An English Auction Protocol for Multi-Attribute Items♣

Esther David[1], Rina Azoulay-Schwartz[1], Sarit Kraus[1]

[1] Bar-Ilan University, Department of Computer Science, Ramat-Gan 52900, Israel
{Saadone, Schwart, Sarit}@macs.biu.ac.il

Abstract. In this paper, we suggest using an English Auction Protocol for a procurement multi-attribute auction in which the item for sale is defined by several attributes, the buyer agent is the auctioneer, and the seller agents are the bidders. Such domains include auctions on task allocation, services, or compound products. At the beginning of the auction the buyer agent announces the required properties of the item, and then various seller agents propose bids, which are composed of specific configurations that match its request. Each proposed bid should be better for the buyer agent than the previous bid, w.r.t. the announced requirements of the buyer agent. Finally, the last suggested bid will win, and the seller agent that suggested this bid will be committed to it. We consider two utility function models for the English auction protocols and provide the optimal bidding strategies for the seller agents and the optimal auction design for the buyer agents regarding both models.

1 Introduction

Auctions are important mechanisms for allocating resources and services among agents [13,14,17]. English auctions are widely used in markets of services (such as cargo deliveries), or unique items, and in on-line auction houses, such as eBay, etc., and it may be very useful in solving resource allocation problems. The widespread research on automated English auctions deals mostly with models where price is the unique strategic dimension [16]. However, in many situations, it is necessary to conduct negotiations on multiple attributes of a deal. For example, in task allocation, the attributes of a deal include the size of a task, starting time, ending deadline, accuracy level, etc. A service can be characterized by its quality, supply time, and the risk involved. English auctions protocols can also be used when the issue to be considered is multi attributed. Several difficulties arise when trying to implement the traditional single-attribute auction protocols for multi-attribute items. In this paper, we provide

♣ This work was supported in part by NSF under Grant No. IIS-0208608 and IIS9820657. Sarit Kraus is also affiliated with UMIACS.

the automated agents (the buyer and the seller agents) that participate in a multi-attribute English auction with stable and efficient strategies to be used in the auction.

We focus on markets in which an agent that wants to buy an item becomes the auctioneer of an auction process and the seller agents are the bidders. At the beginning of the auction, the buyer agent announces the required properties of the item, and then various seller agents propose bids, which are composed of specific configurations that match its request. Each proposed bid should be better for the buyer agent than the previous bid, w.r.t. the announced requirements of the buyer agent. Finally, the last suggested bid will win, and the seller agent that suggested this bid will be committed to it. Such markets appear in several situations, and currently there is no automated mechanism to deal with them. For example, telephone service providers and Internet portals, as well as video-on-demand suppliers, would like to rent extra storage capacity from suppliers over the Internet. The attributes of the required item in this domain are the storage capacity, the access rates to the data, the availability time and deadline, the level of security, etc.

There are several problematic aspects when considering a multi-attribute English auction. Consider a seller participating in a multi-attribute English auction. Preparing a bid that matches the bidder's requirements, which is better than the previous proposed bid and also maximizes the seller agent's utility is a sophisticated task. Another question that the seller faces, is should the composed bid be influenced by the specific values of the proposed bids. If it should, then how should the previous proposed bid influence the next proposed bid?

Some other problematic aspects relate to the buyer, which is the auctioneer in such an auction. For example, what should the announced requirements be in order to achieve the best agreement? Should they be the exact requirements or should certain modifications of the announced requirements be permitted in order to achieve better results.

Another interesting issue concerning the auctioneer and bidders is the multi-attribute offer evaluation process. Calculating the contribution/cost of each attribute's value, accumulating all of these contributions/costs for one value in order to be able to compare the various bids and choosing the best one are certainly not trivial tasks. In this paper, we address these various issues and propose ways to handle them.

In our work, we consider two models for a multi-attribute English auction protocol that differs in the seller agents' utility functions. For both models, we suggest the optimal strategies for the automated agents that participate in such auctions. These results extend our previous work [7] that focused mainly on the multi-attribute sealed bid auction. In [7] we also considered the English auction protocol, and defined four variations of an English auction used for the general case of multi-attribute auctions. We proved under some assumptions that the four variations converge to the same result. However, no optimal strategies were proposed. In this work, we consider one of the English auction's variations, called the *sequential full-information-revelation auction*, for a case where the four variations of the English auction converge to the same result. We provide optimal strategies for agents participating in this auction, and reveal the auction outcome, given the environmental parameters.

2 Related Work

An auction is an efficient protocol for reaching agreements among agents [8, 13, 14, 17]. There are several types of auctions, which are used, including the English auction, first price sealed-bid auction, second-price sealed-bid (Vickery) auction, and the Dutch auction [14]. In this paper, we focus on the English auction. In the English auction, during the bidding process, each bidder can suggest a bid better than or equal to the last bid, and when no more bids are suggested, the last placed bid wins. In real world situations, each auction has its advantages and drawbacks.

Auctions can also be used when the issue to be considered includes more than one attribute. Researchers discuss auctions where a bid is composed of several details, but mostly in the context of a special type of auction, called combinatorial auctions [5, 6, 9, 10, 12, 15]. In this model, a set of available goods is given. Each bidder specifies bundles of goods and the prices it is willing to pay for each specified bundle. The problem that arises is how to determine which agents will obtain the bundles they ask for, since the number of available goods is limited. Although in combinatorial auctions each bid is composed of multiple details, the problems that appear in our model are completely different, since, the auction discussed in our model includes one buyer of one multi-attribute item and several competing sellers. This diversity causes completely different problems when trying to automate the auction mechanism. In this paper, we address these problems, and suggest how to design and implement automated auctions for multi-attribute items. We also provide the automated auctioneer and bidders with stable and efficient strategies to be used in this type of auction

Gimenez-Funes et al. [8] developed trading agents for electronic auctions of multi-attribute items where the seller suggested a given item for sale, and the buyers compete by sending bids that are composed of the price of the item. They suggest the agent apply the Case-Based Decision theory in order to decide which strategy to use in the auction. In this paper, we consider an auction where the buyer is the auctioneer, and the sellers suggest items with multiple attributes. We develop static and stable agent bidding strategies that are based on their beliefs on the environment

Bichler [2] has made an experimental analysis of multi-attribute auctions. He found out that the utility scores achieved in multi-attribute auctions were significantly higher than those of single attribute auctions. The single-attribute auction he refers to is actually a multi-attribute auction in which all the attributes' values were fixed and the bidders actually competed only on one-dimension bids.

