

Outperforming the Competition in Multi-Unit Sealed Bid Auctions

Ioannis A. Vetsikas
School of Electronics and Computer Science
University of Southampton
Southampton SO17 1BJ, UK
iv@ecs.soton.ac.uk

Nicholas R. Jennings
School of Electronics and Computer Science
University of Southampton
Southampton SO17 1BJ, UK
nrj@ecs.soton.ac.uk

ABSTRACT

In this paper, we examine the behavior of bidding agents that are in direct competition with the other participants in an auction setting. Thus the agents are not simply trying to maximize their own utility, rather they wish to maximize a weighted difference of their own gain to that of their competitors. By so doing, this work significantly extends the existing state-of-the-art results on single unit auctions, by generalizing to the multi-unit case. Specifically, our main result is the derivation of symmetric Bayes-Nash equilibria for these agents in both m^{th} and $(m + 1)^{\text{th}}$ price sealed bid auctions. Subsequently, we use these equilibria to examine the profits of different agents and show that aiming to beat the competition is more effective than pure self interest in any competitive setting. Finally, we examine how the auctioneer's revenue is affected and find that the weight that agents place in minimizing the opponents' profit determines whether the m^{th} or the $(m + 1)^{\text{th}}$ price auction yields a higher revenue.

Categories and Subject Descriptors

I.2.11 [ARTIFICIAL INTELLIGENCE]: Multiagent systems; I.2.11 [ARTIFICIAL INTELLIGENCE]: Intelligent Agents

General Terms

Theory, Economics, Experimentation

Keywords

game theory, bidding strategies, equilibrium analysis, revenue, simulation

1. INTRODUCTION

Game theory is widely used in multi-agent systems, as a way to model and predict the interactions between rational agents. Auctions have also become quite popular, especially

in the last decade, particularly with eBay and other similar companies bringing auctions to the Internet and millions of online bidders. Now, in most of the cases examined, either by game theory or by experimental analyses on auction participant behaviour, a key assumption is that bidders are rational and self-interested (i.e. they care only about maximizing their profit in the auction, and not about the profit of other bidders). However, there are observations and scenarios that cannot convincingly be explained by these assumptions alone. These usually involve direct competition between companies or individual agents, where fewer participants mean higher profits for those involved, or cases in which ranking and/or relative profit is important. In such cases, the agents will try to maximize the difference of their utility to that of any other competitor.

In more detail, game theorists have studied a number of examples where self-interest, at least in the short term, is not as important as retaining a hold on the market, especially in scenarios that involve monopolies or oligopolies [6]. In addition, there are examples of real-world scenarios, where myopic self-interest alone cannot explain the observed behaviour. On eBay auctions, for example, bidders sometimes get into bidding wars, which cause those auctions to close much higher than other auctions selling the exact same commodity; this could be attributed partially to a "love for winning" and a total utility that takes into account the fact itself that the bidder lost. In addition, in many European countries, where mobile phone spectrum licenses were sold in various auction settings, it was observed that many telecom companies bid higher than the estimated value of the corresponding licenses, in order to deny them from their competition and to discourage smaller competitors from entering the market. Furthermore, in a competition setting, like the Trading Agent Competition (TAC), it has been observed that participants tend to bid quite high for commodities that are in short demand and which are needed in order to make profit; while part of this bidding behaviour can be attributed to the valuation they have of these items, part of this is also due to the desire to outperform the opposition. In fact, in [10], the authors suggest this competition factor as the main reason why the most successful strategies bid more aggressively than the assumption of self-interested agents alone would predict.

Against this background, this work considers that agents not only wish to maximize their own profit, but at the same time they also wish to minimize the profits of their opponents. Specifically, our assumption for the model used by agents in such settings is that they wish to maximize a

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AAMAS'07 May 14–18 2007, Honolulu, Hawai'i, USA.
Copyright 2007 IFAAMAS .

weighted sum of their profit relatively to the profit of their competitors. The weight of this sum, α , is the spite (or competition) coefficient, that denotes the degree of competition in the setting under examination. It can vary from self-interested agents, when $\alpha = 0$, to completely malicious ones, which only care about minimizing the gain of their opponents, when $\alpha = 1$. Now, there is also a small, but growing, literature of papers that consider cases either equivalent or very similar to this one. For some of the most relevant work see [1, 2, 4, 5, 7, 9]. This paper significantly extends the work presented in these papers and, in particular, in [2]. To be more precise, while these papers examine the strategies and the equilibria for single unit auctions, i.e. second price (in most cases) and/or first price auctions, we examine the case of *multi-unit auctions*. This is an important extension, because in theoretical analyses, as well as in real-world scenarios, it is usually the case that more than one instance of a particular commodity is available.

This paper is organized as follows. In section 2, we formally present the auction setting we examine. Then we provide, for the first time, symmetric Bayes-Nash equilibria for both m^{th} and $(m+1)^{th}$ price sealed bid auctions in sections 3 and 4. We further prove that allowing the bidders to bid for multiple items, which they don't need, in order to deprive them from the competition, does not change the equilibrium strategies.¹ These equilibria, along with the corresponding proofs, are the primary contributions of this paper. From this analysis, we observe a number of differences with the single item case that was examined in related work. For example, for large α , the equilibrium strategy for the m^{th} price auction is to bid more than the true valuation, and we also demonstrate that the equilibrium strategy for the $(m+1)^{th}$ price auction is independent of the number of items sold. In section 5, we examine how the auctioneer's revenue is affected. In particular, we show that revenue equivalence does not hold, and, in fact, the revenue of the $(m+1)^{th}$ price auction is only higher than that for the equivalent m^{th} price auction when α is small; which is not observed in single unit auctions. In section 6, we examine the profit that various strategies yield when pitted against various other strategies, like itself or strategies that do not take the competition into account (as if $\alpha = 0$). From this, we deduce that the most successful strategies in a competition setting are those derived for a coefficient α that is just right to beat the opposition.

2. PROBLEM STATEMENT

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 symmetric² Bayes-Nash equilibria for sealed bid auctions where $m \geq 1$ identical items are being sold. The two most common auction settings in this context are the m^{th} and $(m+1)^{th}$ price auctions, in which the top m bidders win one item

¹This means that the computed equilibria are “*false-name-proof*”, i.e. that no agent may profit by submitting bids under an assumed identity.

