An Automated Agent for Bilateral Negotiation with Bounded Rational Agents with Incomplete Information[1]

Raz Lin[2] and **Sarit Kraus**[2,3] and **Jonathan Wilkenfeld**[3,4] and **James Barry**[5]

Abstract. Many day-to-day tasks require negotiation, mostly under conditions of incomplete information. In particular, the opponent's exact tradeoff between different offers is usually unknown. We propose a model of an automated negotiation agent capable of negotiating with a bounded rational agent (and in particular, against humans) under conditions of incomplete information. Although we test our agent in one specific domain, the agent's architecture is generic; thus it can be adapted to any domain as long as the negotiators' preferences can be expressed in additive utilities. Our results indicate that the agent played significantly better, including reaching a higher proportion of agreements, than human counterparts when playing one of the sides, while when playing the other side there was no significant difference between the results of the agent and the human players.

1 Introduction

Many day-to-day tasks require negotiation, often in a bilateral context. In this setting two agents are negotiating over one or several issues in order to reach consensus [4, 6, 7]. The sides might have conflicting interests, expressed in their utility functions, and they might also cooperate in order to reach beneficial agreements for both sides [1, 12]. While models for bilateral negotiations have been extensively explored, the setting of bilateral negotiation between an automated agent and a bounded rational agent is an open problem for which an adequate solution has not yet been found. Despite this, the advantages of succeeding in presenting an automated agent with such capabilities cannot be understated. Using such an agent could help in training people in negotiations and assist in e-commerce environments by representing humans and other activities in daily life. Our proposed automated agent makes a significant contribution in this respect. An even more difficult problem occurs when incomplete information is added into this environment. That is, there is also lack of information regarding some parameters of the negotiation.

We consider a setting of a finite horizon bilateral negotiation with incomplete information between an automated agent and a bounded rational agent. The incomplete information is expressed as uncertainty regarding the utility preferences of the opponent, and we assume that there is a finite set of different agent types. The negotiation itself consists of a finite set of multi-attribute issues and time-constraints. If no agreement is reached by the given deadline a status-quo outcome is enforced. We propose a model of an automated agent for this type of negotiation by using a qualitative decision making approach [2, 13]. In our experiments we matched our automated agent against human negotiators. By analyzing the results of the experiments, we show that our agent is capable of negotiating efficiently and reaching multi-attribute agreements in such an environment. When playing one of the sides in the negotiation our agent reached significantly better results than the human players, and also allowed both sides to reach an agreement significantly faster. On the other hand, while our agent was playing the other side, though it did not reach significantly better results than the human player, it did not reach worse results. Also, there are no significant differences between the results reached by the human players and the automated agent playing that role.

The rest of the paper is organized as follows. Section 2 provides an overview of bilateral negotiation with incomplete information. Section 3 surveys related work done in the field of negotiation with incomplete information and bounded rational agents. Section 4 presents the design of our automated agent, including its beliefs updating and decision making mechanisms. Section 5 describes the experimental setting and methodology and reviews the results. Finally, Section 6 provides a summary and discusses future work.

2 Problem Description

We consider a bilateral negotiation in which two agents negotiate to reach an agreement on conflicting issues. The negotiation can end either when (a) the negotiators reach a full agreement, (b) one of the agents opts out, thus forcing the termination of the negotiation with an opt-out outcome denoted OPT, or (c) a predefined deadline is reached, denoted dl, in which, if a partial agreement was reached it is implemented or, if no agreement was reached, a status quo outcome, denoted SQ, is implemented. Let I denote the set of issues in the negotiation, o_i the finite set of values for each $i \in I$ and $O = o_1 \times o_2 \times \ldots \times o_{|I|}$ a finite set of values for all issues. Since we allow partial agreements $\emptyset \in o_i$ for each $i \in I$. An offer is then denoted as a vector $\vec{o} \in O$. It is assumed that the agents can take actions during the negotiation process until it terminates. Let **Time** denote the set of time periods in the negotiation, that is **Time** $= \{0, 1, ..., dl\}$. Time also plays a factor and has an influence on the agents' utilities. Each agent is assigned with a time cost which influences his utility as time passes.

