Designing Trading Agents for Real-World Auctions

Ioannis A. Vetsikas and Nicholas R. Jennings

Electronics and Computer Science, Univ. of Southampton, Southampton SO17 1BJ, UK
{iv,nrj}@ecs.soton.ac.uk

Abstract. Online auctions have become a popular method for business transactions. The variety of different auction rules, restrictions in supply or demand, and the agents' combinatorial preferences for sets of goods, mean that realistic scenarios are very complex. Using game theory, we design trading strategies for participating in a single auction or group of similar auctions. A number of concerns need to be considered in order to account for all the relevant features of real-world auctions; these include: budget constraints, uncertainty in the value of the desired goods, the auction reserve prices, the bidders' attitudes towards risk, purchasing multiple units, competition and spitefulness between bidders, the existence of multiple sources for each good. To design a realistic agent, it is necessary to analyze the multi-unit auctions in which a combination of these issues are present together and we have made significant progress towards this goal. Furthermore, we use a principled methodology, utilizing empirical evaluation, to combine these results into the design of agents capable of bidding in the general real-world scenarios.

1 Introduction

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

Game theory is widely used in multi-agent systems as a way to model and predict the interactions between rational agents in auctions. However, the models that are canonically analyzed are rather limited, because the work that incorporates features important in real auctions looks at each feature separately and in most cases analyzes single-unit auctions. Now, while this is useful for economists and perhaps expert bidders, who can integrate the lessons learned using human intuition and imagination, an automated agent cannot immediately benefit. In order to design agents that would represent non-expert humans in real auctions effectively, it is necessary to first analyze the strategic behavior of bidders in auctions that incorporate all the features deemed to be relevant and important in real world auction scenarios; this is the first part of *our agenda towards*

S. Konstantopoulos et al. (Eds.): SETN 2010, LNAI 6040, pp. 275–285, 2010.

the design of trading agents for realistic auctions. To identify these features we looked at a variety of real auctions and at the related literature. We concluded that the following issues capture the essence of almost any real auction[1]: **(i)** *budget constraints* as, in every practical setting, the participating bidders have a certain pot of money that they can spend for purchases and how much they can spend (and thus bid) is limited by this budget [1,2]; **(ii)** *bidders' attitudes towards risk*, meaning whether bidders are conservative or not and how much they are willing to take risks in order to gain a higher profit, which is described by their utility function, which maps profit into utility [3,4]; **(iii)** the *reserve price* of the specific auction, which is the minimum transaction price allowed in this auction [4,5,6]; **(iv)** *uncertainty in the bidders' valuation* for the offered good, as there are many cases when the value of a good or service is not known precisely a priori, i.e. before it is purchased and used, which could be addressed by introspection on the part of the agent [7,8,9]; **(v)** *competition or spite between bidders*, especially in scenarios where one bidder wants to drive a competitor out of the market [10]; **(vi)** *auctions with multiple rounds of bidding*, because they have multiple possible closing times [11]; **(vii)** bidders' desiring to *purchase multiple items* of the good sold, with a different valuation for each one [12,13]; **(viii)** the availability of each good from *multiple sellers* instead of assuming that all of them are sold in a single auction [3,14]; **(ix)** auctions with *asymmetric bidder models*, as in reality it could be known that some bidders have higher valuations or budgets, or different risk attitudes [15,16].

Furthermore, we consider the fact that, in practice, agents (or humans) are rarely interested in a single item, but rather wish to bid in several auctions for multiple interacting goods. Therefore, they must bid intelligently in order to get exactly what they need, as, for example, they may need to acquire a whole set of items.[2] The second part of our agenda is to design agents for this general case. Unlike pre-existing work, e.g. [17,18], which focused mostly on specific cases or used heuristics to tackle this problem, we want to have a principled methodology for design trading agents, which allows us to use the theoretical results obtained from the theoretical analysis we described earlier.