²This 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.

each at a price equal to the m^{th} and $(m+1)^{th}$ highest bid respectively. Specifically, we assume that N bidders (where $N \geq m$) participate in the auction and these agents have a private valuation (utility) u_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) $F(u)$, which is the same for all bidders. Furthermore, let u_l and u_h be the bounds for the possible values of random variables u_i as defined by $F(u)$:³

$$u_l = \max\{u | F(u) = 0\}$$

$$u_h = \min\{u | F(u) = 1\}$$

Finally, we will consider bidding strategies $g(u)$ that are strictly increasing and differentiable over $[u_l, u_h]$, that is the area where $F(u)$ is defined. These are the assumptions commonly made when computing equilibria in the auction theory literature [3].

In this work, we define the objective function (i.e. the total utility) that each agent tries to maximize in the same way as in [2]:

Definition 1. The objective function that each agent wishes to maximize is given by:

$$U_i = (1 - \alpha) \cdot \tilde{u}_i - \alpha \cdot \sum_{j \neq i} \tilde{u}_j$$

where $\alpha \in [0, 1]$ is a parameter called the spite coefficient⁴, \tilde{u}_i is the gain of agent i (i.e. $\tilde{u}_i = 0$, if it does not win any items, and $\tilde{u}_i = u_i - p_i$, if it does) and p_i is the total payment the agent must make to the auctioneer.

We would like to point out that from the point of view of an agent competing against other agents, the following definition of the objective function might look more natural:⁵

$$U_i = \tilde{u}_i - \gamma \cdot \sum_{j \neq i} \tilde{u}_j$$

This is because the agent cares about maximizing its own profit minus the weighted profit of its opponents, where the weight $\gamma \geq 0$. For example, if we have a competition (like TAC), then the agent wishes to maximize the difference of its profit to the expected profit of any opponent, and hence we would use this objective function with $\gamma = \frac{1}{N-1}$. Another parameter that is also interesting in a competition setting, is when $\gamma = 1$, because then the agent maximizes his gain against the sum of the gains of every opponent. However, this objective function is equivalent to the objective function presented in definition 1, when $\gamma = \frac{\alpha}{1-\alpha}$, so definition 1 is more general, and this is the reason why we use it. Moreover, maximizing the difference of the agent's profit to that of any opponent, and to that of all opponents, respectively, can easily be represented using coefficient $\alpha = \frac{1}{N}$, and $\alpha = \frac{1}{2}$, in the objective function.

³For example, if $F(u)$ is the uniform distribution $U[0, 1]$, then $u_l = 0$ and $u_h = 1$. If $F(u)$ is such that u_i can take values in $[0, +\infty)$, then $u_l = 0$ and $u_h = +\infty$. We use these tight bounds to define the boundary conditions for the equilibria we compute.

⁴We also call this the “competition coefficient”, since it determines how much the agents take their competitors into account in the total utility.

⁵In fact this is the total utility function used in [7].

Now that we have defined the general form of the function that each agent tries to maximize, each agent i will do so by submitting a particular bid v_i for the item it needs. Here the “optimal” bids for the agent are those given from the equilibrium strategy $g_\alpha(u)$, assuming that other agents use the same strategy. Thus the bid maximizing the objective function should be $v_i = g_\alpha(u_i)$. To compute the probability distributions of the order statistics of the opponents’ bids and valuations, we use the information provided by the following lemma:

LEMMA 1. *The cdf of the k^{th} order statistic, denoted $U^{(k)}$, of the $(N - 1)$ valuations (utilities) u_i of the opponents of any given agent, is given by:*

$$\text{Prob}[U^{(k)} \leq x] = \Phi_k(x)$$

while the cdf of the k^{th} order statistic, denoted $B^{(k)}$, of the $(N - 1)$ bids of these opponents, when each of them bids $v_i = g_\alpha(u_i)$ according to the equilibrium strategy $g_\alpha(u)$, is given by:

$$\text{Prob}[B^{(k)} \leq x] = \Phi_k(g_\alpha^{-1}(x))$$

where $C(N - 1, i)$ is the total number of possible combinations of i items chosen from $(N - 1)$, and

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

PROOF. The probability distribution from which each utility u_i is drawn has cdf $F(u)$; the probability distribution from which the bids v_i are drawn has cdf $G(x) = F(g_\alpha^{-1}(x))$. We also know that the k^{th} order statistic of a set of $(N - 1)$ i.i.d. random variables drawn from a distribution with cdf $H(x)$ is $\sum_{i=0}^{k-1} C(N - 1, i) \cdot (H(x))^{N-1-i} \cdot (1 - H(x))^i$ (see [8]). Using this equation for $G = F$ and $G = H$, respectively, we get the two equations of this lemma. \square

We will use these formulae extensively in the computations of the Bayes-Nash equilibria in the next sections.

3. M-TH PRICE AUCTIONS

In this section we examine the symmetric Bayes-Nash equilibrium that exists in the case of an m^{th} price auction. Specifically, each agent bids v_i for a unit of the item on sale, and we initially assume that it is only allowed to bid for one item and no more.

THEOREM 1. *In the case of an m^{th} price sealed bid auction with N participating bidders, in which each agent i is interested in purchasing one unit of the good for sale with inherent utility (valuation) for that item equal to u_i , u_i is i.i.d. drawn from $F(u)$, and an α -coefficient for outperforming its competition, the following bidding strategy constitutes a symmetric Bayes-Nash equilibrium:*

$$g_\alpha(u) = u - (F(u))^{-\frac{N-m}{1-\alpha \cdot m}} \cdot \int_{u_l}^u (F(z))^{\frac{N-m}{1-\alpha \cdot m}} \cdot dz \quad (2)$$

if $\alpha \cdot m < 1$,

$$g_\alpha(u) = u \quad (3)$$

if $\alpha \cdot m = 1$, and

$$g_\alpha(u) = u + (F(u))^{-\frac{N-m}{1-\alpha \cdot m}} \cdot \int_u^{u_h} (F(z))^{\frac{N-m}{1-\alpha \cdot m}} \cdot dz \quad (4)$$

if $\alpha \cdot m > 1$.