In each period $t \in$ **Time** of the negotiation, if the negotiation has not terminated earlier, each agent can propose a possible agreement, and the opponent can either accept or reject the offer. Each agent can either propose an agreement which consists of all the issues in the negotiation, or a partial agreement. In contrast to the model of

[1] This research was supported in part by NSF under grant #IIS-0208608.
[2] Department of Computer Science, Bar-Ilan University, Ramat-Gan, Israel 52900. email: {linraz,sarit}@cs.biu.ac.il
[3] Institute of Advanced Computer Studies, University of Maryland.
[4] Department of Government and Politics, University of Maryland, College Park, Maryland. email: jwilkenf@gvpt.umd.edu
[5] Center for International Development and Conflict Management, University of Maryland, College Park, Maryland. email: jbarry6899@aol.com

alternating offers [12], each agent can perform up to M interactions with the opponent agent in a given time period. Thus, an agent must worry that its opponent may opt out in any time period.

Since we deal with incomplete information, we assume that there is a finite set of agent types. These types are associated with different additive utility functions. Formally, we denote the possible types of the agents **Types** $= \{1, \ldots, k\}$. We refer to an agent whose utility function is u_l as an agent of type l, and $u_l : \{(O \cup \{SQ\} \cup \{OPT\}) \times$ **Time**$\} \rightarrow \mathbb{R}$. Each agent is given its exact utility function. The agent, and the subjects in the experiments described later in the paper, are also aware of the set of possible utility functions of the opponent. However, the exact utility function of the rival is private information. Our agent has some probabilistic belief about the type of the other agent. This belief may be updated over time during the negotiation process (for example, using Bayes formula).

3 Related Work

The problem of modeling an automated agent for bilateral negotiation is not new for researchers in the fields of Multi-Agent Systems and Game Theory. However, most research makes simplifying assumptions that do not necessarily apply in genuine negotiations, such as assuming complete information [3, 11] or the rationality of the opponent [3, 4, 5, 7]. None of the above researchers has looked into the negotiation process in which there is both incomplete information and the opponent is bounded rational. While their approaches might be appropriate in their context, they cannot be applied to our settings.

To deal with the opponent's bounded rationality, researchers suggested shifting from quantitative decision theory to qualitative decision theory [13]. In using such a model we do not necessarily assume that the opponent will follow the equilibrium strategy or try to be a utility maximizer. Also, this model is better suited for cases in which the utility or preferences are unknown but can be expressed in ordinal scales or as preference relations [2]. This approach seems appropriate in our settings, and using the maximin criteria, which is generally used in this context, enables our agent to follow a pessimistic approach regarding the probability that an offer will be accepted.

Another way to deal with bounded rationality was suggested by Luce [8], who introduced the notion of *Luce numbers*. A Luce number is a non-negative number that is attached to each offer. The Luce number of an offer $\vec{o^*} \in O$ is calculated using the following formula:

$$lu(o^*) = \frac{u(\vec{o^*})}{\sum_{\vec{o} \in O} u(\vec{o})} \quad (1)$$

That is, the Luce numbers assign probabilistic beliefs regarding the acceptance of the offer by the opponent.

Several methods are proposed when dealing with the incomplete information regarding the preferences of an opponent. For example, Bayes theorem is the core component of the *Bayesian Nash equilibrium* ([12], p. 24-29), and it is used to deduce the current state given a certain signal. It allows us to compensate for incomplete information and enables good adaptation in a negotiation with time-constraints. In finite horizon negotiations there are no past interactions to learn from and not enough time periods to build a complete model. Thus this model supplies a good probabilistic tool to model the opponent, as opposed to neural networks [11] or genetic algorithms [7]. In our settings, Bayes theorem can be used to update the believed type of the opponent and at each time period act as if the opponent is of a certain type.

4 Agent Design

Our agent is built with two mechanisms: (a) a decision making mechanism, and (b) a mechanism for updating beliefs. We describe both mechanisms in the following subsections.

4.1 The Qualitative Decision Making Component (QDM)

The qualitative decision making component takes into account the agent's utility function, as well as the believed type of the opponent. This data is used both for deciding whether to accept or reject an offer and for generating an offer. In our settings, although several offers can be proposed in each time period, we restrict our agent to making a single offer in each period. This is done due to the fact that our mechanism for generating offers only produces one distinct offer at a given time period. The bounded rational agent, on the other hand, is free to propose several offers, and our agent can respond to all the offers, which indeed happened in the simulations. While we provide some theoretical foundations for our approach, we demonstrate its effectiveness using simulations with human subjects in an environment of incomplete information, as described in Section 5.