In order to achieve the goals set by our agenda, we first need to analyze an auction where all the identified relevant features are included in our theoretical analysis. In this paper, we present some selected results we have obtained thus far towards this goal:

(i) the analysis of sealed-bid auctions when budget constraints, reserve prices, varying risk attitudes and valuation uncertainty are examined together [19,20,21] (in section 3); we identify cases where the analysis of two features together produces completely different behaviour than that observed from the analysis of each feature separately. We also present evidence that our strategies work well even when the opponents do not follow the equilibrium strategy, which is not guaranteed by the equilibrium definition.
(ii) the equilibria for auctions with competitive or spiteful bidding [22] (in section 4); there are significant differences between our results and those for single-unit auctions.

---

[1] The citations mostly refer to previous work where the specific feature is examined by itself in the context of single-unit auctions.

[2] For example a person may wish to buy a TV and a VCR, but if she does not have a "flexible plan", she may only end up acquiring the VCR. On the other hand, she bids for VCRs in several auctions, she may end up with more than one. Goods are called complementary (substitutable) if the value of acquiring two of them is more (less) than the sum of their individual values.

(iii) the equilibria for the case of multi-round auctions [11] (in section 5); this a completely novel analysis that has been used in our methodology for general settings.

Thus, we selected to present these results, because not only do they demonstrate the significant progress already made towards our goal, but also show that our analysis provides qualitatively different results to the existing literature looking at simpler versions of the same problems. We have also obtained results on auctions with asymmetric models [23], sequential auctions with partially substitutable goods [24] which extends the standard model found in [3] and some initial results on auctions with bidders having multi-unit demand, which we cannot include in this paper due to space constraints.

In addition, we present our principled methodology for achieving the second part of our agenda. According to this methodology, the problem is decomposed into several subparts, strategies are generated for each subproblem using game-theoretic analysis and empirical knowledge and then systematic experimentation is used to determine the strategies that work best for the whole problem in question. We present this work in section 6. We applied this technique to the Trading Agent Competition (TAC) classic game [25]. The resulting agent WhiteBear is the best agent in the history of the competition, which confirms the validity of our approach.

2 The Multi-unit Auction Model Used in Our Analysis

In this section we formally describe the auction setting to be analyzed and define the objective function that the agents wish to maximize. We also give the notation that we use in the results included in this paper; some additional notation would be required for the model to incorporate the remaining features, which are not presented in this paper.

In particular, we will compute and analyze the symmetric Bayes-Nash equilibria for sealed-bid auctions where $m \geq 1$ identical items are being sold; the Bayes-Nash equilibrium is the standard solution concept used in game theory to analyze auctions and other games. These equilibria are defined by a strategy, which maps the agents' valuations v_i to bids b_i. The two most common settings in this context are the m^{th} and $(m+1)^{th}$ price auctions, in which the top m bidders win one item each at a price equal to the m^{th} and $(m+1)^{th}$ highest bid respectively. We assume that there is a reserve price $r \geq 0$ in our setting; this means that bidders, who wish to participate in the auction, must place bids $b_i \geq r$. N indistinguishable bidders (where $N \geq m$) participate in the auction and they have a private valuation (utility) v_i for acquiring any one of the traded items; these valuations are assumed to be i.i.d. drawn from a distribution with cumulative distribution function (cdf) $F(v)$, which is the same for all bidders.

In the case that there is uncertainty about the valuation v_i, then we need to extend this model. For the $(m+1)^{th}$ price multi-unit auction case, we can use the most general model possible: the agent knows that his valuation v_i is drawn from distribution $G_i()$, but not the precise value. As the valuations v_i are independent, we can assume that any uncertainty that a bidder has about his own valuation is independent of the uncertainty he has about other agents' valuations. For the m^{th} price multi-unit auction case, we use a simpler model, because unlike the previous case, no dominant strategies exist in these cases, therefore the strategy used by opponent bidders affects the strategy used by any specific bidder. Thus, we assume that the true valuation $\overline{v}_i$, which is not known to bidder i, is drawn from distribution $G_{v_i}()$, where v_i is known as being the mean value of distribution $G_{v_i}()$; hence, each bidder i knows approximately his own value as being drawn from distribution $G_{v_i}()$ around the value v_i (and v_i is a known value to bidder i).