Very little theoretical work has been done on multi-attribute auctions. Che [4] considers an auction protocol where a bid is composed of a price and a quality. In addition he assumes that each seller is characterized by only one private cost parameter. In his paper, he proposed a design for first score and second score sealed bid auctions, which are based on the announced scoring rule. A scoring rule is used to define and describe the required item and it associates a score with each possible bid. . In particular the optimal auction design that Che proposed considers a distortion of the scoring rule which is additive. That is, he defines the optimal value of a component that should be reduced from the real score. However, in our model the scoring rule is defined in a way that accumulates the weighted contributions of the various attributes' values. When we developed the optimal auction design, we actually looked for a way to calculate the optimal announced weights that yield the maximum utility for the buyer. In addition, we extended his work by considering more complex models in which the sellers may be characterized by more than one private cost parameter. Also, we consider the English auction protocol, which was not considered by Che [4].

Branco [3] extended the work of Che by assuming that the costs of the firms/bidders are correlated. He considers a governmental procurement auction in which the main goal is to maximize the virtual welfare, which takes into account the private rents that will be given to the firms. Branco uses a method similar to Che's to design the optimal auction considering his model assumptions. In contrast to Branco's work, we assume that the costs of the bidders are independent. In addition, we design the optimal auction from the buyer's point of view and not from the standpoint of the population's welfare.

3 The English Auction with One Cost Parameter per Each Seller

The auction model consists of one buyer, which is the auctioneer, and a fixed number of n sellers, which are the bidders. The English auction protocol is a very useful mechanism in real world decisions about suppliers of services, and in government decision-making processes. In particular, we refer to a version of an English auction called the *full-information revelation protocol.* In this protocol, the buyer agent announces (1) a full scoring-rule function that describes the required item, (2) the *closing interval*, which is the length of the time interval, where if no new bid is made, the auction is closed. The buyer's scoring rule associates a score with each proposed offer and the auction protocol dictates the winner (best scored bid) based on this scoring rule. Each participating seller agent receives a serial number that defines the order of bidding among the agents. In each step, the seller whose turn it is to bid may place a bid, which is better than the previous proposed bid by at least the minimal increment D w.r.t the scoring rule function. In addition, in each step that a bid is proposed any seller agent that wants to bid a bid, which yields the same score, can do so at a predefined interval of time. The buyer chooses one of the bids that yield the same score randomly. The buyer agent tries to derive a scoring rule that maximizes its expected utility in a given auction protocol. Each participant knows its own utility function, and time and bidding are not costly. The auctioneer must be committed to its scoring rule, and the winner agent is required to provide an item with the exact values of the bid it offered (e.g., the exact price, quality, delivery date, etc.).

The utility functions of the agents are similar to the ones we considered in [7]. We recall some of the details here and explain the differences. Each buyer agent and each seller agent is characterized by a utility function that describes its preferences and they use it to evaluate the utility from each given bid. The multi-attribute utility-functions we refer to are based on the Simple Additive Weighting (SAW) method [18], i.e., a utility is obtained by adding the contributions from each attribute. Other methods exist for a multi-attribute utility function (e.g. multiplying the contributions of the various attributes), but the SAW method suits the example of renting storage capacity and the other examples, we consider.

In our model, each seller agent has private information about the costs of improving the quality of the product it sells, or its performance. Each seller agent S_i (bidder) assumes to be characterized by a cost parameter θ_i, which is its private information. As θ_i increases the cost of the seller to produce an item of a higher quality also increases, i.e., the seller is weaker. The buyer (auctioneer) only knows the distribution function of the other sellers' cost parameters, but has no information about the particular value of θ_i for each seller. Similar to the model described by Che [4] and in [7], we assume that the private type θ_i is independently and identically distributed over $[\underline{\theta}, \bar{\theta}]$ ($0<\underline{\theta}<\bar{\theta}<\infty$) according to a distribution function *F* for which a positive, consciously differentiable density *f* exists. In section 4 we will also consider an extended model, which was not considered in [4,7], where each seller is characterized by two private type parameters.

As in [7], for simplicity, we analyze a concrete model, which includes three attributes: the price *p* of the item, and two quality factors, *q1* and *q2* for which the preferences of the buyer and the sellers conflict. In contrast to the model considered in [7], we assume that as *qi* increases, the quality of the item increases, for both *q1* and *q2*. For example, if the announced item is providing a machine, then *q1* can denote the speed of the machine and *q2* can denote its accuracy, or the warranty period for this machine. If the announced service is a video-on-demand supply, then *q1* can denote the storage capacity and *q2* can denote the access rate. There are also quality attributes that denote a higher quality item as the attribute is lower. For example, in a service like a transportation of a cargo, the attributes are the availability time and the path length. As the availability time of the item is shorter, and as the path length is shorter, the service quality is higher. In such domains, we can denote *q1* to be *1/(availability-time),* and *q2* to be *1/(path-length),* in order to have a positive relation between the quality attributes *q1,q2* and the quality of the item.

Using the above influence of q1 and q2, we can conclude that as q1 or q2 increase, the costs of the seller will become higher, since it is more difficult for it to provide a higher quality item, while the utility of the buyer from the provided item will increase. For example, the seller's cost increases if he provides a more accurate machine, but the buyer's utility is higher if it obtains such a machine.

