JAAMAS manuscript No.
(will be inserted by the editor)

Bidding Strategies for Realistic Multi-Unit Sealed-Bid Auctions

Ioannis A. Vetsikas · Nicholas R. Jennings

Received: date / Accepted: date

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. However, 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 and auction format.

Keywords Game Theory · Bidding Strategies · Agent-Mediated e-Commerce · Budget Constraints · Risk Attitudes · Reserve Price · Valuation Uncertainty

1 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,

Ioannis A. Vetsikas
ECS School, University of Southampton, Southampton SO17 1BJ, UK
E-mail: iv@ecs.soton.ac.uk

Nicholas R. Jennings
ECS School, University of Southampton, Southampton SO17 1BJ, UK
E-mail: nrj@ecs.soton.ac.uk

we believe 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. Now, 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 immediately benefit. In order to be able to design autonomous agents that would be able to represent non-expert humans in real auctions, it is therefore necessary to analyze the strategic behavior of agents in multi-unit (m^{th} and $(m+1)^{th}$ price) auction models that incorporate all the relevant features. To identify these relevant features, 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 in the literature.

First, *budget constraints* are very important, whenever businesses and individuals place bids. This is because, in every practical setting, these bidders have a certain pot of money that they can spend in order to purchase the item. Thus, even though they might wish to spend more in order to acquire the desired good or service, they are limited by their available budget. Thus, these budget constraints 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 [4,5], but not for multi-unit ones;[1] it has also been proven that the revenue generated by a 1^{st} price auction is always higher than that of the equivalent 2^{nd} 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, and thus they bid in such a way as to maximize their expected 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. In related work, e.g. in [10], we find the computation of the equilibrium strategy for a risk-averse agent participating in a 1^{st} price auction. Byde[3] also uses agent-based simulations to compute the strategies for a range of sealed-bid auction scenarios in which the participating bidders have a variety of risk attitudes.

Third, setting an appropriate *reserve price* (i.e. a minimum transaction price) in the auction is a way that allows the seller to increase her expected revenue in a number of auction scenarios. By setting this minimum price the seller guarantees that the item will not sell for a very low price, which could happen in cases where all the bidders have low valuations and also it forces the bidders to increase their bids compared to the case where there is no reserve price. Auctions taking place in traditional auction houses typically have a reserve price and this is also the case for many eBay auctions. This case has been examined

[1] The “dual” problem of designing a truthful mechanism for budget constrained bidders has been examined in [1].

for single-unit sealed-bid auctions in [13,16]. In [11], this problem is examined for risk-averse bidders.

Fourth, there may be *uncertainty in the bidders' valuation* of the offered commodity. This is an issue that has received comparatively little attention in the literature and yet it is very important in real scenarios. In particular, in many settings, the value of a commodity is not known precisely a priori. This could be due to a number of factors. 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.[2] Another example is in an electricity market, where a customer or a supplier does not always know precisely the price of energy in the next day as this depends both on the customer demand as well as more unpredictable factors such as the energy production from renewable sources. A final example, in a traditional auction, is that of resale; the buyer may want to resell the item later in order to gain additional profit, however there is uncertainty as to the actual price that it would fetch. 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 [8,17].[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. To this effect, this paper introduces models of valuation uncertainty and presents equilibria that account for this uncertainty, when there is no possibility for introspection.

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) [21]. 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 [18].

Given this background, in this paper we make the following contributions:[4]

- After presenting the model of the problem in section 2, we derive novel equilibria for the multi-unit m^{th} price sealed-bid auction case, when these features are looked at separately. We examine reserve prices in section 3, varying risk attitudes in section 4 and budget constraints in section 5.
- We then combine, for the first time, two of the features that we look at, in section 6. Specifically, we derive the equilibrium strategy for the case of uncertainty in the valuation that bidders have, when the bidders are not risk-neutral, in the settings of the

[2] While some aspect of interdependent valuations is present in some cases of ad auctions, this is not always the case; companies competing for the same ad space sometimes sell different products and/or target different market segments. Thus, the uncertainty in the valuation would often not be removed, even if all the participants private values were known.

[3] There are few notable exceptions where no introspection is assumed, such as [14], which examines how to select the auction that yields the highest revenue and/or efficiency, when bidders have uncertainty in their valuations.

[4] Some of this work has been presented in the conference paper [19]. That paper presented the main result of what happens when all four issues are present together only for the easier case of the $(m+1)^{th}$ price auction. Here, we extend the work and give the equilibrium strategy when all issues are present also for the more complicated case of the m^{th} price auction setting. We also discuss several issues that arise from our analysis, e.g. concerning the various models of valuation uncertainty, and we give unabridged versions of the proofs. Thus in this paper, we give a complete and unified picture of our work.

m^{th} price auction. For the $(m+1)^{th}$ price setting we derive a dominant strategy when examining the same two issues.
– After doing this, we combine all the issues in section 7. We derive the equilibrium strategies both for the m^{th} and $(m+1)^{th}$ price auction settings, in the presence of all four issues that we examine: budget constraints, reserve prices, any bidder risk attitude and uncertainty of bidders' valuation. In the case of the $(m+1)^{th}$ price auction the solution is even stronger in that it is a 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 and auction format.[5] We also examine theoretically and empirically how the valuation uncertainty impacts the bidding strategy of non risk-neutral bidders and discuss various models of valuation uncertainty.

2 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 the end of this section, we give the equilibrium strategies for single-unit cases, when only one issue is present (i.e. reserve price, risk-attitudes and budget constraints).

In particular, we will compute and analyze the symmetric Bayes-Nash equilibria[6] 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. drawn 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, then we need to extend this model. For the $(m+1)^{th}$ price multi-unit auction case, we can use the most general model possible: the agent knows that his valuation v_i is drawn from distribution $G_i()$, 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. For the m^{th} price multi-unit auction case, we use a simpler model, because unlike the previous case, no dominant strategies exist in these cases, therefore the strategy used by opponent bidders affects the strategy used by any specific bidder. Therefore we assume that the true valuation $\overline{v}_i$, which is not known to bidder i, is drawn from distribution $G_{v_i}()$, where v_i is known as being the mean value of distribution $G_{v_i}()$; so each bidder i knows approximately

[5] The use of equilibrium analysis in order to design autonomous, intelligent agents has been done before, e.g. in [10] and [18].

[6] The 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.

his own value as being drawn from distribution $G_{v_i}()$ around the value v_i, which is known to bidder i. We also consider a more general model of valuation uncertainty, in which there can be many more types of uncertainty, so that the distribution from which the true valuation $\overline{v}_i$ is drawn depends on more parameters than just v_i. We discuss this avenue of work in section 7.3.

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 $u()$ 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 $u(x)$ used widely in economics are: $u(x) = x^{\alpha}, \alpha \in (0,1)$ (constant relative risk aversion - CRRA), $u(x) = 1 - exp(-\alpha x), \alpha > 0$ (constant absolute risk aversion - CARA), and $u(x) = -\gamma^x, \gamma \in [0,1]$, all of which indicate risk-averse bidders. We will also use the function used in [10], which is $u(x) = -\gamma^x, \gamma \in [0,1]$; this function is used to indicate risk-averse agents, but here we extend it to also indicate risk-seeking bidders:

$$u(x) = sign(\gamma - 1) \cdot \gamma^x, \forall \gamma \geq 0 \tag{1}$$

where $sign(z)$ is the sign function, which returns $+1$, when $z > 0$, -1 when $z < 0$ and 0 when $z = 0$. These functions indicate risk-neutral bidders for the specific values $\alpha = 1$ for CRRA and as $\gamma \to 1$ for equation 1.

In most scenarios analyzed in auction theory the *default settings* are, that bidders are risk neutral, and thus the utility function is $u(x) = x$, that there is no reserve price, thus $r = 0$, that there are no budget constraints, thus $c_i = \infty$ and that the valuation for the item is known precisely, therefore $G_{v_i}(x) = 0$, when $x < v_i$ and $G_{v_i}(x) = 1$, when $x > v_i$. We will also use these values unless explicitly stated to the contrary.

