

Generating Bayes-Nash Equilibria to Design Autonomous Trading Agents

Ioannis A. Vetsikas and Nicholas R. Jennings

ECS School, University of Southampton,
Southampton SO17 1BJ, UK
{iv,nrj}@ecs.soton.ac.uk

Bart Selman

CS Dept., Cornell University,
Ithaca, NY 14853, USA
selman@cs.cornell.edu

Abstract

This paper presents a methodology for designing trading agents for complex games. We compute, for the first time, Bayes-Nash equilibria for first-price single-unit auctions and m^{th} -price multi-unit auctions, when the auction has a set of possible closing times, one of which is chosen randomly for the auction to end at. To evaluate this approach we used our analysis to generate strategies for the International Trading Agent Competition. One of these was evaluated as the best overall and was subsequently used very successfully by our agent WhiteBear in the 2005 competition.

1 Introduction

Auctions are becoming an increasingly popular method for transacting business, both between individuals over the Internet and between businesses and their suppliers. In this context, auction theory provides us with some simple equilibria, mostly for the case when a single item is being bought or sold. However, in practice, agents (or humans) are rarely interested in a single item. Rather, they often wish to bid in several auctions in parallel for multiple goods. Moreover, the fact that the value of each good to an agent often depends on other goods makes this problem particularly hard. Given this background, the International Trading Agent Competition (TAC) [Wellman *et al.*, 2001] defined a benchmark game that incorporates several elements found in real marketplaces, and different researchers adopted a variety of approaches to tackle the problem (see e.g. [Stone *et al.*, 2002] or [Greenwald and Boyan, 2001]). Now, most of these approaches operate based on some form of heuristics, because it is widely claimed that a principled theoretic approach is not practical. However, in [Vetsikas and Selman, 2003] we outline a general, principled methodology for systematically exploring the space of bidding strategies for complex games with multiple interdependent goods, which are traded in several different auctions, like TAC, where it is not possible to find an equilibrium solution. Specifically, to handle this complexity, we decompose the problem into sub-parts, which are then analyzed separately, before the various resulting strategies are recombined to generate the strategy the agent should use.

Against this background, in this paper we actually compute Bayes-Nash equilibria for some of these sub-problems and

then use them to generate new trading strategies. Even though our general methodology is motivated from the fact that complex games cannot be analyzed using the current game theoretic tools, by concentrating on one sub-problem, we can now analyze it. In particular, we examine auctions that have a set of possible closing times, one of which is chosen randomly for the auction to end at, and compute several novel equilibria. To evaluate our work we applied it to a sub-problem of TAC, the purchase of a hotel room of a particular type. In more detail, in *TAC Classic*, $m = 16$ rooms are available in each category (meaning hotel and day) and they are sold in separate, ascending, multi-unit, m^{th} -price auctions (one per category). These auctions close at randomly determined times and, more specifically, a random auction will close every minute throughout the game. The agents want to buy up to 8 rooms, but for ease of analysis we assume that each such agent is represented by 8 sub-agents, each of which is interested in a single room. In this way, we have $N = 8 \times 8 = 64$ agents each wishing to buy 1 unit (room) in the auction. Between closing times, the agents may place bids, but these are not opened until the next possible closing time; hence each round that takes place between consecutive closing times is effectively a sealed bid auction. The agents' bids are binding, as they cannot subsequently be retracted.¹

This paper makes a number of contributions to the state of the art. In section 2, we analyze the Bayes-Nash equilibria that exist in m^{th} -price auctions with multiple possible closing times and compute several equilibria that have not been seen in the literature before. In particular, the equilibrium that we look for in each case is the unique symmetric equilibrium. Now, some of the initial stages of this work appeared in [Vetsikas and Selman, 2006] where we presented equilibria for the 2-round case, without any implementation details. The current paper extends this very preliminary work by pre-

¹As well as being a benchmark for the international competition, there are also a number of real world auction settings that exhibit these characteristics. The Yahoo! auctions have a "random" closing time, since they are extended every time a new bid is placed. In addition governments sometimes conduct rounds of binding and non-binding bids when privatizing resources (normally the first rounds consist of non-binding bids, and the latter rounds of binding ones). As in TAC, after each round of binding bids, there is a chance that the government will decide, for political or economic reasons, to end the auction (well before the last possible closing time). Furthermore, in many countries, the spectrum auctions for third generation cellular phones involved multiple rounds of bidding.

senting equilibria (with proofs in most cases) for the general multi-round case, as well as for the special case of a first price auction. In section 3, we describe the implementation of this work to various cases, including the TAC game. Finally, in section 4, we generate two new TAC strategies inspired by this analysis, and we determine, via experiments, that one of them is superior. We then used this strategy in the actual competition (where the corresponding agent was highly successful). We conclude this paper, in section 5, by detailing interesting directions for future extensions of this work.

2 Computing the Equilibria

We assume that N risk-neutral agents wish to buy 1 unit of a certain good each. An independent seller sells m units of the desired good in an m^{th} price auction. The agents have valuations (utilities) u_i^r at round r which are i.i.d. drawn from distribution $F_r(u)$. Each agent knows accurately its own valuation u_i^r at the current round and all the distributions $F_r(u)$. 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 . We make two different assumptions about the information each agent i knows about its utility $u_i^{r'}$ ($r' \geq r$) at the start of the round r :

- $u_i^{r'}$ can be assumed to be relatively similar to the current utility u_i^r , and thus $u_i^{r'} \simeq u_i^r$; this is reasonable for TAC because usually there is a correlation between the valuation of the same room over the course of the game. This assumption is used in theorems 3 and 5.
- The more general case is when the agent only knows that $u_i^{r'}$ is drawn from some distribution $G_{r'}(u)$. This distribution could even be $G_{r'}(u) = F_{r'}(u)$, when no information whatsoever exists. We use this assumption in theorem 6.

In addition, *agents may not subtract bids*. In the TAC game, when the valuation of a room drops below the bid price Q_r , it is an optimal behavior for the agent to place a bid equal to Q_r as late as possible, since in this way it might be lucky enough to get rid of the room. If only $k < m$ agents have utilities $u_i^r \geq Q_r$, the rest of the rooms are sold to $(m - k)$ winners of the round $(r - 1)$, that no longer wish to buy; all their bids are equal to Q_r and the TAC rules state that the $(m - k)$ earliest placed such bids will get the rooms, therefore this means that a random selection of those winners will get the rooms, as the order of their bid placements is random.