PROOF. See the appendix, section A.1. \square

It is interesting to notice that for small values of the spite coefficient α (namely $\alpha < \frac{1}{m}$) the equilibrium strategy is to bid less than the agent’s true valuation, for $\alpha = \frac{1}{m}$ to bid truthfully, and for higher values of α ($\alpha > \frac{1}{m}$) to bid more than the true valuation. This is quite different from the first price auction result examined in [2], because in that case, we do not observe the behavior of bidding more than the true valuation. However, in an auction where multiple items are sold, the price paid by all winners is the same and equal to the lowest winning bid. This means that quite often the price paid is significantly lower than one’s bid, especially as m gets bigger, and this leads to agents bidding higher than they would in an identical auction where fewer items are sold. In addition, using small values of α means that the agents primarily care about maximizing their own profit, so they will bid a lot less aggressively, than in the case where they care mostly about minimizing the profit of their opponents (large α). These two facts when taken together explain why in multi-unit auctions there are cases (for large α) when agents should bid more than their true valuation for an item.

COROLLARY 1. *In the case that $F(u)$ is the cdf of a uniform distribution $U[0, 1]$, the equilibrium strategy is:*

$$g_\alpha(u) = \begin{cases} \frac{\beta}{1+\beta} \cdot u & \beta > 0 \\ u & \beta = 0 \\ \frac{\beta}{1+\beta} \cdot u + \frac{u^{-\beta}}{1+\beta} & \beta < 0 \wedge \beta \neq -1 \\ (1 - \ln(u)) \cdot u & \beta = -1 \end{cases} \quad (5)$$

where $\beta = \frac{N-m}{1-\alpha \cdot m}$.

We use this equilibrium strategy derived for $F(u)$ being the uniform distribution $U[0, 1]$, in our analysis of the auctioneer’s revenue in section 5 and the simulations we present in section 6.

Naturally when there are $m > 1$ items for sale, it might be beneficial to attempt to purchase multiple items, which the agent does not use⁶, in order to deny them from its opponents. Now, even if the rules prohibit agents from bidding for multiple items, they can always submit bids under a false identity. Moreover, there are cases when this makes sense, so we will assume that it will happen. As an example consider the following. Assume that $\alpha = 1$ and the opponents bid (erroneously) according to $g_0(u)$ instead of $g_1(u)$. The opponents who win an item make a profit, since $g_0(u) < u$, so the payment is less than their valuations. By bidding for a second item (and winning it), the agent denies its competitors a gain equal to $U^{(m-1)} - p$ (p is the payment), which is a positive value, and forces them to pay a higher new payment $p' > p$ for the rest of the $(m - 2)$ items that they win. However this case arises due to the fact that the opponents do not bid according to the correct equilibrium strategy. In the case that they do bid as per the equilibrium strategy, a bidder i does not gain by bidding for a second item, because the profit from its opponents losing an item and increasing their payments is offset (on expectation) from the fact that the opponents make a payment for one less item, while bidder i pays for one more itself. We examine this case in the next theorem.

⁶We are assuming free disposal, meaning that an agent is allowed to throw away items that it does not wish to use without gain or loss of utility.

THEOREM 2. *In the case of the m^{th} price sealed bid auction described by theorem 1, when each bidder is allowed to place an unrestricted number of bids for multiple items, bidding for exactly one item according to the function $g_\alpha(u)$ given in theorem 1, constitutes a symmetric Bayes-Nash equilibrium.*

PROOF. See the appendix, section A.2. \square

The two theorems presented in this section prove that the strategy $g_\alpha(u)$ that we computed is indeed a Bayes-Nash equilibrium in an m^{th} price auction, even when the agents are allowed to place multiple bids, for items that they don't need.

4. (M+1)-TH PRICE AUCTIONS

In this section we examine the symmetric Bayes-Nash equilibrium that exists in the case of an $(m+1)^{\text{th}}$ price auction. Again we initially assume that each agent bids v_i for a unit of the item on sale, and that each agent is allowed to bid for only one item and no more.

THEOREM 3. *In the case of an $(m+1)^{\text{th}}$ price sealed bid auction with N participating bidders, each agent i interested in purchasing one unit of the good for sale with inherent utility (valuation) for that item equal to u_i , u_i is i.i.d. drawn from $F(u)$, and an α -coefficient for outperforming its competition, the following bidding strategy constitutes a symmetric Bayes-Nash equilibrium:*

$$g_\alpha(u) = u + (1 - F(u))^{-\frac{1}{\alpha}} \cdot \int_u^{u_h} (1 - F(z))^{\frac{1}{\alpha}} \cdot dz \quad (6)$$

PROOF. See the appendix, section A.3. \square

From this, it is interesting to notice that the agent always bids more than its true valuation when $\alpha > 0$ and also that the equilibrium strategy does not depend on the number of agents N or the number of items m for sale, but only on the competition coefficient α . While these observations are consistent with those for the second price auction result examined in [2] (in that case the authors showed that the strategy did not depend on the number of bidders N), we further show that the number of items sold m , also does not matter for the equilibrium strategy.

COROLLARY 2. *In the case that $F(u)$ is the cdf of a uniform distribution $U[0, 1]$, the equilibrium strategy is:*

$$g_\alpha(u) = \frac{u + \alpha}{1 + \alpha} \quad (7)$$

We use this equilibrium strategy derived for $F(u)$ being the uniform distribution $U[0, 1]$, in our analysis of the auctioneer's revenue in section 5 and the simulations we present in section 6.

As before, since there are $m > 1$ items for sale, one might think that it would be beneficial to attempt to purchase multiple items, which the agent does not use, in order to deny them from its opponents. We examine this case in the next theorem.