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

Every agent has a strictly monotonically increasing utility function $u()$ that maps profit into utility and tries to maximize this utility. This function determines the agent's risk attitude. In the case of spiteful bidding, then the objective function that each agent wishes to maximize is given by: $U_i = (1-\alpha)\cdot\widetilde{u}_i - \alpha\cdot\sum_{j\neq i}\widetilde{u}_j$ where $\alpha \in [0,1]$ is a parameter called the spite coefficient, $\widetilde{u}_i$ is the utility-mapped gain of agent i (i.e. $\widetilde{u}_i = u(0)$, if it does not win any items, and $\widetilde{u}_i = u(v_i - p_i)$, if it does) and p_i is the total payment the agent must make to the auctioneer.

Finally, in the case of multiple round auctions, we assume that the valuations v_i^r at round r are i.i.d. drawn from distribution $F_r(u)$, and the probability that round r is the last round is known to be p_r. If more rounds exist, an agent can submit new bids as long as they are greater or equal to the bid price from the end of the previous round; this is the minimum bid allowed at round r which is denoted as Q_r.

3 Equilibria When Budget Constraints, Reserve Prices, Varying Risk Attitudes and Valuation Uncertainty Are Present Together

In this section, we examine the case where a number of issues are taken into account. More specifically, the agents have budget constraints c_i, their risk attitude is described by function $u()$ (not necessarily risk neutral) and the auction has a reserve price r. In addition to these issues, we also consider that the valuations v_i are not known precisely, and we use the models of valuation uncertainty described in the previous section. In [21], we proved the following theorems:

Theorem 1. *In the case of an m^{th} price sealed-bid auction, with reserve price $r \geq 0$, with N participating bidders, in which each bidder i is interested in purchasing one unit of the good for sale with inherent utility (valuation) for that item equal to $\overline{v}_i$ (the exact value unknown to the bidder, and drawn from distribution $G_{v_i}()$), which is approximated by a known variable v_i, the agent's "uncertain valuation", and has a budget constraint c_i, where v_i and c_i are i.i.d. drawn from $F(v)$ and $H(c)$ respectively, and the bidders have a risk attitude which is described by utility function $u()$, the following bidding strategy constitutes a symmetric Bayes-Nash equilibrium:*

$$b_i = min\{g(v_i), c_i\} \tag{1}$$

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

$$g'(v_i) = \frac{(1-H(g(v_i)))F'(v_i)}{\frac{(1-(1-F(v_i))(1-H(g(v_i))))\int_{-\infty}^{\infty} u'(x-g(v_i))G'_{v_i}(x)dx}{(N-m)\left(\int_{-\infty}^{\infty} u(x-g(v_i))G'_{v_i}(x)dx-u(0)\right)} - (1-F(v_i))H'(g(v_i))} \tag{2}$$

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

Theorem 2. *In an $(m+1)^{th}$ price auction, with reserve price r, if a bidder has budget constraint c_i, he knows only imprecisely his own valuation v_i, in that it is drawn from distribution $G_i(v_i)$, and his risk attitude is described by utility function $u_i()$, it is a dominant strategy to bid: $b_i = min\{\beta_i, c_i\}$, if $b_i \geq r$, and not to participate otherwise. The variable β_i is the solution of equation:*

$$\int_{-\infty}^{\infty} u(z-\beta_i)G_i'(z)dz = u(0) \qquad (3)$$

While, sometimes, the bidding behaviour is a combination of those observed when the individual features are present separately, there are cases when this does not occur. For example, it is known that a risk-neutral bidder with uncertainty in his valuation v_i bids as if the value was known and equal to the average of the expected value μ_{G_i}.[8] When bidders are not risk-neutral this is no longer the case. In [21], we prove theoretically, for the case of an $(m+1)^{th}$ price auction, that the more risk-averse (resp. risk-seeking) the bidders become and the higher the uncertainty, meaning the variance of the valuation distribution, the less (resp. more) they will bid.