It should be noted that all the equilibria when $p_1 \neq 1$ are the solutions of differential equations of the form described by theorem 1, which is taken from [Atkinson and Han, 2004].

Theorem 1 Let $f(x, z)$ and $\frac{\partial f(x, z)}{\partial z}$ be continuous functions of x and z at all points (x, z) in some neighborhood of the initial point (x_0, Y_0) . Then there is a unique function $Y(x)$ defined on some interval $[x_0 - \alpha, x_0 + \beta]$, satisfying:

$$Y'(x) = f(x, Y(x)), \quad \forall x : x_0 - \alpha \leq x \leq x_0 + \beta \quad \text{and} \\ Y(x_0) = Y_0$$

This theorem guarantees the existence and uniqueness of the equilibria we compute in the next sections.

2.1 Equilibria For A Single Unit Auction

In this section, $m = 1$ and the single unit is sold to the agent which submitted the highest bid at a price equal to its bid. In theorem 3, we assume that in the second round the utilities are drawn from the same distribution as in the first round, and that $u_i^2 \simeq u_i$. Each agent i submits a bid v_i in the first round. We compute a Bayes-Nash equilibrium $g(u)$ that maps utilities u_i to bids v_i . In the case of $p_1 = 1$ (only one round) and $Q_1 = 0$, we know from standard auction theory that each risk-neutral agent i with valuation u_i should bid:

$$g(u_i) = u_i - \frac{1}{(F_1(u_i))^{N-1}} \cdot \int_0^{u_i} (F_1(\omega))^{N-1} \cdot d\omega$$

Theorem 2 If the starting price is $Q_1 \geq 0$ and the bidding lasts for exactly 1 round ($p_1 = 1$) the equilibrium strategy is

$$g(u_i) = u_i - \frac{\int_{Q_1}^{u_i} (F_1(\omega))^{N-1} \cdot d\omega}{(F_1(u_i))^{N-1}} \quad (1)$$

Proof. Since $Q_1 > 0$, (i) some agents might have stopped participating in the auction, since the current price Q_1 exceeds their private valuation u_i , and (ii) the probability distribution of the valuations $F_1(u)$ has changed, since now we know that the valuation of agents that still participate is $u_i \geq Q_1$. The new probability distribution is:

$$F_{Q_1}(u) = \frac{F_1(u) - F_1(Q_1)}{1 - F_1(Q_1)}, \quad \text{if } u \geq Q_1 \text{ \& } F_{Q_1}(u) = 0, \text{ o.w.} \quad (2)$$

We also know the probability π_k of the event that exactly k agents participate in the auction at price Q_1 ; it is the probability that exactly $(k - 1)$ of the other agents' valuations² u_i are $u_i \geq Q_1$, which is (see e.g. [Rice, 1995]):

$$\pi_k = C(N - 1, k - 1) \cdot (F_1(Q_1))^{N-k} \cdot (1 - F_1(Q_1))^{k-1} \quad (3)$$

Let us assume that it is a Bayes-Nash equilibrium for each agent i to bid $v_i = g(u_i)$. If k agents participate then the utility of agent i is 0, if it does not have the highest bid, and $(u_i - v_i)$, if it does. The expected utility of agent i when it places a bid of v_i (in the case of k agents) is:

$$EU_i(v_i | \#agents = k) = (u_i - v_i) \cdot \text{Prob}[\bigwedge_{j \neq i} v_i \geq v_j] = (u_i - v_i) \cdot \prod_{j \neq i} \text{Prob}[v_i \geq v_j].$$

Since $v_j = g(u_j)$, it is $\text{Prob}[v_i \geq v_j] = F_Q(g^{-1}(v_i))$. Let $\Phi(x) = (F_1(x))^{N-1}$. The expected utility regardless of the number of agents participating is:

$$EU_i(v_i) = \sum_{k=1}^N \pi_k \cdot EU_i(v_i | \#agents = k) \Rightarrow \\ EU_i(v_i) = (u_i - v_i) \cdot \Phi(g^{-1}(v_i)) \quad (4)$$

The bid v_i that maximizes $EU_i(v_i)$ can be found by setting

$$\frac{dEU_i(v_i)}{dv_i} = 0 \Leftrightarrow (u_i - v_i) \cdot \frac{\Phi'(g^{-1}(v_i))}{g'(g^{-1}(v_i))} = \Phi(g^{-1}(v_i)).$$

Since we assumed that the optimal solution is $v_i = g(u_i)$, the previous equation becomes:

$$(u_i - g(u_i)) \cdot \Phi'(u_i) \cdot \frac{1}{g'(u_i)} = \Phi(u_i)$$

The solution of this equation, for boundary condition $g(Q_1) = Q_1$, is equation 1; this can be easily verified. ■

²Because from the point of view of a participating agent it does not know whether the other $(N - 1)$ agents participate.

It should be noted that this theorem is reminiscent of the work in [McAfee and McMillan, 1987] concerning the equilibrium solution in the case that the distribution $F_1(u)$ has a minimum value v_l , presented by equation 5 in that paper. In fact the two equilibria are given by similar equations, where the minimum valuation v_l is substituted by the minimum allowable bid Q_1 . However there are also some important differences between theorem 2 and that work. In particular, in the case examined by theorem 2, the number of agents participating is not known a priori, because some of them may have a private valuation u_i which is smaller than the minimum allowable bid Q_1 ; this creates a fundamental difference compared to the case examined by McAfee and McMillan where there is a minimum valuation v_l . Indeed, dealing with this uncertainty concerning the number of bidders, which participate in the auction, constitutes the core of the proof of theorem 2.