THEOREM 4. *In the case of the $(m+1)^{\text{th}}$ price sealed bid auction described by theorem 3, when each bidder is allowed to place an unrestricted number of bids for multiple items, bidding for exactly one item according to the function $g_\alpha(u)$ given in theorem 3, constitutes a symmetric Bayes-Nash equilibrium.*

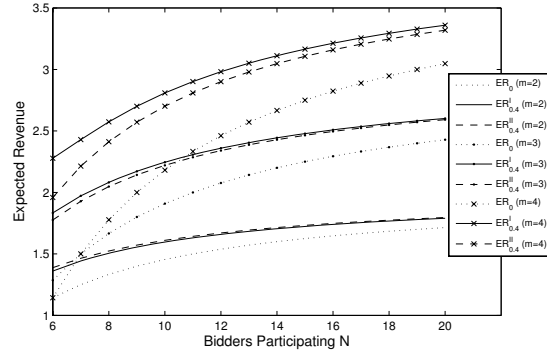


Figure 1: Examination of the expected revenue (ER_α^I and ER_α^{II}) for $N = 6, \dots, 20$ participating bidders. Graphed for $\alpha = 0$ (the revenue is the same for both auctions) and $\alpha = 0.4$.




Figure 2: Examination, for various values of α , of the expected revenue ER_α^I for m^{th} price auctions (solid lines) and ER_α^{II} for $(m+1)^{\text{th}}$ price auctions (dotted lines). From bottom to top the pairs of lines are for $m = 2, m = 3$ and $m = 4$ items sold respectively. The number of agents participating is $N = 2 \cdot m$.

PROOF. (SKETCH) Due to space limitations, we only present a brief sketch of the proof for this case. In a similar way to the proof of theorem 2, we can show that the total utility, when bidder i bids for two items (placing bids v_i and \tilde{v}_i), is a sum of two parts. The first depending on v_i , which is maximized when $v_i = g_\alpha(u_i)$, meaning that the top bid should follow the equilibrium strategy derived in theorem 3. The second depending on \tilde{v}_i , which is maximized when the bidder does not really bid any amount that will allow it to win, i.e. $\tilde{v}_i \leq \frac{u_i + \alpha}{1 + \alpha}$. \square

Of course, given that, in the $(m+1)^{\text{th}}$ price auction, the equilibrium strategy does not depend on the number of opponents, nor the items sold, it only really makes sense to bid for worthless items if the opponents make a serious mistake in the strategy that they use. For example, if $\alpha = 1$, and all the opponents bid less than their true valuation, then, in that case, it does make sense to bid for a second item.

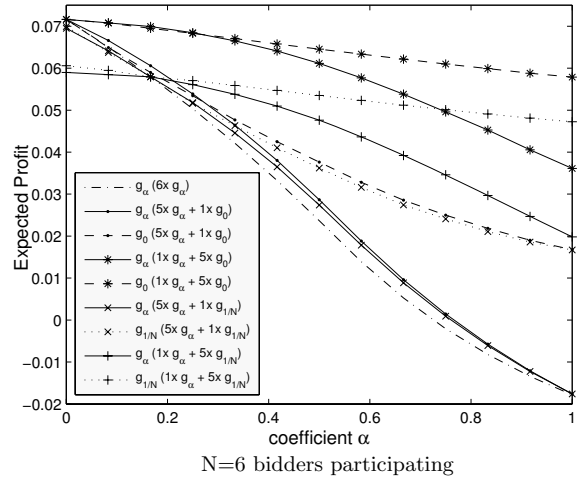



Figure 3: Examination of the expected profit in an m^{th} price auction ($m = 2$) with N participating bidders. For various α coefficients the performance of agents using the $g_\alpha(u)$ strategy is compared against agents using the $g_0(u)$ and $g_{1/N}(u)$ strategies.

5. REVENUE ANALYSIS

In this section we examine the auctioneer's revenue for both the m^{th} and $(m+1)^{\text{th}}$ price auctions. In particular, the expected revenue ER_α^I in an m^{th} price auction with N bidders, whose valuations are drawn from $F(u)$, when they all bid according to function $g_\alpha^I(u)$ is:

$$ER_\alpha^I = m \cdot \int_{u_l}^{u_h} g_\alpha^I(\omega) \cdot \Psi'_m(\omega) \cdot d\omega \quad (8)$$

while the expected revenue ER_α^{II} for the equivalent $(m+1)^{\text{th}}$ price auction, when the bidders bid according to function $g_\alpha^{II}(u)$ is:

$$ER_\alpha^{II} = m \cdot \int_{u_l}^{u_h} g_\alpha^{II}(\omega) \cdot \Psi'_{m+1}(\omega) \cdot d\omega \quad (9)$$

where $\Psi_k(x) = \sum_{i=0}^{k-1} C(N, i) \cdot (F(x))^{N-i} \cdot (1 - F(x))^i$.

Using these equations we examine the expected revenue, when the number N of participating bidders changes (figure 1)⁷, and when the competition coefficient a is varied (figure 2). For $\alpha = 0$ both expected revenues are equal, as we expected given that the Revenue Equivalence Theorem applies in that case. In [2], it is shown that the presence of the coefficient α , when $\alpha \neq 0$ and $\alpha \neq 1$, leads to the second price auction yielding more profit for the auctioneer than the equivalent first price auction ($\forall \alpha, ER_\alpha^I \leq ER_\alpha^{II}$), and also that as α increases, so does the expected revenue. From our experiments, we concur that as α increases, so does the expected revenue. However, it is no longer the case that $ER_\alpha^I \leq ER_\alpha^{II}$. Based on our observations, we know that for α close to 0, it is $ER_\alpha^I \leq ER_\alpha^{II}$, and for α close to 1, it is $ER_\alpha^I \geq ER_\alpha^{II}$. Given this, we believe that the following holds:

⁷It should be noted that the different cases in figure 1 (and this is also the case for the different experiments in figures 3 and 4) are denoted by different symbols ($\times, \cdot, +, *$) in order to make it easy to distinguish them. The expected gain in figures 1 and 2 is computed accurately using numerical integration.

CLAIM 1. For the same coefficient α , when the bidders use the symmetric Bayes-Nash equilibria strategies, the m^{th} and $(m+1)^{\text{th}}$ price auctions yield the same expected revenue, when $\alpha = 0$ or $\alpha = \frac{1}{m}$, while, for $0 < \alpha < \frac{1}{m}$, the $(m+1)^{\text{th}}$ price auction yields more revenue, and for $\alpha > \frac{1}{m}$, the m^{th} price auction yields more revenue.

