Equilibrium Strategies for Multi-unit Sealed-bid Auctions with Multi-unit Demand Bidders

Ioannis A. Vetsikas
Institute of Informatics and Telecommunications,
National Center for Scientific Research (NCSR) "Demokritos"
P. Grigoriou and Neapoleos, Ag. Paraskevi, 15310 Attiki, Greece
ivetsikas@iit.demokritos.gr

ABSTRACT

In most of the existing literature on multi-unit auctions, i.e. auctions selling several identical goods together, it is assumed that bidders demand a single item. Yet this assumption is not valid in most practical auction settings, as often bidders wish to purchase multiple goods. Computing equilibrium strategies in multi-unit uniform-price auctions for bidders with multi-unit demand is an open problem for almost two decades. It is known that they exist in pure strategies, but not how to compute them. Our work addresses this key open problem, when there are no complementarities. More specifically, we examine a model where each bidder's value for the units beyond the first are computed by multiplying the value for the first unit of the good (the most desired one) by preset weights, and then generalize this model by allowing these weights to be different for each participating bidder. We characterize the equilibria and compute equilibrium strategies for both m^{th} and $(m+1)^{th}$ price sealed-bid auctions; then we give some examples examining the properties of these strategies in the process. We conduct experiments that show up to 25% improvement in the performance of trading agents using these strategies as opposed to some heuristic strategies previously used.

Categories and Subject Descriptors

I.2.11 [**ARTIFICIAL INTELLIGENCE**]: Distributed Artificial Intelligence

General Terms

Economics

Keywords

multi-unit sealed-bid auctions, multi-unit demand bidders, game theory, equilibrium strategies, strategic demand reduction

1. INTRODUCTION

Auctions are nowadays used to trade a wide range of goods and have become quite well known to the general public with the advent of online auction houses, such as eBay. A lot of the auction theory literature focuses on single unit auctions, and it is assumed that any results translate to the multi-unit auction case, meaning auctions that sell several identical goods. This is indeed mostly true, if the participating bidders only desire to purchase a single unit of these goods.[11] However, in a number of real world scenarios, bidders want to purchase multiple units of these goods. We say that a bidder has *multi-unit demand* when she desires to buy multiple units of the good sold in a multi-unit auction. Examples of such auctions are the US treasury bills auctions and the FCC spectrum auctions. In fact, the latter have reinvigorated the research on auctions with multi-unit demand bidders in the mid 90's. Before reviewing some of the literature on this problem, we stress that the US treasury bills auctions and the FCC spectrum auctions are sealed-bid uniform-price auctions.[1] In fact, uniform pricing, meaning that all the winners pay the same price, are by far the most common auction pricing model; and even open-cry, ascending auctions have been shown under certain conditions, such as risk neutrality, to be (weakly) strategically equivalent to sealed-bid auctions.[11] This motivates why most of the related work examines uniform-price sealed-bid auctions, and why we examine the same setting in this paper.

Now, a lot of the related work has looked into particular properties of these auctions, and in particular what is called (strategic) *demand reduction*. The idea is that sometimes it can be beneficial to a bidder to reduce her demand and request less units than she desires, because that would reduce the competition for these items and therefore reduce the price paid by everyone: since bidders reduce their demand, the savings in the price that they pay for the remaining items could outweigh the potential loss in valuations from not obtaining the non-requested units. This effect has been studied in the FCC spectrum auctions and other settings, see e.g. [21, 6]. This effect means that the auction is no longer efficient as bidders with lower valuations can win, since some bidders may decide not to bid for their whole demand even if they have relatively high valuations.

Another source of inefficiency [1] is the the fact that bidders will *shade* their bids, which means that they will strategically reduce any bids after the first and this reduction increases as each bidder places bids for lower valued items; this will also cause bidders with lower valuations for their first items to win against bidders with higher valuations for their last items, hence the inefficiency. Why does this strategic re-

Appears in: *Alessio Lomuscio, Paul Scerri, Ana Bazzan, and Michael Huhns (eds.), Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), May 5-9, 2014, Paris, France.*

[1]We do acknowledge that the FCC auctions had multiple rounds which changes the setting somewhat.

duction of bids take place and why is it more pronounced for the last bids placed by each bidder? The answer is that the closing price can be set by any bid and therefore a lower bid placed for one of the least desirable (last) items often sets the price paid for all items that a bidder wins. Therefore, it is in her interest to reduce her last bids. If the bidder does not have market power and cannot affect the final prices then a bidder need not shade her bids.[14, 15] However, in most cases a bidder will affect the prices and therefore bid shading will take place.

Most of the existing literature has provided tools that can analyze simple settings, with usually 2 bidders only or even only one bidder with multi-unit demand and a number of other bidders with single-unit demand, see e.g. [4, 5, 16]. In addition to this limitation, other work has conducted empirical studies as well.[8, 9]. There has been a fair deal more work that we cannot review here due to the space limitations; for a recent survey of multi-unit auctions see [10]. We are aware of no work, though, that has been able to provide tools or characterize the equilibria for more general settings.

Hence, *computing equilibrium strategies in multi-unit uniform-price auctions for bidders with multi-unit demand is an open problem for almost two decades.* McAdams[12, 13] has shown that such equilibria (isotone equilibria to be more specific) exist in pure strategies, but not how to compute them. In fact, their proof is based on fixed point theorems, where continuity would be violated if the equilibrium does not exist; we do not believe that this analysis can be used in any way to compute the equilibria and the fact that no one has been able to compute these equilibria for a decade since reinforces this belief. Therefore, our work addresses this key open problem (using a different analysis), when there are no complementarities between the valuations that each bidder has for the desired items. This is the main contribution of this paper. More specifically:

- We initially examine certain properties of the equilibrium strategies which hold generally, for all the models we will use (Section 3.1), and which will be used in the computation of the equilibria. The only property that we do not prove is that the bids are strictly decreasing in most cases, the exceptions being demand reduction and insufficient competition.
- Then, we examine a model where each bidder's valuation for the units beyond the first are computed by multiplying the value for the first unit of the good (the most desired one) by preset weights. For example, if each bidder wishes to purchase up to 3 units, then if the value for the first one is v then the value for the second and third ones could be 70% and 50% of v respectively. What differentiates the bidders are the values v, while the percentages are the same for all. We give an algorithm for computing the equilibrium strategies for both m^{th} and $(m+1)^{th}$ price auctions (in Sections 3.2 and 3.3 respectively). Our algorithm is able to predict demand reduction as well as bid shading. We then give some examples to clarify certain points of the equilibrium computation.
- Afterwards, in Section 4, we generalize the model to cover all possible valuations (provided that there are no complementarities). To do this, we allow each participating bidder to have different weights for her valuations. This means that one bidder's second valuation might be 50% of her first one and another one might have this be at 100%, etc. We show how to generalize the algorithms to cover this general model as well. We also give examples which describe the properties of the equilibrium. For example, that the bids for the first items are only indirectly affected by the weights for the latter items.
- We conduct experiments that show an improvement of usually 5% to 25% in most cases in the performance of trading agents using these strategies (with only few cases showing a drop in profit of up to 1%); we look at the literature regarding the Trading Agent Competition [22] in order to select heuristic strategies that we compare against.