Theorem 3 *If the starting price is $Q_1 = 0$, a second round of bidding exists with probability $(1 - p_1)$ ($p_1 \neq 0, 1$) and the utility of the agents in the second round is drawn from the same distribution $F_2(u) = F_1(u)$ as in the first round (and each agent i in fact has utility of a similar value to the utility u_i of the first round) then the equilibrium strategy is the solution of the differential equation:*

$$(u_i - g(u_i)) \cdot \frac{\Phi'(u_i)}{g'(u_i)} = \Phi(u_i) \cdot \Psi(g(u_i)) \quad (5)$$

where the boundary condition is $g(0) = 0$,

$$\Phi(x) = (F_1(x))^{N-1} \text{ and } \Psi(x) = 1 + \frac{1-p_1}{p_1} \cdot (F_1(x))^{N-1}.$$

Proof. Due to space considerations we omit this proof. However it has similarities with the proof of theorem 5. ■

As a special case, we can examine this equation when only $N = 2$ agents participate and their valuations $u_i \sim U[0, 1]$, which means that $F_1(u) = u, \forall u \in [0, 1]$. Equation 5 becomes:

$$g_p'(u_i) = \frac{u_i - g_p(u_i)}{u_i \cdot (1 + \frac{1-p}{p} \cdot g_p(u_i))} \quad (6)$$

In figure 1 (left), we graph the solution for various values of $\frac{1-p}{p}$. We provide this figure in order to contrast with the Bayes-Nash equilibrium for the case that $p = 1$, which is $g_1(u_i) = \frac{u_i}{2}$ (see e.g. [Krishna, 2002]). One may notice that as the probability $(1 - p)$ of a second round increases, the equilibrium strategy suggests that the agent should bid less. In section 3, we compare these results with those produced by our algorithm for solving the R -round case.

2.2 Equilibria For Multiple Round Auctions

In this section we present the system of equations for computing the equilibria for the R -round auction ($R \geq 2$).

Let us define $\Phi_r(x)$ and $Y_r(x)$ as follows:

$$\Phi_r(x) = \sum_{i=0}^{m-1} C(N-1, i) \cdot (F_r(x))^{N-1-i} \cdot (1 - F_r(x))^i$$

$$Y_r(x) = \sum_{i=0}^{m-2} C(N-1, i) \cdot (F_r(x))^{N-1-i} \cdot (1 - F_r(x))^i$$

In the short paper [Vetsikas and Selman, 2006],³ we give the following theorem for the 1-round case:

Theorem 4 *If the starting price is $Q_1 \geq 0$ and the bidding*

³This proof and other details that have been omitted from this paper may also be found at Dr. Vetsikas' thesis.[Vetsikas, 2005]

lasts for exactly 1 round ($p_1 = 1$) the equilibrium strategy is

$$g(u) = u - \frac{e^{\int_{Q_1}^u \frac{-Y_1'(\omega)}{\Phi_1(\omega) - Y_1(\omega)} \cdot d\omega}}{\Phi_1(u) - Y_1(u)} \cdot \int_{Q_1}^u \frac{\Phi_1(z) - Y_1(z)}{e^{\int_{Q_1}^z \frac{-Y_1'(\omega)}{\Phi_1(\omega) - Y_1(\omega)} \cdot d\omega}} \cdot dz \quad (7)$$

In the next proofs, we will use the following derivation from the proof of theorem 4, without proving it here; that the expected utility, if the auction closes at this round r , is:

$$U_i^{(r)}(v_i^r) = (u_i^r - v_i^r) \cdot \Phi_r(g_r^{-1}(v_i^r)) + \int_{Q_r}^{v_i^r} Y_r(g_r^{-1}(\omega)) \cdot d\omega \quad (8)$$

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 $F_r(u)$ (and each agent i in fact has utility u_i^r of a similar value to the utility u_i of the first round) then the equilibrium strategy is the solution of the differential equation:*

$$(u_i - g_r(u_i)) \cdot \frac{\Phi_r'(u_i)}{g_r'(u_i)} = (\Phi_r(u_i) - Y_r(u_i)) \cdot \Psi_r(u_i, g_r(u_i)) \quad (9)$$

where $\Psi_r(u, x) = 1 - \frac{1-p_r}{p_r} \cdot \frac{\partial \mathcal{U}_{r+1}(u, x)}{\partial x}$,

and $\mathcal{U}_{r+1}(u_i, Q_{r+1})$ is the expected utility at round $(r + 1)$, when the agent's utility is u_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 $g_r(u_i)$ is then:

$$\mathcal{U}_r(u_i, Q_r) = \quad (10)$$

$$p_r \cdot \left\{ (u_i - g_r(u_i)) \cdot \Phi_r(u_i) + \int_{Q_r}^{u_i} Y_r(\omega) \cdot g_r'(\omega) \cdot d\omega \right\}$$

$$+ (1 - p_r) \cdot \left\{ \int_{Q_r}^{u_i} \mathcal{U}_{r+1}(u_i, g_r(\omega)) \cdot Y_r'(\omega) \cdot d\omega \right.$$

$$+ \mathcal{U}_{r+1}(u_i, g_r(u_i)) \cdot \left\{ \Phi_r(u_i) - Y_r(u_i) \right\}$$

$$\left. + \int_{u_i}^{g_r^{-1}(u_i)} \mathcal{U}_{r+1}(u_i, g_r(\omega)) \cdot \Phi_r'(\omega) \cdot d\omega \right\}$$

Proof. If the auction closes at the first round, then the expected utility is given by equation 8. To compute the expected utility from the other rounds, we need to examine the value of Q_{r+1} , which depends on the bids placed in the current round. The probability distribution of the $(m - 1)^{th}$ and m^{th} highest bids among all other agents at this round, named respectively $B^{(m-1)}$ and $B^{(m)}$, are $Prob[B^{(m-1)} \leq v] = Y_r(g^{-1}(v))$ and $Prob[B^{(m)} \leq v] = \Phi_r(g^{-1}(v))$.

If $B^{(m)} > v_i^r$, then $Q_{r+1} = B^{(m)}$. If $B^{(m-1)} > v_i^r \geq B^{(m)}$, then the agent submitted the m^{th} price, so $Q_{r+1} = v_i^r$. If $v_i^r \geq B^{(m-1)}$, then $Q_{r+1} = B^{(m-1)}$.