A similar result holds for the m^{th} price auction as well; in figure 1(a) we graph the bidding strategies of risk averse bidders, for different degrees of valuation uncertainty. To be more precise, we assume all bidders use utility function $u(x) = (x+1)^{0.01}$, and also that if they have uncertain valuation v_i, then their true valuation for the good sold (which is unknown to them) is uniformly distributed in $[v_i(1-\gamma), v_i(1+\gamma)]$, thus $G'_{v_i}(x) = \frac{1}{2\gamma v_i}$, for $x \in [v_i(1-\gamma), v_i(1+\gamma)]$ and $G'_{v_i}(x) = 0$, otherwise. When $\gamma = 0$, there is no valuation uncertainty, and as γ increases so does the uncertainty. We observe that as γ increases, the bidders will indeed bid lower (for any valuation v_i). The same happens as the bidders' risk averseness increases. As can been seen in the graph, while for small values of the parameter γ the bidders bid higher than $g(v) = \frac{v}{2}$, which is what a risk neutral bidder would bid, on the other hand, for large values of γ, they bid less than that. Therefore, when the valuation uncertainty is significant, the risk averse bidders bid in the same way as risk seeking bidders would (when the latter have no valuation uncertainty). Without our complete analysis, it would not be possible to figure out how these two effects factor in the final equilibrium strategy. This validates our claim that it is important to include all features in our analysis of a realistic auction.

We also examined what happens when several opponents deviate from the equilibrium strategy; there are no theoretical guarantees that our strategies would outperform the other strategies in that case. We simulated an m^{th} price auction with $N = 3$ participating bidders, where $m = 2$ items are sold. The bidders are all risk-averse with utility function $u(x) = x^{\alpha}, \alpha = 0.5$ and they have budget constraints c_i and valuations v_i drawn from uniform distribution $U[0, 1]$. We denote the standard equilibrium strategy (given by equation 2) as *S* and compare it against the following two strategies: (i) *NB* is the strategy when the agent does not take the budget constraint into account, and (ii) *RN* is the strategy when the agent does not take the risk attitudes into account (and assumes that everyone is risk neutral). We choose these two strategies, because they look at less features than strategy *S*. In this sense, they are strategies which do not take advantage of our full analysis, and yet are reasonable, because they do consider some of the desired features. We compare *S* against each of these two strategies by running experiments