2. THE MODEL

In this section we formally describe the auction models used in the paper and present the notation used. We will compute Bayes-Nash equilibria for sealed-bid auctions where $m \geq 1$ identical items are being sold at a uniform price. The two most common auction settings in this context are the m^{th} and $(m+1)^{th}$ price auctions, in which the top m bids placed will be awarded one item each at a price equal to the m^{th} and $(m+1)^{th}$ highest bid respectively; in this paper, we compute the equilibria for both these auction variants. Our model follows in general the models used when computing equilibria in the auction theory literature [11].

More specifically, we assume that N indistinguishable bidders participate in the auction. Each bidder i wants to purchase up to λ units of the item being sold, i.e. has a λ-unit demand. The restriction that $\lambda \leq m$ exists, meaning that a bidder cannot demand more items than the number available for sale and also $\lambda N > m$ as otherwise there are not enough bidders to create competition and the auction is trivial! We also assume that bidders are risk neutral, which means that they only care about maximizing their expected profit, and that their utilities do not have externalities, meaning that they only care about maximizing their own profit without caring at all how their actions will affect the profit of the other bidders they compete against.

We assume that each bidder i has private valuations $v_i^1, \dots, v_i^\lambda$ for these items; these valuations correspond to the marginal profit from obtaining the first up to λ^{th} unit. The marginal profit of the j^{th} items is defined as the profit a bidder obtains by acquiring j items minus the profit from acquiring $(j-1)$ items. In this way we can define any possible case regarding the valuations. For example in the case that the goods are partially substitutable, then $v_i^1 \geq \dots \geq v_i^\lambda$, as the value of getting the first item is higher than the additional value from getting the second etc. They could also be additive (when $v_i^1 = \dots = v_i^\lambda$), which means that the value of getting $\lambda' \leq \lambda$ items is $\lambda' v_i^1$. On the other hand, it could also denote complementarities, when $\exists j : v_i^j > v_i^{j-1}$. For example, if the bidder wants exactly three items (and gets value β from them), and has no desire to obtain less items, then $v_i^1 = v_i^2 = 0$ and $v_i^3 = \beta$. However, in this paper, we do not examine complementarities. Note that a bidder's utility from wining k items in the auction is the sum of the marginal profits $v_i^1, \dots, v_i^k$ for the first k items.

Now, we choose to represent the vector of valuations as a vector of weights $\overline{\alpha_i} = \{\alpha_i^1, \dots, \alpha_i^\lambda\}$, multiplied by valuation v_i, which means that $\forall j : v_i^j = \alpha_i^j v_i$. In the analysis we will present in this paper, this always will be the first valuation v_i^1 w.l.o.g.; this is done so that the results are easier to follow, but it is not an assumption that need to be made for our models and subsequent analysis to work. Using this notation, the profit (or utility) U_i of bidder i when the closing price of the auction is π is:

$$U_i = \begin{cases} (\sum_{j=1}^{k} \alpha_i^j) v_i - k\pi & \text{if bidder } i \text{ wins } k \text{ items,} \\ 0 & \text{if bidder } i \text{ does not win any items.} \end{cases}$$

We make the standard assumption from game theory that the valuations of different bidders are independent of each other which means that the valuations v_i of the bidders are i.i.d. random variables drawn from a known prior distribution with cumulative distribution function (cdf) $F(v)$, which is the same for all bidders. Furthermore, we examine two different models regarding the valuations of the bidders:

- We initially assume that the vector of weights $\overline{\alpha}$ is common for all bidders; this model will be henceforth called the **"common weight vector model"**.

- We then examine the general model, where each bidder i has a different vector of weights $\overline{\alpha_i}$, which is private knowledge of each bidder. We assume that the possible weight vectors $\overline{\alpha}$ have a predetermined, known probability distribution $h(\overline{\alpha})$.[2] This model will be henceforth called the (general) **"asymmetric weight vector model"**.[3]

We will consider bidding strategies $\overline{g}_{\overline{\alpha}}(v)$ defined over the support of distribution $F(v)$, meaning any value v such that $f(v) > 0$, where f is the pdf (probability density function) of distribution F. These strategies return a vector of bids $\overline{g}_{\overline{\alpha}}(v) = \{g_{\overline{\alpha},1}(v), \ldots, g_{\overline{\alpha},\lambda}(v)\}$. Now, each bidder i tries to maximize his profit by submitting a vector of λ bids $b_i^1, \ldots, b_i^\lambda$ for the desired items, where $b_i^1 \geq \ldots \geq b_i^\lambda \geq 0$; these bids are ordered because the auction mechanism will sort all bids and select the highest ones as winners. At the equilibrium, these "optimal" bids for the agent are those given from the equilibrium strategies $\overline{g}_{\overline{\alpha_i}}(v_i)$. Thus the bids maximizing the objective function should be $b_i^j = g_{\overline{\alpha_i},j}(v_i)$.

In the next sections, we characterize the equilibria first for the common weight vector model and then for the general model.

3. COMMON WEIGHT VECTOR MODEL

In this section, we compute the symmetric equilibrium strategies that the bidders adopt in the setting where all bidders have the same vector of weights, which will be denoted as $\overline{\alpha}$. Note that as this is the same for all bidders we drop the subscript. These equilibrium strategies are symmetric meaning that bidders with the same valuation v will submit the same bids $\overline{g}_{\overline{\alpha}}(v)$. Initially, we examine some properties of the equilibrium strategies and compute the probability distributions which will subsequently be used to characterize the equilibrium strategies of the m^{th} and $(m+1)^{th}$ price auctions in Sections 3.2 and 3.3 respectively.