In fact, we expect the proof for the case when $\alpha \in [0, \frac{1}{m}]$ to be very similar to the proof of the theorem for the single unit case (when the expected revenue of the first price auction is no more than that of the second price auction), presented in [2]. To prove the claim when $\alpha > \frac{1}{m}$, will require using the fact that the two revenues are equal for $\alpha = \frac{1}{m}$ and the equilibrium strategies that we computed, and is left for future work.

6. SIMULATIONS

In this section we present the results of simulations that were conducted to empirically verify the validity of the computed equilibria and to analyze the profits derived by various agents when pitted against other agents. In figure 3 we present the expected gain⁸ for a simulation of an m^{th} price auction with $m = 2$ items being sold, for N participating bidders ($N = 3, 6$) for varying values of the coefficient α . In each case we performed five simulations: (i) strategy $g_\alpha(u)$ against itself, (ii) strategy $g_\alpha(u)$ against strategy $g_0(u)$ (which is a self-interested agent), and (iii) strategy $g_\alpha(u)$ against strategy $g_{1/N}(u)$ (which is an agent maximizing the relative difference between its own profit to that of any competitor). For the latter two cases, we conducted two simulations, one in which $(N - 1)$ of the agents used strategy $g_\alpha(u)$ and one in which only one agent used strategy $g_\alpha(u)$. We also did the same simulations for the case of an

⁸In figures 3 and 4, the expected gain is computed by discretizing the space of possible valuations for all agents (some billions of cases to examine) and then computing a near accurate (estimated error is well under 0.01%) estimate of the expected gain as the average gain from all these cases.




Figure 4: Examination of the expected profit in an $(m + 1)^{th}$ price auction ($m = 2$) with $N = 4$ participating bidders. For various α coefficients the performance of agents using the $g_\alpha(u)$ strategy is compared against agents using the $g_0(u)$ and $g_{1/N}(u)$ strategies.

$(m + 1)^{th}$ price auction, with $N = 4$ participating bidders and the results are presented in figure 4.

There are a number of observations that can be made from these simulations. When more aggressive agents⁹ participate in the auction, the profits for all agents are decreased. Furthermore, we verify that each strategy $g_\alpha(u)$ does constitute a best response to itself, when the objective function is equal to the weighted difference of the profits with weight equal to α . Even for $\alpha = 1$, when the profit is lowest, the agent does what is best, because the opponents' profit is also lowest. However, we further observe that in cases where $(N - 1)$ of these aggressive agents using strategy $g_\alpha(u)$ participate, their profit is almost the same as when all of them use $g_\alpha(u)$, while the opponent which is using a less aggressive strategy $g_{\alpha'}(u)$ ($\alpha' < \alpha$), obtains a much higher profit. Our observations are different in the case that there is only one aggressive agent. In this case, the aggressive agents beat the agents using $g_0(u)$ quite often, although this depends on the relative number of bidders N to the number m of items sold, with higher N or lower m reducing this benefit. Moreover, they even manage to achieve a score which is reasonably close to that of the agents using strategy $g_{1/N}(u)$. These latter agents (using strategy $g_{1/N}(u)$) always obtain a higher profit than their opponents, which is not entirely surprising given that they try to maximize the profit difference between themselves and any opponent. They also tend to get relatively good profit, no matter what agents they compete against, which indicates that their performance is relatively robust under most (if not all) possible opposition. This observation and the fact that, in direct competition, which is the most common case in real world scenarios, they

⁹These are agents using strategy $g_\alpha(u)$, where the coefficient α takes high values (i.e. close to 1), and thus they place bids at higher prices.

always outperform their opponents, suggest that we should consider using this strategy unless there is a strong reason to believe that the coefficient α has some different value.

7. CONCLUSIONS

We examined the behavior of bidding agents that wish to maximize a weighted difference of their own gain to that of their competitors. We provided, for the first time, the symmetric Bayes-Nash equilibria that exist for the standard multi-unit auctions, namely both the m^{th} and $(m + 1)^{th}$ price sealed bid auctions, thus extending the existing state-of-the-art results, which were analyzing only single unit auctions. We also observed a number of differences between the single and multi-unit cases; the most notable of which is that for $\alpha > \frac{1}{m}$, the equilibrium strategy for the m^{th} price auction becomes to bid more than the true valuation. We further proved that allowing the bidders to bid for multiple items, that they don't need, in order to deprive them from the competition, does not change the equilibrium strategies. Then we examined how the auctioneer's revenue is affected. We showed that revenue equivalence does not hold, and in fact the expected revenue of the $(m + 1)^{th}$ price auction is higher than that of the equivalent m^{th} price auction for certain values of α , while for others the opposite is true. Finally, we conducted simulations to empirically verify the validity of our computed strategies. Indeed we observed that the agent using $a = \frac{1}{N}$ (i.e. trying to maximize the profit difference between itself and any opponent) outperforms its competitors, when conducting a comparison of the relative profits.

There are a number of unresolved issues in this paper. In particular we plan to examine the revenue analysis and prove the claim that we made in section 5. We are also currently working towards generating equilibria for multi-demand auctions when $\alpha = 0$ and then will plan to extend these results to any value of α .

8. ACKNOWLEDGMENTS

We would like to thank Alex Rogers and Radjeep Dash at the University of Southampton for comments on this work. 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).

