

## NRC Publications Archive Archives des publications du CNRC

### TAC-REM - The Real Estate Market Game: A Proposal for the Trading Agent Competition

Buffett, Scott; Fasli, M.

This publication could be one of several versions: author's original, accepted manuscript or the publisher's version.  
/ La version de cette publication peut être l'une des suivantes : la version prépublication de l'auteur, la version acceptée du manuscrit ou la version de l'éditeur.

#### **Publisher's version / Version de l'éditeur:**

*Trading Agent Design and Analysis and Agent Mediated Electronic Commerce VIII [Proceedings], 2006*

**NRC Publications Archive Record / Notice des Archives des publications du CNRC :**  
<https://nrc-publications.canada.ca/eng/view/object/?id=d16831d9-7c36-418d-982c-de9cd5acf819>  
<https://publications-cnrc.canada.ca/fra/voir/objet/?id=d16831d9-7c36-418d-982c-de9cd5acf819>

Access and use of this website and the material on it are subject to the Terms and Conditions set forth at  
<https://nrc-publications.canada.ca/eng/copyright>

READ THESE TERMS AND CONDITIONS CAREFULLY BEFORE USING THIS WEBSITE.

L'accès à ce site Web et l'utilisation de son contenu sont assujettis aux conditions présentées dans le site  
<https://publications-cnrc.canada.ca/fra/droits>

LISEZ CES CONDITIONS ATTENTIVEMENT AVANT D'UTILISER CE SITE WEB.

**Questions?** Contact the NRC Publications Archive team at  
PublicationsArchive-ArchivesPublications@nrc-cnrc.gc.ca. If you wish to email the authors directly, please see the first page of the publication for their contact information.

**Vous avez des questions?** Nous pouvons vous aider. Pour communiquer directement avec un auteur, consultez la première page de la revue dans laquelle son article a été publié afin de trouver ses coordonnées. Si vous n'arrivez pas à les repérer, communiquez avec nous à PublicationsArchive-ArchivesPublications@nrc-cnrc.gc.ca.

# TAC-REM - The Real Estate Market Game: A Proposal for the Trading Agent Competition

Scott Buffett<sup>1</sup> and Maria Fasli<sup>2</sup>

<sup>1</sup> Institute for Information Technology - e-Business, National Research Council Canada, Fredericton, New Brunswick, Canada, E3B 9W4, scott.buffett@nrc.gc.ca

<sup>2</sup> University of Essex, Department of Computer Science, Wivenhoe Park, Colchester CO4 3SQ, UK, mfasli@essex.ac.uk

**Abstract.** In this game, agents will face-off against each other in the ultra-competitive real estate market. Each competitor will act as a real estate agent, working on behalf of clients who need to move into new homes. These clients need to buy a new home as well as sell their current home. The game will test competitors' technology in two main research areas: preference elicitation and multi-issue negotiation. Each time an agent acquires a new client, it must query the client about its various preferences (e.g. price range, number of bedrooms, etc.) for their new home. Agents then search the listings of the other agents, seeking a possible match. Once found, the agent then engages in negotiations with the selling agent, haggling over various aspects of the deal. Once a house has been purchased, the client's old house needs to be sold. The objective of the game is to earn the most money. Selling agents earn commissions from sales. Buying agents do not earn commissions, but instead need to maximize the utility of their clients by obtaining a good deal. Satisfied clients are more likely to keep their agent to sell their old house.

## 1 Introduction

Research in automated negotiation and preference elicitation has been gaining momentum in recent years. While the traditional storefront e-commerce model with take-it-or-leave-it pricing has been the norm because of its relative ease, in purchase scenarios where several attributes related to the transaction (other than just price) need to be decided upon, a one-size-fits-all method of determining such values becomes too restrictive. Attributes must instead be tailored to the needs of the customer as well as to the availability of the seller. In such situations, negotiation becomes a valuable tool for searching the space of possible agreements for a mutually acceptable transaction. Given the speed with which transactions can be negotiated and executed through various electronic services today, research in intelligent agent technology has been focusing increasingly on automated negotiation [3, 5, 6]. Two of the main focuses of research in automated negotiation include the specification of protocols (i.e. the set of rules that the agents must follow during negotiation), and the construction of effective strategies (i.e. the rules agents use to determine how to operate within the protocols to achieve a successful or satisfactory outcome).