As the strategy of each agent-bidder depends on the strategies of the opponents, we need to first compute the probability distributions of top m order statistics of all bids placed by her opponents. In simple terms, how many items (if any) are won by the agent depends on whether his first, second, ..., and λ^{th} bids are respectively higher than the m^{th}, $(m-1)^{th}$, ..., and $(m-\lambda+1)^{th}$ highest of the bids placed by his opponents. Let us denote the k^{th} highest bid placed by opponents as $B^{(k)}$. Then, assuming that the opponents follow the equilibrium strategies, we can compute the probability distribution of this bid $\Phi_k(x) = Prob[B^{(k)} \leq x]$ as follows:

$$\Phi_k(x) = \sum_{j=1}^{k} \sum_{\substack{d_1,\ldots,d_{\min\{j,\lambda\}} \\ s.t.:\sum_l l d_l = j-1}} \frac{(N-1)!}{(N-1-\sum_l d_l)! \prod_l d_l!} (1-F(g^{-1}_{\overline{\alpha},\lambda}(x)))^{d_\lambda}$$

$$\cdot \prod_{l<\lambda} (F(g^{-1}_{\overline{\alpha},l+1}(x)) - F(g^{-1}_{\overline{\alpha},l}(x)))^{d_l} (F(g^{-1}_{\overline{\alpha},1}(x)))^{N-1-\sum_l d_l} \quad (1)$$

To explain how this complex formula is derived, consider the following: to have $B^{(1)} \leq x$ it must be that all the top bids of all opponents are also $\leq x$, and as their distribution has cdf $F(g^{-1}_{\overline{\alpha},1}(x))$, this means that $\Phi_1(x) = Prob[B^{(1)} \leq x] = F(g^{-1}_{\overline{\alpha},1}(x)))^{N-1}$. In order to have a bid x be between the $(j-1)^{th}$ and j^{th} highest opponent bids, it must be the case that it is lower than the lowest (i.e. λ^{th}) bid of d_λ opponents, between the l^{th} and $(l+1)^{th}$ bid of d_l opponents and higher than the top bid of the remaining $(N-1-\sum_l d_l)$ opponents, where $\sum_l l d_l = j-1$. Note that, since it is $d_l \geq 0, \forall l$, it must be $d_l = 0, \forall l \geq j$. There are $\binom{N-1}{d_1,\ldots,d_\lambda}$ ways to select combinations of opponents that fit this pattern for each particular combination of the values $\{d_1, \ldots, d_\lambda\}$. From these facts, we compute the probability $Prob[B^{(j)} \leq x \leq B^{(j-1)}]$ and since $Prob[B^{(j)} \leq x \leq B^{(j-1)}] = Prob[B^{(j)} \leq x] - Prob[B^{(j-1)} \leq x] = \Phi_j(x) - \Phi_{j-1}(x)$, we derive equation 1.

The derivative of equation 1 is:

$$\Phi'_k(x) = \sum_{j=1}^{k} \sum_{\substack{d_1,\ldots,d_{\min\{j,\lambda\}} \\ s.t.:\sum_l l d_l = j-1}} \frac{(N-1)!}{(N-1-\sum_l d_l)! \prod_l d_l!} (1-F(g^{-1}_{\overline{\alpha},\lambda}(x)))^{d_\lambda}$$

$$\cdot \prod_{l<\lambda} (F(g^{-1}_{\overline{\alpha},l+1}(x)) - F(g^{-1}_{\overline{\alpha},l}(x)))^{d_l} (F(g^{-1}_{\overline{\alpha},1}(x)))^{N-1-\sum_l d_l}$$

$$\Big((N-1-\sum_l d_l) \frac{F'(g^{-1}_{\overline{\alpha},1}(x)) \frac{dg^{-1}_{\overline{\alpha},1}(x)}{dx}}{F(g^{-1}_{\overline{\alpha},1}(x))} - d_\lambda \frac{F'(g^{-1}_{\overline{\alpha},\lambda}(x)) \frac{dg^{-1}_{\overline{\alpha},\lambda}(x)}{dx}}{1 - F(g^{-1}_{\overline{\alpha},\lambda}(x))}$$

$$+ \sum_{l<\lambda} d_l \frac{(F'(g^{-1}_{\overline{\alpha},l+1}(x)) \frac{dg^{-1}_{\overline{\alpha},l+1}(x)}{dx} - F'(g^{-1}_{\overline{\alpha},l}(x))) \frac{dg^{-1}_{\overline{\alpha},l}(x)}{dx}}{F'(g^{-1}_{\overline{\alpha},l+1}(x)) - F(g^{-1}_{\overline{\alpha},l}(x))} \Big) \quad (2)$$

3.1 Properties of the Equilibrium Strategies

We now present *properties of the best response strategies* for this problem. As the equilibrium strategies are the best response strategies to the equilibrium strategies themselves, these properties will be shared by the equilibrium strategies

[2]We will present the equilibria for this model for $h(\overline{\alpha})$ being a discrete distribution, i.e. that the set of possible weight vectors $\overline{\alpha}$ has a finite number of elements. However, these results can be extended in a straight forward way to continuous distributions corresponding to infinite many possible weight vectors $\overline{\alpha}$; in fact, we present such an example later.

[3]It is also possible to examine this model when the models of each bidder are known. This is in fact a little harder to solve and we will detail this in Example 5.

themselves. Note that these properties hold for the general model as well.

Property 1: The j^{th} bid b_i^j depends only on the marginal profit from obtaining the j^{th} item as well as the rank (j) of the bid; it does not depend at all on the other valuations for the other items that the bidder wishes to obtain.

To prove this property think as follows: obviously it depends on the marginal profit $v_i^j = \alpha_i^j v_i$ as, by bidding higher, the bidder increases the probability of winning j items rather than $(j-1)$, thus increasing his profit by an amount equal to the marginal profit v_i^j. At the same time, this bid may also set the price for the top j items she wins (or $(j-1)$ top items in the case of the $(m+1)^{th}$ price auction), therefore the bid b_i^j depends on this number j as well. By this last argument, the bid does not depend on the marginal profits $v_i^{j'}, j' < j$, as the actual valuations for these items do not matter (only how many they are). The bid cannot set the payment for any item after the j^{th}, therefore the bid does not depend on the marginal profits $b_i^{j'}, j' > j$.

Property 2: The bids must be ordered.

This follows immediately from the fact that the auction will order all bids from all bidders and select the m highest ones as winning bids. Each bidder will in turn get a number of items - let it be j - and her profit will be the sum of the marginal profits for the first j items. These j items were won by the j^{th} highest bid placed by this bidder, therefore whichever orders the bids are submitted, in the end they are sorted by the auction from highest to lowest. Therefore, w.l.o.g. we place the restriction that $b_i^1 \geq \ldots \geq b_i^\lambda \geq 0$, where 0 is assumed to be the price indicating that a bidder is not willing to participate in the auction.

In general, we make the following observation about the bids:

CLAIM 1. Provided that there are no complementarities and sufficient competition between the bidders, the bids from a bidder are strictly decreasing (unless they are zero): $\exists l \leq \lambda : b_i^1 > \ldots > b_i^l$ and, if $l < \lambda$, then $b_i^{l+1} = \ldots = b_i^\lambda = 0$.