9. REFERENCES

- [1] S. Brainov. The role and the impact of preferences on multiagent interaction. In *ATAL'99*, pages 349–363, 1999.
- [2] F. Brandt, T. Sandholm, and Y. Shoham. Spiteful bidding in sealed-bid auctions. In *IJCAI-07*, pages 1207–1214, 2007.
- [3] V. Krishna. *Auction theory*. Academic Press, 2002.
- [4] D. K. Levine. Modeling altruism and spitefulness in experiments. *Review of Economic Dynamics*, 1:593–622, 1998.
- [5] E. Maasland and S. Onderstal. Auctions with financial externalities. Working paper, 2003.
- [6] A. Mas-Colell, M. D. Whinston, and J. R. Green. *Microeconomic theory*. Oxford University Press, 1995.
- [7] J. Morgan, K. Steiglitz, and G. Reis. The spite motive and equilibrium behavior in auctions. *Contributions to Economic Analysis & Policy: Vol. 2(1)*, 2003.
- [8] J. A. Rice. *Mathematical Statistics and Data Analysis*. Duxbury Press, California, 1995.
- [9] J. Sobel. Interdependent preferences and reciprocity. *Journal of Economic Literature: Vol. XLIII*, pages 392–436, June 2005.
- [10] I. A. Vetsikas, N. R. Jennings, and B. Selman. Generating Bayes-Nash equilibria to design autonomous trading agents. In *IJCAI-07*, pages 1543–1550, 2007.

APPENDIX

A. PROOFS OF MAIN THEOREMS

We use the following lemma to simplify the differential equations in the proofs:

LEMMA 2. For all N, m , such that $N \geq m$ the following equations hold:

$$\begin{aligned} \frac{\Phi'_m(u_i)}{(\Phi_m(u_i) - \Phi_{m-1}(u_i))} &= (N - m) \cdot \frac{F'(u_i)}{F(u_i)} \\ \frac{\Phi'_m(u_i)}{(\Phi_{m+1}(u_i) - \Phi_m(u_i))} &= m \cdot \frac{F'(u_i)}{1 - F(u_i)} \end{aligned}$$

PROOF. Computing the derivative of $\Phi_m(u_i)$ using equation 1 leaves only one term at the end and then it's easy to prove the first equation. Then to prove the second one, observe that $\frac{\Phi_m(u_i) - \Phi_{m-1}(u_i)}{\Phi_{m+1}(u_i) - \Phi_m(u_i)} = \frac{m}{N-m} \cdot \frac{F(u_i)}{1-F(u_i)}$. \square

A.1 Proof of Theorem 1

We assume that all bidders except i follow the strategy $g_\alpha(u)$. Bidder i bids v_i , the bid that maximizes his objective function on expectation. Let C be the sum (on expectation) of the top $(m-1)$ highest opponent valuations. Since in all cases that we will examine, whether bidder i wins or not, we know that the opponents with the top $(m-1)$ valuations will win an item, we know that they will gain this amount C from doing so. Now:

$$C = E\left(\sum_{j=1}^{m-1} U^{(j)}\right) = \sum_{j=1}^{m-1} E(U^{(j)}) = \sum_{j=1}^{m-1} \int_{u_1}^{u_h} \omega \cdot \Phi'_j(\omega) \cdot d\omega$$

is a constant and does not depend on the bid v_i . We need to consider the following three cases:

(i) When $B^{(m)} > v_i$, bidder i does not win any item and the closing price is $B^{(m)}$. Therefore bidder i 's gain is 0 and the opponents make a gain from gaining an extra item (the m^{th}), in addition to the $(m-1)$ items that they always win (this was counted in the computation of constant C), but they also must make total payments of $m \cdot B^{(m)}$. The total additional¹⁰ expected utility for bidder i in this case is hence:

$$\Delta U_1 = \alpha \cdot \int_{v_i}^{\infty} (m \cdot \omega - g_\alpha^{-1}(\omega)) \cdot \frac{d}{d\omega} (\Phi_m(g_\alpha^{-1}(\omega))) \cdot d\omega$$

(ii) When $B^{(m-1)} > v_i \geq B^{(m)}$, bidder i wins an item and the closing price is v_i . Therefore bidder i 's gain is $u_i - v_i$ and the opponents pay $(m-1) \cdot v_i$ for the items that they win. The total additional expected utility for bidder i is:

$$\Delta U_2 = ((1 - \alpha) \cdot (u_i - v_i) + \alpha \cdot (m - 1) \cdot v_i) \cdot (\Phi_m(g_\alpha^{-1}(v_i)) - \Phi_{m-1}(g_\alpha^{-1}(v_i)))$$

(iii) When $v_i \geq B^{(m-1)}$, bidder i wins an item and the closing price is $B^{(m-1)}$. Therefore bidder i 's gain is $u_i - B^{(m-1)}$ and the opponents must pay $(m-1) \cdot B^{(m-1)}$ for the items that they purchase. The total additional expected utility for bidder i in this case is:

$$\Delta U_3 = \int_0^{v_i} ((1 - \alpha) \cdot (u_i - \omega) + \alpha \cdot (m - 1) \cdot \omega) \cdot \frac{d}{d\omega} (\Phi_{m-1}(g_\alpha^{-1}(\omega))) \cdot d\omega$$

¹⁰We mean additional to the fact that the agent always loses utility $\alpha \cdot C$, since its opponents always gain a value C from the top $(m-1)$ items.

The total expected utility for bidder i when considering all possibilities is therefore $EU_i = -\alpha \cdot C + \Delta U_1 + \Delta U_2 + \Delta U_3$. This implies:

$$\begin{aligned} EU_i &= \alpha \cdot \int_{v_i}^{\infty} (m \cdot \omega - g_\alpha^{-1}(\omega)) \cdot \frac{\Phi'_m(g_\alpha^{-1}(\omega))}{g'_\alpha(g_\alpha^{-1}(\omega))} \cdot d\omega \\ &+ ((1 - \alpha) \cdot u_i + (\alpha \cdot m - 1) \cdot v_i) \cdot (\Phi_m(g_\alpha^{-1}(v_i)) - \Phi_{m-1}(g_\alpha^{-1}(v_i))) \\ &+ \int_0^{v_i} ((1 - \alpha) \cdot u_i + (\alpha \cdot m - 1) \cdot \omega) \cdot \frac{\Phi'_{m-1}(g_\alpha^{-1}(\omega))}{g'_\alpha(g_\alpha^{-1}(\omega))} \cdot d\omega \\ &- \alpha \cdot \sum_{j=1}^{m-1} \int_{u_1}^{u_h} \omega \cdot \Phi'_j(\omega) \cdot d\omega \end{aligned} \quad (10)$$