Note that $Prob[Q_{r+1} = v_i^r] =$

$$Prob[B^{(m-1)} > v_i^r \geq B^{(m)}] = \Phi_r(g_r^{-1}(v_i^r)) - Y_r(g_r^{-1}(v_i^r)).$$

As a result, we can now compute the expected utility from the rest of the rounds U_i^* as follows:

$$U_i^* = \int_{Q_r}^{u_i} \mathcal{U}_{r+1}(u_i, \omega) \cdot Prob[Q_{r+1} = \omega] \cdot d\omega \Rightarrow$$

$$U_i^* = \int_{Q_r}^{v_i^r} \mathcal{U}_{r+1}(u_i, \omega) \cdot \frac{d}{d\omega} Y_r(g_r^{-1}(\omega)) \cdot d\omega \quad (11)$$

$$+ \mathcal{U}_{r+1}(u_i, v_i^r) \cdot \left\{ \Phi_r(g_r^{-1}(v_i^r)) - Y_r(g_r^{-1}(v_i^r)) \right\}$$

$$+ \int_{v_i^r}^{u_i} \mathcal{U}_{r+1}(u_i, \omega) \cdot \frac{d}{d\omega} \Phi_r(g_r^{-1}(\omega)) \cdot d\omega$$

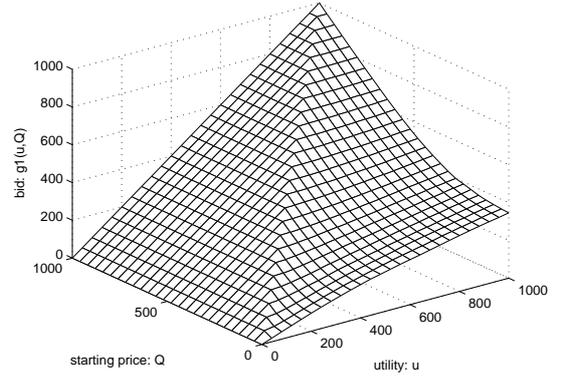
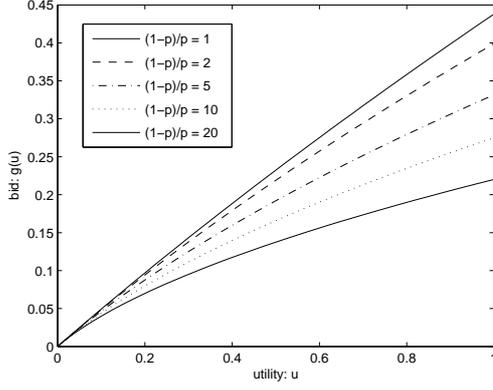


Figure 1: Equilibrium strategies for case $N = 2$, $m = 1$ and utility drawn from uniform distribution $U[0, 1]$, (left) using equation 6 for the case when $R = 2$ rounds exist, and $Q_1 = 0$, and (right) using the algorithm of figure 2 for the 1st round of 8 and probabilities p_r equal to those of TAC.

The expected utility for both rounds is:

$$EU_i^r(v_i^r) = p_r \cdot U_i^{(r)} + (1 - p_r) \cdot U_i^* \quad (12)$$

The bid v_i that maximizes $EU_i^r(v_i^r)$ can be found by setting $\frac{dEU_i^r(v_i^r)}{dv_i^r} = 0$. Then we set $v_i^r = g_r(u_i)$ and eventually we derive the differential equation 9.

If $u_i = Q_r$ then the agent must bid $v_i^r = Q_r$, hence the boundary condition.

From equations 11 and 12, by substituting $v_i^r = g_r(u_i)$, we derive equation 10. ■

In the following theorem the utility in the later rounds is assumed not to be precisely known and to be drawn from distribution $G_r(u)$.

Theorem 6 *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 $F_r(u)$ (and each agent i knows more accurately that, in fact, its utility u_i^r is drawn from distribution $G_r(u)$), then the equilibrium strategy $g_r(u_i^r, Q_r)$ is the solution of the differential equation:*

$$(u_i^r - g_r(u_i^r, Q_r) + \frac{1-p_r}{p_r} \cdot U_L^r(g_r(u_i^r, Q_r))) \cdot \frac{\Phi_r'(u_i)}{\partial u_i^r} = (\Phi_r(u_i^r) - Y_r(u_i^r)) \cdot \Psi_r(g_r(u_i^r, Q_r)) \quad (13)$$

where $\Psi_r(x) = 1 - \frac{1-p_r}{p_r} \cdot \frac{d}{dx}(\tilde{U}^{r+1}(x) + U_L^{r+1}(x))$ and

$$U_L^{r+1}(x) = - \sum_{k=0}^{m-1} \left\{ \frac{N \cdot (m-k)}{m \cdot (N-k)} \cdot C(N-1, k) \cdot (F_{r+1}(x))^{N-1-k} \cdot (1 - F_{r+1}(x))^k \cdot \int_0^x G_{r+1}(\omega) \cdot d\omega \right\}$$

and the expected utility at round r when u_i^r is drawn from $G_r(u)$ is:

$$\tilde{U}^r(Q_r) = \int_{Q_r}^{+\infty} U_r(\omega, Q_r) \cdot \frac{d}{d\omega} G_r(\omega) \cdot d\omega \quad (14)$$

The boundary condition is $g(u_i^r, Q_r) = Q_r$, when $u_i^r = Q_r$. $U_r(u_i^r, Q_r)$ is the expected utility at round r , when the agent's

utility is u_i^r and the starting price is Q_r and given that $g_r(u_i^r, Q_r)$ has been computed it is:

$$U_r(u_i^r, Q_r) = (1-p_r) \cdot \left\{ \int_{Q_r}^{u_i^r} \tilde{U}^{r+1}(g_r(\omega, Q_r)) \cdot Y_r'(\omega) \cdot d\omega + \tilde{U}^{r+1}(g_r(u_i^r, Q_r)) \cdot \{\Phi_r(u_i^r) - Y_r(u_i^r)\} \right. \\ \left. + \int_{u_i^r}^{g_r^{-1}(u_i^r, Q_r)} \tilde{U}^{r+1}(g_r(\omega, Q_r)) \cdot \Phi_r'(\omega) \cdot d\omega \right\} \\ + p_r \cdot \left\{ (u_i^r - g_r(u_i^r, Q_r)) \cdot \Phi_r(u_i^r) + \int_{Q_r}^{u_i^r} Y_r(\omega) \cdot g_r'(\omega, Q_r) \cdot d\omega \right\} \quad (15)$$