In order to construct an effective strategy, an agent must have a utility model over the set of possible outcomes of the negotiation. That is, the agent must have an idea as to the degree of preference over outcomes, so that it can work toward achieving an outcome that is highly preferred. If an agent is working on behalf of a human user (which is most often the case in electronic commerce), then it must have some idea of the user's preferences. Preference elicitation (also referred to as utility elicitation here) [2] is the field of research dedicated to studying the problem of effectively extracting these preferences or utilities from the user. The main issues here involve determining how to query the user, how to infer information from responses to queries, and how to determine when sufficient information has been elicited.

The Trading Agent Competition (TAC) [9] is an international competition where players submit autonomous trading agents that participate in a simulated market. In the original version of the game, referred to as "TAC Classic", agents participate in auctions to purchase components of potential vacation packages. Each agent works on behalf of several clients who have utilities for various travel packages. Agents attempt to satisfy their clients as much as possible while minimizing costs.

In 2003 the Supply Chain Management game (TAC-SCM) [1] was developed and added to the TAC event. In this game, agents procure computer components from suppliers, assemble the computers and sell them to customers. Supplier prices are set according to supply and demand, and computers are sold via auctions. The object of this game is to make the most money.

In this paper, we propose a new game to the series referred to as the Real Estate Market game (TAC-REM). In this game, agents work on behalf of clients to buy and sell houses. To determine which houses may be potential matches for their clients, or what transaction conditions may or may not be acceptable, agents must elicit their clients' preferences. Once a potential match is found, agents representing the potential buyer and seller enter into negotiations in order to find a mutually acceptable agreement. Commissions are earned on houses sold. The winner at the end of a game is the agent with the most money. While the first two TAC installments focused on auction and supply-chain technology, this new game will capitalize on agents' capabilities in preference elicitation and automated negotiation.

In addition to a new TAC game, the ideas given here have the potential to break new ground in the real estate business. Such technology can help pave the way for real estate companies and agents to explore the use of information technology and electronic commerce. Gibler and Nelson's [4] demonstration of the descriptive theories of how home buyers make choices shows that there is a gap between decision making in practice, and the more normative theory of choice as dictated by the axioms of utility theory [7]. This is largely because of the high complexity of quantifying a buyer's preference over the large number of attribute values that need to be considered when comparing houses. Instead, buyers often take an "elimination by aspects" approach [8]. Here, the most important attribute is chosen (in this case usually price), and an acceptable range of values

is chosen. All alternatives falling outside of this range are then eliminated. Next the second-most important attribute is chosen, and so on. The decision-maker may iterate through this process several times, adjusting these ranges each time until a suitably-sized set of alternatives is found. Naturally, this decision-making strategy can provide sub-optimal results. By using our proposed system as simulation software, a real-estate agent can devise and test new intelligent methods for eliciting preference information from a client, and for using that information to make better decisions on which homes will make better matches.

In this preliminary paper, we present a general description of the game in section 2, followed by the game specification in section 3. This section gives details on such matters as daily activities, preference elicitation rules and procedures, utility functions and negotiation protocols, among others. Section 4 discusses a few issues that need to be resolved during the implementation and testing phases of the game production. In section 5 we close the paper with some final remarks.