This claim is based on our observations from the cases we examined while conducting this research. We do not prove this fact, as it is quite difficult for all possible opponent bids. An argument why this holds is that the competition for winning the j^{th} desired unit is less than that for winning one more, i.e. the $(j+1)^{th}$ one. Furthermore, when there are no complementarities the value from obtaining the j^{th} unit is higher than that of the next one. Therefore, it is reasonable that the bid $b_j > b_{j+1}$.

There are cases which are not covered by this claim. These are the following:

Property 3: *Demand reduction*, which means bidding 0 for the last items.This covers the case that the last bids are zero (without the valuations being zero): $b_i^{l+1} = \ldots = b_i^\lambda = 0$.

Property 4: *Insufficient competition*, which is defined as a case when a bidder is certain to win one or more units no matter the bids from the other bidders. This only happens when the number of units $m > \lambda(N-1)$, meaning that a bidder will win a unit even if all opponents win λ units (the maximum number). In some cases the bids for the first units do not matter at all, but in some cases they may set the winning price, therefore to be safe in this case, we can assume that all the top bids are equal: $b_i^1 = \ldots = b_i^{l+1}$, where $l = m - \lambda(N-1)$.

Property 5: *Complementarities* can lead to bids being equal. More specifically, a complementarity exists when $\exists k', k : k' < k \wedge \frac{\sum_{j=k'}^{k-1} v_i^j}{k-k'} < v_i^k$, because in that case obtaining the k^{th} unit will give more profit than the average marginal profit from obtaining any of the previous $(k-k')$ units. In such a case it is possible that $b_i^{k'} = \ldots = b_i^k$.

These properties are important because they will be used in the algorithms for computing the equilibrium strategies. More specifically we can now compute the necessary conditions that give the equilibrium strategies for this model. We next analyze the m^{th} and $(m+1)^{th}$ price auctions respectively in the next two sections.

3.2 m-th Price Auctions

In this section, we characterize the equilibrium strategies for the m^{th} price auction. We will start by characterizing the equilibria when there is no demand reduction (Theorem 1) and then describe the algorithm for the general case, which will use that earlier theorem. We also give a number of examples.

THEOREM 1. *The equilibrium strategy for the setting that we examine (when the bids are not equal to each other) is given by the system of differential equations:*

$$(\alpha_k v_i - g_{\overline{\alpha},k}(v_i))\Phi'_{m-k+1}(g_{\overline{\alpha},k}(v_i)) = \tag{3}$$

$$k(\Phi_{m-k+1}(g_{\overline{\alpha},k}(v_i)) - \Phi_{m-k}(g_{\overline{\alpha},k}(v_i)))$$

where Φ and Phi' are computed from equations 1 and 2. The boundary condition is $g_{\overline{\alpha},k}(v^L) = v^L$, where v^L is the lower valuation allowed, meaning the lower bound of the support of the prior distribution F.

PROOF. The bids placed by agent i are sorted from highest to lowest and thus w.l.o.g. it is $b_i^1 \geq \ldots \geq b_i^\lambda \geq 0$. The last constraint follows from the fact that the bids must be non-negative in order to be valid. Given these bids, the expected profit of the agent is:

$$EP_i(b_i^1, \ldots, b_i^\lambda) = \sum_{k=1}^{\lambda} (v_i \sum_{j=1}^{k} \alpha_j - k b_i^k)(\Phi_{m-k+1}(b_i^k) - \Phi_{m-k}(b_i^k))$$

$$+ \sum_{k=1}^{\lambda-1} \int_{b_i^{k+1}}^{b_i^k} (v_i \sum_{j=1}^{k} \alpha_j - k\omega)\Phi'_{m-k}(\omega)d\omega$$

$$+ \int_0^{b_i^\lambda} (v_i \sum_{j=1}^{\lambda} \alpha_j - \lambda\omega)\Phi'_{m-\lambda}(\omega)d\omega \tag{4}$$

The partial derivatives of EP_i are:

$$\frac{\vartheta EP_i}{\vartheta b_i^k} = (\alpha_k v_i - b_i^k)\Phi'_{m-k+1}(b_i^k) - k(\Phi_{m-k+1}(b_i^k) - \Phi_{m-k}(b_i^k)) \tag{5}$$

To find the bids that maximize the expected profit EP_i we use Lagrange multipliers. By introducing factors δ_k to convert the constraints on the bids from inequalities to equalities, these constraints become: $b_i^k - b_i^{k+1} = \delta_k^2$. Then, the Lagrange equation becomes:

$$\Lambda(b_i^1, \ldots, b_i^\lambda, \mu_1, \ldots, \mu_\lambda, \delta_1, \ldots, \delta_\lambda) =$$

$$-EP_i(b_i^1, \ldots, b_i^\lambda) + \sum_{k=1}^{\lambda} \mu_k (b_i^k - b_i^{k+1} - \delta_k^2)$$

The values of the variables which maximize this equation are found by setting the partial derivative for all the dependent variables to 0:

$$\frac{\vartheta\Lambda}{\vartheta\mu_k} = 0 \Leftrightarrow b_i^k - b_i^{k+1} = \delta_k^2$$
$$\frac{\vartheta\Lambda}{\vartheta\delta_k} = 0 \Leftrightarrow \mu_k\delta_k = 0$$
$$\frac{\vartheta\Lambda}{\vartheta b_i^k} = 0 \Leftrightarrow \frac{\vartheta EP_i}{\vartheta b_i^k} = \begin{cases} \mu_k - \mu_{k-1} & \text{if k>1,} \\ \mu_1 & \text{if k=1.} \end{cases}$$

However, as we assumed that there is no demand reduction and no bids are equal to each other, this means that all $\mu_i = 0$,[4] hence the optimal bids are given by:

$$(\alpha_k v_i - b_i^k)\Phi'_{m-k+1}(b_i^k) = k(\Phi_{m-k+1}(b_i^k) - \Phi_{m-k}(b_i^k))$$

and since at the equilibrium it should be: $b_i^k = g_{\overline{\alpha},k}(v_i)$, this leads to the system of differential equations 3.