We also use the following additional notation in the proofs: $Z(x)$ is the probability distribution (cdf) of any opponent's bid b_j. Thus $Z(x) = Prob[b_j \leq x]$, 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 [15]:

$$\Phi_k(x) = \sum_{i=0}^{k-1} C(N-1,i)(Z(x))^{N-1-i}(1-Z(x))^i \tag{2}$$

where the notation $C(n,k)$ is the total number of possible combinations of k items chosen from n. As shown in [18], $\forall N \geq m$ the following holds:

$$\Phi'_m(x) = (N-m)\big(\Phi_m(x) - \Phi_{m-1}(x)\big)\frac{Z'(x)}{Z(x)} \tag{3}$$

Now that we presented the model and the notation used, we give the Bayes-Nash equilibrium strategies for the case when each issue is examined separately for the case of a 1^{st} price auction. (see e.g. [7]) These results are being extended by this paper.

- *The presence of a Reserve Price* r:

$$g(v) = v - (F(v))^{-(N-1)} \cdot \int_r^v (F(z))^{N-1} \cdot dz$$

- *Varying Risk-Attitudes*: When the utility function is $u()$, the equilibrium is the solution of the differential equation:

$$g'(v_i) = \frac{u(v_i - g(v_i)) - u(0)}{u'(v_i - g(v_i))} \cdot (N-1) \cdot \frac{F'(v_i)}{F(v_i)}$$

– *The presence of Budget Constraints*: In this case the equilibrium strategy is to bid the minimum of the budget constraint c_i and the solution of the differential equation[7]

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

3 Equilibria in the Presence of Reserve Prices

In this section we examine the equilibria that exist in the case that the reserve price of the auction is $r \geq 0$. For the remaining issues, we use the default settings; thus, we assume the bidders have no budget constraints, they are risk-neutral and there is no uncertainty about their valuations.

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 the valuations v_i are i.i.d. variables drawn from $F(v)$, the following bidding strategy constitutes a symmetric Bayes-Nash equilibrium:*

$$g(v) = v - (F(v))^{-(N-m)} \cdot \int_r^v (F(z))^{N-m} \cdot dz \tag{4}$$

PROOF. Because the main elements of this proof are also used in the proof of theorem 9 we omit it here. ■

From equation 4, it is evident that, for the same valuation v_i, the bid b_i increases with each increase of the reserve price.

In the case of an $(m + 1)^{th}$ price auction, the optimal strategy is [7]:

Theorem 2 *In an $(m + 1)^{th}$ price auction, with reserve price r, where the bidders are risk-neutral, have valuations v_i and no budget constraints, it is a dominant strategy to bid truthfully: $b_i = v_i$, if $v_i \geq r$, and not to participate otherwise.*

4 Equilibria in the Case of Varying Risk Attitudes

In this section we examine the equilibria that exist in the case that agents are not risk-neutral, but rather have a utility function $u()$ that maps profit into utility. If this function is concave, the agents are risk-averse; if it is convex, they are risk-seeking. For the remaining issues, we use the default settings.

Theorem 3 *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 $F(v)$, the bidders have no budget constraints and they have a risk attitude which is described by utility function*

[7] Actually the case examined in [4] is more complex in the sense that the budget constraints need not be hard, meaning that the utility loss for overspending is not ∞, and is characterized by a penalty function.

$u()$, the bidding strategy $g(v)$, which constitutes a symmetric Bayes-Nash equilibrium, is the solution of the differential equation:

$$g'(v_i) = \frac{u(v_i - g(v_i)) - u(0)}{u'(v_i - g(v_i))} \cdot (N - m) \cdot \frac{F'(v_i)}{F(v_i)} \tag{5}$$

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

PROOF. Let us assume that the equilibrium strategy is described by a function $g()$ which maps valuations to bids. We consider any bidder i, who places a bid b_i in the auction. The distribution $Z(x)$ of the bid b_j, that any opponent j ($j \neq i$) of agent i places, is:

$$Z(x) = F(g^{-1}(x)) \tag{6}$$

The k^{th} order statistic of these bids $B^{(k)}$ is drawn from distribution $\Phi_k(x)$, described by equation 2.

Depending on the value of b_i, the following three cases are possible:

- If $b_i < B^{(m)}$, then bidder i is outbid and doesn't win any items, therefore his utility is $u_i = u(0)$.[8] The probability of this case happening is: $Prob[b_i \leq B^{(m)}] = 1 - \Phi_m(b_i)$.
- If $B^{(m)} \leq b_i \leq B^{(m-1)}$, then bidder i has placed the last winning bid. Thus the payment equals his bid, his profit is $(v_i - b_i)$, and his utility is $u_i = u(v_i - b_i)$. The probability of this case happening is: $Prob[B^{(m)} \leq b_i \leq B^{(m-1)}] = \Phi_m(b_i) - \Phi_{m-1}(b_i)$.
- If $B^{(m-1)} < b_i$, then bidder i is a winner, the payment is equal to bid $B^{(m-1)}$, his profit is equal to $(v_i - B^{(m-1)})$ and his utility is $u_i = u(v_i - B^{(m-1)})$. Note that the probability: $Prob[B^{(m-1)} \leq \omega] = \Phi_{m-1}(\omega)$.

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

$$\begin{aligned} Eu_i(b_i) =& u(0) \cdot (1 - \Phi_m(b_i)) + u(v_i - b_i) \cdot (\Phi_m(b_i) - \Phi_{m-1}(b_i)) \\ &+ \int_0^{b_i} u(v_i - \omega) \cdot \frac{d}{d\omega} \Phi_{m-1}(\omega) \cdot d\omega \end{aligned} \tag{7}$$

The bid which maximizes this expected utility, is found by setting: $\frac{dEu_i}{db_i} = 0$. This becomes:

$$(u(v_i - b_i) - u(0)) \cdot \Phi'_m(b_i) = u'(v_i - b_i) \cdot (\Phi_m(b_i) - \Phi_{m-1}(b_i))$$

Using equation 3, to simplify this equation, we derive:

$$(u(v_i - b_i) - u(0)) \cdot (N - m) \cdot \frac{F'(g^{-1}(b_i))}{g'(g^{-1}(b_i)) \cdot F(g^{-1}(b_i))} = u'(v_i - b_i)$$

This value b_i is equal to $b_i = g(v_i)$, since it maximizes the expected utility $Eu_i(b_i)$. Using this substitution, we derive differential equation 5.

The boundary condition is $g(0) = 0$, because an agent with valuation $v_i = 0$ will bid $b_i = 0$. ■

If we use the function $u(x)$ from equation 1, we can solve equation 5, to get the following equilibrium strategy:

$$g(v_i) = v_i - \log_\gamma \left[1 + \frac{\ln \gamma}{F(v_i)^{N-m}} \cdot \int_0^{v_i} F(\omega)^{N-m} \cdot \gamma^{v_i - \omega} \cdot d\omega\right] \tag{8}$$

[8] Note that profit 0 does not necessarily mean that the utility is 0; it depends on the form of the utility function $u()$.

In [10], the authors take the limits of $g(v_i)$ as γ approaches 1 and 0, which represent the cases when the agent becomes risk-neutral and very risk-averse respectively. Here, we do the same, and also compute the limit as γ approaches ∞, which represents the case when the agent becomes very risk-seeking:

$$\lim_{\gamma \to 0} g(v_i) = v_i \tag{9}$$

$$\lim_{\gamma \to 1} g(v_i) = v_i - \frac{\int_0^{v_i} F(\omega)^{N-m} \cdot d\omega}{F(v_i)^{N-m}} \tag{10}$$

$$\lim_{\gamma \to \infty} g(v_i) = 0 \tag{11}$$

We observe that, when $\gamma \to 1$, i.e. the agents tend to become risk-neutral, equation 8 gives the same solution as the one known from the literature, for the case when agents just maximize their profit (risk-neutral agents) [7]. When $\gamma \to 0$, i.e. the agents become very risk-averse, they bid truthfully, because they are worried too much about losing no matter how small this possibility is.[9] When $\gamma \to \infty$, i.e. the agents become very risk-seeking, they bid 0 (or $\epsilon > 0$, if zero bids are not allowed), gambling on the unlikely chance that there is no competition and they receive the item for free. This reasoning as to why very risk-averse and very risk-seeking agents would bid in this way holds for any family of utility functions, and it is not important to give a formal proof for this fact. As for the case when an agent becomes risk-neutral, we can see that in this case equation 5, becomes: $g'(v_i) = (v_i - g(v_i))(N-m)\frac{F'(v_i)}{F(v_i)}$, whose solution is equation 10. Thus, these results can be generalized to any family of utility functions; the agents bid according to equations 9, 10 and 11, when they are respectively very risk-averse, risk-neutral and very risk-seeking.

