


Provided for non-commercial research and education use.
Not for reproduction, distribution or commercial use.



This article was published in an Elsevier journal. The attached copy is furnished to the author for non-commercial research and education use, including for instruction at the author's institution, sharing with colleagues and providing to institution administration.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.

In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

<http://www.elsevier.com/copyright>



ELSEVIER

Available online at www.sciencedirect.com
 ScienceDirect

Artificial Intelligence 172 (2008) 823–851

**Artificial
Intelligence**

www.elsevier.com/locate/artint

Negotiating with bounded rational agents in environments with incomplete information using an automated agent [☆]

Raz Lin ^{a,*}, Sarit Kraus ^{a,c}, Jonathan Wilkenfeld ^{b,c,d}, James Barry ^d^a Department of Computer Science, Bar-Ilan University, Ramat-Gan, Israel 52900^b Department of Government and Politics, University of Maryland, College Park, Maryland, USA^c Institute of Advanced Computer Studies, University of Maryland, USA^d Center for International Development and Conflict Management, University of Maryland, College Park, Maryland, USA

Received 6 March 2007; received in revised form 24 September 2007; accepted 24 September 2007

Available online 17 October 2007

Abstract

Many tasks in day-to-day life involve interactions among several people. Many of these interactions involve negotiating over a desired outcome. Negotiation in and of itself is not an easy task, and it becomes more complex under conditions of incomplete information. For example, the parties do not know in advance the exact tradeoff of their counterparts between different outcomes. Furthermore information regarding the preferences of counterparts might only be elicited during the negotiation process itself. In this paper we propose a model for an automated negotiation agent capable of negotiating with bounded rational agents under conditions of incomplete information. We test this agent against people in two distinct domains, in order to verify that its model is generic, and thus can be adapted to any domain as long as the negotiators' preferences can be expressed in additive utilities. Our results indicate that the automated agent reaches more agreements and plays more effectively than its human counterparts. Moreover, in most of the cases, the automated agent achieves significantly better agreements, in terms of individual utility, than the human counterparts playing the same role.

© 2007 Elsevier B.V. All rights reserved.

Keywords: Bilateral negotiation; Bounded rationality; Incomplete information; Automated agent

1. Introduction

Various tasks in day-to day life require negotiation capabilities. These can be as simple and ordinary as haggling over a price in the market, through deciding what show to watch on TV. On the other hand, it can also involve issues over which millions of lives are at stake, such as interstate disputes [25] and nuclear disarmament [9]. No matter what the domain, the negotiation process itself is not an easy task. The parties will have conflicting interests with reference to some aspects of the negotiation. On the other hand, both sides might also have the incentive to cooperate with each

[☆] This research was supported in part by NSF under grant #IIS-0208608. A preliminary version of this paper was published at ECAI 2006, Italy. This paper was selected for the Fast Track submission.

* Corresponding author.

E-mail addresses: linraz@cs.biu.ac.il (R. Lin), sarit@cs.biu.ac.il (S. Kraus), jwilkenf@gvpt.umd.edu (J. Wilkenfeld), jbarry6899@aol.com (J. Barry).

other, as reaching an agreement could be more beneficial for them than walking away without any agreement ([20], Chapter 7).

Negotiation is a profession, yet on many occasions, ordinary people need to become involved in negotiation tasks. Thus, success in modeling such an agent has great advantages and implications: from using it for training people in negotiations, to assisting in e-commerce environments, as well as for the development of tools for modeling negotiation behavior in general. We propose a model of such an automated agent. Thus, we make a significant contribution in this respect.

With regard to the negotiation model, we consider a setting of a finite horizon bilateral negotiation with incomplete information between an automated agent and a human counterpart. The negotiation is said to have a finite horizon if the length of every possible history of the negotiation is finite ([20], p. 90). 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. The negotiation consists of multi-attribute issues if the parties have to negotiate an agreement which involves several attributes for each issue. This can help in making complex decisions while taking into account multiple factors [10]. Costs are assigned to each agent, such that during the negotiation process, the agents might gain or lose utility over time. If no agreement is reached by the given deadline a status quo outcome is enforced.