To prove the boundary condition: when a bidder has the lowest possible valuation v^L, then she would never bid higher than v^L, as she would potentially lose profit by winning. On the other hand, it cannot be optimal to bid less than v^L for any valuations, because an opponent would take advantage of this to win by bidding at least v^L even if he had a low valuation too. Therefore, that bidder needs to bid v^L. □

To accommodate the possibility of some bids being equal to each other, we extend the above theorem as follows:

Theorem 2. *The equilibrium strategy for this setting is given by the following algorithm:*
1. $b_i^1 = b_i^2 = \ldots = b_i^{k^*}$*, where* $k^* = \max\{1, m+1-(N-1)\lambda\}$
2. Remove from the system of equations 3 the equations having $\forall k < k^*$ *and replace them with:* $g_{\overline{\alpha},k}(x) = g_{\overline{\alpha},k^*}(x)$
3. For $\Lambda = k^* \ldots \lambda$ *do:*
4. Solve the system of equations 3 setting $\lambda := \Lambda$
5. If $\Lambda < \lambda$ *and* $\left.\frac{\vartheta EP_i}{\vartheta b_i^{\Lambda+1}}\right|_{b_i^{\Lambda+1}=v^L} > 0$*, then continue loop (use equation 5 and compute* Φ *from the solutions of step 4)*
6. else stop loop
7. Return the last found solution from step 4

Proof. The first step of the algorithm checks whether a bidder will win some items for certain due to lack of competition. These are $k^* - 1 = m - (N-1)\lambda$ items, so the bids for these must be as low as possible; however they can never become lower than the next bids (only equal to these), hence they are set equal to $b_i^{k^*}$. In step 2, because of these equalities, this means that when maximizing the Lagrange equation, we should set $\delta_k = 0, \forall k < k^*$, which means that $\mu_k \neq 0$ and therefore the corresponding equations in 3 are no longer valid: they need to be removed and any reference to bidding functions $g_{\overline{\alpha},k}(x)$ must be replaced by$g_{\overline{\alpha},k^*}(x)$, as $b_i^k = b_i^{k^*}$.

The loop which follows examines whether there is demand reduction, meaning that bidders will prefer to bid for Λ items (where $\Lambda < \lambda$). They will bid for at least k^* items though, so the loop starts from that value. In each iteration, first the equilibrium is found (conditional on reduced demand Λ) in step 4, and then step 5 checks whether it is optimal to bid for one more item, if all the opponents bid for Λ items: this happens exactly when $\frac{\vartheta EP_i}{\vartheta b_i^{\Lambda+1}} > 0$ at the point when $b_i^{\Lambda+1} = v^L$ and for some $v_i > v^L$, as this means that bidding $b_i^{\Lambda+1} = v^L + \epsilon, \epsilon > 0$ would increase the expected profit. If it is optimal to get more items, then our assumption that bidders would limit themselves at Λ items is incorrect, and the loop is continued. Of course we do not need to perform this check if $\Lambda = \lambda$.

Finally, when the loop ends or $\Lambda = \lambda$, then value of Λ has the correct value of how many items a bidder should limit herself to, and therefore that solution gives the equilibrium strategies. □

Now we proceed to examine a number of examples:

Example 1. *Assume that* $N \geq 2$ *bidders participate in the auction, where* $m = 2$ *units are for sale. Each bidder wishes to purchase* $\lambda = 2$ *units with valuation* v_i *for the first unit drawn from the uniform distribution on* $[0, 1]$*, i.e.* $F(x) = x, x \in [0, 1]$*, and the value of the second unit is* αv_i*, i.e. the weight* $\alpha^2 = \alpha$*, where* $\alpha \leq 1$*.*
In this case, $k^* = 1$*, so there are no guaranteed items to be won.*
We set $\Lambda = 1$ *(step 3 of the algorithm) and compute the solution:* $g_{\overline{\alpha},1}(v_i) = \frac{1}{2}v_i$*. (step 4)*
We verify that $\left.\frac{\vartheta EP_i}{\vartheta b_i^2}\right|_{b_i^2=0} > 0$ *when the opponents bid* $b_i^1 = \frac{1}{2}v_i$ *and* $b_i^2 = 0$*, therefore the loop is continued. (step 5)*
Setting $\Lambda = 2$*, meaning that there is no demand reduction, we compute the solution:*

$$g_{\overline{\alpha},1}(v_i) = \begin{cases} \frac{1}{2}\left(\frac{N-1}{N+1}\alpha + 1\right)v_i & \text{when } v_i \leq 2\frac{\frac{N-1}{N+1}\alpha}{\frac{N-1}{N+1}\alpha+1} \\ \frac{N-1}{N}v_i + \left(\frac{N-1}{N+1}\alpha - \frac{N-1}{N}2\frac{\frac{N-1}{N+1}\alpha}{\frac{N-1}{N+1}\alpha+1}\right)\left(\frac{2\frac{\frac{N-1}{N+1}\alpha}{\frac{N-1}{N+1}\alpha+1}}{v_i}\right)^{N-2} & o.w. \end{cases}$$

$$g_{\overline{\alpha},2}(v_i) = \frac{N-1}{N+1}\alpha v_i$$

This is the equilibrium bidding strategy, as $\Lambda = \lambda$ *and the loop of the algorithm terminates.*

Notice that there are two cases for the first bid $g_{\overline{\alpha},1}(v_i)$ because the distribution F has a maximum possible valuation $v^H = 1$. This means that bids from the second bid can only go as much as high as $\frac{N-1}{N+1}\alpha$, which affect Φ_2, which in turn changes the function for $g_{\overline{\alpha},1}(v_i)$ for values $v_i > 2\frac{\frac{N-1}{N+1}\alpha}{\frac{N-1}{N+1}\alpha+1}$.

Furthermore, notice that the bidding strategy for the first bid does depend on α. It might seem that this contradicts property 1, however, as we will see more clearly in the examples of Section 4, this is not so. What happens is that the parameter α enters the equation (indirectly) through the computation of the distributions of the opponent bids Φ; this α refers to how this parameter changes their second value of all bidders (and hence their second bids).

3.3 (m+1)-th Price Auctions

Theorem 3. *The equilibrium strategy for the setting that we examine is given by the algorithm of Theorem 2, with the change that in step 4 the following system of differential equations needs to be solved:*

$$(\alpha_k v_i - g_{\overline{\alpha},k}(v_i))\Phi'_{m-k+1}(g_{\overline{\alpha},k}(v_i)) = \tag{6}$$
$$(k-1)(\Phi_{m-k+2}(g_{\overline{\alpha},k}(v_i)) - \Phi_{m-k+1}(g_{\overline{\alpha},k}(v_i)))$$

[4] In the next theorem we will deal with the other cases, i.e. bids equal to each other