## 2 Game Overview

The environment in which the agents compete is a typical real estate market. Each agent represents several clients, some of which are buyers and some of which are sellers. Buyers have preferences over the various attributes associated with potential houses. Such attributes include the number of bedrooms, number of bathrooms, square footage, type of neighbourhood, and whether or not it has a garage. Preferences for the combinations of various attribute values are modeled by a multi-attribute utility function. The agent representing the buyer must attempt to determine this utility function as accurately as possible by asking elicitation queries. There is little time to spend with clients, and each agent will likely have multiple clients, so the number of queries that can be made is limited. The agent must then find houses for sale that best meet the client's preferences.

Each selling client has a house for sale which is described using the values of the house attributes given above, with the addition that a price range is given as well. Each day, agents publish a listing of all of their clients' houses, accessible to all other agents.

A transaction is also made up of several attributes, including price, closing date, whether or not landscaping work will be done, whether or not maintenance will be performed on the heating/cooling system, and whether or not plumbing work will be done.

When an agent finds a house for sale that it believes will satisfy one of its buyers, it sets up a "showing". At this point, the client observes the attributes of the house and, based on its satisfaction with the house, computes and reveals its multi-attribute utility for the house, and indicates how likely it is that a potential transaction will meet or exceed its utility acceptability threshold. The seller of the house also has such a utility function and utility threshold for the various potential agreements. The two agents then enter into negotiations. Negotiations are bilateral where participants take turns exchanging one offer at a time, starting with the buyer. Each time an offer is received, the receiving

agent can choose to either accept the offer, reject and counteroffer, or reject and quit.

If the agents are successful in achieving a deal, the transaction is made with the selling agent receiving a commission on the purchase price. The selling client then happily disappears from the market, and the agent representing the seller takes on a new buyer client. The client that bought the house now needs to sell its current home, and thus becomes a seller. It may or may not, however, retain the services of its current real estate agent. The client will choose to leave the agent with some probability depending on its utility of the negotiated transaction. The higher the utility, the higher the probability that it will stay with the current agent. If it chooses to leave the agent, it will choose another agent in the game randomly with uniform probability.

The objective of the game is to earn the most money. Money can only be earned through commission on selling houses. The exact percentage of the commission is computed based on how well the various aspects of the eventual deal satisfy the client. Thus the selling agents have an incentive to find acceptable deals not only with high prices, but also with other attribute values that will satisfy its client. Buyer agents receive no commission from transactions, but rather retain the buyer as a client based on the client's satisfaction. If the buyer is retained then it becomes a seller, making it a potential money-maker. So there is an incentive for buyer agents to find deals that satisfy their clients as well.

### 3 Game Specifications

#### 3.1 Initial Setup

Six agents compete in a game. Initially each agent has 20 clients: 10 buyers and 10 sellers. Each buyer owns a home and wants to purchase a new one. Each seller simply has a home for sale (and presumably another one that they live in, but not for sale).

#### 3.2 Summary of Daily Activities

1. *Obtaining new clients.* Each day begins with the registration of new clients. An agent receives new buyer clients in one of two ways: 1) by selling houses: each agent obtains a number of new buyer clients equal to the number of houses sold the previous day (i.e. each satisfied seller is replaced with a new buyer the next day), and 2) by obtaining other agents' unsatisfied buyer clients who were unable to buy a home after 14 days. An agent obtains new seller clients in one of three ways: 1) by purchasing a house for a buyer the previous day and convincing it to stay on as a seller client, 2) by obtaining other agents' buyer clients who left them after buying a home the previous day, and 3) by obtaining other agents' unsatisfied seller clients who were unable to sell their home after 14 days. See section 3.9 for more specific details.