Our automated agent is capable of negotiating efficiently in such environments, as our experimental results show. Nonetheless, in order to allow our agent to negotiate efficiently in these settings, we had to decide how to allow it to deal with the uncertainty both regarding the environment and the opponent. To achieve this, we incorporated two mechanisms in the automated agent. The first deals with incomplete information regarding the opponent by using a practical learning algorithm based on Bayes' theorem which updates the agent's beliefs regarding its opponent. The second mechanism deals with the bounded rationality of the opponent. While our model applies utility functions, it is based on a non-classical decision making method. Instead of focusing on maximizing the expected utility, we are motivated by qualitative decision making approaches [4,24] and we use the maximin function and the qualitative valuation of offers in our model. Using these methods our automated agent generates offers and decides whether to accept or reject proposals it has received.

We conducted three sets of experiments in which we matched our automated agent against (a) human negotiators, (b) an automated agent following an equilibrium strategy, and (c) against itself—that is, the same models of an agent playing both sides. The experiments were run on two distinct domains. In the first domain, England and Zimbabwe negotiate in order to reach an agreement evolving from the World Health Organization's Framework Convention on Tobacco Control, the world's first public health treaty. In the second domain a negotiation takes place after a successful job interview between an employer and a job candidate.

By analyzing the results of the experiments we conducted, we show that our automated agent is capable of negotiating efficiently and reaching multi-attribute agreements in such environments.

When playing one of the sides in the negotiation (England in the first domain and the job candidate in the second domain) our agent achieved significantly higher utility values than the human players, and agreements were reached faster than when an agent was not present in the negotiation. On the other hand, when the agent played the other side, though it reached higher utility values than the human player, these results were not significantly higher than the humans' results.

This paper contributes to research on automated negotiation in several ways. First, it tackles the problem of multi-attribute negotiation with incomplete information. Given the importance of negotiating efficiently in such an environment we provide a generic mechanism that achieves just that. Second, we present an automated negotiation agent which is domain independent and allows to experiment with almost any real-life domain. Together, the automated negotiation environment will enable exploration of future research directions and thereafter it can be used to better understand behavioral and cognitive aspects of negotiations undertaken by human negotiators.

The remainder of the paper is organized as follows. Section 2 provides an overview of bilateral negotiation with incomplete information. Section 3 surveys related work in the field of negotiation with incomplete information and bounded rational agents. Section 4 presents the design of the 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 lists recommendations for future work in this area.

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 , whereby, if a partial agreement is reached it is implemented or, if no agreement is reached, a status quo outcome, denoted SQ , is implemented. Since no agreement is worse than any agreement, and a status quo is implemented if the deadline is reached, we assume that default values are assigned to each attribute. Thus, if both sides agree only on a subset of the issues and the deadline is reached, the unresolved issues are assigned with their default value and thus a partial agreement can be 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 a finite set of values for all issues ($O_1 \times O_2 \times \dots \times O_{|I|}$). We allow partial agreements, $\perp \in O_i$ 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 $\mathbf{Time} = \{0, 1, \dots, dl\}$. Time also impacts the agents' utilities. Each agent is assigned a time cost which influences its utility as time passes.

In each period $t \in \mathbf{Time}$ of the negotiation, if the negotiation has not terminated earlier, each agent can propose a possible agreement, and the opponent can either accept the offer, reject it or opt out. 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 alternating offers ([20], pp. 118–121), each agent can perform up to $M > 0$ interactions with the opponent agent in each time period. Thus, an agent must take into account 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 (e.g., one type might have a long term orientation regarding the final agreement, while the other type might have a more constrained orientation). Formally, we denote the possible types of the agents $\mathbf{Types} = \{1, \dots, k\}$. Given $l \in \mathbf{Types}$, $1 \leq l \leq k$, we refer to the utility of an agent of type l as u_l , and $u_l: ((O \cup \{SQ\}) \cup \{OPT\}) \times \mathbf{Time} \rightarrow \mathbb{R}$. Each agent is given its exact utility function. Also, each agent of type $l \in \mathbf{Types}$ is given a certain reservation price, denoted r_l , which is held private. The reservation price states the minimum value of the utility of an offer under which the agent is unwilling to accept the offer. The agent, and the subjects in the experiments described later in the paper, are also aware of the set of possible types 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' rule).

