) = u(0) = F_j(b_j)$ and thus that $b_i > b_j$. ■

Putting It All Together

In this section we derive the equilibria for the case when agents have budget constraints c_i , their risk attitude is described by function $u(\cdot)$ (not necessarily risk neutral) and the auction has a reserve price r in the setting of an m^{th} price auction. In the setting of an $(m+1)^{\text{th}}$ price auction, in addition to these issues, we also consider that the valuations v_i are not known precisely.

Theorem 6 In the case of an m^{th} price sealed-bid auction, with reserve price $r \geq 0$, with N participating bidders, in which each bidder i is interested in purchasing one unit of the good for sale with inherent utility (valuation) for that item equal to v_i , and has a budget constraint c_i , where v_i

and c_i are i.i.d. drawn from $F(v)$ and $H(c)$ respectively, and the bidders have a risk attitude which is described by utility function $u(\cdot)$, the following bidding strategy constitutes a symmetric Bayes-Nash equilibrium:

$$b_i = \min\{g(v_i), c_i\} \quad (9)$$

where $g(v)$ is the solution of the differential equation:

$$g'(v_i) = \frac{(1 - H(g(v_i)))F'(v_i)}{\frac{u'(v_i - g(v_i))(1 - (1 - F(v_i))(1 - H(g(v_i))))}{(N-m)(u(v_i - g(v_i)) - u(0))} - (1 - F(v_i))H'(g(v_i))} \quad (10)$$

with boundary condition $g(r) = r$.

PROOF.(shortened) Using the same reasoning as that presented in the proof of theorem 3, we assume that $b_i = \min\{g(v_i), c_i\}$ are the bids placed by the agents.

Because of the reserve price r , there is a chance that an agent will not be able to participate in the auction, either because its budget is $c_i < r$, or because its valuation for the item is $v_i < r$. We therefore begin by analyzing the case when exactly $n \leq N$ agents can participate in the auction; these agents have $c_i \geq r$ and $v_i \geq r$. The probability that a particular agent participates in the auction is equal to:

$$\pi(r) = \text{Prob}[c_i \geq r \wedge v_i \geq r] = (1 - F(r))(1 - H(r))$$

The probability that exactly n (out of the N total) agents participate in this auction is thus:

$$\pi_n = C(N - 1, n - 1)(\pi(r))^{n-1} \cdot (1 - \pi(r))^{N-n} \quad (11)$$

The distribution of the opponents' bids b_j is:

$$Z_r(x) = 1 - \frac{1 - F(g^{-1}(x))}{1 - F(r)} \frac{1 - H(x)}{1 - H(r)} \quad (12)$$

The distribution of the k^{th} highest opponent bid $B^{(k)}$ is:

$$\Phi_k^{n,r}(x) = \sum_{i=0}^{k-1} C(n - 1, i)(Z_r(x))^{n-1-i}(1 - Z_r(x))^i \quad (13)$$

The expected utility of bidder i , when he places bid b_i , is:

$$Eu_i^{n,r}(b_i) = u(0)(1 - \Phi_m^{n,r}(b_i)) + u(v_i - b_i)\Phi_m^{n,r}(b_i)$$

$$+ \int_r^{b_i} u'(v_i - \omega)\Phi_{m-1}^{n,r}(\omega)d\omega \quad (14)$$

Using Bayes' rule, we compute the expected utility that bidder i gets, by placing bid b_i , for any possible numbers of total participating agents:

$$Eu_i(b_i) = u(0)(1 - \Phi_m(b_i)) + u(v_i - b_i)\Phi_m(b_i)$$

$$+ \int_r^{b_i} u'(v_i - \omega)\Phi_{m-1}(\omega)d\omega \quad (15)$$

where $\Phi_k(x) = \sum_{i=0}^{k-1} C(N - 1, i)(Z(x))^{N-1-i}(1 - Z(x))^i$ and $Z(x) = 1 - (1 - F(g^{-1}(x)))(1 - H(x))$.