To find the value of v_i that maximizes equation 10, we set $\frac{dEU_i}{dv_i} = 0$. If strategy $g_\alpha(u)$ gives the equilibrium strategy, then it must be the case that the value of v_i that maximizes the total utility is given by $g_\alpha(u)$ (i.e. that $v_i = g_\alpha(u_i)$). By substituting this into the equation derived from setting $\frac{dEU_i}{dv_i} = 0$, we finally get the differential equation:

$$(u_i - g_\alpha(u_i)) \cdot \frac{\Phi'_m(u_i)}{g'_\alpha(u_i)} = (1 - \alpha \cdot m) \cdot (\Phi_m(u_i) - \Phi_{m-1}(u_i)) \quad (11)$$

To simplify this equation we use lemma 2 to get:

$$g'_\alpha(u_i) = (u_i - g_\alpha(u_i)) \cdot \frac{N - m}{1 - \alpha \cdot m} \cdot \frac{F'(u_i)}{F(u_i)} \quad (12)$$

The solution of this equation is:

$$g_\alpha(u) = u - (F(u))^{-\frac{N-m}{1-\alpha \cdot m}} \cdot \int_c^u (F(z))^{-\frac{N-m}{1-\alpha \cdot m}} \cdot dz \quad (13)$$

where c depends on the boundary condition.

To select the appropriate boundary condition one should note that terms $(u_i - g_\alpha(u_i))$ and $(1 - \alpha \cdot m)$ have the same sign, since all other terms in equation 12 are positive. Therefore:

(i) if $\alpha \cdot m < 1$, then $u_i - g_\alpha(u_i) > 0$ and the boundary condition is $g_\alpha(u_i) = u_i$, and the resulting strategy is given by equation 2;

(ii) if $\alpha \cdot m = 1$, then $u_i - g_\alpha(u_i) = 0 \Rightarrow g_\alpha(u_i) = u_i$ (equation 3), and

(iii) if $\alpha \cdot m > 1$, then $u_i - g_\alpha(u_i) < 0$ and the boundary condition is $g_\alpha(u_h) = u_h$, and the resulting strategy is given by equation 4.

A.2 Proof of Theorem 2

We assume that all bidders except i follow the strategy $g_\alpha(u)$. In this case, bidder i , is allowed to bid for any number of items. In order to show a contradiction we assume that he bids for 2 items, and the bids are equal to v_i and \tilde{v}_i . Without loss of generality, we assume that $v_i \geq \tilde{v}_i$.

By considering several possible cases that can occur depending on the relative values of v_i and \tilde{v}_i to the bids of the opponents, we conclude that the total expected utility in this case is:

$$\widetilde{EU}_i = EU_i + \Delta EU_i$$

where EU_i depends only on v_i (not \tilde{v}_i) and is given by equation 10 and ΔEU_i depends only on \tilde{v}_i (not v_i). Therefore we need to maximize both terms in order to maximize the total utility. The value that maximizes the first term is $v_i = g_\alpha(u_i)$. The second term is equal to:

$$\begin{aligned} \Delta EU_i = & - \int_0^{\tilde{v}_i} ((1-\alpha) \cdot u_i + (\alpha \cdot m - 1) \cdot \omega) \cdot \frac{\Phi'_{m-1}(g_\alpha^{-1}(\omega))}{g'_\alpha(g_\alpha^{-1}(\omega))} \cdot d\omega \\ & + ((1-\alpha) \cdot u_i + (\alpha \cdot m - 2) \cdot \tilde{v}_i) \cdot (\Phi_{m-1}(g_\alpha^{-1}(\tilde{v}_i)) - \Phi_{m-2}(g_\alpha^{-1}(\tilde{v}_i))) \\ & + \int_0^{\tilde{v}_i} ((1-\alpha) \cdot u_i + (\alpha \cdot m - 2) \cdot \omega) \cdot \frac{\Phi'_{m-2}(g_\alpha^{-1}(\omega))}{g'_\alpha(g_\alpha^{-1}(\omega))} \cdot d\omega \\ & + \alpha \int_0^{\tilde{v}_i} g_\alpha^{-1}(\omega) \cdot \frac{\Phi'_{m-1}(g_\alpha^{-1}(\omega))}{g'_\alpha(g_\alpha^{-1}(\omega))} \cdot d\omega \end{aligned}$$

To simplify this equation, let us set $\tau = g_\alpha^{-1}(\tilde{v}_i)$. Then $u_i \leq \tau \leq u_i$. Hence ΔEU_i can be expressed as a function of τ and α :

$$\Delta EU_i(\tau, \alpha) = \int_{u_i}^\tau (\alpha \cdot \omega - g_\alpha(\omega)) \cdot \Phi'_{m-1}(\omega) \cdot d\omega \quad (14)$$

$$+ (\alpha \cdot m - 2) \cdot \int_{u_i}^\tau g'_\alpha(\omega) \cdot (\Phi_{m-1}(\omega) - \Phi_{m-2}(\omega)) \cdot d\omega$$

If $\alpha \cdot m \leq 2$ it is easy to see that both integrals are negative, when $\tau > u_i$, so the maximizing value is $\tau = u_i$, which means that one should not bid for the second item.¹¹

If $\alpha \cdot m > 2$, then using lemma 2 and equation 12, equation 14 can be rewritten as:

$$\Delta EU_i(\tau, \alpha) = \int_{u_i}^\tau \left(-(N - m + 1)(g_\alpha(\omega) - \alpha \cdot \omega) \right) \quad (15)$$

$$+ (N - m) \cdot \frac{\alpha m - 2}{\alpha m - 1} (g_\alpha(\omega) - \omega) \frac{F'(\omega)}{F(\omega)} (\Phi_{m-1} - \Phi_{m-2})(\omega) \cdot d\omega$$

So even in this case $\forall \alpha > \frac{2}{m}, \tau$:

$$\Delta EU_i(\tau, \alpha) \leq \int_{u_i}^\tau \left(-(N - m + 1)(g_\alpha(\omega) - \omega) \right)$$

$$+ (N - m) \cdot \frac{\alpha m - 2}{\alpha m - 1} (g_\alpha(\omega) - \omega) \frac{F'(\omega)}{F(\omega)} (\Phi_{m-1} - \Phi_{m-2})(\omega) \cdot d\omega$$

$$\Rightarrow \Delta EU_i(\tau, \alpha) \leq 0$$

This shows that it is not beneficial to bid for a second item,¹² when the opponents use the equilibrium strategy $g_\alpha(u)$ of theorem 1. Thus the best response to strategy $g_\alpha(u)$ is bidding according to $g_\alpha(u)$ for exactly one item.