Consider the cost function of the sellers. As in [7], we start by assuming that there are fixed coefficients for each of the quality dimensions, which are identical for all the sellers. Namely, 'a' is the coefficient of q_1, and 'b' is the coefficient of q_2. The seller's utility function, is:

$$U_{s_i}(\theta) = p - \theta(a \cdot q_1 + b \cdot q_2). \qquad \textbf{(1)}$$

where a,b>0. Notice that, the utility function of the seller is the difference between the payment (profit) and the costs, and as the payment it obtains increases, the utility increases. The influence of q1 and q2 is assumed to be independent and linear: as q1 increases by one unit, the cost of the seller will increase by θa, and as q2 increases by one unit, the cost will increase by θb. In Section 4 we will consider a more complex utility function, where the cost of the seller depends on two private cost parameters.

The utility function of the buyer agent (the auctioneer) from the item or service is:

$$U_{buyer}(p,q_1,q_2) = -p + W_1 \cdot \sqrt{q_1} + W_2 \cdot \sqrt{q_2} \quad \textbf{(2)}$$

where W_1 and W_2 are the weights the buyer assigns to q_1 and q_2, respectively. In fact, as the price decreases the buyer's utility increases. It is clear that as q1 and q2 increase, the utility of the buyer increases. We assume that the influence of q1 and q2 is independent, but, in contrast to the model of [7], we consider a non-linear influence: as *q1* increases, the influence of one additional unit of *qi* becomes smaller. This assumption is valid in many domains. For example, enlarging the speed of a machine from 100 Mhz to 200 Mhz will have a higher influence than enlarging the speed from 200 Mhz to 300 Mhz. The effect of q_1 and q_2 is weighted by W_1 and W_2, respectively, where *Wi* can be smaller or larger than 1. As *Wi* increases, the importance of attribute *qi* increases w.r.t. the price and the second attribute. The price is not weighted since we assume that the function is normalized by the weight of the price and it is easy to extend it to the case in which the price is weighted.

The buyer is required to announce a scoring rule function at the beginning of an auction, which is used for choosing among the bids. The scoring rule published by the buyer can be different than its real utility function in a sense that the published weights w_1, w_2 may be different than the actual weights W_1, W_2. In particular, the scoring rule is of the form

$$S(p,q_1,q_2) = -p + w_1 \cdot \sqrt{q_1} + w_2 \cdot \sqrt{q_2} \quad \textbf{(3)}$$

In the following, we will analyze the above model. We first present the optimal bids to be suggested by each bidder, and then we prove which bidder will win, and what will be its winning bid. We proceed by analyzing the expected revenue of the auctioneer and its optimal scoring rule.

3.1 Optimal Bidders' Strategies

One may believe that by announcing the scoring rule at the beginning of the auction the problem of the multi-attribute case reduces to the case of a single attribute auction in which the bidders compete only over the score. Also, each bidder maps the score using his cost functions for the qualities' value, which maximizes its utility. However, in the following Lemma, we show that the bidders decide about the qualities that maximize their utility independently of the exact score they want to achieve, but only with regard to the announced scoring rule. Thus, only the price is calculated with respect to the desired score and the predefined qualities' values. Namely, the bidder saves much computation effort since only the price should be recalculated in each step. In addition, we show that this is the optimal strategy and that every bidder would have no interest in using a different strategy. Another important issue is that by knowing the bidding strategy the auctioneer may estimate its expected revenue from the auction and may decide about the scoring rule, as explained in section 3.2. We begin by considering the optimal bid to be offered by each bidder in each step. We show that the qualities *q1,q2* will be chosen independently of the current selected bid, and we find the optimal price to be offered in each step of the auction, given the bidder's properties and given the current selected bid, that is denoted by *selected*. Calculating the price is different from [4,7] because of the English protocol that is applied. For space reasons we do not present the proofs here.

Lemma 1

Given the scoring rule and the seller's utility functions, in a ***sequential full-information revelation*** *English auction, and given the last selected bid, the seller's best strategy is to bid*

$$p^*(\theta, selected) = \frac{w_1^2}{2 \cdot a \cdot \theta} + \frac{w_2^2}{2 \cdot b \cdot \theta} - S(selected) - D \tag{4}$$

$$q_1^*(\theta) = \left(\frac{w_1}{2 \cdot a\theta}\right)^2, q_2^*(\theta) = \left(\frac{w_2}{2 \cdot b\theta}\right)^2. \tag{5}$$

Sketch of proof: Suppose to the contrary an equilibrium bid *(p,q1,q2)* in which, q1≠q1* or/and q2 ≠q2* at least for one seller with $\theta < \overline{\theta}$. A contradiction is derived by showing that the bid is dominated by an alternative bid *(p1,q1',q2')* where q1'=q1*, q2'=q2*, and p' is chosen in a way that the scores of the two bids, (p,q1,q2) and *(p',q1',q2')*, are equal. We can show that the utility of the seller for the alternative bid is higher than in *(p,q1,q2)*. It remains to find the optimal p* to be offered, given the current selected bid. According to the defined protocol the next bid must be better than the selected bid in terms of the scoring value at least in an increment of *D*. That is: *S(p*,q1*,q2*)=S(selected)+D*. By substituting the values of *q1** and *q2**, we can find the value of p* for which the above equation holds. ■

Lemma 1 specifies the details of the next bid to be announced by a seller. Notice that in the sequential full-information-revelation English auction the strategy of choosing the price to offer does not include the seller's beliefs about the types of the other sellers. In the sequential full-information-revelation auction, the price depends on the bidder's type and on the previous best bid. The qualities offered in the bid are chosen independently of the price. As the bidder's efficiency decreases (has a higher θ), the price it requires decreases, in order to be able to compete in the auction.

The next question is, given the auction participants utility functions, the range of the sellers types (i.e., $[\underline{\theta}, \overline{\theta}]$), the announced scoring rule, and the sellers optimal bidding strategy, can the buyer agent estimate which of the sellers will win and can it estimate its expected revenue? We will answer these questions in the next section.