4.1.1 Generating Offers

The motivation behind the mechanism for generating offers is that the automated agent would like to propose an offer, which yields him the highest utility value. However, due to conflict of interests, there is a high probability that this agreement will be rejected by the opponent. To overcome this, our agent uses a qualitative decision strategy. Basically, the agent evaluates all possible offers based on its utility and the probability that the rival will accept them. *Luce numbers* are used to estimate this probability. We note that for every two offers x and y and agent j, where $lu_j(x)$ and $lu_j(y)$ denote the Luce numbers associated with offers x and y respectively for agent j, if $u_j(x) \geq u_j(y)$ then $lu_j(x) \geq lu_j(y)$. That is, the higher the Luce number of an offer, the greater the probability of it being accepted.

Since the opponent himself also tries to reason whether an offer will be accepted by our agent, we take the Luce numbers of both our agent and the opponent into account. That is, our agent tries to estimate, from the opponent's point of view, whether the opponent will accept the offer. This is done by calculating the sum of the the Luce numbers of the agent and the opponent. This sum is used as an estimation for the acceptance of the offer, and is multiplied by the utility value of the opponent from that offer. Finally, our agent compares those values with its own utility values. Similarly to the qualitative decision theory, which uses the maximin value [2, 13], our agent selects the minimum value between those two values, under the pessimistic assumption that the probability that an offer is accepted is based on the agent that favors the offer the least. After calculating the minimum value between all the offers, our agent selects the offer with the maximum value among all the minima, in order to also try and maximize its own utility. Thus, our qualitative offer generation mechanism selects, intuitively, the best offer among the offers that the agent believes that might be accepted. In the rest of this section we describe this process formally.

We assume that, at each time period t, the automated agent has a belief about the type of its opponent. This believed type, denoted by $BT(t)$, is the one that the agent presumes to be the most probable for the opponent. The agent uses the utility function associated with that type in all of the calculations in that time period. In Section 4.2 we describe in detail the elicitation of this belief. We denote by $u_{opp}^{BT(t)}$ the utility function associated with the believed type of the opponent at time t. From this utility function, our agent derives the *Luce numbers* [8]. Since the Luce number is calculated based on a given utility function, we denote the Luce number of an offer derived from the opponent's believed utility at time t, $BT(t)$, by $lu_{opp}(\text{offer} \mid BT(t))$, and the Luce number for an offer derived by the agent's own utility simply as $lu(\text{offer})$. We denote our function by $QO(t)$ (standing for Qualitative Offer), where $t \in$ **Time** is the time of the offer. If the current agent is j, the strategy selects an offer o in time t such that:

$$
\begin{aligned}
QO(t) &= \arg\max_{o \in O} \min\{\alpha, \beta\} \qquad (2)\\
\text{where} \quad \alpha &= u_j(o,t)\\
\text{and} \quad \beta &= [lu_{opp}(o \mid BT(t)) + lu(o)] \cdot u_{opp}^{BT(t)}(o,t)
\end{aligned}
$$

Seemingly, our QO function is a non-classical method for generating offers. However, not only were we able to show its efficacy by empirical experiments, in which it was used in negotiations with bounded ration agents, as we describe in Section 5, we also showed (Section 4.1.3) that it also conforms to some properties from classical negotiation theory, which are mainly used by mediators. Note that the formula does not build on the bounded rationality of the opponent.

The next subsection deals with the question of when the agent should accept an incoming offer or reject it.