A.3 Proof of Theorem 3

We assume that all bidders except i follow the strategy $g_\alpha(u)$. Bidder i bids v_i , the bid that maximizes his objective function on expectation. Let C be the sum (on expectation) of the top m highest opponent valuations. In most cases that we will examine (when bidder i does not win), the opponents with the top m valuations will win an item, and they will gain this amount C from doing so. In the case that bidder i outbids the competition (case iii below), then we will subtract the m^{th} valuation from the opponents' gain to compensate. Now:

¹¹To be more precise, $\tau = u_i$ means to bid as if it had the minimum valuation for the item, which in practice is equivalent to not bidding.

¹²And since bidding for a second is not beneficial, bidding for a second and third and/or more, is not beneficial either.

$$C = E\left(\sum_{j=1}^m U^{(j)}\right) = \sum_{j=1}^m E(U^{(j)}) = \sum_{j=1}^m \int_{u_i}^{u_h} \omega \cdot \Phi'_j(\omega) \cdot d\omega$$

is a constant and does not depend on the bid v_i . We need to consider the following three cases:

(i) When $B^{(m+1)} > v_i$, bidder i does not win any item and the closing price is $B^{(m+1)}$. Therefore bidder i 's gain is 0 and the opponents must make total payments of $m \cdot B^{(m+1)}$. The total additional expected utility for bidder i in this case is hence:

$$\Delta U_1 = \alpha \cdot \int_{v_i}^\infty m \cdot \omega \cdot \frac{d}{d\omega} (\Phi_{m+1}(g_\alpha^{-1}(\omega))) \cdot d\omega$$

(ii) When $B^{(m)} > v_i \geq B^{(m+1)}$, bidder i does not win and the closing price is v_i . Therefore bidder i 's gain is 0 and the opponents must pay $m \cdot v_i$ for the items that they purchase. The total additional expected utility for bidder i in this case is:

$$\Delta U_2 = \alpha \cdot m \cdot v_i \cdot (\Phi_{m+1}(g_\alpha^{-1}(v_i)) - \Phi_m(g_\alpha^{-1}(v_i)))$$

(iii) When $v_i \geq B^{(m)}$, bidder i wins an item and the closing price is $B^{(m)}$. Therefore bidder i 's gain is $u_i - B^{(m)}$ and the opponents must pay $(m - 1) \cdot B^{(m)}$ for the items that they purchase. We also need to subtract the m^{th} highest valuation $U^{(m)} = g^{-1}(B^{(m)})$ from the total gain of the opponents, since they only won $(m - 1)$ items. The total additional expected utility for bidder i in this case is:

$$\begin{aligned} \Delta U_3 = & \int_0^{v_i} ((1-\alpha) \cdot (u_i - \omega) + \alpha \cdot ((m-1) \cdot \omega + g_\alpha^{-1}(\omega))) \cdot \\ & \frac{d}{d\omega} (\Phi_m(g_\alpha^{-1}(\omega))) \cdot d\omega \end{aligned}$$

The total expected utility for bidder i when considering all possibilities is therefore $EU_i = -\alpha \cdot C + \Delta U_1 + \Delta U_2 + \Delta U_3$. This implies:

$$\begin{aligned} EU_i = & \alpha \cdot \int_{v_i}^\infty m \cdot \omega \cdot \frac{\Phi'_{m+1}(g_\alpha^{-1}(\omega))}{g'_\alpha(g_\alpha^{-1}(\omega))} \cdot d\omega \\ & + \alpha \cdot m \cdot v_i \cdot (\Phi_{m+1}(g_\alpha^{-1}(v_i)) - \Phi_m(g_\alpha^{-1}(v_i))) \\ & + \int_0^{v_i} ((1-\alpha) \cdot u_i + (\alpha \cdot m - 1) \cdot \omega + \alpha \cdot g_\alpha^{-1}(\omega)) \cdot \frac{\Phi'_m(g_\alpha^{-1}(\omega))}{g'_\alpha(g_\alpha^{-1}(\omega))} \cdot d\omega \\ & - \alpha \cdot \sum_{j=1}^m \int_{u_i}^{u_h} \omega \cdot \Phi'_j(\omega) \cdot d\omega \quad (16) \end{aligned}$$

To find the value of v_i that maximizes equation 16, we set $\frac{dEU_i}{dv_i} = 0$. If strategy $g_\alpha(u)$ gives the equilibrium strategy, then it must be the case that the value v_i that maximizes the total utility is given by $g_\alpha(u)$ (i.e. that $v_i = g_\alpha(u_i)$). By substituting this into the equation derived from setting $\frac{dEU_i}{dv_i} = 0$, we finally get the differential equation:

$$(u_i - g_\alpha(u_i)) \cdot \frac{\Phi'_m(u_i)}{g'_\alpha(u_i)} = -\alpha \cdot m \cdot (\Phi_{m+1}(u_i) - \Phi_m(u_i)) \quad (17)$$

To simplify this equation we use lemma 2 to get:

$$g'_\alpha(u_i) = -\frac{1}{\alpha} \cdot (u_i - g_\alpha(u_i)) \cdot \frac{F'(u_i)}{1 - F(u_i)} \quad (18)$$

The solution of this equation is:

$$g_\alpha(u) = u - (1 - F(u))^{-\frac{1}{\alpha}} \cdot \int_c^u (1 - F(z))^{\frac{1}{\alpha}} \cdot dz \quad (19)$$

where c depends on the boundary condition.

To select the appropriate boundary condition one should note that term $(u_i - g_\alpha(u_i))$ is negative since $-\alpha \cdot m$ is negative and all the other terms in equation 12 are positive. Therefore the boundary condition is $g_\alpha(u_h) = u_h$, and the resulting strategy is given by equation 6.