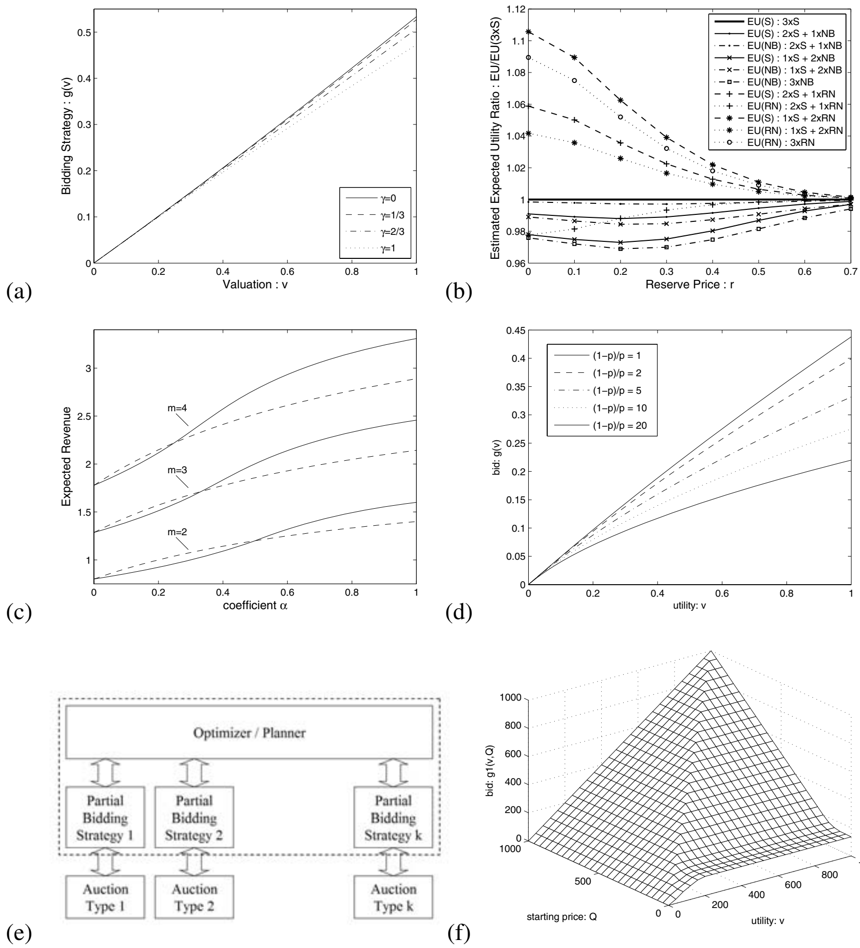

Fig. 1. Graphs of the experimental results, as well as the diagram of our methodology for decomposing the general problem into subproblems, each of which can then be analyzed separately

in which some agents bid according to *S* and some according to *NB* (according to *RN* in the second comparison we did), for various values of the reserve price r. The results are presented in figure 1(b); they are presented as the ratio of the corresponding utility divided by the utility of the case when all agents use strategy *S* (experiment "3xS"). From this figure we can observe that, in every single instance, an agent using strategy *NB* (or *RN*) would always obtain a higher utility by switching to strategy *S*. This means that strategies *NB* and *RN* are *dominated* by *S*, when we consider agents who play either *NB*/*RN* or *S*. This fact that *the trading strategies we compute analytically perform well even when opponents deviate from the equilibrium strategies* has also been observed in experimental simulations we conducted in many other settings.

4 Equilibria in the Case of Competitive or Spiteful Bidding

In [22], we examine the case of spiteful bidding. We define v_l and v_h to be respectively the lowest and highest value that the bidders' valuations can take. We proved that:

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

$$g_\alpha(v) = \begin{cases} v - (F(v))^{-\frac{N-m}{1-\alpha m}} \int_{v_l}^{v} (F(z))^{\frac{N-m}{1-\alpha m}} dz & \alpha m < 1 \\ v & \alpha m = 1 \\ v + (F(v))^{-\frac{N-m}{1-\alpha m}} \int_{v}^{v_h} (F(z))^{\frac{N-m}{1-\alpha m}} dz & \alpha m > 1 \end{cases} \tag{4}$$

Theorem 4. *In the case of the equivalent $(m+1)^{th}$ price sealed bid auction, the symmetric Bayes-Nash equilibrium is given by the strategy:*

$$g_\alpha(v) = v + (1 - F(v))^{-\frac{1}{\alpha}} \cdot \int_{v}^{v_h} (1 - F(z))^{\frac{1}{\alpha}} \cdot dz \tag{5}$$

We note some notable differences when analyzing the multi-unit auction case compared to the existing results for the single-unit auction case, which were presented in [10]. Firstly, bidders in an m^{th} price auction would now bid higher than their true value v_i, when $\alpha > \frac{1}{m}$. An even more interesting results is obtained when examining the revenue of the seller under the two auction formats. In the single-unit case, the $(m+1)^{th}$ price auction yields more revenue then the equivalent m^{th} price one, therefore the seller should always use this auction. In the extended (multi-unit) setting, for the same coefficient α, the m^{th} and $(m+1)^{th}$ price auctions yield the same expected revenue, at equilibrium, when $\alpha = 0$ or $\alpha = \frac{1}{m}$, while, for $0 < \alpha < \frac{1}{m}$, the $(m+1)^{th}$ price auction yields more revenue, and for $\alpha > \frac{1}{m}$, the m^{th} price auction yields more revenue. This is shown in figure 1(c), where we graph the revenue for the two auction types when the spite coefficient a is varied; the number of items is $m = 2, 3, 4$, the number of bidders $N = 2m$ and the distribution $F = U[0, 1]$ (uniform). This validates our claim that it is important to analyze the multi-unit setting of each auction.