4.1.2 Accepting Offers

The agent needs to decide what to do when it receives an offer from its opponent, offer_{opp}, at time $t-1$. If we refer to the automated agent as agent 1 and the bounded rational agent as agent 2, if $u_1(\text{offer}_{opp}) \geq u_1(QO(t))$ then our agent accepts the offer. Otherwise, our agent should not immediately rule out accepting the offer it has just received. Instead, it should take into consideration the probability that its counter-offer will be accepted or rejected by the opponent. This is done by comparing the believed utility of the opponent from the original offer as compared with the opponent's utility from our offer. If the difference is lower than a given threshold[6] T, that is $\mid u_2(QO(t)) - u_2(\text{offer}_{opp}) \mid \leq T$, then there is a high probability that the opponent will be indifferent between its original offer and our counter-offer, so our agent will reject the offer and propose a counter-offer (taking a risk that the offer will be rejected), since the counter-offer has a better utility value for our agent. If the difference is greater than the threshold, i.e., there is a higher probability that the opponent will not accept our counter-offer, our agent will accept the opponent's offer with a given probability, which is attached to each outcome. To this end we define the rank number, which is associated with each offer and a given utility function u, denoted $rank(\text{offer})$. The rank number of an offer is calculated by ordering all offers on an ordinal scale between 1 and $|O|$ according to their utility values, and dividing the offer's ordering number by $|O|$. That is, the agent will accept an offer with a probability $rank(\text{offer}_{opp})$ and reject and make a counter-offer $QO(t)$ with probability $1 - rank(\text{offer}_{opp})$. The intuition behind this is to enable the agent also to accept agreements based on their relative values, on an ordinal scale of $[0..1]$, and not based on their absolute values.

In the next subsection we demonstrate that our proposed solution also conforms to some properties of the *Nash bargaining solution*. This gives us the theoretical basis for the usage of our technique in bilateral negotiation, and for the assumption that offers proposed by our agent will also be considered to be accepted by the opponent.

[6] In the simulations, T was set to 0.05.

4.1.3 QO: An Alternative to Nash Bargaining Solution

We employ from Luce and Raiffa [9] the definitions of a bargaining problem and the Nash bargaining solution. We denote by $B = (u_1(\cdot), u_2(\cdot))$ the bargaining problem with two utilities, u_1 and u_2. The Nash bargaining solution (note that the solution is not the offer itself, but rather the payoff of the offer) is defined by several characteristics and is usually designed for a mediator in an environment with complete information. A bargaining (or a negotiation) solution f should satisfy symmetry, efficiency, invariance and independence of irrelevant alternatives, wherein *symmetry* states that if both players have the same bargaining power, then neither player will have any reason to accept an agreement which yields a lower payoff for it than for its opponent. For example, for the solution to be symmetric, it should not depend on the agent who started the negotiation process. *Efficiency* states that two rational agents will not agree on an agreement if its utility is lower for both of them than another possible agreement. This solution is said to be Pareto-optimal. *Invariance* states that for all equivalent problems B and B', that is $B' = (\alpha_1 + \beta_1 \cdot u_1(\cdot),\ \alpha_2 + \beta_2 \cdot u_2(\cdot)),\ \alpha_1,\ \alpha_2 \in \mathbb{R},\ \beta_1,\ \beta_2 \in \mathbb{R}^+$, the solution is also the same, $f(B) = f(B')$. That is, two positive affine transformations can be applied on the utility functions of both agents and the solution will remain the same. Finally, *independence of irrelevant alternatives* asserts that the solution $f(B) = f(B')$ whenever $B' \subseteq B$ and $f(B) \subseteq B'$. That is, if new agreements are added to the problem in such a manner that the status quo remains unchanged, either the original solution is unchanged or it becomes one of the new agreements.

It was shown by Nash [10] that the only solution that satisfies all of these properties is the product maximizing of the agents' utilities. However, as we stated, the Nash solution is usually designed for a mediator. Since we propose a model for an automated agent which negotiates with bounded rational agents following the QO function (Equation 2), our solution cannot satisfy all of these properties. To this end, we modified the independence of irrelevant alternatives property to allow for a set of possible solutions instead of one unique solution:

- PROPERTY 4A (Independence of irrelevant alternatives solutions) A negotiation solution f satisfies independence of irrelevant alternatives solutions if the set of all possible solutions of $f(B)$ is equal to the set of all possible solutions of $f(B')$ whenever $B' \subseteq B$ and $f(B) \subseteq B'$.

Proving that our agent's strategy for proposing offers conforms to those properties is important since although the agent should maximize its own utility, it should also find agreements that would be acceptable to its opponent.

Theorem 1 *The QO function satisfies the properties of symmetry, efficiency and independence of irrelevant alternatives solutions.*

Due to space limitations, we do not present the proof of the theorem. We also showed that in certain cases QO generates agreements which are better for the automated agent than agreements that would have been generated by following the Nash solution.