2. *Update Listings.* Agents send the specifications of the houses that their clients have for sale to the central Multiple Listing Service. The central listing is updated once daily, early in the day. Entries can be submitted at any time, but those not submitted in time for the update will not appear until the next day. See section 3.3 for more specific details.
3. *Utility Elicitation.* Throughout the remainder of the day, agents can elicit utility information from their buyer clients to determine their preferences for different types of houses. Queries can ask about house and transaction preferences, and can offer two possible responses (e.g. “Would you prefer A or B?”). To further ascertain utility, standard gambles may be included in the queries (e.g. “Would you prefer house type A or a gamble where you would receive house type B with 0.6 likelihood or house type C with 0.4 likelihood?”). See section 3.6 for more specific details.
4. *Showings.* A “showing” takes place when an agent allows a buyer client to assess its utility for a house that is for sale. At this point the client reveals its utility for the house (with average utilities used for the transaction attributes), and indicates the likelihood of being interested in buying the house, given its preferences for the likely transaction outcomes. An agent may choose to show its client a house belonging to one of its own selling clients, however the agent will incur penalties if it ultimately represents both clients in a transaction. See section 3.7 for more specific details.
5. *Negotiation.* If, after a showing, it is determined that there is an acceptable likelihood that an agreement on a transaction can be reached, then the agent may choose to enter in negotiations with the selling agent. Negotiations are bilateral, with the agent representing the buyer making the first offer. Each offer consists of a value for each transaction attribute. Negotiation continues until an agreement is reached or one of the agents quits. See section 3.8 for more specific details.
6. *Client Shuffle.* A client will find a new agent to work with if they are unsatisfied with their current one. Buyers and seller clients will stay with an agent for a maximum of two weeks (14 days). If the current agent has been unsuccessful in selling their current home or buying them a new one, they will leave the agent and enlist with a new one (randomly chosen uniformly). Also, a buyer client may leave its agent immediately after buying a new home, if it is somewhat unsatisfied with the purchase. The higher the utility of the purchase, the higher the probability that it will stay with its current agent. If it leaves, it will randomly choose a new agent. This shuffle phase takes place at the end of each day. No more of the above activities will take place once this phase begins. See section 3.9 for more specific details.

### 3.3 House Attributes, Listings

The house attributes are the attributes of a house that are non-negotiable, such as the size or the number of bedrooms. The house attributes and their domains are given in Table 1.

Attribute	Domain
House type	1-story, 2-story, split-level
Number of bedrooms	integer 1-5
Number of bathrooms	integer 1-4
Size	1000, 1500, 2000, 2500, 3000, 4000 sq ft
Neighbourhood type	urban, suburban, rural
Garage?	yes, no

**Table 1.** Attributes and domains for houses

The Multiple Listing Service (MLS) provides a public directory of all houses up for sale. In addition to the attribute values of each house, the listing provides information on the selling agent, the number of days on the market, and the price range. Example entries in the listing are given in Table 2.

Listing	Agent	Attributes	Price Range	Days on Market
0881	1	232421	[225000, 250000]	4
1021	6	343511	[250000, 275000]	2
1088	5	122230	[150000, 175000]	1

**Table 2.** Example listing of available houses

For example, the first entry indicates there is a house for sale with reference number 0881, agent number 1 is the selling agent, and the values for the attributes are as follows: the first attribute (type) has value 2 (2-story), the second attribute (# of bedrooms) has value 3 (3 bedrooms), the third attribute (# of bathrooms) has value 2 (2 bathrooms), the fourth attribute (size) has value 4 (2500 square feet), the fifth attribute (neighbourhood type) has value 2 (suburban) and the sixth attribute (garage?) has value 1 (yes). The price range is 225,000 - 250,000, and it has been on the market for 4 days. Note that all price ranges have size 25,000 with upper- and lower-limits being multiples of 25,000.

### 3.4 Transaction Attributes

The transaction attributes are those negotiable issues and conditions pertaining to the agreement of a sale, such as the price or whether landscaping work will be done by the current owners as a condition of sale. The transaction attributes and their domains are given in Table 3.