3.2 Auction Results

In Lemma 2 we will show that the winning seller is the seller with the lowest type θ. The intuition behind this is that for any bid that another seller offers the seller with the lowest type can offer a better bid from the buyer's point of view, which yields a higher score. This means that it can overcome any of the other sellers along the sequential full-information revelation (English) auction process.

Lemma 2

Denote by Si the seller with the lowest value of θ, and by Sj the seller with the second lowest θ. In the sequential full-information revelation English auction protocol, seller Si will be the winner of the auction whenever

$$D \le (\sqrt{a \cdot w_1} + \sqrt{b \cdot w_2}) \cdot (2\sqrt{\theta_j} - 2\sqrt{\theta_i}) \cdot \quad \textbf{(6)}$$

Sketch of proof: Denote the best bid that seller s_j can offer by $(p(\theta j),q1(\theta j),q2(\theta j))$. In order to suggest a higher bid, the scoring value of the bid $(p(\theta i),q1(\theta i),q2(\theta i))$ of seller *Si* should be at least $(p(\theta j),q1(\theta j),q2(\theta j))+D$. We proceed by substituting the values of the best bid of seller Sj with the optimal q1 and q2, and with the lowest p for which the utility of Sj is non-negative, and then we find the scoring rule of this bid. Similarly, we find the scoring rule of the best bid seller Si can make. In order to enable seller Si to make a bid after the best one of seller Sj, D should be less or equal to the difference between the scoring rule of the best bids of Si and Sj. ■

The next question to address is what should the winning bid be? In other words, the question is, at which price will the seller win since the qualities are determined independently of the other sellers' types and bids following Lemma 1? The assumption in an English auction is that a seller bids while its profit is positive. Suppose that seller s_i is the seller with the lowest type θ_i, and seller s_j is the seller with type θ_j that is the second lowest type among the set of bidders. Then, seller s_i actually competes with seller s_j, which is its strongest competitor. Therefore, the prices that seller s_i will offer will decrease until seller s_j quits the auction and this happens when its utility becomes non-positive. From this point, seller s_i has to reduce the price in such a way that will increase its score in D which is the minimal increment allowed in the auction protocol we discuss. In Lemma 3 we define the winning bid considering these assumptions.

Lemma 3

Given a sequential full-information revelation auction protocol, assuming the model described in section 3, if seller s_i has the lowest type θ_i, seller s_j has the second lowest type θ_j, and equation (6) holds, then the ***winning bid*** *is:*

$$\left\{ p = \left(\frac{w_1^2}{2\cdot a} + \frac{w_2^2}{2\cdot b} \right) \cdot \left(\frac{1}{\theta_i} - \frac{1}{2\cdot\theta_j} \right) - D; q_1 = \left(\frac{w_1}{2\cdot a\cdot\theta_i} \right)^2; q_2 = \left(\frac{w_2}{2\cdot b\cdot\theta_i} \right)^2 \right\}. \tag{7}$$

Sketch of proof: Suppose seller s_i has the lowest type θi, and seller Sj has the second lowest type θj. The qualities q1(θj) and q2(θj) are calculated using lemma 1. p(θj) is the minimal value that j can suggest while its utility is still non negative. Given that the last possible bid of seller Sj is (q1(θj),q2(θj), p(θj)), by lemma 2, the winning bidder will be θi, and the qualities of its bid are determined by lemma 1. The price p of the winning bid should cause the scoring of the winning bid to exceed S(q1(θj),q2(θj), p(θj)) by D. Solving the above equation, we obtain the value of p. ■

The qualities of the winning bid are defined following Lemma 1. However, the price depends implicitly on the cost function of the second best seller, denoted by j. In particular, as the expenses of the second best seller increase, i.e., θ_j increases, the winning price offered by seller i increases, but as the expenses of the winning seller increases, its winning price decreases. Intuitively, the reason for this result is that as θ_i increases, the qualities *q1, q2* decrease, so the seller should compensate this by suggesting a lower price, in order to compete with the other bidders. In addition, as θ_j decreases, the second best bidder is stronger, that is, its type is higher, and seller s_i should suggest a more competitive price in order to win against agent j, and thus, the winning price decreases.

Given the information about the auction, we can try to analyze it from the auctioneer's (buyer's) point of view. That is, to calculate the buyer's expected revenue given the environment details. In order to estimate the buyer's expected revenue, we actually have to estimate the best bid and in which probability the buyer may receive it. This brings us to the following theorem that explicitly calculates the buyer's expected revenue ER^E (Expected Revenue for the version of an English auction).

Theorem 1

In a sequential full-information revelation auction protocol of one buyer and n sellers with types independently and identically distributed over $[\underline{\theta}, \overline{\theta}]$*, given equation (6), and given the real weights* W_1, W_2 *of the buyer's utility and the published weights* w_1*,* w_2 *of the scoring, the buyer's expected revenue* ER^E *is:*

$$ER^{E}(\underline{\theta},\overline{\theta}) = \frac{1}{\left(\overline{\theta}-\underline{\theta}\right)^{n}} \cdot n \cdot (n-1) \cdot \Bigg(\left(\frac{w_1^2}{4 \cdot a} + \frac{w_2^2}{4 \cdot b} \right) \cdot \int_{\underline{\theta}}^{\overline{\theta}} \int_{\theta_i}^{\overline{\theta}} \frac{\left(\overline{\theta}-\theta_j\right)^{n-2}}{\theta_j} d\theta_j d\theta_i \tag{8}$$

$$\left(-\left(\frac{w_1^2}{2 \cdot a} + \frac{w_2^2}{2 \cdot b} \right) + \frac{W_1 \cdot w_1}{2 \cdot a} + \frac{W_2 \cdot w_2}{2 \cdot b} \right) \cdot \int_{\underline{\theta}}^{\overline{\theta}} \frac{1}{\theta_i} \cdot \int_{\theta_i}^{\overline{\theta}} \left(\overline{\theta}-\theta_j\right)^{n-2} d\theta_j d\theta_i$$

$$+ D \cdot \int_{\underline{\theta}}^{\overline{\theta}} \int_{\theta_i}^{\overline{\theta}} \left(\overline{\theta}-\theta_j\right)^{n-2} d\theta_j d\theta_i \Bigg)$$