Proof. Due to space limitations and because the proof of this theorem shares a number of common steps with the proof of theorem 5, we will give a sketch of the proof. One should notice that equation 15 is derived in almost the same way as equation 10. However since the utility u_i^r in round r is not precisely known, and is drawn from distribution $G_r(u)$, we must use equation 14 in order to compute the expected utility at a round r , based on all possible utility values. The other important difference between the two proofs, is that we must also compute the expected gain of utility (actually it's negative, so it's a loss) $U_L^{r+1}(Q_{r+1})$ if the agent is a winner in the round r and in the next round ($r+1$) its utility is $u_i^{r+1} < Q_{r+1}$. Initially we compute the utility difference $U_L^{r+1}(u_i^{r+1}, Q_{r+1})$, and then it is

$$U_L^{r+1}(Q_{r+1}) = \int_0^{Q_{r+1}} U_L^{r+1}(u_i^{r+1}, Q_{r+1}) \cdot \text{Prob}[u_i^{r+1} = \omega] \cdot d\omega$$

The final equation for $U_L^{r+1}(Q_{r+1})$ is:

$$U_L^{r+1}(Q_{r+1}) = - \sum_{k=0}^{m-1} \left\{ \frac{N \cdot (m-k)}{m \cdot (N-k)} \cdot C(N-1, k) \cdot (F_{r+1}(Q_{r+1}))^{N-1-k} \cdot (1 - F_{r+1}(Q_{r+1}))^k \cdot \int_0^{Q_{r+1}} G_{r+1}(\omega) \cdot d\omega \right\}$$

The rest of the proof proceeds along the same lines as the proof of theorem 5. ■

```

Set  $\mathcal{U}_{R+1}(u, Q) = 0, \forall u, Q$ 
for  $r = R$  to 1 do
   $\forall u, Q$  set  $\Psi_r(u, Q) = 1 - \frac{1-p_r}{p_r} \cdot \frac{\partial}{\partial Q} \mathcal{U}_{r+1}(u, Q)$ 
   $\forall Q$  compute  $g_r(u, Q)$  by solving Differential Eq. 9
   $\forall u, Q$  compute  $\mathcal{U}_r(u, Q_r)$  using equation 10
end for

```

Figure 2: Algorithm for solving the system of Theorem 5.

3 Implementation

Figure 2 provides the algorithm that we used in order to compute the equilibrium described by theorem 5 using Matlab. At each round r , starting from the last round (R), the equilibrium strategy and expected utility from round $(r + 1)$ are used to compute the equilibrium strategy for round r , using the equations of theorem 5. To solve equation 9, we used the differential equation solver *ode45*. R is set to $R = 8$ rounds in all cases, since TAC hotel auctions can have up to 8 rounds.

Initially we computed the equilibria for a *uniform distribution*. The simplest case is when $N = 2$ agents are buying $m = 1$ items; this case was used in order to compare the results of the algorithm with the known solution as shown in figure 1. We only present the equilibrium strategy for the first round of 8 rounds here in figure 1 (right); the solutions of all the other rounds have similar graphs with different curvatures. One notices that in earlier rounds the agent bids less than in the later rounds as was expected. By comparing figure 1 (left) to the solution for the 7th round, we verify that they are identical for the case that the current price $Q = 0$. However we observed that the solution at each round is very similar to the solution that we would get if, instead of several later rounds (with probability $(1 - p)$ that the auction will close later), we only have one possible later round (again with probability $(1 - p)$ that the auction does not close at this round, but at the next one). This indicates that, *from the point-of view of the agent, the number of possible rounds does not matter, as long as the probability of having more rounds is the same*. For example, the equilibrium strategies for all the rounds are almost identical to the one presented in figure 1 (left) when $Q = 0$. This would mean that we only need to solve one differential equation instead of a system when the utility distribution is the same in every round.

Still using a *uniform distribution*, we computed the equilibria when $N = 64$ agents are buying $m = 16$ items; this uses the correct number of agents and items for a real TAC scenario. In every round the equilibrium strategy is to bid closer to each agent’s utility; that is 94% or more in round 1 and almost 100% in the last round. This happens because of the fact that many more agents participate, compared to the number of items offered. In this case, the agent makes a profit by the fact that its bid is likely not to be the m^{th} price, and thus, usually pays less than its actual bid, if it wins.

Finally, we computed the equilibria using an *experiment derived distribution*. Since our stated goal was to apply these equilibria to a TAC agent in order to generate more candidate strategies, we collected the utilities of the hotel rooms from a large number of actual games and used these to create the distributions $F_r(u)$ and $G_r(u)$. The cdf of the distribution $F_r(u)$ that we used is presented in figure 3. We make sure that

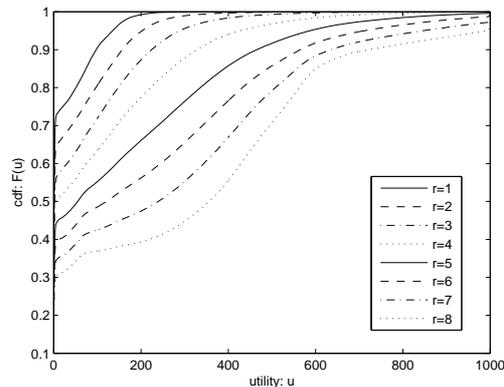


Figure 3: Experimental distribution $F_r(u)$, which is used to generate the strategies for our TAC agent.

these functions and their derivatives are continuous, so that we can apply *ode45* and guarantee uniqueness. To do this, we group all the samples in buckets of size 2, and then to make it continuous, we expand each bucket to a normal distribution with $\sigma = 2$ and center μ equal to the center of the bucket; this is the most common way, which is used in statistics, for turning sampled discrete distributions into continuous ones. In figure 4, we present the equilibrium strategy for rounds 1, 5 and 8. The fact that the utility distribution $F_r(u)$ is different in every round causes the bid function to vary significantly between different rounds. In the later rounds, it is advisable to bid close to the marginal utility u_i that the agent gets if it buys a unit of the commodity sold in the auction, whereas in the early rounds, it is preferable to bid closer to the current price Q of the auction.