### 3.5 Utility Functions

Each buyer client  $b$  has a utility function  $u^b$  for each house and transaction attribute. These utility functions are additively used to make up a buyer's mul-

Attribute	Domain
Price	integer 100000-500000
Preferred closing date	buyer's, seller's
Landscaping?	yes, no
Heating system maintenance?	yes, no
Plumbing maintenance?	yes, no

**Table 3.** Attributes and domains for transactions

tiattribute utility function that assigns a utility to each potential combination of house and transaction attribute values:

$$u^b(\text{house}, \text{transaction}) = \sum_{a \in H \cup T} w_a^b u_a^b(x_a) \quad (1)$$

where  $H \cup T$  holds the house and transaction attributes,  $w_a^b$  is the weight of attribute  $a$ ,  $u_a^b$  is the utility function for  $a$ , and  $x_a$  is the value for  $a$  given the house and transaction.

During the house-hunting process, a buyer's utility for a house is computed using the average utility for each transaction attribute, since the transaction attribute values are unknown and do not yet come into play. The particular values for these attributes are then used during the negotiation phase to determine the buyer's overall utility of a deal.

Each seller client  $s$  has a utility function  $u^s$  for each transaction attribute. These utility functions are additively used to make up a seller's multiattribute utility function that assigns a utility to each potential combination of transaction attribute values:

$$u^s(\text{transaction}) = \sum_{a \in T} w_a^s u_a^s(x_a) \quad (2)$$

where  $T$  holds the transaction attributes,  $w_a^s$  is the weight of attribute  $a$ ,  $u_a^s$  is the utility function for  $a$ , and  $x_a$  is the value for  $a$  given the transaction.

For each type of client, the range of utilities for each attribute is  $[0,1]$ , with the most favourable value yielding a utility of 1 and the least favourable yielding a utility of 0. All weights sum to 1, and thus the multiattribute utilities lie in the  $[0,1]$  range. When  $a$  represents price,  $u_a$  is not necessarily linear, but can rather be shaped to model the client's marginally decreasing (or increasing) utility.

Many attributes have a natural ordering in the preference of their values. For example, buyers always prefer smaller values for price while sellers always prefer larger values. As another example, buyers always prefer larger values for size than smaller values. It may be the case that buyers with limited funds prefer houses of smaller size, but this is likely because smaller houses are usually cheaper. All things being equal, a buyer will prefer a larger house to a smaller one. Values for other attributes such as neighborhood type do not have such a natural ordering.



Some buyers may prefer urban living, while others prefer rural. Table 4 gives a summary of the orientation of each attribute.

Attribute	Orientation
House type	none
Number of bedrooms	buyer prefers more
Number of bathrooms	buyer prefers more
Size	buyer prefers more
Neighbourhood type	none
Garage?	buyer prefers yes
Price	buyer prefers less, seller prefers more
Preferred closing date	each prefer their own
Landscaping?	buyer prefers yes, seller prefers no
Heating system maintenance?	buyer prefers yes, seller prefers no
Plumbing maintenance?	buyer prefers yes, seller prefers no

**Table 4.** Orientation of attribute values

### 3.6 Preference Elicitation

In each elicitation query, the agent asks for the client’s preference over two standard gambles  $g_1$  and  $g_2$ . Each gamble  $g_i$  involves two potential deals  $d_i$  and  $d'_i$ , each made up of values for the house and transaction attributes, and a probability  $p_i$ . The uncertain outcome of  $g_i$  is  $d_i$  with probability  $p_i$  or  $d'_i$  with probability  $1 - p_i$ . So the query asks the following: “Would you prefer to take gamble  $g_1$  where you would receive deal  $d_1$  with probability  $p_1$  or  $d'_1$  with probability  $1 - p_1$ , or would you prefer to take gamble  $g_2$  where you would receive deal  $d_2$  with probability  $p_2$  or  $d'_2$  with probability  $1 - p_2$ ?” Values for all attributes are not required in the descriptions of  $d_1$  and  $d_2$ ; attributes with missing values will be ignored by the client. Agents can query about preferences over deals rather than gambles simply by making  $p = 1$  and leaving  $d'$  blank.