5 Equilibria for Multi-round Auctions

In this section, we present the case when an auction does not close at one preset time; we assume that there is a set of possible closing times, each one having a known probability of being selected. We analyze this case in [11]:

Theorem 5. *If the starting price of the current round r is $Q_r \geq 0$, the next round of bidding $(r+1)$ exists with probability $(1-p_r)$ ($p_r \neq 0, 1$) and the utility of the agents in round r is drawn from the distribution $F_r(v)$ (and each agent i in fact has utility v_i^r of a similar value to the utility v_i of the first round) then the equilibrium strategy is the solution of the differential equation:*

$$\big(v_i - g_r(v_i)\big) \cdot \frac{\Phi_r'(v_i)}{g_r'(v_i)} = \big(\Phi_r(v_i) - Y_r(v_i)\big) \cdot \Psi_r(v_i, g_r(v_i)) \tag{6}$$

where $\Psi_r(v,x) = 1 - \frac{1-p_r}{p_r} \cdot \frac{\partial \mathcal{U}_{r+1}(v,x)}{\partial x}$, *and* $\mathcal{U}_{r+1}(v_i, Q_{r+1})$ *is the expected utility at round* $(r+1)$*, when the agent's valuation is* v_i *and the starting price is* Q_{r+1}.
The boundary condition is $g(Q_r) = Q_r$.
In addition, the expected utility at round r *given this strategy* $g_r(v_i)$ *is then:*

$$\mathcal{U}_r(v_i, Q_r) = p_r \cdot \big\{(v_i - g_r(v_i)) \cdot \Phi_r(v_i) + \int_{Q_r}^{v_i} Y_r(\omega) \cdot g_r'(\omega) \cdot d\omega\big\} + (1-p_r) \cdot \quad (7)$$

$$\big\{ \int_{Q_r}^{v_i} \mathcal{U}_{r+1}(v_i, g_r(\omega)) Y_r'(\omega) d\omega + \mathcal{U}_{r+1}(v_i, g_r(v_i)) \big\{\Phi_r(v_i) - Y_r(v_i)\big\}$$
$$+ \int_{v_i}^{g_r^{-1}(v_i)} \mathcal{U}_{r+1}(v_i, g_r(\omega)) \Phi_r'(\omega) d\omega \big\}$$

We present this result, because we have used this equilibrium strategy in our methodology, which is discussed in the next section. In figure 1(d), we graph the equilibrium strategy for $R = 2$ rounds and various values of $\frac{1-p}{p}$, where $p = p_1$ is the probability of the auction closing after the first round of bids. As the probability $(1-p)$ of a second round increases, the equilibrium strategy is that the agent should bid progressively less.

6 A Principled Methodology for Designing Trading Agents

In the previous sections, we demonstrated that equilibrium strategies can be quite useful in designing trading agents. However, the game-theoretic analysis has its limitations. It is not possible to fully analyze a setting with various different auction types, each with different rules, selling different commodities, when these commodities are complementary or substitutable to each other. In order to design an agent for such a setting, in [26], we presented a methodology that allows us to use the equilibrium strategies computed in the design of our agents. The high-level description of the proposed methodology is:

A. Decompose the problem into subproblems as shown in figure 1(e):
  1. OPTIMIZER: This component decides the quantities to buy assuming that everything will be bought at current or predicted prices (*optimize utility*)
  2. CANDIDATE STRATEGIES: For each different auction type (and good) do:
    a. Determine boundary "partial strategies" for this auction
    b. Generate "intermediate" strategies as follows:
      - combine the boundary strategies, or
      - modify them using empirical knowledge from the domain, or
      - Use equilibrium strategies (game-theoretic analysis for that auction type)

B. Use rigorous experimentation to explore the space of candidate strategies:
  1. Select one auction type and fix other partial strategies for all agents
  2. Experiment to find the best partial strategy for the specific auction as follows:
    a. Fix the agents using some intermediate candidate strategies
    b. Vary the number of agents using the boundary strategies
  3. Find the best candidate strategy and use it for all agents in the experiment
  4. Repeat step 1 for a different auction type, until the best strategies do not change