To find the bid which maximizes the expected utility, we set $\frac{dEu_i}{db_i} = 0$. Using equation 2, to simplify this equation, and the fact that the bid b_i that maximizes the expected utility is $b_i = g(v_i)$, we get equation 10.

When the bidder's valuation is $v_i = r$, he bids $b_i = r$, no matter what c_i is, hence the boundary condition. ■

To illustrate the optimal bidding function, we again consider the special case when the valuations v_i and budget constraints c_i are drawn from uniform distribution $U[0, 1]$:




Figure 2: Equilibrium strategy $g(v)$ for the cases when bidders are risk-neutral ($\alpha = 1$), and risk-averse ($\alpha = 0.5$). The valuations v and the budget constraints c are both drawn from the uniform distribution $U[0, 1]$. The auction has a reserve price, which takes values $r = 0, 0.25, 0.5$. The number of bidders, N , and the number of items being sold, m , have values $N - m = 1$. The dotted lines represent the valuations for which no bids are placed due to the reserve price.

Corollary 2 *In the case that $F(v)$ and $H(c)$ are uniform distributions $U[0, 1]$, the equilibrium strategy is: $b_i = \min\{g(v_i), c_i\}$ where $g(v)$ is the solution of d.e.:*

$$g'(v_i) = \frac{1 - g(v_i)}{\frac{u'(v_i - g(v_i))(v_i + g(v_i) - v_i g(v_i))}{(N-m)(u(v_i - g(v_i)) - u(0))} - (1 - v_i)} \quad (16)$$

with boundary condition $g(r) = r$.

We choose to use the CRRA utility function $u(x) = x^\alpha$, $\alpha \in (0, 1)$. By substituting it into equation 16, we get:

$$g'(v_i) = \frac{1 - g(v_i)}{\frac{\alpha(v_i + g(v_i) - v_i g(v_i))}{(N-m)(v_i - g(v_i))} - (1 - v_i)}$$

Comparing this equation with equation 7, we observe that they are almost identical. Indeed, *the fact that the agent is now risk-averse is strategically equivalent to having more opponents*. In particular, a risk-averse agent using the CRRA utility function with parameter α , who enters an auction with N participating bidders in total, will bid in exactly the same way as a risk-neutral agent, who enters an auction with $N' = \frac{N + (1-\alpha)m}{\alpha}$ participating bidders.

Now that we are aware of the effect that varying risk attitudes produce in relation to the number of participating bidders, we examine the effect that reserve prices have in this setting. In figure 2, we vary the reserve price r between values 0, 0.25 and 0.5, as well as the parameter α of the utility function $u(x) = x^\alpha$, which takes values $\alpha = 0.5$ (risk-averse bidder) and $\alpha = 1$ (risk-neutral bidder). We fix the number of participating bidders, N , and items sold, m , so that $N - m = 1$. We observe that for relatively small values

of the reserve price (when $r = 0.25$), the effect that it has on increasing the bids of the agents, is smaller than the effect of the varying risk attitude. However, as r increases, and it becomes equal to $r = 0.5$, this effect is strengthened; in fact, for all possible valuations, the risk-neutral bidder will bid more when the reserve price is $r = 0.5$, than the risk-averse bidder when the reserve price is $r = 0.25$. This can potentially generate much more revenue for the seller, but there is a higher risk, now, of items not being sold.