Sketch of proof: The expected revenue is calculated by:

$$\int_{\underline{\theta}}^{\overline{\theta}} \left[\int_{\theta_i}^{\overline{\theta}} \left(U_{buyer}(wining\_bid) \cdot (1-F(\theta_j))^{n-2} \cdot f(\theta_j) \cdot f(\theta_i) \cdot n \cdot (n-1) \right) d\theta_j \right] d\theta_i$$

The expected revenue actually describes the average utility of the buyer from each possible winning bid weighted by the probability of this winning bid, where the winning bid is the bid calculated in Lemma 3. In this double integral, we run all the probabilities of the lowest type θ_i and all the probabilities of the second lowest type θj, multiplying by the buyer's utility from the best bid in the given types. The probability of a certain lowest type θi and second lowest type θj is actually the probability that all the other n-2 sellers have higher types than θ_j, multiplied by the probability of having one seller with type θ_j and having one seller with type θ_i, and this probability is f(θ_i)*f(θ_j)*n*(n-1). By substituting all the explicit functions and simplifying them, we receive the required formula. ■

Notice that the only information that the buyer should get is the sellers' types range and the number of sellers, in order to estimate its expected revenue. It is interesting to notice that if we multiply $\underline{\theta}$ and $\overline{\theta}$ by a given value, and we divide a and b by the same value, the result of ER^E remains the same. This means that the expected revenue does not depend on the exact values of the parameters, but only on the relations between them.

We proceed by finding the optimal scoring rule to be announced by the auctioneer. The ability of predicting the buyer's expected revenue leads to the most interesting phase of the auction design, which is searching for the optimal scoring rule that will maximize the buyer's expected revenue. In other words, finding the optimal weights w_1 and w_2 to be announced. These optimal values of the scoring function can be found by differentiating the expected revenue function once by w_1 and once by w_2.

Theorem 2

Suppose a sequential full-information revelation auction protocol of one buyer and n sellers with types independently and identically distributed over $[\underline{\theta},\overline{\theta}]$, and equation (6) holds. The optimal values of the weights w1 and w2 included in the scoring rule given the real weights W_1 and W_2, of the buyer's utility function, are:

$$w_i = W_i \cdot \frac{\int_{\underline{\theta}}^{\overline{\theta}} \frac{1}{\theta_i} \cdot \int_{\theta_i}^{\overline{\theta}} \left(\overline{\theta} - \theta_j\right)^{n-2} d\theta_j d\theta_i}{2 \cdot \int_{\underline{\theta}}^{\overline{\theta}} \frac{1}{\theta_i} \cdot \int_{\theta_i}^{\overline{\theta}} \left(\overline{\theta} - \theta_j\right)^{n-2} d\theta_j d\theta_i - \int_{\underline{\theta}}^{\overline{\theta}} \int_{\theta_i}^{\overline{\theta}} \frac{\left(\overline{\theta} - \theta_j\right)^{n-2}}{\theta_j} d\theta_j d\theta_i} \tag{9}$$

Sketch of proof: In order to find the weight w_1 of the scoring rule that maximizes the buyer's expected revenue ER^E, we differentiate the function ER^E by w_1.

$$\frac{\partial ER^E(\underline{\theta},\overline{\theta})}{\partial w_1} = \frac{1}{(\overline{\theta}-\underline{\theta})^n} \cdot (n-1) \cdot n \cdot \left(\left[-\frac{2 \cdot w_1}{2 \cdot a} + \cdot \frac{W_1}{2 \cdot a} \right] \cdot \int_{\underline{\theta}}^{\overline{\theta}} \left(\frac{1}{\theta_i} \cdot \int_{\theta_i}^{\overline{\theta}} (\overline{\theta} - \theta_j)^{n-2} d\theta_j \right) d\theta_i \right.$$
$$\left. + \frac{2 \cdot w_1}{4 \cdot a} \cdot \int_{\underline{\theta}}^{\overline{\theta}} \left(\int_{\theta_i}^{\overline{\theta}} \frac{(\overline{\theta} - \theta_j)^{n-2}}{\theta_j} d\theta_j \right) d\theta_i \right)$$

By comparing the differentiation of ER^E to zero the maximum value of w_1 is identified. Then, we differentiate this function by w_2 in a similar manner since the weights w_1 *and* w_2 are independent. ■

Given the above results, we can check the influence of different parameters on the ratio w1/W1, and symmetrically, on *w2/W2*. First, it can be shown from the formula of *w1/W1* that the ratio depends on the ratio $\overline{\theta}/\underline{\theta}$ and on n, and no other factors influence this ratio. That is, if $\underline{\theta}$ and $\overline{\theta}$ are multiplied by the same number, the announced weight will remain unchanged. Notice that the relation $\overline{\theta}/\underline{\theta}$ actually defines the relative distribution of the sellers' types. Namely, if the relative distribution is very low the sellers are more or less of the same type, namely, they are homogenous. However if the relative distribution is high the sellers are heterogeneous.

Figure 1 shows the influence of the number of bidders *n* (2..20) on *w1/W1*, given $\underline{\theta}$=0.5, $\overline{\theta}$=1 (circles), and given $\underline{\theta}$=0.2, $\overline{\theta}$=1 (line). We can see that the announced weight is lower than the real weight of the utility function, but as *n* increases, the announced weight approaches the real one. However, as the ratio between $\overline{\theta}/\underline{\theta}$ increases, the initial ratio of w1/W1 is lower, and its convergence is slower. This is shown by comparing the circles graph, where the ratio $\overline{\theta}/\underline{\theta}$ is 2, with the lined graph, where the ratio $\overline{\theta}/\underline{\theta}$ is 5. Finally, Figure 2 shows that as the ratio $\overline{\theta}/\underline{\theta}$ increases, given a particular *n*, the announced weights decrease, and their difference from the real weights increase.