In the case of an $(m+1)^{th}$ price auction, the agents submit truthful bids [7]:

Theorem 4 *In an $(m+1)^{th}$ price auction, where the bidders have valuations v_i, they have no budget constraints and they have a risk attitude described by utility function $u()$, it is a dominant strategy to bid truthfully: $b_i = v_i$*

5 Equilibria in the Case of Budget Constraints

In this section we examine the equilibria for the case when agents have budget constraints only. For the remaining issues, we use the default settings.

Theorem 5 *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:*

$$b_i = min\{g(v_i), c_i\} \tag{12}$$

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

$$g'(v_i) = \frac{(1 - H(g(v_i)))F'(v_i)}{\frac{1-(1-H(g(v_i)))(1-F(v_i))}{(N-m)(v_i - g(v_i))} - (1 - F(v_i))H'(g(v_i))} \tag{13}$$

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

9 Both these results are consistent with those reported in [10] for the case of single-unit auctions.

PROOF. Consider any bidder i. We assume that each opponent j ($j \neq i$) bids $b_j = min\{g(v_j), c_j\}$. As the variables v_j and c_j are independent of each other and are drawn from distributions $F(v)$ and $H(c)$ respectively, the distribution $Z(x)$ of the opponent's bid b_j can be computed as follows:

$$\begin{aligned} Z(x) &= 1 - Prob[min\{g(v_j), c_j\} > x] \\ &= 1 - Prob[g(v_j) > x \wedge c_j > x] \\ &= 1 - (1 - F(g^{-1}(x))) \cdot (1 - H(x)) \end{aligned} \tag{14}$$

The k^{th} order statistic of these bids $B^{(k)}$ is drawn from distribution $\Phi_k(x)$, described by equation 2.

We can now analyze the expected profit of bidder i. Let b_i be the bid that he places in the auction. We are going to assume at first that his bid b_i is not limited by the budget constraint c_i and impose this restriction later on. We distinguish the following cases:

- If $b_i < B^{(m)}$, then bidder i is outbid and doesn't win any items, therefore his utility is $u_i = 0$.
- If $B^{(m)} \leq b_i \leq B^{(m-1)}$, then bidder i has placed the last winning bid. Thus the payment equals his bid and his utility is $u_i = v_i - b_i$. The probability of this case happening is: $Prob[B^{(m)} \leq b_i \leq B^{(m-1)}] = \Phi_m(b_i) - \Phi_{m-1}(b_i)$.
- If $B^{(m-1)} < b_i$, then bidder i is a winner, the payment is equal to bid $B^{(m-1)}$ and his utility is $u_i = v_i - B^{(m-1)}$. Note that: $Prob[B^{(m-1)} \leq \omega] = \Phi_{m-1}(\omega)$.

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

$$Eu_i(b_i) = (v_i - b_i) \cdot (\Phi_m(b_i) - \Phi_{m-1}(b_i)) + \int_0^{b_i} (v_i - \omega) \cdot \frac{d}{d\omega}\Phi_{m-1}(\omega) \cdot d\omega \tag{15}$$

The bid b_i which maximizes this utility is found by setting:

$$\frac{dEu_i(b_i)}{db_i} = 0 \Leftrightarrow -(\Phi_m(b_i) - \Phi_{m-1}(b_i)) + (v_i - b_i) \cdot \Phi'_m(b_i) = 0$$

Using equation 3, we get:

$$\Phi_m(b_i) - \Phi_{m-1}(b_i) = (v_i - b_i) \cdot (N - m) \cdot (\Phi_m(b_i) - \Phi_{m-1}(b_i)) \cdot \frac{Z'(x)}{Z(x)}$$

Since this bid maximizes the agent's utility, it must be equal to $b_i = g(v_i)$. Using this fact and equation 14 we get differential equation 13.

The boundary condition is $g(0) = 0$, no matter the value of c_i, because when the bidder's valuation $v_i = 0$, he must bid $b_i = 0$.

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 needs to bid $b_i = g(v_i)$, as this maximizes the expected profit $Eu_i(b_i)$. In the case that $g(v_i) > c_i$, we need to select a bid $b_i \leq c_i$. Since the only point at which $\frac{dEu_i}{db_i} = 0$ is when $b_i = g(v_i)$ and $\frac{dEu_i}{db_i} > 0$ for small values of b_i, while $\frac{dEu_i}{db_i} < 0$ for large values of b_i, this means that for values $b_i < g(v_i)$, it is $\frac{dEu_i}{db_i} > 0$. Therefore the value of $b_i \leq c_i$ that maximizes Eu_i is $b_i = c_i$. This means that the bid is always $b_i = min\{g(v_i), c_i\}$, which was our initial assumption (about how the agents bid). ■

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:

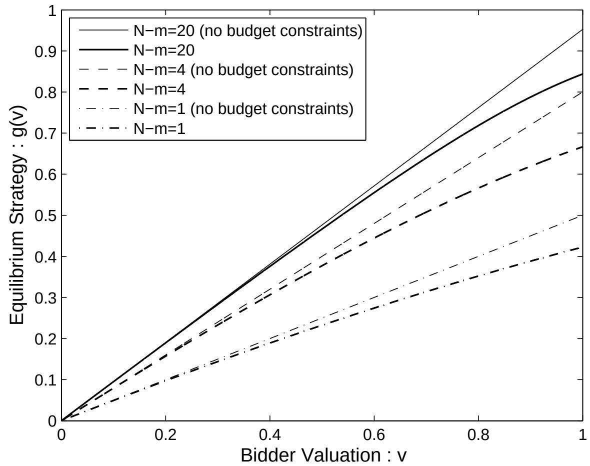

Fig. 1 Equilibrium strategy $g(v)$ for the case when the valuations v and the budget constraints c are both drawn from the uniform distribution $U[0,1]$; the equilibrium bid is $b = min\{g(v), c\}$. The number of bidders, N, and the number of items being sold, m, take values $N - m = 1, 4, 20$. The equilibrium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

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.:*

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

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 [7]).

It is also interesting to note that function $g()$ deviates from $g_\infty()$ most in the case when $N - m = 4$. This is due to two 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.[10]

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 [7]:

Theorem 6 *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: $b_i = min\{v_i, c_i\}$.*

[10] We 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.

6 Equilibria in the Case of Uncertainty in Agents' Valuations

In this section we examine the equilibria for the case when the bidders do not precisely know their own valuations. In the case that bidders are risk-neutral, it is an equilibrium strategy to bid as if they had value $v_i = \mu_G$ and the valuation uncertainty does not matter.[11] This is the reason why we choose to examine this issue when the agents are not necessarily risk-neutral. For the remaining issues, we use the default settings.

6.1 The $(m+1)^{th}$-price Auction Case

In this section, we examine the case when the bidders do not precisely know their own valuations in an $(m+1)^{th}$-price auction. 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 this section.

Theorem 7 *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:*

$$\int_{-\infty}^{\infty} u(z - b_i) G_i'(z) dz = u(0) \tag{17}$$

PROOF. From the point of view of each bidder i, he knows that his 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 he doesn't win, which happens with probability $(1 - \Omega_i(b_i))$. Therefore the expected utility of bidder i, when he 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:

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

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 17. The solution of this equation is the bid b_i which maximizes the expected utility of bidder i. ■

From equation 17, 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 made these observations empirically for various functions of valuation uncertainty and

[11] The fact that risk-neutral agents participating in a second price auction bid the mean value μ_G that they have for their valuation has been observed in some of the related work, e.g. in [8,17,14].

risk attitude functions. However, we can prove theoretically these statements for the special case when the distribution $G_i()$ is symmetric. More specifically, we prove that these statements always hold, in the following two 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).

PROOF. Equation 17 becomes $b_i - \mu_{G_i} = 0$, for linear function $u_i(x)$, which is the case of risk-neutral bidders.