In the case of an $(m + 1)^{th}$ price auction, we can prove the following theorem:

Theorem 7 *In an $(m + 1)^{th}$ price auction, with reserve price r , where the bidders have valuations v_i and budget constraints c_i , and they have a risk attitude described by utility function $u(\cdot)$, it is a dominant strategy to bid: $b_i = \min\{v_i, c_i\}$, if $b_i \geq r$, and not to participate otherwise.*

We now extend this theorem to the case when the bidders do not precisely know their own valuations. The following theorem is a generalization of theorems 5 and 7:

Theorem 8 *In an $(m + 1)^{th}$ price auction, with reserve price r , if a bidder has budget constraint c_i , he knows only imprecisely his own valuation v_i , in that it is drawn from distribution $G_i(v_i)$, and his risk attitude is described by utility function $u_i(\cdot)$, it is a dominant strategy to bid: $b_i = \min\{\beta_i, c_i\}$, if $b_i \geq r$, and not to participate otherwise. The variable β_i is the solution of equation:*

$$\int_{-\infty}^{\infty} u(z - \beta_i) G'_i(z) dz = u(0) \quad (17)$$

PROOF. Using the exact same reasoning as in theorem 5, the expected utility of an agent bidding β_i is:

$$Eu_i = \int_{-\infty}^{\beta_i} \left(\int_{-\infty}^{\infty} u(z - x) G'_i(z) dz \right) \Omega'_i(x) dx + (1 - \Omega_i(\beta_i)) u(0)$$

Note that this formula does not change as a result of the reserve price, nor the budget constraint, because these two parameters limit the bid b_i that the agent is allowed to place. The unconstrained bid β_i which maximizes the expected revenue is found by setting: $\frac{dEu_i}{d\beta_i} = 0$, which leads to equation 17. As $\frac{dEu_i}{dx} > 0, \forall x < \beta_i$, the agent should try to place a bid b_i as close as possible to β_i from below, as the budget constraint c_i will allow, so $b_i = \min\{\beta_i, c_i\}$. It also should never bid above β_i , because then $Eu_i \leq u(0)$; thus if $b_i < r$, it should not participate. ■

Experimental Evaluation

In this section, we validate experimentally the usefulness in practice of the theoretical analysis. The strategy computed for the case when valuation uncertainty exists in an $(m+1)^{th}$ price auction, is a (weakly) dominant strategy, and therefore we know for certain, without the need of simulations, that this is always going to yield the highest utility. However, the strategy given by equation 10 is an equilibrium strategy and therefore there are no theoretical guarantees that this one will always yield the highest revenue, especially if the opponents don't bid according to the equilibrium strategy.

We simulated the case when $N = 3$ bidders participate in an m^{th} price auction, where $m = 2$ items are sold; this is a




Figure 3: Experimental comparison of the equilibrium strategy S against strategies NB (budget constraints are ignored) and RN (risk attitudes are ignored). In all experiments some agents used strategy S and all the rest use the same strategy (either NB or RN).

simple, yet representative case of a multi-unit auction. The bidders are all risk-averse (using the CRRA utility function with $\alpha = 0.5$) and they have budget constraints c_i and valuations v_i drawn from uniform distribution $U[0, 1]$. We denote this standard equilibrium strategy (given by equation 10) as S and compare it against the following two strategies: (i) NB is the strategy when the agent does not take the budget constraint into account, and (ii) RN is the strategy when the agent does not take the risk attitudes into account (and assumes that everyone is risk neutral).¹⁰ We compare S against each of these two strategies by running experiments in which some agents bid according to S and some according to NB (according to RN in the second comparison we did), for various values of the reserve price r . The results are presented in figure 3; they are presented as the ratio of the corresponding utility divided by the utility of the case when all agents use strategy S (experiment “3xS”).¹¹ From this figure we can observe that, in every single instance, an agent using strategy NB (or RN) would always obtain a higher utility by switching to strategy S . For example, in the case that all agents use RN , any one of them would get a higher expected utility by switching to strategy S , and this is true for all other possible cases, in which some agent uses a strategy other than S .

¹⁰The reason why these two strategies were selected, is because they look at less features than strategy S . In this sense, they are strategies which don’t take advantage of the full analysis presented in this paper, and yet are reasonable, because they do consider some of the desired features. It should be pointed that the pre-existing state-of-the-art equilibrium strategies are less advanced than even these strategies (NB and RN).