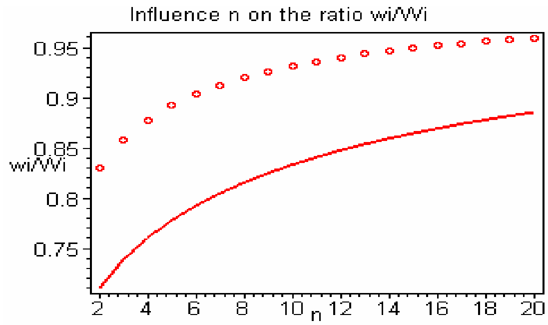

Fig. 1. The influence of n on w1/W1. Circles: $\underline{\theta}$=0.5, $\overline{\theta}$=1. Line: $\underline{\theta}$=0.2, $\overline{\theta}$=1

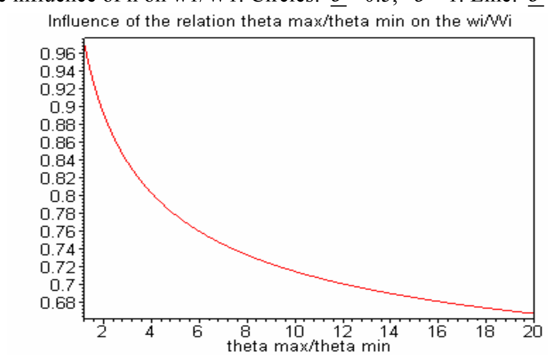

Fig. 2. The influence of $\overline{\theta}/\underline{\theta}$ on w1/W1, where n=4

These two results can be explained as follows. First, it is clear that as *n* increases, the bidders' natural competition is higher, and thus the revenue of the auctioneer is higher. (This can be inferred from the expected revenue formulas.) Thus, as *n* increases, the gains of the auctioneer become higher and the additional gains expected from manipulating the scoring rule decrease, so the optimal weights it publishes are closer to the real ones. Second, we should explain why increasing the ratio $\overline{\theta}/\underline{\theta}$, which is the relative distribution of the sellers' types, causes the announced weight to decrease. As the relative distribution increases, the sellers become more and more heterogeneous. That is the competition among the sellers decreases and therefore the natural revenue for the auctioneer/buyer decreases. In this case manipulating the scoring rule can increase the expected revenue. As shown in Figure 2 as the relative distribution increases the auctioneer will be motivated to publish weights with a larger deviation from the real ones.

In an ongoing research, we have been comparing the expected revenue of the auctioneer when using the English auction with its expected revenue in the first price auction that was presented in [7]. We revealed that a constant difference of *D* exists between the expected revenue of the auctioneer when using the English auction and when using the first price sealed bid auction. We also found that the announced weights are the same for both models. For space reasons we do not present these results here.

4 The English Auction with Two Cost Parameters per Each Seller

In this section, we consider an extension of the model discussed in Section 3. In particular, we consider a model where the cost of the seller depends differently on q_1 and q_2, and there are two private cost parameters per each seller. This extension is very interesting since another dimension of uncertainty is added to the model which the bidders and the auctioneer should take it into consideration in order to behave optimally in such an auction.

The cost function of each seller is:

$$U_s(q_1, q_2, \theta_i^1, \theta_i^2) = p - \theta_i^1 \cdot q_1 - \theta_i^2 \cdot q_2 \qquad \textbf{(10)}$$

We assume that both θ_i^1 and θ_i^2 of seller *i* are unknown by the other participating sellers in the auction. Since there are two private cost parameters, the buyer should consider both when evaluating its expected utility and when deciding about its optimal scoring function. The idea behind this model is that different sellers may have different abilities in producing items: it can be easier for one seller to provide a higher quality of attribute 1, while it is easier for another seller to increase the second quality attribute. For example, in task allocation, suppose that each bid is composed of the time when the task will be completed and the accuracy level of the task performance. Then, it may be easier for one machine to perform the task earlier, while it may be easier for another machine to return a higher accuracy result. In the service domain, suppose each bid includes the service quality and its supply time, and again, different companies may have different abilities regarding these properties. Thus, it seems to be more reasonable to assume that the cost to supply a bid is based on different private cost parameters.

The utility function of the buyer remains as in the previous section, and the format of the scoring rule published by the buyer also remains the same. In the following, we will reveal the optimal strategies to be used by the sellers given their types. Based on these strategies, we will show the expected revenue of the auctioneer, and then we will be able to describe how the auctioneer will find the optimal scoring rule to be published at the beginning of the auction. Similarly to the model presented in Section 3, we start by considering the optimal bid to be made by each bidder at each step of the auction, given the announced scoring rule, and given the private cost parameters θ_i^1 and θ_i^2.

Lemma 5

Suppose that there are n sellers including seller i, and j, where i is the seller that can offer the best bid from the scoring function's point of view without a loss, *and seller j is the seller that can offer the second best bid after seller i, without a loss. Then the winning bid is:*

$$\begin{Bmatrix} p = \frac{w_1^2}{2 \cdot \theta_i^1} + \frac{w_2^2}{2 \cdot \theta_i^2} - \frac{w_1^2}{4 \cdot \theta_j^1} - \frac{w_2^2}{4 \cdot \theta_j^2} - D; \\ q_1^*(\theta_i^1) = \left(\frac{w_1}{2 \cdot \theta_i^1} \right)^2, q_2^*(\theta_i^2) = \left(\frac{w_2}{2 \cdot \theta_i^2} \right)^2 \end{Bmatrix} \quad \textbf{(11)}$$