Let us now examine what happens for risk-averse bidders:
$\int_{-\infty}^{\infty} u(z-\mu_{G_i})G_i'(z)dz = \int_{-\infty}^{\infty} u(x)G_i'(x+\mu_{G_i})dx$
$= \int_{-\infty}^{0} u(x)G_i'(x+\mu_{G_i})dx + \int_{0}^{\infty} u(x)G_i'(x+\mu_{G_i})dx$
Since $G()$ is symmetric, it is $G_i'(\mu_{G_i}-x) = G_i'(\mu_{G_i}+x), \forall x$ and
it is: $\int_{-\infty}^{0} u(x)G_i'(x+\mu_{G_i})dx = \int_{0}^{\infty} u(-x)G_i'(-x+\mu_{G_i})dx$, thus:
$\int_{-\infty}^{\infty} u(z-\mu_{G_i})G_i'(z)dz = \int_{0}^{\infty} u(-x)G_i'(-x+\mu_{G_i})dx + \int_{0}^{\infty} u(x)G_i'(x+\mu_{G_i})dx$
$= \int_{0}^{\infty}(u(-x)+u(x))G_i'(x+\mu_{G_i})dx$
Because $u()$ is concave, as the bidders are risk-averse, it is $u(-x)+u(x) < 2u(0), \forall x \neq 0$. Thus:
$\int_{-\infty}^{\infty} u(z-\mu_{G_i})G_i'(z)dz < \int_{0}^{\infty} 2u(0)G_i'(x+\mu_{G_i})dx = 2u(0)\cdot\frac{1}{2} \Rightarrow$

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

This means that if we set $b_i = \mu_{G_i}$, the left hand side of equation 17 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: $b_i < \mu_{G_i}$.

When the bidders are risk-seeking, $u()$ is convex, so it is $u(-x)+u(x) > 2u(0), \forall x \neq 0$ and therefore we prove in the same way that:

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

This implies that it should be: $b_i > \mu_{G_i}$. ■

One other issue that should be noted from this proof is the fact that the difference between $u(-x)+u(x)$ and $2u(0)$ increases with the concavity (resp. convexity) of $u()$, which occurs as the bidders become more risk averse (resp. risk seeking). This however implies that we need to decrease (resp. increase) the bid b_i even further in order to satisfy equation 17. This implies that *the more risk-averse bidders are, the lower they will bid*, while *the more risk-seeking bidders are, the higher they will bid*, in settings with exactly the same valuation uncertainty.

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 $u()$: if $u()$ is concave (i.e. risk-averse bidders), then the bidders' bids are: $b_i > b_j$, and if $u()$ is convex (i.e. risk-seeking bidders), then the bidders' bids are: $b_i < b_j$.*

PROOF. We give the proof for concave function $u()$. Let
$F_i(x) = \int_{-\infty}^{\infty} u(z-x)G_i'(z)dz \Leftrightarrow$
$F_i(x) = \int_0^1 u(G_i^{-1}(p)-x)dp \Leftrightarrow$
$F_i(x) = \int_0^{\frac{1}{2}} u(G_i^{-1}(p)-x)dp + \int_{\frac{1}{2}}^1 u(G_i^{-1}(p)-x)dp \Leftrightarrow$
$F_i(x) = \int_0^{\frac{1}{2}} \Big(u(G_i^{-1}(p)-x) + u(G_i^{-1}(1-p)-x)\Big)dp$
and define $F_j(x)$ in the same way. It is then: $F_i(b_i) = u(0) = F_j(b_j)$. Since $G_i()$ and $G_j()$ 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 $u()$ is concave, we get:

$$u(G_i^{-1}(p)-b_i) - u(G_j^{-1}(p)-b_i) > u(G_j^{-1}(1-p)-b_i) - u(G_i^{-1}(1-p)-b_i) \Leftrightarrow$$
$$u(G_j^{-1}(p)-b_i) + u(G_j^{-1}(1-p)-b_i) < u(G_i^{-1}(p)-b_i) + u(G_i^{-1}(1-p)-b_i)$$

From this equation, given how we defined $F_i()$, we get that: $F_j(b_i) < F_i(b_i)$. And since, $F_i(b_i) = u(0) = F_j(b_j)$ this means that $F_j(b_i) < F_j(b_j)$ and therefore $b_i > b_j$.

For convex $u()$ the proof is similar; we show in the same way that $F_j(b_i) > F_j(b_j)$ and therefore $b_i < b_j$. ■

6.2 The m^{th}-price Auction Case

In this section, we examine the case when the bidders do not precisely know their own valuations in an m^{th}-price auction. Like in the previous setting, the agents are not risk-neutral, but rather have a utility function $u()$,[12] there are no budget constraints and the reserve price of the auction is $r = 0$.

The model that we assume is slightly different (somewhat more restricted) from that used in the $(m+1)^{th}$-price auction case.[13] More specifically, the bidders have uncertainty about their true valuation, which we will denote $\overline{v}_i$. Each agent knows that this valuation is drawn from distribution $G_{v_i}()$, where v_i is a known value (which we will call "the agent's uncertain valuation") around which the agent's uncertainty is distributed. v_i is drawn from distribution $F()$ and is known to each agent. The equilibrium strategy of each agent depends on v_i and gives the bid $b_i = g(v_i)$ of each agent.

It is appropriate to give an example of what the practical meaning of having a distribution $G_{v_i}()$ is in conjunction to variables v_i and $\overline{v}_i$. v_i is known to bidder i, because each bidder has an idea of what his true value is. For example, a bidder knows that the value of the item is approximately $10\$$, so $v_i = 10$, but also that his actual true valuation is, for example, somewhere within 20% of this value, so $G_{v_i}()$ is a uniform distribution $U[8, 12]$, and thus $\overline{v}_i$ is drawn from this distribution, but is not precisely known at the time of bidding. If another bidder has $v_i = 7$, then $G_{v_i}()$ is $U[5.6, 8.4]$ etc. It might also be the case that the uncertainty is the same for everyone, for example that the true value is within $[-1, 1]$ of

[12] Note that since the model used could allow for negative profit as well, $u()$ should be such (or extended) that it allows for negative profit as well.

[13] We do this because a bidder's best strategy depends on the strategies of his opponents, unlike in the $(m+1)^{th}$-price auction case. We discuss how to extend this model in section 7.3.

the mean value v_i, so $G_{v_i}()$ is $U[v_i - 1, v_i + 1]$, or, more generally, the distribution could be any function (not necessarily a uniform distribution). The model we use allows that the distribution from which the true valuation is drawn should depend only on one parameter, the agent's "uncertain valuation" v_i, which is known to the agent a priori, even though the true valuation is not.

Theorem 8 *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 $\overline{v}_i$ (where $\overline{v}_i$ is not precisely known and is drawn from a distribution $G_{v_i}()$ and the parameters v_i are i.i.d. drawn from a known distribution $F(v)$), the bidders have no budget constraints and they have a risk attitude which is described by utility function $u()$, the bidding strategy $g(v)$, which constitutes a symmetric Bayes-Nash equilibrium, is the solution of the differential equation:*

$$g'(v_i) = \frac{\int_{-\infty}^{\infty} u(x - g(v_i))G'_{v_i}(x)dx - u(0)}{\int_{-\infty}^{\infty} u'(x - g(v_i))G'_{v_i}(x)dx}(N - m)\frac{F'(v_i)}{F(v_i)} \tag{18}$$

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

PROOF. Using the same reasoning as in the case when the valuation is known precisely (see derivation of equation 7), we compute the expected utility of bidder i, who places bid b_i, when his true valuation is $\overline{v}_i$, to be:

$$\begin{aligned} Eu_i(b_i, \overline{v}_i) =& u(0)\cdot(1 - \Phi_m(b_i)) + u(\overline{v}_i - b_i)\cdot(\Phi_m(b_i) - \Phi_{m-1}(b_i)) \\ &+ \int_0^{b_i} u(\overline{v}_i - \omega)\cdot \frac{d}{d\omega}\Phi_{m-1}(\omega)\cdot d\omega \end{aligned} \tag{19}$$

The distribution of the opponent bid is given by equation 2, where $Z(x) = F(g^{-1}(x))$, because the bids of each bidder j depends on his "uncertain valuation" v_i, which is known to each bidder.