Another type of query is an acceptability query. This is used when an agent wants to know whether a proposed agreement meets or exceeds the client’s acceptability threshold. Such queries will typically be made during a negotiation. Here, the agent simply leaves  $g_2$  blank. If the client indicates that the offer (or gamble) given in  $g_1$  is preferred, then this means that  $g_1$  is deemed acceptable to the client.

Table 5 shows how to formulate queries to elicit particular types of information, other than the general query for the client’s preference over two standard gambles.

To determine a response to the query, the client uses its own secret utility function (described above) to compute the utility of each gamble:

$$u(g_i) = u(d_i)p_i + u(d'_i)(1 - p_i) \quad (3)$$

Query Type	$g_1$			$g_2$		
	$d_1$	$d'_1$	$p_1$	$d_2$	$d'_2$	$p_2$
Determine whether utility of deal $d$ is $>$ or $<$ $u$	$d^+$	$d^-$	$u$	$d$	-	-
Determine which of deals $d^A$ and $d^B$ is preferred	$d^A$	-	1	$d^B$	-	1
Determine whether deal $d$ is acceptable	$d$	-	-	-	-	-

**Table 5.** Formulation of query types (where  $d^+$  and  $d^-$  are known to be the most and least preferred deals, respectively).

where  $u(d_i)$  is computed using equation 1 for buyers and equation 2 for sellers. The client then indicates which of the two gambles has higher utility. In the case of a tie,  $g_1$  is selected.

When an attribute value is missing, the client computes the overall utility of a deal using the average utility over all possible values for this attribute. When querying about specific houses, the agent can enter the house reference number instead of entering all of the attribute values into the query.

### 3.7 Showings

When an agent believes that a house is a potential match for a buyer client, it may choose to “show” the house to the client. At this point, the client reveals its utility for the house (with average utilities used for the transaction attributes), and indicates the likelihood of being interested in buying the house. This likelihood is computed as the number of joint outcomes for transaction attribute values that meet its acceptability threshold (given the utilities of the house attribute values), divided by the total number of joint outcomes for transaction attribute values. Given this likelihood, the agent can decide whether to put in an offer (i.e. enter into negotiations with the seller agent), or to keep looking.

It may be the case from time to time that the best match for a buyer is a house that is being sold by a client represented by the same agent. An agent may represent both the buyer and the seller in a transaction, but will incur the following penalties:

1. 50% the commission earned from the sale
2. The buyer will leave the agent and choose to sell with another agent with 100% certainty

Thus such transactions are permissible but discouraged, and are likely only favourable as a last resort.

As showings take valuable time, each agent may show no more than 5 houses in a day, regardless of the number of clients it has.

### 3.8 Negotiating

A negotiation session begins with an agent representing a buyer sending an offer on a house to an agent representing the seller of that house. Each message in the

negotiation consists of a house (given by the reference number), the intended buyer (given by the buyer number), and values for all transaction attributes. Each time an offer is received, the receiving agent may 1) accept, 2) reject and send a counteroffer or 3) reject and quit. An offer is valid for a fixed length of time.


Each time an offer is sent, the agent must first ensure that the client's utility for the offer meets or exceeds its acceptability threshold. This can be done using an acceptability query, or else can be concluded using past information, since a client's preferences do not change over time. If an agent ever accepts an offer that does not meet its client's acceptability threshold, or submits such an offer which is subsequently accepted, this agreement is still binding. The offending agent(s) in such a situation is (are) penalized an amount equal to the amount of money that could be offered to the client to make the deal acceptable. For example, consider an agent who agrees on a purchase of a house with price 200,000 on behalf of a buyer client with acceptability threshold  $u'$ . If the deal has utility for the buyer less than  $u'$ , but a price of 190,000 (all else equal) would raise the utility of the purchase to  $u'$ , then the buyer would be penalized  $200,000 - 190,000 = 10,000$ . The same holds true for sellers. In this case, sellers could effectively forgo some of their commission in order to help a client sell their house. Agents may have negative balances.