We also computed the equilibrium strategies described by theorem 6. We set the distribution of the later rounds $G_r(u)$ to be equal to $G_r(u) = F_r(u|u \geq u_{r-1})$, which means that we disregard any value of u that is smaller than the utility u_{r-1} in the previous round; this is reasonable for TAC, since utilities rarely decrease between rounds. We only need to add one step to the algorithm of figure 2 in order to implement this: before finishing each iteration, we add a step (see equation 14) in order to compute the utility of this round when the utility u_i^r is only known to be drawn from distribution $G_r(u)$. Even though, due to space limitations, we can’t show any figures of the equilibrium strategies, they all have the same general form as the strategies generated by theorem 5. However, for any values of the valuation u and starting price Q and at every round r these strategies bid less than the corresponding strategies given by theorem 5; this is something that we expected, because of the form of distribution $G_r(u)$, which only allows for higher than the current utility u_i^r .

4 Strategy Generation and Experiments

Each candidate strategy for the hotel auctions determines the prices offered in each round; they are differentiated by how aggressively the agent bids. In [Vetsikas and Selman, 2003], we present three strategies. Two of them are the boundary strategies, that is the lowest and highest possible prices respectively, that a rational agent would ever consider bidding at:

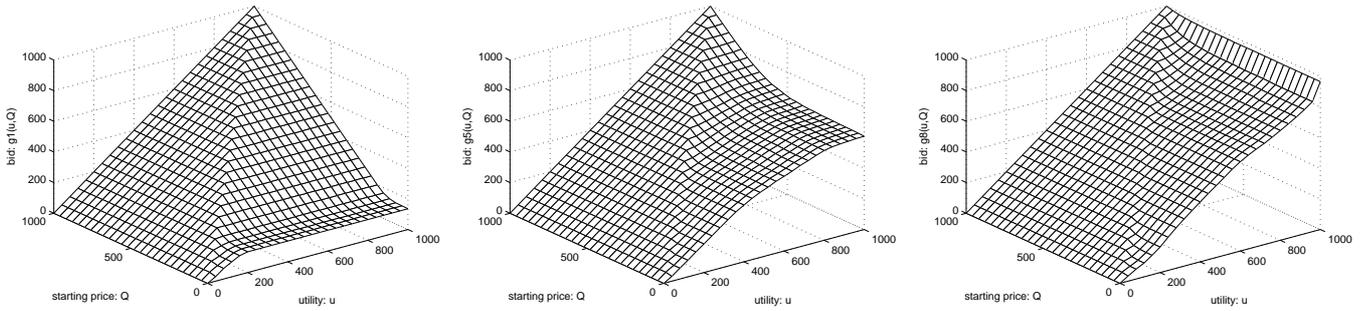


Figure 4: The equilibrium strategies $g_1(u, Q)$, $g_5(u, Q)$ and $g_8(u, Q)$, are presented (from left to right), for rounds 1, 5 and 8 respectively. These strategies are generated for the experimental distribution $F_r(u)$ when $N = 64$ and $m = 16$. It should be noted that when $u_r \leq Q_r$ (left half of each diagram), the agent should bid u_r , but this bid will be rejected by the auction.

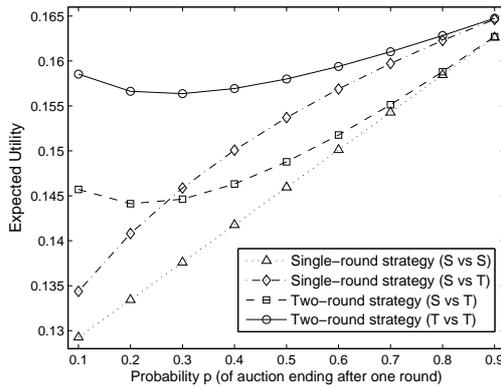


Figure 5: Expected utilities for a game, where $N = 2$ agents participate and $m = 1$ item is sold, for all possible combinations of agents. The agents use either (i) the single-round (S) equilibrium strategy $g(u_i) = \frac{u_i}{2}$, or (ii) the two-round (T) equilibrium strategy described by equation 6.

- (i) the *L*(owest aggressiveness) strategy is to bid reasonably close to the current bid price Q of the auction, which is the lowest price that one can offer for its bid to be valid, and
- (ii) the *H*(ighest aggressiveness) strategy is to bid progressively closer to the actual (marginal) utility it has for the item, which is the highest price that a rational agent should bid at if it doesn't wish to lose money.

Then we must generate a number of intermediate strategies, which would bid between the lowest and highest reasonable prices (these are the bids of the boundary strategies) at each instance. In [Vetsikas and Selman, 2003], we proposed one such strategy which worked very well in practice (iii) the *M*(oderate aggressiveness) strategy, which combines the behaviour of the boundary strategies (*L* and *H*) depending on the marginal utility of the desired good: if the marginal utility is low, then it bids similarly to the *L* strategy, whereas if it is high, it bids similarly to the *H* strategy.

In this paper we expand this methodology to use strategies based on the equilibria presented in section 2.2:

- (iv) the actual *E*(quilibrium) strategy $g_r(u, Q)$, which was computed in section 3 by applying the equations of theorem 5 for the TAC (experiment derived) distribution, is used.

However, the basic premise for generating this equilibrium was that each agent is interested in maximizing its own profit. *In contrast, in a competition, like TAC, the actual goal of each agent is to beat the competition; that is, maximize the difference of its profit to that of the other agents, rather than to just maximize its profit.* This change can have a significant impact on the equilibrium strategy. To explore this issue we looked at the single item ($m = 1$) first price auction with $N = 2$ agents and $R = 2$ rounds when the valuations are drawn from uniform distribution $U[0, 1]$.