However, the true valuation $\overline{v}_i$ is not known. Bidder i knows that $\overline{v}_i$ is drawn from $G_{v_i}()$. Therefore it is necessary to use Bayes rule to compute the expected utility of bidder i when he places bid b_i given that $\overline{v}_i$ is unknown, but the "uncertain valuation" v_i is. Therefore, the expected utility is:

$Eu_i(b_i) = \int_{-\infty}^{\infty} Eu_i(b_i, x) \cdot G'_{v_i}(x) \cdot dx \Rightarrow$

$$\begin{aligned} Eu_i(b_i) =& u(0) \cdot (1 - \Phi_m(b_i)) + (\Phi_m(b_i) - \Phi_{m-1}(b_i)) \int_{-\infty}^{\infty} u(x - b_i)G'_{v_i}(x)dx \\ &+ \int_0^{b_i} \Big(\int_{-\infty}^{\infty} u(x - \omega)G'_{v_i}(x)dx\Big) \frac{d}{d\omega}\Phi_{m-1}(\omega)d\omega \end{aligned} \tag{20}$$

The bid which maximizes this expected utility, is found by setting: $\frac{\vartheta Eu_i}{\vartheta b_i} = 0$. This becomes:

$$\Big(\int_{-\infty}^{\infty} u(x - b_i)G'_{v_i}(x)dx - u(0)\Big)\Phi'_m(b_i) =$$

$$(\Phi_m(b_i) - \Phi_{m-1}(b_i)) \int_{-\infty}^{\infty} u'(x - b_i)G'_{v_i}(x)dx$$

Using equation 3, to simplify this equation, we derive:

$$\Big(\int_{-\infty}^{\infty} u(x - b_i)G'_{v_i}(x)dx - u(0)\Big) \frac{(N - m)F'(g^{-1}(b_i))}{g'(g^{-1}(b_i))F(g^{-1}(b_i))} = \int_{-\infty}^{\infty} u'(x - b_i)G'_{v_i}(x)dx$$

This value b_i is equal to $b_i = g(v_i)$, since it maximizes the expected utility $Eu_i(b_i)$. Using this substitution, we derive the differential equation 18.

The boundary condition is $g(0) = 0$, because an agent with valuation $v_i = 0$ will bid $b_i = 0$. This is actually also a result of the fact that v_i are assumed to have non negative values. For a more general boundary condition see the result of the general theorem 9. ■

Now that we have computed the equilibrium strategy for the m^{th} price auction, we examine how the bidding strategy is affected by the valuation uncertainty. In the $(m+1)^{th}$ price auction, we proved that that risk-averse (resp. risk-seeking) agents bid less (resp. more) than the mean μ_G of their valuation distribution $G()$ and that this effect increases as the variance of distribution $G()$ increases, i.e. the more uncertain they get about their valuation and also as the bidders become more risk-averse (resp. risk-seeking). However, in section 4, we showed that as bidders become more risk-averse (resp. risk-seeking), when there is no valuation uncertainty, they bid more (resp. less). Therefore we have two opposite effects. For this reason, in section 8.1, we present an experiment where we explore these effects. What we find out is that, when the valuation uncertainty becomes significant, then risk-averse bidders will bid much less, even to the point of bidding like risk seeking bidders would (when the latter don't have valuation uncertainty).

7 Equilibria in the Presence of All Issues

In the previous sections, we considered each issue independently of the others, in order to examine the behavior of the bidders in the presence of each one; this makes it easier to understand the behavior when all issue are considered. This is exactly what is done in this section. In particular, we examine the general setting where all the issues are taken into account. Thus, the agents have budget constraints c_i, their risk attitude is described by function $u()$ (not necessarily risk neutral) and the auction has a reserve price r. In addition to these issues, we also consider that the valuations v_i are not known precisely, and we use the models of valuation uncertainty that we utilized in the previous section.

7.1 The m^{th}-price Auction Case

Here, we derive the equilibria in the setting of an m^{th} price auction.

Theorem 9 *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 $\overline{v}_i$ (the exact value unknown to the bidder, and drawn from distribution $G_{v_i}()$), which is approximated by a known variable v_i, the agent's "uncertain valuation", 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 $u()$, the following bidding strategy constitutes a symmetric Bayes-Nash equilibrium:*

$$b_i = min\{g(v_i), c_i\} \tag{21}$$

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

$$g'(v_i) = \frac{(1-H(g(v_i)))F'(v_i)}{\frac{(1-(1-F(v_i))(1-H(g(v_i)))) \int_{-\infty}^{\infty} u'(x-g(v_i))G'_{v_i}(x)dx}{(N-m)\left(\int_{-\infty}^{\infty} u(x-g(v_i))G'_{v_i}(x)dx-u(0)\right)} - (1-F(v_i))H'(g(v_i))} \tag{22}$$

with boundary condition $g(\overline{r}) = r$, where $\overline{r}$ satisfies the equation $\int_{-\infty}^{\infty} u(x-r)G'_{\overline{r}}(x)dx = u(0)$.

PROOF. Using the same reasoning as that presented in the proof of theorem 5, 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 his budget is $c_i < r$, or because his valuation for the item is $v_i < \overline{r}$, where $\overline{r}$ is the solution of equation $\int_{-\infty}^{\infty} u(x-r)G_{\overline{r}}'(x)dx = u(0)$. Therefore $\overline{r}$ is the smallest valuation for which a bidder would participate in this auction. If a bidder has uncertain valuation v_i smaller than this price then since he needs to bid at least r in the auction, we would make negative expected utility.[14]

We therefore begin by analyzing the case when exactly $n \le N$ agents can participate in the auction; these agents have $c_i \ge r$ and $v_i \ge \overline{r}$. The probability that a particular agent participates in the auction is equal to:

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

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

$$\pi_n = C(N-1, n-1)(\pi(r))^{n-1} \cdot (1 - \pi(r))^{N-n} \tag{23}$$

The distributions $H_r(c)$ and $F_r(v)$ from which the participating agents' c_i and v_i are drawn, are the initial distributions $H()$ and $F()$ respectively, conditional on the fact that $c_i \ge r$ and $v_i \ge \overline{r}$. Thus it is:

$H_r(c) = \frac{H(c)-H(r)}{1-H(r)},\ if\ c \ge r$ & $H_r(c) = 0,\ if\ c < r$ and

$F_r(v) = \frac{F(v)-F(\overline{r})}{1-F(\overline{r})},\ if\ v \ge \overline{r}$ & $F_r(v) = 0,\ if\ v < \overline{r}$.

The distribution, from which the opponents' bids b_j are drawn, is:

$$\begin{aligned} Z_r(x) &= Prob[b_i \le x] = 1 - (1 - F_r(g^{-1}(x))) \cdot (1 - H_r(x)) \\ &= 1 - \frac{1 - F(g^{-1}(x))}{1 - F(\overline{r})} \cdot \frac{1 - H(x)}{1 - H(r)} \end{aligned} \tag{24}$$

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

$$\Phi_k^{n,r}(x) = \sum_{i=0}^{k-1} C(n-1, i) \cdot (Z_r(x))^{n-1-i} \cdot (1 - Z_r(x))^i \tag{25}$$

To analyze the expected profit of a bidder who places a bid b_i in the auction, we distinguish the following cases:

- If $b_i < B^{(m)}$, then bidder i is outbid and doesn't win any items, therefore his utility is $u_i = u(0)$. This happens with probability: $Prob[b_i \le B^{(m)}] = 1 - \Phi_m^{n,r}(b_i)$.
- If $B^{(m)} \le b_i \le B^{(m-1)}$, then bidder i has placed the last winning bid. Thus the payment equals his bid and his utility is $u_i = u(\overline{v}_i - b_i)$. This happens with probability: $Prob[B^{(m)} \le b_i \le B^{(m-1)}] = \Phi_m^{n,r}(b_i) - \Phi_{m-1}^{n,r}(b_i)$.
- If $B^{(m-1)} < b_i$, then bidder i is a winner, the payment is equal to bid $B^{(m-1)}$ and his utility is $u_i = u(\overline{v}_i - B^{(m-1)})$. Note that: $Prob[B^{(m-1)} \le \omega] = \Phi_{m-1}^{n,r}(\omega)$.

Therefore the expected utility of bidder i, who has true valuation $\overline{v}_i$, when he places a bid equal to b_i, is equal to:

$$\begin{aligned} Eu_i^{n,r}(b_i, \overline{v}_i) =& u(0) \cdot (1 - \Phi_m^{n,r}(b_i)) + u(\overline{v}_i - b_i) \cdot \Phi_m^{n,r}(b_i) \\ &+ \int_r^{b_i} u'(\overline{v}_i - \omega) \cdot \Phi_{m-1}^{n,r}(\omega) \cdot d\omega \end{aligned} \tag{26}$$

[14] Note that when there is no uncertainty in the valuation, or the bidders are risk neutral, then $\overline{r} = r$.