PROOF. Similar to the m^{th} price auction, the bids placed by agent i are $b_i^1 \geq \ldots \geq b_i^\lambda \geq 0$. Given these bids, the expected profit of the agent is:

$$EP_i(b_i^1, \ldots, b_i^\lambda) = \sum_{k=1}^{\lambda-1} \int_{b_i^{k+1}}^{b_i^k} (v_i \sum_{j=1}^{k} \alpha_j - k\omega)\Phi'_{m-k+1}(\omega)d\omega$$

$$+\sum_{k=1}^{\lambda-1}(v_i \sum_{j=1}^{k} \alpha_j - kb_i^{k+1})(\Phi_{m-k+1}(b_i^{k+1}) - \Phi_{m-k}(b_i^{k+1}))$$

$$+\int_0^{b_i^\lambda} (v_i \sum_{j=1}^{\lambda} \alpha_j - \lambda\omega)\Phi'_{m-\lambda}(\omega)d\omega \qquad (7)$$

The partial derivatives of EP_i are:

$$\frac{\vartheta EP_i}{\vartheta b_i^k} = (\alpha_k v_i - b_i^k)\Phi'_{m-k+1}(b_i^k) - (k-1)(\Phi_{m-k+2}(b_i^k) - \Phi_{m-k+1}(b_i^k)) \qquad (8)$$

If an assumption is made that no bids are equal to each other, then we need to take the first order conditions: $\frac{\vartheta EP_i}{\vartheta b_i^k} = 0$ and since at the equilibrium it must be: $b_i^k = g_{\overline{\alpha},k}(v_i)$, this leads to the system of differential equations 6. The boundary conditions remain the same.

To account for the possibility of bids being equal to neighboring bids, we use the same reasoning as in Theorem 2, hence why the rest of the algorithm remains the same. □

We can now give two more properties of the equilibrium:
Property 6: Bidder bid truthfully for the first item in this auction.
Indeed, this is known and it is trivial to prove, by setting $k = 1$ in equation 6.
Property 7: It is easier to solve the system of this theorem than than of Theorem 1.
Since we know that $b_i^1 = v_i$ in this theorem, the size of the problem reduces by 1, and the system of differential equation has $(\lambda-1)$ equations and $(\lambda-1)$ unknown functions. The one from Theorem 1 has λ equations and λ unknown functions, so it is harder to solve.

We will give two examples; one where there is demand reduction and another with insufficient competition:

EXAMPLE 2. *Assume that $N \geq 2$ bidders participate in the auction, where $m = 2$ units are for sale. Each bidder wishes to purchase $\lambda = 2$ units with valuation v_i for the first unit drawn from the uniform distribution on $[0, 1]$, i.e. $F(x) = x, x \in [0, 1]$, and the value of the second unit is αv_i, i.e. the weight $\alpha^2 = \alpha$, where $\alpha \leq 1$.*
In this case, applying the algorithm we can check that when $\alpha < 1$, the bidder will bid $g_{\overline{\alpha},1}(v_i) = v_i$ and $g_{\overline{\alpha},2}(v_i) = 0$: there is demand reduction to one unit.

EXAMPLE 3. *Assume that $N = 2$ bidders participate in the auction, where $m = 4$ units are for sale. Each bidder wishes to purchase $\lambda = 4$ units with valuation v_i for the first unit drawn from the uniform distribution on $[0, 1]$. The value for the other units are $\alpha_2 v_i \geq \alpha_3 v_i \geq \alpha_4 v_i$.*
In this case, applying the algorithm, we find that $b_i^1 = b_i^2$. However, now that system of equations can have an infinite number of solutions, as $\frac{\vartheta EP_i}{\vartheta b_i^2} = 0, \forall b_i^2$; we can choose $g_{\overline{\alpha},1}(v_i) = g_{\overline{\alpha},2}(v_i) = \alpha_1 v_i$ w.l.o.g. Then we check that no bidder will want to bid for three items, as $\frac{\vartheta EP_i}{\vartheta b_i^3} < 0$, hence there is demand reduction to two units.

4. THE GENERAL MODEL: BIDDERS WITH ASYMMETRIC WEIGHT VECTORS

In this section, we extend the results of the previous section to the general model where bidders can have different weight vectors. In this way, it is possible to simulate any possible combination of valuations for each bidder, provided that there are no complementarities meaning as long as $v_i^1 \geq \ldots \geq v_i^\lambda$.

THEOREM 4. *The algorithm of Theorem 2 still gives the equilibrium strategy for a m^{th} price auction in this setting, with the change that equation 1 is replaced by:*

$$\Phi_k(x) = \sum_{j=1}^{k} \sum_{\substack{d_1,\ldots,d_{\min\{j,\lambda\}} \\ s.t.: \sum_l ld_l = j-1}} \frac{(N-1)!}{(N-1-\sum_l d_l)! \prod_l d_l!} (1 - \widetilde{F}_\lambda(x))^{d_\lambda}$$

$$\cdot \prod_{l<\lambda} (\widetilde{F}_{l+1}(x) - \widetilde{F}_l(x))^{d_l} (\widetilde{F}_1(x))^{N-1-\sum_l d_l} \qquad (9)$$

where:

$$\widetilde{F}_l(x) = \sum_{\overline{\alpha}} h(\overline{\alpha}) F(g_{\overline{\alpha},l}^{-1}(x))$$

PROOF. The main change compared to Theorems 2 and 1 is that the opponents can have a number of weight vectors $\overline{\alpha}$ each with probability $h(\overline{\alpha})$. This affects the distributions of the opponent bids in the following ways:

1. The first bid of each opponent is given by distribution $\sum_{\overline{\alpha}} h(\overline{\alpha}) F(g_{\overline{\alpha},1}^{-1}(x)) = \widetilde{F}_1(x)$, which is obtained by using Bayes' law.

2. Similarly using Bayes' law, the probability of being between the l^{th} and $(l+1)^{th}$ bids of one opponent is: $\sum_{\overline{\alpha}} h(\overline{\alpha}) \big(F(g_{\overline{\alpha},l+1}^{-1}(x)) - F(g_{\overline{\alpha},l}^{-1}(x))\big) =$ $\sum_{\overline{\alpha}} h(\overline{\alpha}) F(g_{\overline{\alpha},l+1}^{-1}(x)) - \sum_{\overline{\alpha}} h(\overline{\alpha}) F(g_{\overline{\alpha},l}^{-1}(x)) =$ $\widetilde{F}_{l+1}(x) - \widetilde{F}_l(x)$.

3. Finally, the probability of having a lower value than the lowest bid is: $\sum_{\overline{\alpha}} h(\overline{\alpha}) \big(1 - F(g_{\overline{\alpha},\lambda}^{-1}(x))\big) = \sum_{\overline{\alpha}} h(\overline{\alpha}) -$ $\sum_{\overline{\alpha}} h(\overline{\alpha}) F(g_{\overline{\alpha},\lambda}^{-1}(x)) = 1 - \widetilde{F}_\lambda(x)$.