The equilibrium strategy for a (S)ingle round is $g(u_i) = \frac{u_i}{2}$, and for (T)wo rounds is given by equation 6. Figure 5 presents the expected utilities in all possible combinations of the two agents for varying values of the probability p of the auction closing after the first round. We observe that, when only the S strategy is used, the scores are the lowest, whereas when only the T strategy is used, the scores are highest. However, when the two agents use different strategies, one uses S and the other T, the S strategy usually performs better than the T strategy. This verifies that in a competition the more aggressive S strategy, which bids higher than T, would win, whereas if we care only about maximizing profit, the S strategy is dominated by the T strategy; we expected this since T is an *equilibrium strategy* for the two-round auction.

We also computed the equilibrium strategy for the “competition objective function” (that is to maximize the difference of profit rather than the actual profit). This strategy would beat the $g(u_i) = \frac{u_i}{2}$ strategy in direct competition for a single-round auction: its expected utility is 0.136 vs 0.125.⁴ **Theorem 7** *If we care about maximizing the difference in profit from the opponent, then the Bayes-Nash equilibrium strategy for the first-price, single-unit, single-round auction when two agents participate with i.i.d. utilities drawn from the uniform distribution $U[0, 1]$ is $g^*(u_i) = \frac{2}{3} \cdot u_i$*

⁴In [Brandt *et al.*, 2007] the authors present an analysis of the equilibria that exist in 1st and 2nd price sealed bid auctions, for a variety of objective functions (i.e. for varying significance of the effect that the opponents' gain has on the agent's utility). In [Vetsikas and Jennings, 2007], we extend these results and present the equilibria for multi-unit auctions, i.e. the m^{th} and $(m + 1)^{\text{th}}$ price sealed bid auctions. We prove that, in some cases, the agent should increase significantly its bid if it wishes to outperform its competition; this can lead to bidding more than its true valuation even in an m^{th} price auction.

This demonstrates that the E strategy's bids are lower than they should be in order to beat the opponents. By looking at the form of $g_r(u, Q)$ in all rounds (see figure 4), we observe that, at round 1 the bids are a bit higher than the current price Q , at round 8 closer to the marginal utility, and in the intermediate rounds, the bids are approximately equal to a weighted sum of the values provided by these extreme strategies with the weight equal to the probability p_r of the current round r being the one when the auction closes. Since the problem of the E strategy are these bid levels, we modify it:

(v) the P strategy, which uses the boundary strategies L and H (rather than E) as the strategies that give the two extreme values (low and high), and mixes these values in the same way as E , namely by placing bids at a price equal to the weighted sum of those bids with weights p_r and $(1 - p_r)$ respectively.

We would like to point out that bidding more aggressively, i.e. higher, than suggested by the strategies we present in this paper, is actually going to produce inferior results, even in the case that the agent is interested in outperforming the competition. The reason for this is that, if the agent bids higher than that point, it will end up paying too much for whatever it buys, and this amount will be higher than the benefit from the opportunities, which are stolen from its competitors.

We now have 3 intermediate strategies to test, in addition to the two boundary ones. According to our methodology, we must organize a tournament among the intermediate strategies in order to determine the best strategy, as we should not use more than two of the agents employing intermediate strategies. The rest of the agents use the boundary strategies and the number is varied, so that the whole spectrum of possible strategic combinations might be explored. To determine which is the best strategy we use paired t tests⁵ on the scores of the agents; and, generally, we consider scores to be statistically significant if the test is below 5% (and definitely less than 10% in a small number of cases only).⁶ In the first phase, we use intermediate strategies M and E . The results are presented in table 1. The M agents outperform the E agents by a small margin (about 2.5%). This is mainly due to the fact that in competition there is a different objective than simply maximizing one's profit, as we have demonstrated earlier in this section. Note that the difference between the M and H agents is not quite statistically significant, but, with more experiments, it probably would be and thus is denoted by $\checkmark?$.⁷ This is not important, however, as we found out what we wanted from this phase, namely that we should promote the M agent

⁵This test was used in similar work, e.g. [Stone *et al.*, 2002], to evaluate performance differences between agents.

⁶To make this more precise, let us consider the first experiment in table 1 where the average score difference between the L and M agents is 1%. There are two M and four L agents participating in that experiment, with average scores 3359 and 3364 for each of the individual M agents and 3326, 3327, 3322 and 3335 for each of the L agents. We calculate p-values for all eight possible pairings of M with L agents. Six of these p-values are between 2% and 5% (all of them except for the values of the L agent with average score 3335 to both the M agents). These last two p-values are 6% and 9%. Based on our criteria, we consider this to be an indication of statistically different scores.

⁷The p-values, in this case, for all eight agent combinations that we examine, range from 3% to 16%, so it does not meet our stated criteria for statistical significance.

#H agents	Average Scores				Stat. Significant Diff.?		
	E	M	L	H	L/M	M/H	L/H
0 (600)	3292	3361	3327	N/A	✓		
2 (599)	3240	3323	3264	3339	✓	×	✓
4 (626)	2954	3046	N/A	3011		✓?	

Table 1: Scores for agents of low (L), moderate (M) and high (H) bidding aggressiveness, as well as the equilibrium (E) agent, as the number of aggressive agents participating increases. In each experiment, we use two E and two M agents. \checkmark indicates statistically significant difference in the scores of the selected agents, and \times statistically similar scores. (In parenthesis we give the number of games ran for each experiment)

#H agents	Average Scores				Stat. Significant Diff.?		
	P	M	L	H	P/M	M/H	P/H
0 (545)	3403	3417	3373	N/A	×		
2 (573)	3275	3242	3235	3241	×?	×	×?
4 (603)	2999	2981	N/A	3006	×	×	×

Table 2: Scores for agents of low (L), moderate (M) and high (H) bidding aggressiveness as well as the P agent, as the number of aggressive agents participating increases. In each experiment, two P and two M agents participate.

to the next phase of the tournament. Thus, in the last phase, we use intermediate strategies M and P . The results are presented in table 2. Here the M and the P agents outperform the L agents. The other agents perform reasonably similarly between themselves, and, in fact, the performances were statistically similar for most cases (as denoted in the table). The only exceptions to this rule were the performance of the P agents (about +1%) as opposed to that of the M and H agents in the case that two agents of each type participated. Based on this, and empirical data from the early rounds of TAC 2005, we decided that the P agent is performing best, which is the reason it was chosen as the strategy we used in most of the seeding round and the finals of TAC 2005.