4.2 The Bayesian Updating Rule Component

The Bayesian updating rule (*BUR*) is based on Bayes Theorem described above and it provides a practical learning algorithm. We assert that there are several types (or clusters) of possible agents. The bounded rational agent should be matched to one such type. In each time period, our agent consults the BUR component in order to update its belief regarding the opponent's type.

Recall that there are k possible types for an agent. At time $t = 0$ the prior probability of each type is equal, that is, $P(\text{type}) = \frac{1}{k}, \forall\, \text{type} \in \textbf{Types}$. Then, for each time period t we calculate the posteriori probability for each of the possible types, taking into account the history of the negotiation. This is done incrementally after each offer is received or accepted. Then, this value is assigned to $P(\text{type})$. Using the calculated probabilities, the agent selects the type whose probability is the highest and proposes an offer as if this is the opponent's type. Formally, at each time period $t \in \textbf{Time}$ and for each type $\in$ **Types** and offer$_t$ (the offer at time period t) we compute:

$$P(\text{type}|\text{offer}_t) = \frac{P(\text{offer}_t|\text{type})P(\text{type})}{P(\text{offer}_t)} \tag{3}$$

where $P(\text{offer}_t) = \sum_{i=1}^{|\textbf{Types}|} P(\text{offer}_t|\text{type}_i) \cdot P(\text{type}_i)$. Since the Luce numbers actually assign probabilities to each offer, $P(\text{offer}_t|\text{type})$ is computed using the Luce numbers.

Now we can deduce the believed type of the opponent for each time period t, denoted as $BT(t)$, using the following equation:

$$BT(t) = \underset{\text{type}\in\textbf{Types}}{\arg\max}\ P(\text{type}|\text{offer}_t), \quad \forall t \in \textbf{Time} \tag{4}$$

Using this updating mechanism enables our updating component to conform to the following conditions, which are generally imposed on an agent's system of beliefs, and which are part of the conditions for a sequential Bayesian equilibrium [12]: (a) *consistency* and (b) *never dissuaded once convinced*. *Consistency* implies that agent j's belief should be consistent with its initial belief and with the possible strategies of its opponents. These beliefs are updated whenever possible, while *never dissuaded once convinced* implies that once an agent is convinced of its opponent's type with a probability of 1, or convinced that its opponent cannot be of a specific type, it is never dissuaded from its view. The results of the simulation indeed show that in more than 70% of the simulations our agent believed that its opponent is of the correct type with a probability of 1 or with the highest probability amongst the other possible types.

5 Experiments

We developed a simulation environment which is adaptable such that any scenario and utility functions, expressed as multi-issue attributes, can be used, with no additional changes in the configuration of the interface of the simulations or the automated agent. Our agent can play either role in the negotiation, while the human counterpart accesses the negotiation interface via a web address. The negotiation itself is conducted using a semi-formal language. Each agent constructs an offer by choosing the different values constituting the offers. Then, the offer is constructed and sent in plain English to its counterpart.

To test the efficiency of our proposed agent, we have conducted experiments on a specific negotiation domain[7]. In the following subsections we describe our domain, the experimental methodology and review the results.

[7] To show that our agent is capable of negotiating in other domains as well, we loaded another domain and tested the agent. As expected, the agent performs as well as in the specified domain. That is, only the utility functions play a role, and not the scenario or the domain.

5.1 Experimental Domain

The experimental domain adheres to the problem definitions described in Section 2. In our scenario England and Zimbabwe are negotiating in order to reach an agreement growing out of the World Health Organization's Framework Convention on Tobacco Control, the world's first public health treaty. The principal goal of the convention is "to protect present and future generations from the devastating health, social, environmental and economic consequences of tobacco consumption and exposure to tobacco smoke." There are three possible agent types, and thus a set of six different utility functions was created. This set describes the different types or approaches towards the negotiation process and the other party. For example, type (a) has a long term orientation regarding the final agreement, type (b) has a short term orientation, and type (c) has a compromise orientation.