That means that the terms F from the original equation need to be replaced with the equivalent terms $\widetilde{F}$, which generates equation 9. □

The same process extends Theorem 3:

THEOREM 5. *The algorithm of Theorem 3 still gives the equilibrium strategy for a $(m+1)^{th}$ price auction in this setting, with the change that equation 1 is replaced by equation 9.*

Now, it is time to re-examine Example 1:

EXAMPLE 4. *Assume the same setting as in Example 1, with the difference that the weight $\alpha^2 = \alpha$, is no longer the same for all bidders, but rather the weights of every bidder α^2 are i.i.d. random variables drawn from the uniform distribution $U[0, 1]$.*
There is no demand reduction and the equilibrium is computed as:

$$g_{\overline{\alpha},2}(v_i) = \frac{N-1}{N+1} \alpha v_i$$

$$g_{\overline{\alpha},1}(v_i) = \frac{1}{2}\big(\frac{N-1}{N+1}\alpha^* + 1\big)v_i$$

where

$$a^* = \frac{1}{\sum_\alpha \frac{h(\alpha)}{\alpha}}$$

when $v_i \leq 2\frac{\frac{N-1}{N+1}\min\{\alpha\}}{\frac{N-1}{N+1}\alpha^*+1}$.

We do not detail here the equilibrium strategies for values $v_i > 2\frac{\frac{N-1}{N+1}\min\{\alpha\}}{\frac{N-1}{N+1}\alpha^*+1}$; this is a mathematical exercise. We want to focus instead and point out certain properties of the equilibrium strategies. One can notice that each bidder's weight α only affects directly her second bid. The first bid in not affected, which is consistent with the properties of the equilibrium that we detailed earlier. In addition, the first bid placed by each bidder does not depend on her individual weight α, but rather on an aggregate weight α^* which is equal to the harmonic mean of all the possible values that α can take, weighted by the probability of each. This is the indirect effect of the weights α to the first bids placed by the bidders.

We can also examine this example for the case when the weight α_i of each bidder are known, i.e. they are common knowledge to all bidders:

Example 5. *Assume the same setting as in Example 4, with the difference that the weights α_i are known.*
There is no demand reduction and the equilibrium is computed as in Example 4, with the difference that now:

$$a^* = \frac{N}{\sum_{\alpha_i}\frac{1}{\alpha_i}}$$

Again, we make the exact same observations, with the difference that now the aggregate weight α^* is equal to the harmonic mean of all the bidders' weights α_i. This parameter indeed summarizes the indirect effect of the weights α_i (which in turn give the bidders' valuations for getting a second item) to the first bids placed by the agents.

5. EXPERIMENTAL EVALUATION

In this section, we conduct experiments to test the equilibrium strategies against heuristic strategies used in such an auction setting in previous work.

First, we examined the literature on trading agent design and the Trading Agent Competition (TAC): in TAC Classic, the first benchmark game that was created for TAC, some of the auctions used were m^{th} price sealed bid auctions. These were the auctions selling the hotel rooms, where $m = 16$ hotel rooms of each type were available.[22] We examined the strategies used in order to bid specifically in these auctions. Some agents, see [7, 17], placed bids based on the historical closing prices that were learnt from previous games, usually some weighted average of the current price (which is zero at the beginning) and the predicted price. Similarly, Walverine[3] used a similar reasoning, however, the price predictions were partly based on a principled competitive analysis which attempted to stabilize price predictions based on the information that has been observed in the current game; as such it could be less sensitive to changes in the participants strategies. Finally, in [19], the authors simply ignore the effect of bid shading and split each bidder into multiple subbidders where each one bids independently for each item; the authors note that because of the competition effect, the bid shading is essentially canceled, and examine the effect of competition in [18] basing their work on [2]. However, in all the TAC agents, we notice two common elements: (a) demand reduction to 3, 4 or 5 units (out of the possible 8 desired) and (b) most agents are willing to place small bids in order to grab a hotel cheaply even if they do think that they can use it.

Based on this review, we chose the following strategies to test:

1. Heuristic Strategy S_I: this ignored the bid shading altogether and assumed that there are $N' = N\lambda$ bidders, as in [19].

2. Heuristic Strategy S_R: this strategy uses the same strategy as S_I for b_1 and bids a small percentage of the valuation for the remaining ones. Essentially this is inspired from the demand reductions strategies, but does not go all the way to completely bidding zero, thinking of the way that most TAC agents are willing to pay something small to grab a good for a bargain price (if they are lucky). [3, 7, 17, 19] To this end we set this function to bid 25% of the second value in the experiment.

3. Equilibrium Strategy S_E: this is the equilibrium strategy computed in this paper.

We conducted our experiments on the setting of Example 1, for $N = 4$ bidders and values of $\alpha = 0.1, \ldots, 0.9, 1$. ($\alpha = 0$ has no meaning because there would be no multi-unit demand) We will assume that between 1 and 3 bidders using strategy S_I and S_R respectively compete against the remaining agents who use strategy S_E. The results of this experiment are presented in Figure 1. Note that the errorbars would very small and that is why there are not depicted in this Figure. What we notice is that with the exception of the case when there is only one agent using strategy S_I (or S_R resp.) and then only for small values of $\alpha <= 0.5$), we see a significant improvement in the expected profit generated from using strategy S_E, which varies between 5% and 25%. In the few cases, that this is not the case, the drop in performance (in those few cases) is no more than 1%. Furthermore, even in this case, when there is one bidder not playing the equilibrium strategy, this bidder might be able to do slightly better, however, if she were to switch to strategy S_E, then the profit of all bidders would improve; this is theoretically guaranteed after all from the fact that S_E is an equilibrium strategy and no one bidder can gain by deviating from it! And in fact, the explanation why S_I (or S_R resp.) manages to outperform slightly S_E in these few cases stems exactly from the fact that we are now looking at the relative performance of the strategies (not the absolute value), therefore we face the issues described in [2, 18].

6. DISCUSSION AND CONCLUSIONS

In this paper, we characterized and computed the equilibrium strategies for both m^{th} and $(m+1)^{th}$ price sealed-bid multi-unit auctions, where the participating bidders have multi-unit demand; this was an open problem for almost two decades. We examined a number of properties of the equilibrium and then examined two models, the second of which is general enough to capture any set of bidder valuations provided that there are no complementarities in these.