Strategy M was used as our TAC entry since the semi-finals of TAC 2002 thru the finals of TAC 2004. In 2005 we used the P strategy mainly. The performance in the seeding, semi-final and final rounds of TAC show this (in parenthesis we give the difference from the top competing agent):

- 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%).

Throughout this time, our agent has consistently achieved the highest scores in the competition. In 2004 our agent was first in all rounds with statistically significant differences in the scores, which has never happened before or since in TAC Classic. In 2005 the top three finalists *Mertacor*, *WhiteBear05*, and *Walverine* were very close to each other. The reason for this extremely competent performance by all agents seems to have been (at least partly) the fact that the other two top scoring agents adopted some of our techniques; in particular, the founding principle of our methodology that trading commodities should be based entirely on maximiz-

ing the agent's utility. Walverine used this especially in re-designing their procurement of entertainment tickets. Moreover, the Walverine team also used an experimental methodology [Reeves *et al.*, 2005] which is probably the closest approach to the one we outlined in [Vetsikas and Selman, 2003]. However, no one else thought that it was feasible to look for an equilibrium strategy, even for a sub-problem of a large problem like TAC, and thus our approach as outlined in this paper is truly unique.

5 Conclusions and Future Work

This paper analyzed auctions that have a set of possible closing times, one of which is chosen randomly for the auction to end at. These auctions can be decomposed into one or more rounds, each of which is defined by the interval between possible closing times, and during which the auction is treated as sealed bid. Specifically, we analyzed the Bayes-Nash equilibria that exist in such cases and computed several novel solutions for these auctions. In so doing, we deal with the tradeoff that exists in this case, namely the price offered in the agent's bids. We also explained how to apply these results in the design of a TAC agent; we generated two new trading strategies (*E* and *P*) based on the form of the equilibria we computed, one of which (*P*) was successfully used in our top-scoring agent during the 2005 competition.

We are currently working towards generating equilibria for multi-demand auctions; this would allow us to remove the current restriction that each agent must bid for a single item. While this improvement is not expected to change the results of our analysis for the TAC game significantly,⁸ success in this endeavor will lead to a more generally applicable result, which is important for a wide range of game theoretic situations.

We are also working on recomputing the equilibria for the "competition objective function", namely when the agent wishes both to maximize its own profit and minimize that of its opponents, which is more realistic for a competition setting. We have been able to generate equilibria for the m^{th} and $(m+1)^{th}$ price auctions with one round of bidding for a range of objective functions. We plan to incorporate this work in the analysis presented in this paper in order to generate strategies that are more appropriate to a competition setting like TAC.

Acknowledgements

We would like to thank prof. Robin Mason at the University of Southampton, and prof. Steven H. Strogatz at Cornell University, as well as the reviewers of this paper, for comments on this work.

Most of this research was undertaken as part of the first author's thesis. The latter parts, as well as ongoing related work,

⁸The reason why this happens is as follows. Our analysis shows that the top bid of each agent is the same under both approaches and the other bids, from highest to lowest, are affected progressively more. This means that the bids for the top-valued three or four rooms are not affected (the top one) or are affected relatively little. As these are the highest bids placed by the agents, they are the ones likely to affect the closing price of the auction; the rest are for rooms with minimal marginal utilities, as observed in the simulations, and thus cannot in practice affect the final result.

is undertaken as part of the *ALADDIN* (Autonomous Learning Agents for Decentralised Data and Information Systems) project, which is jointly funded by a BAE Systems and EP-SRC (Engineering and Physical Research Council) strategic partnership (EP/C548051/1).

References

- [Atkinson and Han, 2004] Kendall Atkinson and Weimin Han. *Elementary Numerical Analysis*. John Wiley & Sons Inc., 2004.
- [Brandt *et al.*, 2007] F. Brandt, T. Sandholm, and Y. Shoham. Spiteful bidding in sealed-bid auctions. In *Proc of the 20th International Joint Conference on Artificial Intelligence*, 2007.
- [Greenwald and Boyan, 2001] Greenwald and Boyan. Bid determination for simultaneous auctions. In *EC-01*, October 2001.
- [Krishna, 2002] Vijay Krishna. *Auction theory*. Academic Press, 2002.
- [McAfee and McMillan, 1987] R. Preston McAfee and John McMillan. Auctions and bidding. *Journal of Economic Literature*, 25:699–738, June 1987.
- [Reeves *et al.*, 2005] D. M. Reeves, K. M. Lochner, S. Cheng, and R. Suri. Approximate strategic reasoning through hierarchical reduction of large symmetric games. In *AAAI-05*, July 2005.
- [Rice, 1995] John A. Rice. *Mathematical Statistics and Data Analysis*. Duxbury Press, California, 1995.
- [Stone *et al.*, 2002] P. Stone, R. Schapire, M. Littman, J. Csirik, and D. McAllester. ATTac-2001: A learning, autonomous bidding agent. In *AMEC IV. LNCS, vol. 2531*. Springer Verlag, 2002.
- [Vetsikas and Jennings, 2007] I. A. Vetsikas and N. R. Jennings. Outperforming the competition in multi-unit sealed bid auctions. In *(forthcoming)*, 2007.
- [Vetsikas and Selman, 2003] I. A. Vetsikas and B. Selman. A principled study of the design tradeoffs for autonomous trading agents. In *Proceedings of the 2nd International Joint Conference on Autonomous Agents and Multi-Agent Systems*, pages 473–480, 2003.
- [Vetsikas and Selman, 2006] I. A. Vetsikas and B. Selman. Bayes-Nash equilibria for m -th price auctions with multiple closing times. In *SigEcom Exchanges*, 2006.
- [Vetsikas, 2005] Ioannis A. Vetsikas. *A principled methodology for the design of autonomous trading agents with combinatorial preferences in the presense of tradeoffs*. PhD thesis, Cornell University, August 2005.
- [Wellman *et al.*, 2001] M. P. Wellman, P. R. Wurman, K. O'Malley, R. Bangerla, S. Lin, D. Reeves, and W. E. Walsh. Designing the market game for TAC. *IEEE Internet Computing*, April, 2001.