Note that even if $n \leq m$, this equation holds, as it is $\Phi_{m-1}^{n,r}(x) = \Phi_m^{n,r}(x) = 1$.

From Bayes' rule, we know that the expected utility that bidder i gets, by placing bid b_i, for any possible numbers of total participating agents, is: $Eu_i(b_i, \overline{v}_i) = \sum_{n=1}^{N} \pi_n \cdot Eu_i^{n,r}(b_i, \overline{v}_i)$. Then using equations 23, 24, 25 and 26, this becomes:

$$Eu_i(b_i, \overline{v}_i) = u(0) \cdot (1 - \Phi_m(b_i)) + u(\overline{v}_i - b_i) \cdot \Phi_m(b_i) + \int_r^{b_i} u'(\overline{v}_i - \omega) \cdot \Phi_{m-1}(\omega) \cdot d\omega \tag{27}$$

The terms $\Phi_m(x)$ and $\Phi_{m-1}(x)$ after a number of mathematical transformations[15] can be shown to be equal to:

$$\Phi_m(x) = \sum_{i=0}^{m-1} C(N-1, i) \cdot (Z(x))^{N-1-i} \cdot (1 - Z(x))^i \tag{28}$$

$$\Phi_{m-1}(x) = \sum_{i=0}^{m-2} C(N-1, i) \cdot (Z(x))^{N-1-i} \cdot (1 - Z(x))^i \tag{29}$$

where $Z(x) = 1 - (1 - F(g^{-1}(x))) \cdot (1 - H(x))$.

As the true valuation $\overline{v}_i$ is not known and bidder i knows that $\overline{v}_i$ is drawn from $G_{v_i}()$, we use Bayes rule to compute the expected utility of bidder i when he places bid b_i given that $\overline{v}_i$ is unknown, but the "uncertain valuation" v_i is. Therefore, the expected utility is: $Eu_i(b_i) = \int_{-\infty}^{\infty} Eu_i(b_i, x) \cdot G_{v_i}^{'}(x) \cdot dx \Rightarrow$

$$Eu_i(b_i) = u(0)(1 - \Phi_m(b_i)) + \Phi_m(b_i) \int_{-\infty}^{\infty} u(x - b_i) G_{v_i}^{'}(x) dx + \int_r^{b_i} \Big(\int_{-\infty}^{\infty} u'(x - \omega) G_{v_i}^{'}(x) dx \Big) \Phi_{m-1}(\omega) d\omega \tag{30}$$

To find the bid which maximizes the expected utility, we set $\frac{\vartheta Eu_i}{\vartheta b_i} = 0$. Using equation 3, to simplify this equation, we get to the following equation:

$(N-m) \cdot (\int_{-\infty}^{\infty} u(x-b_i) G_{v_i}^{'}(x) dx - u(0)) \cdot \Big((1 - F(g^{-1}(b_i))) \cdot H'(b_i) + (1 - H(b_i)) \cdot \frac{F'(g^{-1}(b_i))}{g'(g^{-1}(b_i))} \Big) = \int_{-\infty}^{\infty} u'(x - b_i) G_{v_i}^{'}(x) dx \cdot \Big(1 - (1 - F(g^{-1}(b_i))) \cdot (1 - H(b_i)) \Big)$

Since the bid b_i that maximizes the expected utility is $b_i = g(v_i)$, we substitute this in the equation to get equation 22.

When the bidder's valuation is $v_i = \overline{r}$, he bids $b_i = r$, provided the $c_i \geq r$ is, hence the boundary condition. ■

We would like to point out that the boundary condition changes because of the valuation uncertainty. More specifically, it is no longer necessarily the case that a bidder will only participate if his valuation is higher than r. Since they are not risk neutral and they don't know their true valuation, we need to compute value $\overline{r}$, which is the lowest valuation that

[15] We explain briefly how to simplify the term for $\Phi_m(b_i)$:
$\Phi_m(x) = \sum_{n=1}^{N} C(N-1, n-1) \cdot ((1 - F(\overline{r})) \cdot (1 - H(r)))^{n-1} \cdot (1 - (1 - F(\overline{r})) \cdot (1 - H(r)))^{N-n} \cdot \big(\sum_{i=0}^{m-1} C(n-1, i) \cdot (1 - \frac{1 - F(g^{-1}(x))}{1 - F(\overline{r})} \cdot \frac{1 - H(x)}{1 - H(r)})^{n-1-i} \cdot (\frac{1 - F(g^{-1}(x))}{1 - F(\overline{r})} \cdot \frac{1 - H(x)}{1 - H(r)})^i \big) \Leftrightarrow$
$\Phi_m(x) = \sum_{i=0}^{m-1} C(N-1, i) \cdot ((1 - F(g^{-1}(x))) \cdot (1 - H(x)))^i \cdot \big(\sum_{n=1}^{N} C(N-1-i, N-n) \cdot (1 - (1 - F(\overline{r})) \cdot (1 - H(r)))^{N-n} \cdot ((1 - F(\overline{r})) \cdot (1 - H(r)) - (1 - F(g^{-1}(x))) \cdot (1 - H(x)))^{n-1-i} \big)$
This gives us the desired formula. A similar proof simplifies term Φ_{m-1}.

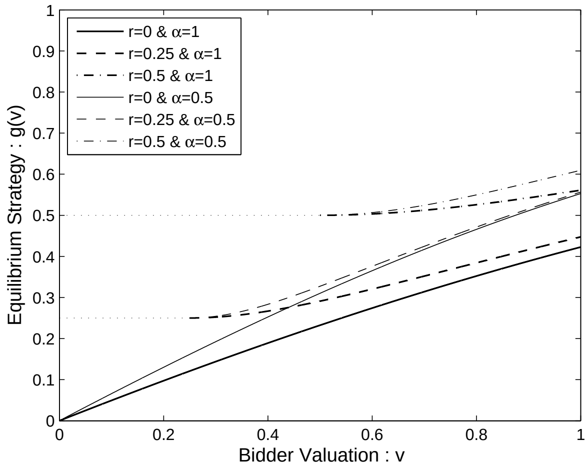

Fig. 2 Equilibrium strategy $g(v)$ for the cases when bidders are risk-neutral ($\alpha = 1$), and risk-averse ($\alpha = 0.5$). The valuations v and the budget constraints c are both drawn from the uniform distribution $U[0,1]$. The auction has a reserve price, which takes values $r = 0, 0.25, 0.5$. 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.

gives expected utility equal to $u(0)$, i.e. what the agent makes by not participating in the auction, when he is forced to make a payment equal to r. Any agent with uncertain valuation $v_i > \overline{r}$ will participate in the auction since he will make a higher expected utility if he pays r and any agent with $v_i < \overline{r}$ will not, since he makes less than $u(0)$.

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 $U[0,1]$ and there is no uncertainty about the valuation:[16]

Corollary 2 *In the case that $F(v)$ and $H(c)$ are uniform distributions $U[0,1]$, and there is no uncertainty about the valuation v_i (i.e. $\overline{v}_i = v_i$), the equilibrium strategy is: $b_i = min\{g(v_i), c_i\}$ where $g(v)$ is the solution of d.e.:*

$$g'(v_i) = \frac{1 - g(v_i)}{\frac{u'(v_i - g(v_i))(v_i + g(v_i) - v_i g(v_i))}{(N-m)(u(v_i - g(v_i)) - u(0))} - (1 - v_i)} \tag{31}$$

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

We choose to use the CRRA utility function $u(x) = x^{\alpha}, \alpha \in (0,1)$, as this is a standard utility function used in the literature. By substituting it into equation 31, 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 16, 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,

[16] As we have already examined, in the previous section, how valuation uncertainty affects the bidding strategy, we decided to simplify the example in this section, so as to be able to concentrate on the other three issues and how those, in turn, affect the bidding strategy.

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.

7.2 The $(m+1)^{th}$-price Auction Case

Here, we examine the same problem for the setting of an $(m+1)^{th}$ price auction.