Each negotiator was assigned a utility function at the beginning of the negotiation but had incomplete information regarding the opponent's utility. That is, the different possible types of the opponent were public knowledge, but the exact type of the opponent was unknown. The negotiation lasts at most 14 time periods, each with a duration of two minutes. If an agreement is not reached by the deadline then the negotiation terminates with a status quo outcome. Each party can also opt out of the negotiation if it decides that the negotiation is not proceeding in a favorable way. Opting out by England means trade sanctions imposed by England on Zimbabwe (including the ban on the import of tobacco from Zimbabwe), while if Zimbabwe opts out then it will boycott all British imports.

A total of 576 agreements exist, consisting of the following issues: (a) the total amount to be deposited into the Tobacco Fund to aid countries seeking to rid themselves of economic dependence on tobacco production, (b) impact on other aid programs, (c) trade issues, and (d) creation of a forum to explore comparable arrangements for other long-term health issues. While on the first two issues there are contradicting preferences for England and Zimbabwe, for the last two issues there are options which might be preferred by both sides.

5.2 Experimental Methodology

We tested our agent against human subjects, all of whom are computer science undergraduates. The experiment involved 44 simulations with human subjects, divided into 22 pairs. Each simulation was divided into two parts: (i) negotiating against another human subject, and (ii) negotiating against the automated agent. The subjects did not know in advance against whom they played. Also, in order not to bias the results as a consequence of the subjects getting familiar with the domain and the simulation, for exactly half of the subjects the first part of the simulation consisted of negotiating with a human opponent, while the other half negotiated first with the automated agent. The outcome of each negotiation is either the reaching of a full agreement, opting out, or reaching the deadline.

5.3 Experimental Results

The main goal of our proposed method was to enable the automated agent to achieve higher utility values at the end of the negotiation than the human negotiator, playing the same role, achieved, and to allow the negotiation process to end faster as compared to negotiations without our agent. Another secondary goal was to test whether our agent increased the social welfare of the outcome.

Table 1 summarizes the average utility values of all the negotiations and the average of the sums of utility values in all the experiments. HZ and HE denote the utility value gained by the human playing the role of Zimbabwe or England, respectively, and QZ and QE denote the utility value gained by the QO agent playing either role. The utility values ranged from -575 to 895 for the England role and from -680 to 830 for the Zimbabwe role. The Status-Quo value in the beginning of the negotiation was 150 for England and -610 for Zimbabwe. England had a fixed gain of 12 points per time period, while Zimbabwe had a fixed loss of -16 points.

Table 1. Final negotiations utility values and sums of utility values

Parameter	Avg	Stdev
QZ vs. HE	-27.4	287.2
HZ vs. HE	-260	395.6
QE vs. HZ	150.5	270.8
HE vs. HZ	288.0	237.1
HZ vs. QE	113.09	353.23
HE vs. QZ	349.14	299.36
Sum - HE vs. QZ	321.7	223.4
Sum - HZ vs. QE	263.6	270.5
Sum - HE vs. HZ	27.86	449.9

First, we examined the final utility values of all the negotiations for each player, and the sums of the final utility values. The results show that when the automated agent played the role of Zimbabwe, it achieved significantly higher utility values as opposed to a human agent playing the same role (using paired *t-test*: $t(22) = 2.23, p < 0.03$). On the other hand, when playing the role of England, there is no significant difference between the utility values of our agent and the human player, though the average utility value for the human was higher than for the automated agent. The results also show that the average utility values when the human played against another human are lower than the average utility values when he played against our agent. However, these results are significant only when the human players played the role of Zimbabwe ($t(22) = -3.43, p < 0.003$). One explanation for the higher values achieved by the QO agent is that the QO agent is more eager to accept agreements than humans. When playing Zimbabwe side, which has a negative time cost, accepting agreements sooner, rather than later, allowed the agent to gain higher utility values than the human playing the same side.

Comparing the sum of utility values of both negotiators, based on the role our agent played, we show that this sum is significantly higher when the negotiations involved our agent (either when our agent played the role of Zimbabwe or England), as opposed to negotiations involving only human players (using 2-sample *t-test*: $t(22) = 2.74, p < 0.009$ and $t(22) = 2.11, p < 0.04$).

Another important aspect of the negotiation is the outcome - whether a full agreement was reached or whether the negotiation ended with no agreement (either status-quo or opting out) or with a partial agreement. While only 50% of the negotiations involving only humans ended with a full agreement, 80% of the negotiations involving the automated agent ended with a full agreement. Using *Fishers Exact test* we determined that there is a correlation between the kind of the opponent agent (be it our agent or the human) and the form of the final agreement (full, partial or none). The results show that there is a significantly higher probability of reaching a full agreement when playing against our agent ($p < 0.006$).