Since accepted offers are considered binding agreements, an agent may make only one offer at a time for any given client. That is, a seller agent may make an offer to only one potential buyer for a house that it is selling. No other offers may be made for that house until 1) the first offer expires, or 2) the agent that received the first offer rejects it or submits a counteroffer. The agent may entertain several offers from buyers for a house, thus effectively participating in several simultaneous negotiations for the house, but may only have one outstanding offer from itself at any given time. The same goes for buyer agents who represent a client that is interested in more than one house. At any time, the agent may submit at most one offer on behalf of the buyer. Naturally, an agent can represent several buyers that are interested in the same house.

### 3.9 Obtaining/Losing Clients

Each time an agent successfully sells a house for a selling client, that client disappears from the game, and the agent receives a newly-produced buyer client the next day. When an agent successfully purchases a house for a buyer client, that client turns into a seller client the next day. However, if the client was not particularly satisfied, it may choose to leave the agent and find a new one. Let  $u'$  be the minimum utility deemed acceptable by the client and  $u(d)$  be the utility achieved by the sale given by  $d$ . If  $u(d) = u'$  (the worst case), the agent has a 50% chance of losing the client. If  $u(d) = 1$  (the best case), the agent has a 0% chance of losing the client. The probability for all other  $u(d)$  is computed by equation 4. The function from equation 4 is depicted in Figure 1. Note that  $u(d)$  will never really be less than  $u'$ . Whenever an agent makes a deal where this is the case, the agent is penalized the amount that effectively brings  $u(d)$  up to  $u'$ .

$$p(\text{lose}) = \frac{(\frac{u(d)-u'}{1-u'} - 1)^2}{2} \quad (4)$$



**Fig. 1.** Probability of a buyer client leaving its agent

Using this type of function (as opposed to a more linear function) ensures that the probability of losing a client is relatively low in most cases, but increases sharply if the agent exerts minimal effort and barely satisfies the client. This discourages the agent to simply agree to the first acceptable offer so it will have time to move on to other tasks.

Both buyer and seller clients will stay with the same agent for a maximum of 14 days. At the end of the 14th day, if the agent has not either purchased (in the case of the buyer) or sold (in the case of the seller) a house for the client, the client will leave the agent. A client that leaves an agent will enlist with a randomly chosen agent at the beginning of the next day. This client retains its original preferences, and thus the client's former agent may retain information learned about these preferences for possible later use. After an additional 14 days, if the client has still not completed a transaction, it is deemed to be not satisfiable and is removed from the system. It is then replaced with a new client of the same type and is assigned to an agent (possibly the current one) randomly.

When a buyer converts to a seller and stays with the same agent, this day counter is reset, with the next day (the first day that the agent can sell for them) being day 1.

### 3.10 Commission

The selling agent receives a commission from each sale. This is the only way to earn money in the game. The seller earns a base commission of 5% of the sale price for making the sale, plus a bonus of up to 5% more based on client satisfaction. The total commission  $c$  earned for deal  $d$  with sale price  $s$  is computed by

$$c = (0.05 + 0.05(\frac{u(d) - u'}{1 - u'}))s \quad (5)$$

where  $u'$  is the minimum utility deemed acceptable by the client and  $u(d)$  is the client's utility for  $d$ .

### 3.11 Duration

The game lasts 365 days. At the end the winner is chosen to be the agent that made the most money.

## 4 Future Work

At this point in the game development, there are several issues that remain unresolved. For the most part, such resolutions can reasonably be made only after implementation and experimentation. We mention some of these issues here.