We start by examining the case, where all the issues except valuation uncertainty are present; we can prove the following theorem:

Theorem 10 *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 $u()$, it is a dominant strategy to bid: $b_i = min\{v_i, c_i\}$, if $b_i \geq r$, and not to participate otherwise.*

PROOF. We know that truthful bidding is the dominant strategy when bidders have any risk attitude. The presence of the budget constraint c_i serves to restrict this bid to be less that c_i, and the reserve price r means that the bidder will only participate if it would be higher than r. ■

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 7 and 10:

Theorem 11 *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 $u_i()$, 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:*

$$\int_{-\infty}^{\infty} u(z - \beta_i) G_i'(z) dz = u(0) \tag{32}$$

PROOF. Using the exact same reasoning as in theorem 7, the expected utility of an agent bidding β_i is:

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

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 32. As $\frac{dEu_i(x)}{dx} > 0, \forall x < \beta_i$, 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 $b_i = min\{\beta_i, c_i\}$. He also should never bid above β_i, because then $Eu_i \leq u(0)$; thus if $b_i < r$, he should not participate. ■

Let us now examine briefly how the strategy is affected by the four issues. We have already examined this, in section 6, for the issues of risk attitudes and valuation uncertainty. In this section we also considered reserve prices r and budget constraints c_i. However, the only way that these two parameters affect the strategy is that they restrict the bid to be less than c_i and more than r. Therefore, it is the risk attitudes and the valuation uncertainty that really determine the agent's bid and as we have already examined these we will not discuss this issue further.

7.3 Discussion: Unifying the Models of Valuation Uncertainty

The model of valuation uncertainty that we considered in the $(m+1)^{th}$ price auction is the most general form of valuation uncertainty possible. That is because each bidder can have his own distribution from which his true valuation is drawn. On the other hand, the model we used for the m^{th} price auction is more restricted; essentially all bidders have the same model of uncertainty in the sense that the distribution $G()$ from which the true valuation is drawn depends on the parameter v_i, the approximate (or uncertain) valuation of the bidder. Now, there is a way to extend this setting so that bidders can have different models of uncertainty. Specifically, we can parameterize the distribution $G()$ of each bidder by a second parameter α. Let $G_{\alpha,v_i}()$ be this distribution. For example, a bidder knows approximately that his value is $v_i = 10$, and that his true valuation is within $\alpha\%$ of that value, so that his true valuation is drawn between $[v_i(1-\alpha), v_i(1+\alpha)]$. We could have a whole family of functions $G_{v_i}()$ parameterized by α. If we do this then essentially the m^{th} price model becomes identical to the $(m+1)^{th}$ price model we used.

Now, this setting is similar to the settings examined in [12,9], where the bidders are asymmetric in the sense that their valuations are drawn from different prior distributions. However, there is still a small difference between the two models because there is no uncertainty about what models the other bidders are using. In [20], we have looked at asymmetric bidder models under uncertainty (i.e. where the bidders don't know their opponent models, unlike in [9]), where we examine bidders with different risk attitudes and competitiveness. From these papers, we can observe that in all cases (both in the settings without and with uncertainty) where there are a number of possible different models for the asymmetric bidders, the equilibria are computed by solving a system of differential equations of the type described in [20], and we have given in that paper an algorithm for solving these systems of differential equations. Therefore unifying the uncertainty models requires to consider asymmetric bidder models, which is a novel direction of research that we are currently undertaking.

8 Experimental Evaluation

In this section, we present three different examples of experiments that we performed. The first one examines how the bidding strategy changes in an m^{th} price auction in the presence

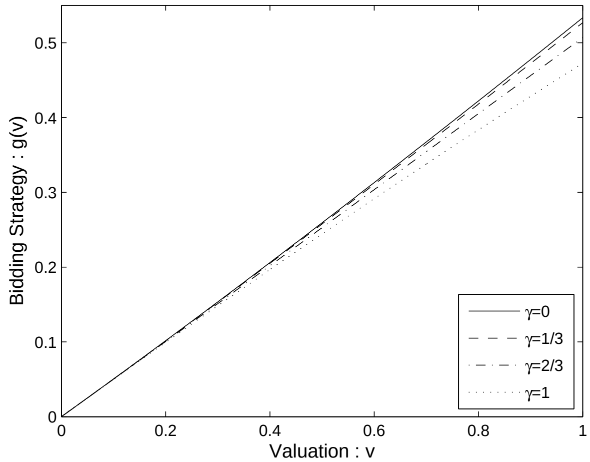

Fig. 3 Bidding strategies for the case of an m^{th} price auction, in the presence of valuation uncertainty and risk averse bidders. $m = 2$ items are auctioned to $N = 3$ participating bidders with uncertain valuations v_i drawn from uniform distribution $U[0, 1]$. The utility function used is $u(x) = (x + 1)^{\alpha}$ with $\alpha = 0.01$. The bidders true valuation are within γ of the uncertain valuation v_i, so they are drawn uniformly from $[v_i(1 - \gamma), v_i(1 + \gamma)]$.

of valuation uncertainty. This complements the theoretical results about how valuation uncertainty affects the bidding in the $(m + 1)^{th}$ price auction, by showing similar behavior for the m^{th} price auction setting. The remaining two experiments are intended to demonstrate the usefulness of our theoretical analysis in practice, by showing the potential benefits for both bidders and seller obtained by using our analysis. It should be noted that we examine cases where there is no valuation uncertainty in these two experiments so as to be able to concentrate on the other issues and not overly complicate our explanations, especially given that the first experiment already accounted for how valuation uncertainty affects the bidding in conjunction with varying risk attitudes.

8.1 Effect of Valuation Uncertainty on Bidding

This example is designed to examine how the presence of valuation uncertainty affects the equilibrium bidding strategy in the m^{th} price auction. We have already proven theoretically that the more risk averse (resp. risk seeking) the bidders become and the higher the uncertainty (i.e. variance) of the valuation uncertainty, the less (resp. more) they will bid in the case of the $(m + 1)^{th}$ price auction. This example shows that a similar result will hold in general. We graph in figure 3 the bidding strategies of risk averse bidders, for different degrees of valuation uncertainty. To be more precise, we assume all bidders use utility function $u(x) = (x + 1)^{0.01}$,[17] and also that if they have uncertain valuation v_i, then their true valuation for the good sold (which is unknown to them) is uniformly distributed in

[17] This is a modified CRRA utility function with $\alpha = 0.01$, but we add $+1$ to x so as to make sure that a utility is assigned to cases where the profit is negative as well, which is not covered by the original CRRA utility function. The reason for using $\alpha = 0.01$, is so that the difference becomes sufficient to be visible on the graph; smaller values of α give similar results, but the bidding strategies for the different values of γ show smaller differences.

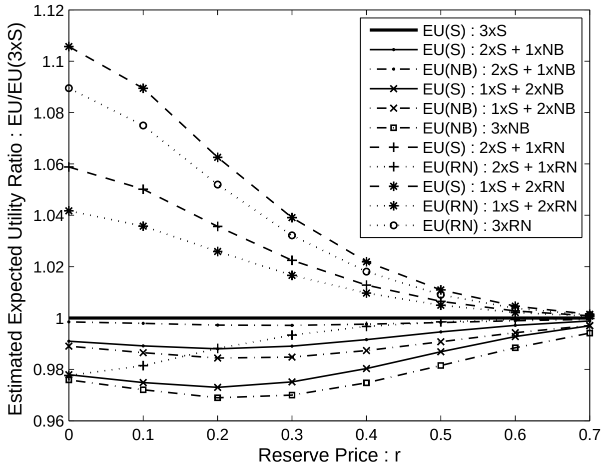

Fig. 4 Experimental comparison of the equilibrium strategy *S* against strategies *NB* (budget constraints are ignored) and *NR* (risk attitudes are ignored). In all experiments some agents used strategy *S* and all the rest use the same strategy (either *NB* or *RN*).

$[v_i(1-\gamma), v_i(1+\gamma)]$, thus $G'_{v_i}(x) = \frac{1}{2\gamma v_i}$, for $x \in [v_i(1-\gamma), v_i(1+\gamma)]$ and $G'_{v_i}(x) = 0$, otherwise. When $\gamma = 0$, there is no valuation uncertainty, and as γ increases so does the uncertainty. From figure 3, we verify that as γ increases, the bidders will indeed bid lower (for any valuation v_i).