¹¹The notation “2xS + 1xNB” means that two agents using strategy S and one using strategy NB participate in the experiment, etc.




Figure 4: Expected revenue both for m^{th} and the $(m + 1)^{th}$ price auctions, in the presence of varying reserve prices and bidder risk attitudes. $m = 2$ items are auctioned to $N = 3$ participating bidders with valuations v_i and budget constraints c_i drawn from uniform distribution $U[0, 1]$. The utility function used is $u(x) = x^\alpha$.

This means that strategies NB and RN are *dominated* by S , when we consider agents who play either NB/RN or S , and therefore *using our novel analysis does lead to higher bidder utility compared to cases where some feature is not taken into account*.

Given the fact that bidders will use this new equilibrium strategy (as shown by the first experiment), it makes sense for the seller to also use the theoretical results of this paper, in order to maximize her revenue, by selecting the best reserve price and the correct auction type (m^{th} or $(m + 1)^{th}$).¹² The expected revenue of the seller ER_k in a k^{th} ($k = m, m + 1$) price auction, when the bidders bid according to function $g(v)$, is: $ER_k = m \int_r^\infty \omega \cdot d\Psi_k(\omega)$, where $\Psi_k(x) = \sum_{i=0}^{k-1} C(N, i) (Z(x))^{N-i} (1 - Z(x))^i$ and $Z(x) = 1 - (1 - F(g^{-1}(x)))(1 - H(x))$. We assume the same values as in the previous experiment, i.e. $N = 3, m = 2$ and the bidders are risk-averse with $\alpha = 0.5$. In figure 4, we graph the seller’s revenue, when she sets a reserve price $r \in [0, 1]$. The correct reserve price is $r = 0$, for the case when $a = 0.5$ (in this case the bidders actually use strategy S). However, if the seller assumes (erroneously) that the bidders are risk-neutral ($\alpha = 1$), and thus they use strategy RN , she would select $r = 0.13$ which would lead to a 3.95% loss of revenue; if she assumes (again erroneously) that the bidders have ∞ budget, and thus they use strategy NB , she would select $r = 0.02$ which would lead to a 0.02% loss

¹²Note that the auctions no longer assign the items to the bidders with the top m valuations, due to the budget constraints, and furthermore, the winner can be different between the m^{th} and the corresponding $(m + 1)^{th}$ price auction. Thus, the revenue equivalence theorem does not apply here.

of revenue.¹³ Once more we see that *our analysis is necessary to determine the correct reserve price that maximizes the sellers revenue*. In figure 4, we also graph two additional cases: when the bidders are risk-averse to the extreme ($\alpha \rightarrow 0$), and when an $(m + 1)^{th}$ price auction is used (the revenue doesn't depend on α in this case). As was expected, when $\alpha \rightarrow 0$, the expected revenue is maximized. Here for risk-neutral and risk-averse bidders, the $(m + 1)^{th}$ price auction yields a lower revenue than the m^{th} price one. This is not the case though for risk-seeking bidders. Our analysis allows to compute the optimal reserve price for both auctions and then we can select the correct auction type.

Conclusions

We examined the behavior of agents participating in multi-unit sealed-bid auctions, when budget constraints, reserve prices, varying risk attitudes and valuation uncertainty exist. We provided a number of novel equilibria. First, we derived equilibria for each case separately. Second, we derived the dominant strategy for the case of uncertainty in the valuation that bidders have, when the bidders are not risk-neutral, in the setting of the $(m + 1)^{th}$ price auction. We also showed the the variance in this uncertainty and the risk attitude of a bidder determine the deviation of the bid from the expected value of his valuation. Third, we combined all the features in our analysis. We derived the equilibrium strategies for both the m^{th} and the $(m + 1)^{th}$ price auction, in the presence of budget constraints, reserve prices and any possible bidder risk attitude. Then we also included the uncertainty of bidders' valuation for the case of the $(m + 1)^{th}$ price auction and derived the dominant strategy. Fourth, we used simulations to show that this analysis is useful, in practice, both for the bidding agents in order to maximize their utility, and also for the seller in order to select the correct reserve price and thus maximize her revenue.