- *Creating clients.* New clients are created randomly, but must be done so as to ensure fairness among the competitors. Clearly certain agents will have an advantage if they always receive buyer clients that are easy to please and seller clients with premium houses.
- *Reputation.* Under the current game specification, when a client leaves an agent it selects a new agent randomly with uniform probability. To make the game more realistic, we hope to introduce the concept of reputation. Agents' reputations will be based on how well they are able to satisfy their clients. Those agents with high reputation would then be more likely to land new clients.
- *Running time.* The game should not run longer than 60 to 70 minutes, which is comparable to the TAC Classic and TAC-SCM games. Otherwise it will be impractical to run in a competition setting. The length of a day would then be  $1/365$  of the game time.
- *Number of queries per day.* In order to be more realistic, and encourage the use of intelligent methods for determining optimal queries and inferring information from query responses, the number of times an agent is allowed to query a client should be limited. It may be the case that number of possible queries is limited by the short time that is available for doing so each day. Otherwise a hard cap of, say two queries per day per client, may have to be imposed.
- *Stages for each activity.* It may be best that each of the major activities that take place during the course of a day are divided into stages. For example, the day could begin with the listing stage, followed by the elicitation stage, followed by the showing stage, followed by the negotiation stage. Only the specified activity can be done during each stage. This will likely make things easier from an implementation perspective, while imposing a reasonable workflow on the agents' daily activities.
- *Offer expiry.* Offers cannot be valid indefinitely. An agent might not get an immediate response from the receiver of the offer, and cannot be expected to wait very long before looking elsewhere. This issue will be best resolved during the testing phase of the project.

## 5 Final Remarks

This paper gives a proposal for a third edition to the TAC series, called the Real Estate Market game (TAC-REM). Agents work on behalf of clients to buy and sell houses. To determine which houses may be potential matches for their clients, or what transaction conditions may or not be acceptable, agents must elicit their clients' preferences. Once a potential match is found, agents representing the potential buyer and seller enter into negotiations in order to find a mutually acceptable agreement. Commissions are earned on houses sold. The objective of the game is to make the most money over a 365-day period.

This paper serves as an invitation to the TAC community to consider the game presented in this paper, and provide us with feedback and suggestions on how to improve it. The game should provide an excellent forum for researchers in the areas of automated negotiation and preference elicitation to test their techniques and solutions. It may also pave the way for real-world real estate companies to explore the use of information technology and electronic commerce to conduct business. Above all, the game promises to be interesting and engaging, and thus has high potential to promote research in these areas by attracting and encouraging new researchers and students to participate.

## References

1. Raghu Arunachalam and Norman Sadeh. The 2003 supply chain management trading agent competition. In *Proc. of the 6th International Conference on Electronic Commerce (ICEC2004)*, pages 113–120, Delft, The Netherlands, 2004.
2. U. Chajewska, D. Koller, and R. Parr. Making rational decisions using adaptive utility elicitation. In *AAAI-00*, pages 363–369, Austin, Texas, USA, 2000.
3. S.S. Fatima, M. Wooldridge, and N. R. Jennings. An agenda-based framework for multi-issue negotiation. *Artificial Intelligence*, 152(1):1–45, 2004.
4. K. M. Gibler and S. L. Nelson. Consumer behavior applications to real estate education. *Journal of Real Estate Practice and Education*, 6(1):63–84, 2003.
5. N. R. Jennings, P. Faratin, A. Lomuscio, S. Parsons, C. Sierra, and M. Wooldridge. Automated negotiation: prospects, methods and challenges. *Int. J. of Group Decision and Negotiation*, 10(2):199–215, 2001.
6. N. R. Jennings, S. Parsons, C. Sierra, and P. Faratin. Automated negotiation. In *5th International Conference on the Practical Application of Intelligent Agents and Multiagent Systems (PAAM-2000)*, pages 23–30, Manchester, UK, 2000.
7. R. L. Keeney and H. Raiffa. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. John Wiley and Sons, Inc., 1976.
8. A. Tversky. Elimination by aspects: a theory of choice. *Psychological Review*, 79:281–299, 1972.
9. MP Wellman, PR Wurman, K O'Malley, R Bangera, S d Lin, D Reeves, and WE Walsh. Designing the market game for a trading agent competition. *IEEE Internet Computing*, 5:43–51, 2001.