The first part of our methodology requires the decomposition of the general problem into several subproblems. The quantities placed in each bid are determined independently by maximizing the utility of the agent assuming that all the goods are bought (or sold) at some predicted prices and that every unit will be bought instantly. A dedicated module called "the planner" is doing this task. How to bid for these goods is determined by the "partial strategies"; one such strategy is used per auction type (single auction or group of similar auctions). For each such auction type, we first compute the boundary strategies that are possible. We then combine parts of the boundary strategies or modify some of their parts to form intermediate strategies that behave between the extreme bounds (e.g. if the one boundary strategy will place a bid at price p_{low} and the other at price p_{high} in a certain case, then the intermediate strategy should place its bid at price $p : p_{low} \leq p \leq p_{high}$). Another way to generate these intermediate strategies, which is of particular interest, is *to use the equilibrium strategy computed for the particular subproblem*. Having done this, the second part of our methodology advocates the use of experimentation in order to select the best combination of partial strategies, which is then implemented in the final design of the agent. For additional details, see [26].

Aplication: designing an agent for the TAC Classic game
We selected the Trading Agent Competition as an application domain for our methodology and designed agent WhiteBear. The equilibrium strategy derived from the game-theoretic analysis of multi-unit, multi-round auctions was computed and incorporated into our agent design. In figure 1(f), we graph this equilibrium bidding strategy for the first round; the priors $F_r()$ we used in theorem 5 to generate this strategy were obtained by sampling the valuations of the goods from the actual simulations. We then used rigorous experimentation to generate the final agent. (more details are given in [11])

The performance in the seeding, semi-final and final rounds of TAC show that our agent was the best performing agent in the history of the competition (in parenthesis we give the difference from the top competing agent), which validates our decision to use a principled methodology that included the use of game-theoretic analysis:
2002: Semi-final **1st** ($+0.55\%$), Final **1st** ($+1.85\%$).
2003: Seeding **1st** ($+1.63\%$), Semi-final **1st** ($+5.37\%$), Final **3rd** (-1.81%).
2004: Seeding **1st** ($+3.12\%$), Semi-final **1st** ($+6.57\%$), Final **1st** ($+7.10\%$).
2005: Seeding **1st** ($+2.63\%$), Semi-final **1st** ($+1.61\%$), Final **2nd** (-0.50%).

7 Conclusions

In this paper, we described our agenda for designing trading agent which would be able to participate in a large number of auctions, bidding for sets of goods. We identified the relevant features of real-world auctions which are important and we presented several theoretical results on analysis of auctions, which lead towards accomplishing this goal. We showed that including all the features in the analysis is necessary, because the behaviour when two features are present can be quite different to the cases when each one is examined independently. For a similar reason, we need to examine multi-unit auctions. In addition, we demonstrated that playing the equilibrium strategy outperforms other strategies even in settings where opponents deviate from the equilibrium.

We also presented our methodology that allows the design of trading agents for the general setting, which involves bidding in many different auctions while trying to acquire several goods with combinatorial valuations. The methodology allows the incorporation of theoretic results obtained for a single auction or group of similar auctions.

We showed the usefulness of our methodology by applying it to the design of our agent WhiteBear, which was the most successful agent in the history of TAC Classic.

We are currently continuing the theoretical work to analyze cases with bidders with multi-unit demand, asymmetric bidder cases and the presence of multiple auctions (either in sequence or in parallel). Furthermore, we continue our analysis and combine our results towards the final goal of an analytically computed strategy that would incorporate all the important features together.

Acknowledgements. This research was undertaken as part of the ALADDIN project and is jointly funded by a BAE Systems and EPSRC strategic partnership (EP/C548051/1).