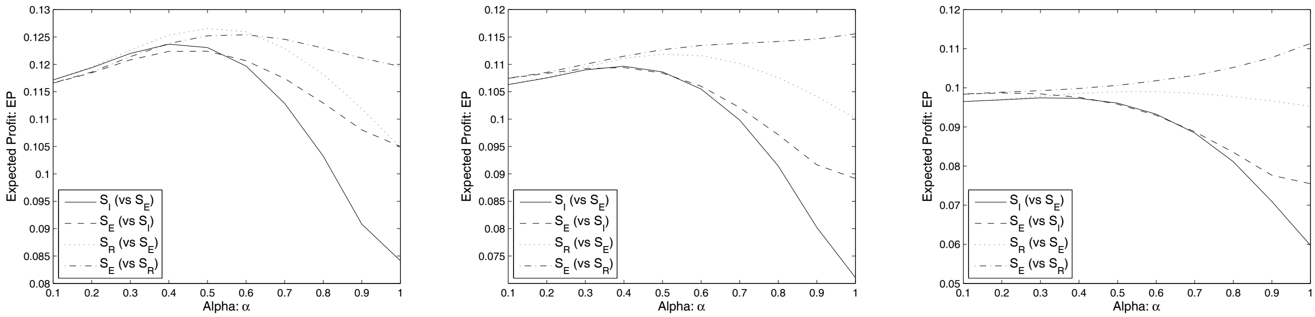

Figure 1: Expected profit for: (left) 1 bidder using S_I (resp. S_R), (center) 2 bidders using S_I (resp. S_R), and (right) 3 bidders using S_I (resp. S_R).

To the best of our knowledge, we are the first to give an algorithm for computing the equilibrium strategies for this problem and for any values of bidder N, items for sale m and demand λ. In the end, we conducted experiments to examine what happens when comparing the equilibrium strategies with other heuristic ones, like those used previously in the literature; our experiments show an improvement of usually 5% to 25% in most cases when using the equilibrium strategies (with only few cases showing a drop in profit of up to 1%).

Now, this work is far from over. The main issue is that the systems of differential equations that need to be solved are notoriously unstable, as we have observed in our previous work, when looking at asymmetric auctions.[20] Because of this, in all the examples we gave, the equilibrium strategies were computed analytically rather than computationally. This does not mean that our theorems are incorrect; far from it in fact. However, in order to be able to automatically generate these strategies for any input, the parameters N, m, λ, the weights and prior distributions for the valuations and the weights, we need to resolve this issue.

Furthermore, as future work, we plan to examine ways in which we might be able to extend our model to cover also complementary valuations. Of course, this is not an easy problem as even in the case of multiple auctions selling complementary items, because in addition to the bid shading the agent would have to face also the well-known exposure problem.[11]

7. ACKNOWLEDGMENTS

This work was supported by a grant from the Institute of Informatics and Telecommunications of NCSR "Demokritos" and from the SKEL lab.

8. REFERENCES

[1] L. M. Ausubel and P. Cramton. Demand reduction and inefficiency in multi-unit auctions. 2002.

[2] F. Brandt, T. Sandholm, and Y. Shoham. Spiteful bidding in sealed-bid auctions. In *Proceedings of the 20th international joint conference on Artifical intelligence*, pages 1207–1214. Morgan Kaufmann Publishers Inc., 2007.

[3] S.-F. Cheng, E. Leung, K. M. Lochner, K. O'Malley, D. M. Reeves, J. L. Schvartzman, and M. P. Wellman. Walverine: A walrasian trading agent. *Decision Support Systems*, 39(2):169–184, 2005.

[4] R. Engelbrecht-Wiggans and C. M. Kahn. Multi-unit auctions with uniform prices. *Economic Theory*, 12(2):227–258, 1998.

[5] R. Engelbrecht-Wiggans, J. A. List, and D. H. Reiley. Demand reduction in multi-unit auctions with varying numbers of bidders: Theory and evidence from a field experiment*. *International Economic Review*, 47(1):203–231, 2006.

[6] J. K. Goeree, T. Offerman, and R. Sloof. Demand reduction and preemptive bidding in multi-unit license auctions. *Experimental Economics*, 16(1):52–87, 2013.

[7] A. Greenwald and J. Boyan. Bidding algorithms for simultaneous auctions. In *Proceedings of 3rd ACM Conf. on Electronic Commerce*, pages 115–124. ACM, 2001.

[8] J. H. Kagel and D. Levin. Behavior in multi-unit demand auctions: Experiments with uniform price and dynamic vickrey auctions. *Econometrica*, 69(2):413–454, 2001.

[9] J. H. Kagel and D. Levin. Multi-unit demand auctions with synergies: behavior in sealed-bid versus ascending-bid uniform-price auctions. *Games and Economic Behavior*, 53(2):170–207, 2005.

[10] J. H. Kagel and D. Levin. Auctions: A survey of experimental research, 1995-2008. *Handbook of experimental economics*, 2, 2008.

[11] V. Krishna. *Auction theory*. Academic press, 2009.

[12] D. McAdams. Isotone equilibrium in games of incomplete information. *Econometrica*, 71(4):1191–1214, 2003.

[13] D. McAdams. Monotone equilibrium in multi-unit auctions. *The Review of Economic Studies*, 73(4):1039–1056, 2006.

[14] D. Nautz. Optimal bidding in multi-unit auctions with many bidders. *Economics Letters*, 48(3):301–306, 1995.

[15] D. Nautz and E. Wolfstetter. Bid shading and risk aversion in multi-unit auctions with many bidders. *Economics Letters*, 56(2):195–200, 1997.

[16] C. Noussair. Equilibria in a multi-object uniform price sealed bid auction with multi-unit demands. *Economic Theory*, 5(2):337–351, 1995.

[17] P. Stone, M. L. Littman, S. Singh, and M. Kearns. Attac-2000: An adaptive autonomous bidding agent. In *Proceedings of the fifth international conference on Autonomous agents*, pages 238–245. ACM, 2001.

[18] I. A. Vetsikas and N. R. Jennings. Outperforming the competition in multi-unit sealed bid auctions. In *Proceedings of the 6th international joint conference on Autonomous agents and multiagent systems*, page 103. ACM, 2007.

[19] I. A. Vetsikas, N. R. Jennings, and B. Selman. Generating bayes-nash equilibria to design autonomous trading agents. In *IJCAI*, pages 1543–1550, 2007.

[20] I. A. Vetsikas, S. Stein, and N. R. Jennings. Multi-unit auctions with a stochastic number of asymmetric bidders. In *Proceedings of the 20th European Conference on Artificial Intelligence*, pages 816–821, 2012.

[21] R. J. Weber. Making more from less: Strategic demand reduction in the fcc spectrum auctions. *Journal of Economics & Management Strategy*, 6(3):529–548, 1997.

[22] M. P. Wellman, A. Greenwald, and P. Stone. *Autonomous bidding agents: Strategies and lessons from the trading agent competition*. MIT Press, 2007.