There are a number of unresolved issues in this paper. We plan to enrich our model in order to analyze the m^{th} price auction equilibria, in the presence of valuation uncertainty. In addition, we would like to examine the case of identical items being sold in *multiple concurrent auctions* (Gerding et al. 2007); in this case it is necessary to place bids in all the auctions. Finally, there are settings in which competition between agents negates the traditional assumption that agents are self-interested, i.e. maximizing their profit (Vetsikas and Jennings 2007); this leads to more aggressive bidding.

Acknowledgments

This research was undertaken as part of the ALADDIN (Autonomous Learning Agents for Decentralised Data and Information Systems) project and is jointly funded by a BAE Systems and EPSRC (Engineering and Physical Research Council) strategic partnership (EP/C548051/1).

¹³This difference is small in this case, but if a different α had been selected, this error could be more significant. E.g. for $\alpha = \frac{2}{3}$, the seller would select $r = 0.07$, for a 0.26% loss of revenue, if the budget is ignored, and she would select $r = 0.13$, for a 2.19% loss of revenue, if the risk attitudes are ignored.

References

- Borgs, C.; Chayes, J. T.; Immorlica, N.; Mahdian, M.; and Sabar, A. 2005. Multi-unit auctions with budget-constrained bidders. In *Proceedings of the 6th ACM Conference on Electronic Commerce (ACM EC'05)*.
- Che, Y., and Gale, I. 1998. Standard auctions with financially constrained bidders. *Review of Economic Studies* 65(1):1–21.
- Gerding, E. H.; Dash, R. K.; Yuen, D. C. K.; and Jennings, N. R. 2007. Bidding optimally in concurrent second-price auctions of perfectly substitutable goods. In *Proceedings of the 6th International Joint Conference on Autonomous Agents and Multi-Agent Systems*, 267–274.
- Krishna, V. 2002. *Auction theory*. Academic Press.
- Larson, K., and Sandholm, T. 2001. Costly valuation computation in auctions. In *8th Conference of Theoretical Aspects of Knowledge and Rationality (TARK VIII)*.
- Liu, Y.; Goodwin, R.; and Koenig, S. 2003. Risk-averse auction agents. In *Proceedings of the 2nd International Joint Conference on Autonomous Agents and Multi-Agent Systems*, 353–360.
- Maskin, E., and Riley, J. 1984. Optimal auctions with risk averse buyers. *Econometrica* 52(6):1473–1518.
- Myerson, R. B. 1981. Optimal auction design. *Mathematics of Operations Research* 6:58–73.
- Parkes, D. C. 1999. Optimal auction design for agents with hard valuation problems. In *Proc. IJCAI-99 Workshop on Agent Mediated Electronic Commerce*, 206–219.
- Rice, J. A. 1995. *Mathematical Statistics and Data Analysis*. California: Duxbury Press.
- Riley, J. G., and Samuelson, W. F. 1981. Optimal auctions. *The American Economic Review* 71(3):381–392.
- Thompson, D., and Leyton-Brown, K. 2007. Valuation uncertainty and imperfect introspection in second-price auctions. In *Proceedings of the 22nd Conference on Artificial Intelligence (AAAI)*.
- Vetsikas, I. A., and Jennings, N. R. 2007. Outperforming the competition in multi-unit sealed bid auctions. In *Proceedings of the 6th International Joint Conference on Autonomous Agents and Multi-Agent Systems*, 702–709.
- Vetsikas, I. A., and Jennings, N. R. 2008. Towards agents participating in realistic multi-unit sealed-bid auctions. In *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multi-Agent Systems (to appear)*.
- Weber, R. 1997. Making more from less: Strategic demand reduction in the fcc spectrum auctions. *Journal of Economics and Management Strategy* 6(3):529–548.