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 assumptions that do not necessarily apply in genuine negotiations with humans, such as assuming complete information [5,19] or the rationality of the human negotiator [5–7,13]. In this sense, they assume that both parties are rational in their behavior (e.g., describing the decisions made by the agents as rational and that they are utility maximizing agents that cannot deviate from their prescribed behavior). For example, Faratin et al. [5] assume that the agents are motivated by maximizing the joint utility of the outcome, that is, the agents are utility maximizers that seek Pareto-optimal agreements. In a similar manner, [6,7,13] assert that the agent's strategy must follow the equilibrium's strategy such that it should be the best response to the opponent's strategy. In other words, no agent will have an incentive to deviate from the strategy. None of the above researchers has looked into the negotiation process involving both incomplete information and negotiations against humans. While their approaches might be appropriate in their context, they cannot be applied to our settings.

Dealing only with the bounded rationality of the opponent, several researchers suggested new notions of equilibria (e.g., the *trembling hand equilibrium* described in Rasmusen [21] (p. 139)) or other probability models. For example, Capra et al. [3] use what is called a "*standard logit model*". In this model probabilities are assigned to the decisions. Those probabilities are proportional to exponential functions of expected payoffs. They use this model in order to enable the players to update their beliefs about other players. This model is equivalent to assuming expected payoffs are subjected to deviations with an extreme value distribution. That is, the logit model assumes that the decisions are

not perfect and may have some noise. It also tries to deal with such situations. These noisy decisions can be interpreted either as unobserved random changes in preferences or as errors in responding to expected payoffs. Similar to Capra et al., our agent also assigns probability to the believed type of the opponent. However, we try to avoid the need of adding a special mechanism that assumes that the actions of the opponent are characterized by noise.

In addition, in a recent paper, Kraus et al. [11] describe an automated agent that negotiates efficiently with humans. Although they also deal with negotiation against humans, in their settings there is complete information. That is, the agents know exactly what the world state parameters are and how changing them affects the other agent. First, they identified the perfect equilibrium strategies in their model. Then, however, they observed, that the human players do not necessarily follow these equilibrium strategies. Thus, they added heuristics and argumentation tailored to their specific settings to enable effective negotiation by their perfect equilibrium agent. We, on the other hand, propose a general model of an agent, who can negotiate efficiently in multi-attribute negotiations against bounded rational agents with incomplete information.

Other researchers suggested shifting from quantitative decision theory to qualitative decision theory [24]. 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 [4]. 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 [15], who introduced the *choice axiom*. The choice axiom, in relation to negotiations, states that the probability of selecting one offer over another from a pool of offers, is not affected by the presence or absence of other items in that pool. The axiom introduces the notion of *Luce numbers*. A Luce number is a non-negative number that is associated with each offer. The Luce number of an offer $\vec{o} \in O$ is calculated using the following formula:

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

From the mathematical definition the following property follows:

Property 3.1 (*Luce Number Relation*). For every offer \vec{x} and \vec{y} , if $u(\vec{x}) \geq u(\vec{y})$ then $lu(\vec{x}) \geq lu(\vec{y})$ when $lu(\vec{x})$ and $lu(\vec{y})$ denote the Luce number associated with offer \vec{x} and \vec{y} respectively.

Mostly in economics, this model is used to assign probabilistic beliefs regarding the tendency of the consumer's to choose one offer over another (e.g., see Luce's survey of the choice axiom [16]). As such, we believe that this model can be used as an estimation of the acceptance rate of the opponent's offer.

Several methods are proposed when dealing with incomplete information regarding the preferences of an opponent. For example, Bayes' theorem is the core component of the *Bayesian Nash equilibrium* ([20], pp. 24–29), and it is used to deduce the current state given a certain signal. One motivation for using this notion of equilibrium is that it allows one to compensate for incomplete information and enables good adaptation in negotiations 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 provides a good probabilistic tool to model the opponent, as opposed to using feed-forward neural networks [19] or genetic algorithms [13], both of which require considerable time to facilitate adequate learning and are more sensitive to the domain in which they are run.

Zeng and Sycara [26] also use the Bayesian analysis as a learning mechanism in negotiations. And like them, we also use Bayes' theorem to update the believed type of the opponent. Thus, we allow the negotiator, at each time period, to act as if the opponent is of a certain type.