Sketch of proof: Similarly to Lemma 1, we calculate the last bid of seller j, which includes the lowest possible price that seller j can bid without a loss. Then, we find the winning bid, which is offered by seller i. We first show that the optimal qualities *q1* and *q2* of the winning bid are determined independent of the price and the current bid, and we find their optimal values. In order to be able to offer the winning bid, its score should be higher than the score of seller j's last bid by *D,* Thus, the price of the winning bid is determined as the value for which the score of the bid is higher by *D* than the score achieved by the last bid of seller j. ■

The optimal bid is similar to that of Lemma 1, but instead of considering the cost coefficients $a\theta$ and $b\theta$, as in the one θ model, we consider the cost coefficients of the 2-θ model, namely, θ^1 and θ^2. Note that each optimal quality attributes *q1* and *q2* depends only on the private cost parameters θ_i^1 and θ_i^2 and on the announced weights w1 and w2, respectively. This makes it easy to generalize the model in order to also consider more than two quality attributes, since the values of the attributes are selected only with regard to their coefficients in both the sellers' utility function and in the buyer's scoring function. The optimal price to be offered in each step of the auction depends on the last proposed bid, since it must be better than the previous bid by at least *D*.

Given the optimal bid to be offered in each step of the auction, it is possible to find the best two sellers, given their private information types, and given the best two sellers, it is easy to calculate the winning bid. To find the winner seller, we calculate for each seller the price p for which its utility becomes 0, and we calculate the scoring function for the bid including price p for which the seller utility is 0, and qualities q1* and q2*, as evaluated by Lemma 5. The seller that is able to suggest a bid with the highest score is the strongest seller, since no other seller is able to suggest a better bid regarding the scoring function, without losing. Similarly, the second-highest seller can be found, by choosing the seller that achieves the second highest scoring value when suggesting the lowest p for which it does not lose.

Assume that seller i is the strongest seller, and seller j is the second best seller, with regard to the announced scoring function. Then, we can compute exactly what the utility of the winning bid suggested by seller i, is using the fact that seller i's bid comes after the best bid that seller j could suggest. We denote this utility by *Ubuyer*. Using *Ubuyer*, we can calculate the expected revenue from the auction, by considering all possible cost parameters for both best sellers, seller i and seller j.

Theorem 3

Given the utility of the buyer from the winning bid

$$Ubuyer = -\left(\frac{w_1^2}{2\cdot\theta_i^1}+\frac{w_2^2}{2\cdot\theta_i^2}-\frac{w_1^2}{4\cdot\theta_j^1}-\frac{w_2^2}{4\cdot\theta_j^2}-D\right)+W_1\cdot\left(\frac{w_1}{2\cdot\theta_i^1}\right)^2+W_2\cdot\left(\frac{w_2}{2\cdot\theta_i^2}\right)^2 \quad \textbf{(12)}$$

and given the probability of an arbitrary seller to be lower than seller j,

$$F(\theta_j^1,\theta_j^2)=\frac{1}{\left(\overline{\theta}-\underline{\theta}\right)^2}\cdot\int_{\underline{\theta}}^{\overline{\theta}}\int_{\frac{1}{\frac{w1^2}{w2^2}\cdot\left(\frac{1}{\theta_j^1}-\frac{1}{\theta_a^1}\right)+\frac{1}{\theta_j^2}}}^{\overline{\theta}} d\theta_a^2\cdot d\theta_a^1 \quad \textbf{(13)}$$

The expected revenue of the buyer agent from this auction is:

$$Expected\,Revenue(\underline{\theta},\overline{\theta})= \quad \textbf{(14)}$$

$$\int_{\underline{\theta}}^{\overline{\theta}}\int_{\underline{\theta}}^{\overline{\theta}}\int_{\underline{\theta}}^{\overline{\theta}}\int_{\frac{1}{\frac{w1^2}{w2^2}\cdot\left(\frac{1}{\theta_i^1}-\frac{1}{\theta_j^1}\right)+\frac{1}{\theta_i^2}}}^{\overline{\theta}} Ubuyer\cdot\left(F(\theta_j^1,\theta_j^2)\right)^{n-2}\cdot\frac{1}{\left(\overline{\theta}-\underline{\theta}\right)^4}\cdot n\cdot(n-1)d_j^2\cdot d\theta_j^1\cdot d\theta_i^2\, d\theta_i^1$$

Sketch of proof: We run all the possible types of seller i and seller j, and for each pair, we calculate the probability for such a pair to exist, and to be the best sellers, multiplied by the utility of the buyer in this case. The utility *Ubuyer* is calculated by substituting the values p*,q1* and q2* of the winning bid of seller i, as calculated in Lemma 5, given seller i and seller j, in the formula of the buyer's utility. The probability for the pair, seller i and seller j, to be the best bidders is equal to the probability of all the other n-2 sellers to be "weaker", w.r.t. the scoring function.

We denote the probability for one arbitrary seller to be weaker than agent j (and thus, weaker also than agent i) by $F(\theta_j^1, \theta_j^2)$. In order for a given seller a to be less efficient than seller j, it should hold that the score of the best bid seller a can suggest will be lower than the score of the best bid of seller j. The best bid a seller can suggest is evaluated by assigning the optimal qualities values from lemma 5, and finding the lowest possible price by comparing the utility of the seller to 0. We obtain the following condition: $w_1^2/(4\theta_j^1)+w_2^2/(4\theta_j^2)>w_1^2/(4\theta_a^1)+w_2^2/(4\theta_a^2)$, and using this equation, we can find the condition of θ_a^2 as a function of θ_j^1,θ_j^2 and θ_a^1.