The same will happen as the bidders' risk averseness increases. However, since for an m^{th} price auction, when there is no valuation uncertainty, as bidders become more risk averse, they increase their bids, the valuation uncertainty produces the opposite effect from just the risk averseness when there is no valuation uncertainty. As can been seen in the graph, while for small values of the parameter γ the bidders bid higher than $g(v) = \frac{v}{2}$, which is what a risk neutral bidder would bid, on the other hand, for large values of γ, they bid less than that. Therefore, when the valuation uncertainty is significant, the risk averse bidders bid in the same way as risk seeking bidders would (when the latter have no valuation uncertainty). Without our complete analysis, it would not be possible to figure out how these two effects factor in the final equilibrium strategy.

8.2 Out of Equilibrium Play : Validating Our Analysis

Using this example, we further validate experimentally the usefulness in practice of our 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 22 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.

To this end, we simulated the case when $N = 3$ bidders participate in an m^{th} price auction, where $m = 2$ items are sold; this is a 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 $U[0, 1]$.

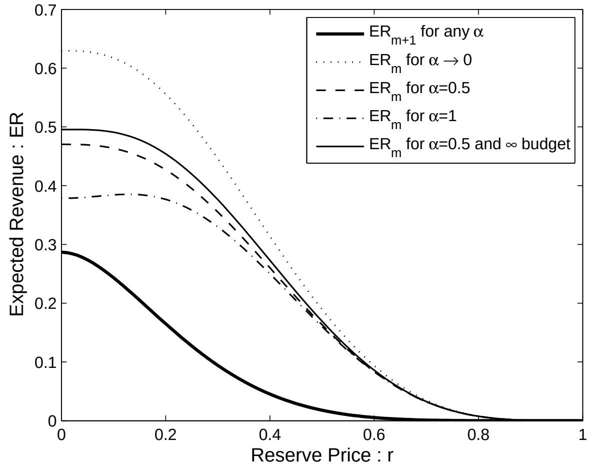

Fig. 5 Expected revenue both for m^{th} and the $(m+1)^{th}$ price auctions, in the presence of varying reserve prices and bidder risk attitudes. $m = 2$ items are auctioned to $N = 3$ participating bidders with valuations v_i and budget constraints c_i drawn from uniform distribution $U[0, 1]$. The utility function used is $u(x) = x^{\alpha}$.

We denote this standard equilibrium strategy (given by equation 22) 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).[18] 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 4; 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").[19] 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*. This means that strategies *NB* and *RN* are *dominated* by *S*, when we consider agents who play either *NB*/*RN* or *S*, and therefore *the use of our novel analysis does lead to higher bidder utility compared to cases where some feature is not taken into account*.

[18] The 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*).

[19] We have used a sufficiently large number of samples spanning all possible combinations of valuations when we computed the expected revenue in each case, therefore the error is very small (less than 0.001%), and this is the reason why we don't put error bars in the graph. Additionally, the notation "2xS + 1xNB" means that two agents using strategy *S* and one using strategy *NB* participate in the experiment, etc.

8.3 Examining the Seller Revenue

Given the fact that bidders will use this new equilibrium strategy (as shown by the previous 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 $(m+1)^{th}$).[20] 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 $\Psi_k(x) = \sum_{i=0}^{k-1} C(N,i)\big(Z(x)\big)^{N-i}\big(1-Z(x)\big)^i$ 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 5, 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 of revenue.[21] Once more we see that *our analysis is necessary to determine the correct reserve price that maximizes the seller's revenue*. In figure 5, we also graph two additional cases: when the bidders are risk-averse to the extreme ($\alpha \to 0$), and when an $(m+1)^{th}$ price auction is used (the revenue doesn't depend on α in this case). As was expected, when $\alpha \to 0$, 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.

9 Conclusions

In this paper, 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 them for the case of uncertainty in the valuation that bidders have, when the bidders are not risk-neutral, for both the m^{th} and the $(m+1)^{th}$ price auction settings. We also showed theoretically and experimentally that the variance in this uncertainty and the risk attitude of a bidder determine the deviation of the equilibrium bidding strategy from bidding 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 and we also included the uncertainty of bidders' valuation. We then discussed about extending these results for various models of valuation uncertainty. Fourth, we used simulations to show that this analysis is useful 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 auction format and thus maximize her revenue.

[20] Note 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.

[21] This 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.

As future work, we would like to examine the case of identical items being sold in *multiple concurrent auctions* [6]; in this case it is necessary to place bids in all the auctions. Furthermore, there are settings in which competition between agents negates the traditional assumption that agents are self-interested, i.e. maximizing their profit, which leads to more aggressive bidding [18, 2]; we intend to incorporate this issue into our results. Finally, as we discussed in section 7.3, we are currently looking at asymmetric bidder models, a setting which extends and unifies our models of valuation uncertainty.

Acknowledgements 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).

References

1. Borgs, C., Chayes, J.T., Immorlica, N., Mahdian, M., Saber, A.: Multi-unit auctions with budget-constrained bidders. In: Proceedings of the 6th ACM Conference on Electronic Commerce (ACM EC'05), pp. 44–51 (2005)
2. Brandt, F., Sandholm, T., Shoham, Y.: Spiteful bidding in sealed-bid auctions. In: IJCAI-07, pp. 1207–1214 (2007)
3. Byde, A.: Applying evolutionary game theory to auction mechanism design. HP Laboratories Bristol, Tech Report HPL-2002-321 (2002)
4. Che, Y., Gale, I.: Standard auctions with financially constrained bidders. Review of Economic Studies **65**(1), 1–21 (1998)
5. Che, Y.K., Gale, I.: Expected revenue of all-pay auctions and first-price sealed-bid auctions with budget constraints. Economics Letters **50**(3), 373 – 379 (1996)
6. Gerding, E.H., Dash, R.K., Yuen, D.C.K., Jennings, N.R.: 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, pp. 267–274. Hawaii, USA (2007)
7. Krishna, V.: Auction theory. Academic Press (2002)
8. Larson, K., Sandholm, T.: Costly valuation computation in auctions. In: 8th Conference of Theoretical Aspects of Knowledge and Rationality (TARK VIII), pp. 169–182 (2001)
9. Lebrun, B.: First price auctions and the asymmetric n bidder case. Int. Econ. Rev. **40**(1), 125–142 (1999)
10. Liu, Y., Goodwin, R., Koenig, S.: Risk-averse auction agents. In: Proceedings of the 2nd International Joint Conference on Autonomous Agents and Multi-Agent Systems, pp. 353–360 (2003)
11. Maskin, E., Riley, J.: Optimal auctions with risk averse buyers. Econometrica **52**(6), 1473–1518 (1984)
12. Maskin, E., Riley, J.: Asymmetric auctions. Review of Economic Studies **67**, 413–438 (2000)
13. Myerson, R.B.: Optimal auction design. Mathematics of Operations Research **6**, 58–73 (1981)
14. Parkes, D.C.: Optimal auction design for agents with hard valuation problems. In: Proc. IJCAI-99 Workshop on Agent Mediated Electronic Commerce, pp. 206–219 (1999)
15. Rice, J.A.: Mathematical Statistics and Data Analysis. Duxbury Press, California (1995)
16. Riley, J.G., Samuelson, W.F.: Optimal auctions. The American Economic Review **71**(3), 381–392 (1981)
17. Thompson, D., Leyton-Brown, K.: Valuation uncertainty and imperfect introspection in second-price auctions. In: Proceedings of the 22nd Conference on Artificial Intelligence (AAAI), pp. 148–153 (2007)
18. Vetsikas, I.A., Jennings, N.R.: Outperforming the competition in multi-unit sealed bid auctions. In: Proceedings of the 6th International Joint Conference on Autonomous Agents and Multi-Agent Systems, pp. 702–709 (2007)
19. Vetsikas, I.A., Jennings, N.R.: Bidding strategies for realistic multi-unit sealed-bid auctions. In: Proceedings of AAAI-08, pp. 182–189 (2008)
20. Vetsikas, I.A., Jennings, N.R.: Considering assymmetric opponents in multi-unit sealed-bid auctions. In: Proceedings of the 1st International Workshop on Market Based Control (MBC) (2008)
21. Weber, R.: Making more from less: Strategic demand reduction in the fcc spectrum auctions. Journal of Economics and Management Strategy **6**(3), 529–548 (1997)