Then we examined the final time period of the negotiations. All of the test results showed that in negotiations in which our agent was involved the final time period was significantly lower than the final time period in negotiations in which only humans were involved. The average final time period with two human negotiators was 11.36, while the average final time period against our agent playing the role of Zimbabwe and the role of England was 6.36 and 6.27 respectively. We used the 2-sample Wilcoxon test to compare between the final time periods of our agent and the human agent when playing against the same opponent (for England role $p < 0.001$ and for Zimbabwe role $p < 0.002$). In addition, we used the Wilcoxon signed rank test, to examine the final time period achieved by the same human player in the two simulations in which he participated - once involving another human player, and another involving an automated agent (for England role $p < 0.0004$, for Zimbabwe role $p < 0.001$).

We also examined the total number of offers made and received in each negotiation. The average number of offers made and received in negotiations involving our agent was 10.86 and 10.59, while in negotiation involving only human players the average number was 18.09. Using 2-sample Wilcoxon test we show that significantly fewer offers were made when our agent was involved ($p < 0.007$ and $p < 0.004$).

6 Conclusions

This paper outlines an automated agent design for bilateral negotiation with bounded rational agents where there is incomplete information regarding the opponent's utility preferences. The automated agent incorporated a qualitative decision making mechanism. The results showed that our agent is indeed capable of negotiating successfully with human counterparts and reaching efficient agreements. In addition, the results demonstrated that the agent played at least as well as, and in the case of one of the two roles, gained significantly higher utility values, than the human player. Though the experiments were conducted on a specific domain, it is quite straightforward to adapt the simulation environment to any other scenario.

Although much time was spent on designing the mechanism for generating an offer, the results showed that most of the agreements reached were offered by the human counterpart. Our agent accepted the offers based on a probabilistic heuristic. A future research direction would be to improve this mechanism and the probabilistic heuristic and then test it by conducting the same sets of experiments.

REFERENCES

[1] C. Boutilier, R. Das, J. O. Kephart, G. Tesauro, and W. E. Walsh, 'Cooperative negotiation in autonomic systems using incremental utility elicitation', in *Proceedings of UAI-03*, pp. 89–97, (2003).

[2] D. Dubois, H. Prade, and R. Sabbadin, 'Decision-theoretic foundations of qualitative possibility theory', *European Journal of Operational Research*, **128**, 459–478, (2001).

[3] P. Faratin, C. Sierra, and N. R. Jennings, 'Using similarity criteria to make issue trade-offs in automated negotiations', *AIJ*, **142**(2), 205–237, (2002).

[4] S. Fatima and M. Wooldridge, 'An agenda based framework for multi-issue negotiation', *AIJ*, **152**, 1–45, (2004).

[5] S. Fatima, M. Wooldridge, and N. R. Jennings, 'Bargaining with incomplete information', *AMAI*, **44**(3), 207–232, (2005).

[6] P. Hoz-Weiss, S. Kraus, J. Wilkenfeld, D. R. Andersen, and A. Pate, 'An automated agent for bilateral negotiations with humans', in *Proc. of AAAI/IAAI-02*, pp. 1000–1001, (2002).

[7] R. J. Lin and S. T. Chou, 'Bilateral multi-issue negotiations in a dynamic environment', in *Workshop on AMEC V*, (2003).

[8] R. D. Luce, *Individual Choice Behavior: A Theoretical Analysis*, John Wiley & Sons, NY, 1959.

[9] R. D. Luce and H. Raiffa, *Games and Decisions - introduction and critical survey*, John Wiley & Sons, NY, 1957.

[10] J. F. Nash, 'The bargaining problem', *Econ.*, **18**, 155–162, (1950).

[11] M. Oprea, 'An adaptive negotiation model for agent-based electronic commerce', *Studies in Informatics and Control*, **11**(3), 271–279, (2002).

[12] M. J. Osborne and A. Rubinstein, *A Course In Game Theory*, MIT Press, Cambridge MA, 1994.

[13] M. Tennenholtz, 'On stable social laws and qualitative equilibrium for risk-averse agents', in *KR*, pp. 553–561, (1996).