We will now to explain the expected revenue formula. We should consider the values of θ^1 and θ^2 for both seller i and seller j. For this purpose, we run all possible values of θ^1 and θ^2 of both, and thus we need 4 levels of integration. Since the winner seller is i, seller j should be weaker than seller i. Thus, $\theta_j{}^2$ runs only from a particular value, which depends on the scoring function, on the θ values of seller i, and on θ_j^1. The particular limit for the value of $\theta_j{}^2$ is found similarly to the limit for θ_a^2, inside the formula $F(\theta_j^1, \theta_j^2)$. ■

In the model of one cost parameter, it was simple to find out the probability of an arbitrary seller to be worse than seller j, since it only requires t that θ of this seller will be higher than θj. In this model, the strength of a seller depends on θ^1 and θ^2, and also on the announced scoring function, since only given the scoring function, we can evaluate which seller can make a better bid without losing. Thus, also calculating the probability for a seller to be weaker than some other sellers becomes more complex, as we stated in the last theorem. Consider the possible values for θ_j^2 for which j is still "weaker" than i. Intuitively, we can see that as θ_j^1 decreases, there are more possible values of θ_j^2 for which seller j is still "weaker" than seller i. Whenever θ_i^1 or θ_i^2 increase, i.e., agent i is weaker, agent j must also be weaker, i.e., θ_j^2 starts from a higher limit. Finally, as w1 increases w.r.t. w2, q1 becomes more important, so θ_j^2 should be smaller in order for j to be weaker.

Given the expected revenue of the buyer, the buyer agent is now able to calculate the optimal values of w1 and w2, in order to maximize its expected revenue. These values can be found by differentiating the expected revenue formula according to w1 and w2, and by finding when the differentiation is equal to zero. The optimal values of w1 and w2 will be used by the buyer when publicizing its optimal scoring function. Differentiating the optimal scoring function is difficult in general, but given a particular auction, with known values of n and the utility function, and given the buyer's beliefs about the seller's distribution, it becomes easier for the buyer to differentiate the optimal *w1* and *w2* to be announced.

5. Conclusion

In this paper, we consider a version of the English auction protocol that suits the case of multi-attribute auctions, where the auctioneer is the buyer of an item or a service, and various sellers bid and offer diverse configurations of the item or service they were asked for, and we suggest stable and beneficial strategies for the buyer agent and for the seller agents participating in the auction. We also consider a situation where two private cost parameters, θ_1 and θ_2, are associated with each seller. We show how to analyze this model and suggest how to derive the optimal scoring function. The optimal scoring rule ensures that the buyer will maximize its expected revenue assuming the sellers' types range. We show that as the number of sellers increases, the buyer has no motivation to manipulate since the natural competition among the sellers will achieve the maximum gain for the buyer. However, in case of small number of sellers we show that the buyer is motivated to manipulates the scoring rule in such a way that he announces weights will be lower than the real weights, in order to achieve the best result.

In future work, we intend to consider more general utility functions and the distribution function, and check how it will influence the auction results. The extension to more than two attributes can follow the same steps we took to extend from one attribute to two attributes. We also intend to consider various English auction formats under the constraints of a time deadline, and to suggest stable and efficient strategies for this case.

References

1. Andersson, A. Tenhunen, M. and Ygge, F. Integer programming for combinatorial auction winner determination. In ICMAS-00, Boston, MA (2000) 39-46.
2. Bichler, M. An experimental analysis of multi-attribute auction. *Decision Support Systems* Vol. 29, 2000, p. 249-268.
3. Branco. F. The Design of Multidimensional Auctions. Rand Jurnal of Economics Vol. 28 N0. 1 Spring 1997, pp. 63-81.
4. Che, Y. K., Design competition through multidimensional auctions, RAND Journal of Economics, Vol. 24, (1993) 668-680.
5. Collins, J. Demir, G. and Gini, M. Bidtree Ordering in IDA* Combinatorial Auctionb Winner Determination with Side Constraints. Proceedings of AMEC-IV LNCS No. 2531 (2002).
6. Conen, W. and Sandholm, T. Differential-Revelation VCG Mechanisms for Combinatorial auctions, Proceedings of AMEC-IV LNCS No. 2531(2002).
7. David, E., Azoulay-Schwartz, R., Kraus, S., Protocols and Strategies for Automated Multi-Attributes Auctions, in Proc. of the 1st conference on autonomous agents and multi-agent systems, Bologna, Italy (2002), to appear.
8. Gimenez-Funes, E., Godo, L., Rodriguez-Aguilar, J. A. and Garcia-Calves, P. Designing bidding strategies for trading agents in electronic auctions. In ICMAS-98, Paris, France (1998) 136-143.
9. Hudson, B. and Sandholm, T, Effectiveness of Preference Elicitation in Combinatorial Auctions. , Proceedings of AMEC-IV LNCS No. 2531 (2002).
10. Nisan, N., Bidding and allocation in combinatorial auctions. In Proc. of the ACM Conference on Electronic Commerce (ACM-EC), Minneapolis, MN (2000) 1-12.
11. Parkes, D. C. Price-Based Information Certificates for Minimal-Revelation Combinatorial Auctions. Proceedings of AMEC-IV LNCS No. 2531 (2002).
12. Parkes, D. C., and Ungar, L. H. 2000b. Iterative combinatorial auctions: Theory and practice. In proc. AAAI-2000 (2000) 74-81.
13. Rosenschein, J., and Zlotkin, G. Rules of encounter: designing conventions for automated negotiation among computers. Cambridge, Mass. MIT Press. (1994)
14. Sandholm, T.W. Limitations of the Vickrey Auction in Computational Multiagent Systems. In proc. of ICMAS-96 (1996) 299-306.
15. Sandholm, T.W. An algorithm for optimal winner determination in combinatorial auctions. In IJCAI-99 (1999) 542-547.
16. Stone, P. Schapire, R. E. and Csirik, J. A.Attac-2001:A Learning Autonomous Bidding Agent. Proceedings of AMEC-IV LNCS No. 2531 (2002).
17. Walsh, W.E., Wellman, M. P., Wurman, P.R. and MacKie-Mason, J. K. Auction protocols for decentralized scheduling. In Proc. 18th International Conference on Distributed Computing Systems (1998).
18. Yoon, K. and C. Hwang. Multiple attribute decision making: an introduction. Thousand Oaks: Sage (1995).