References

1. Che, Y., Gale, I.: Standard auctions with financially constrained bidders. Review of Economic Studies 65(1), 1–21 (1998)
2. Che, Y.K., Gale, I.: Expected revenue of all-pay auctions and first-price sealed-bid auctions with budget constraints. Economics Letters 50(3), 373–379 (1996)
3. Krishna, V.: Auction theory. Academic Press, London (2002)
4. Maskin, E., Riley, J.: Optimal auctions with risk averse buyers. Econometrica 52(6), 1473–1518 (1984)
5. Myerson, R.B.: Optimal auction design. Maths of Operations Research 6, 58–73 (1981)
6. Riley, J.G., Samuelson, W.F.: Optimal auctions. Amer. Econ. Review 71(3), 381–392 (1981)
7. Larson, K., Sandholm, T.: Costly valuation computation in auctions. In: 8th Conference of Theoretical Aspects of Knowledge and Rationality (TARK VIII), pp. 169–182 (2001)
8. Thompson, D., Leyton-Brown, K.: Valuation uncertainty and imperfect introspection in second-price auctions. In: AAAI 2007, pp. 148–153 (2007)
9. Parkes, D.C.: Optimal auction design for agents with hard valuation problems. In: Proc. Agent Mediated Electronic Commerce (AMEC), pp. 206–219 (1999)
10. Brandt, F., Sandholm, T., Shoham, Y.: Spiteful bidding in sealed-bid auctions. In: IJCAI 2007, pp. 1207–1214 (2007)
11. Vetsikas, I.A., Jennings, N.R., Selman, B.: Generating Bayes-Nash equilibria to design autonomous trading agents. In: IJCAI 2007, pp. 1543–1550 (2007)
12. McAdams, D.: Isotone equilibria in games of incomplete information. Econometrica 71, 1191–1214 (2003)
13. McAdams, D.: Monotone equilibrium in multi-unit auctions. Review of Economic Studies 73(4), 1039–1056 (2006)
14. Gerding, E.H., Dash, R.K., Yuen, D.C.K., Jennings, N.R.: Bidding optimally in concurrent second-price auctions of perfectly substitutable goods. In: AAMAS 2007, pp. 267–274 (2007)
15. Maskin, E., Riley, J.: Asymmetric auctions. Review of Economic Studies 67, 413–438 (2000)
16. Lebrun, B.: First price auctions and the asymmetric n bidder case. International Economic Review 40(1), 125–142 (1999)
17. Anthony, P., Hall, W., Dang, V., Jennings, N.R.: Autonomous agents for participating in multiple on-line auctions. In: IJCAI workshop on e-Business & Intelligent Web, pp. 54–64 (2001)
18. Preist, C., Bartolini, C., Phillips, I.: Algorithm design for agents which participate in multiple simultaneous auctions. In: Agent Mediated e-Commerce (AMEC) III, pp. 139–154 (2001)
19. Vetsikas, I.A., Jennings, N.R.: Towards agents participating in realistic multi-unit sealed-bid auctions. In: AAMAS 2008, pp. 1621–1624 (2008)
20. Vetsikas, I.A., Jennings, N.R.: Bidding strategies for realistic multi-unit sealed-bid auctions. In: AAAI 2008, pp. 182–189 (2008)

21. Vetsikas, I.A., Jennings, N.R.: Bidding strategies for realistic multi-unit sealed-bid auctions. Journal of Autonomous Agents and Multi-Agent Systems, JAAMAS (2009) (in press) doi : 10.1007/s10458–009–9109–6
22. Vetsikas, I.A., Jennings, N.R.: Outperforming the competition in multi-unit sealed bid auctions. In: AAMAS 2007, pp. 702–709 (2007)
23. Vetsikas, I.A., Jennings, N.R.: Considering assymmetric opponents in multi-unit sealed-bid auctions. In: Proc. of the 1st Int. Workshop on Market Based Control, MBC (2008)
24. Vetsikas, I.A., Jennings, N.R.: Sequential auctions with partially substitutable goods. In: Proc. IJCAI Workshop on Trading Agent Design and Analysis (TADA), pp. 123–130 (2009)
25. Wellman, M.P., Wurman, P.R., O'Malley, K., Bangera, R., Lin, S., Reeves, D., Walsh, W.E.: Designing the market game for TAC. IEEE Internet Computing (April 2001)
26. Vetsikas, I.A., Selman, B.: A principled study of the design tradeoffs for autonomous trading agents. In: AAMAS 2003, pp. 473–480 (2003)