Since we aim to design an agent that can negotiate efficiently against bounded rational agents in conditions of incomplete information, we should design it not only with a good mechanism to deal with the bounded rationality of the opponent, but also with a mechanism which can help overcome the incomplete information settings. Following the surveyed related work, we believe that implementing a non-classical valuation mechanism, using the Luce numbers, for generating and responding to offers, while also incorporating a Bayesian belief update component to deal with the incomplete information, can lead to an agent's efficient design. We elaborate on this design in the next section.

4. Agent design

Due to the unique settings in which we operate—incomplete information and bounded rationality of the opponent—the agent is built with two mechanisms: (a) a learning mechanism, and (b) a decision making mechanism.

For the learning mechanism we use the formal model of the Bayesian updating rule and for the decision making mechanism we incorporate a non-classical model of offers' valuation, rather than the traditional quantitative decision making model. We describe both mechanisms in the following subsections.

4.1. The decision making valuation component

As Brafman and Tennenholtz [1] state, there are extensions to qualitative decision theory yet to be pursued. We contend that one such extension should be in negotiation under uncertainty against a bounded rational agent. One reason is due to the fact that, as in real life, we do not impose restrictions on the bounded rational agent and we do not assert that it will follow an equilibrium strategy. We propose a unique decision making model. While our model is quantitative in spirit, we shift from the model of expected utility maximization and try to evaluate the offers in a more qualitative way, using the maximin method and the ranking of offers, described below.

The decision making valuation component takes into account the agent's utility function, as well as the believed type of the opponent (note that the believed type of the opponent is also influenced by the offers proposed by the opponent, as described in Section 4.2). 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 the 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 opponent, on the other hand, is free to propose several offers, and the agent can respond to all the offers, which indeed happened in the experiments. While we provide some theoretical foundations for our approach, we demonstrate its effectiveness using experiments with people 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 it the highest utility value. However, due to conflicting interests, there is a high probability that this offer will be rejected by the opponent. To overcome this, the agent evaluates all possible offers based on their utility and the probability that the rival will accept them.

An initial version of this offer generating mechanism was presented in [14]. Based on the preliminary experimental results, we have decided to improve this mechanism. Thus, instead of directly using the utility value of an offer, our mechanism uses the ranking value of an offer, which is associated with each offer and a given utility function u , denoted $rank(\cdot)$. The rank number of an offer $\vec{o} \in O$ is calculated using the following formula:

$$rank(\vec{o}) = \frac{order(\vec{o}, O)}{|O|} \quad (2)$$

where $order(\cdot, \cdot)$ is the ordering function which orders the offer \vec{o} in an ordinal scale between 1 and $|O|$ according to its utility value compared to all other offers in O . In order to facilitate the computations we also divide the offer's ordering number by $|O|$ to obtain a range between $[0, 1]$. Note that a certain offer might be ranked differently when using the utility value of the agent for that offer and the believed utility value of the opponent. In addition, the Luce numbers are used to estimate the probability of the agent accepting an offer, following Luce's choice axiom models, which assign a probability to each offer [15]. Based on Property 3.1, we get that the higher the utility of the offer the higher its Luce number as well. As the choice axioms is usually used to assign probabilistic beliefs regarding the tendency of an agent to propose an offer, we also believe that using the Luce numbers can provide a good estimation of an acceptance of offers.

Since the opponent also tries to estimate whether an offer will be accepted by the agent, we take the Luce numbers of both into account. Then, the agent tries to estimate, from the opponent's point of view, whether the opponent will accept the offer. This is done by estimating the sum of the Luce numbers of the agent and the presumed type of the opponent. This sum is used as an estimation of the acceptance of the offer, and is multiplied by the ranking value of that offer of the opponent. The sum of both agents is used in the calculation as a method for estimating the *social*

welfare of both sides. Finally, the agent compares these values with its own estimation of the offer, which is based on its ranking of the offer and the probability that it will accept the offer.

Similar to qualitative decision theory, which uses the maximin value [4,24], the agent selects the minimum value between these two values (the agent's own estimation of the offer and the agent's estimation from the opponent's point of view), 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 of all the offers, the 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 might be accepted. In the rest of this section we describe this process formally.

First we will demonstrate this using the following example.

Example 1. Bob (*b*) and Alice (*a*) are negotiating about what to do over the weekend. They have to decide about the activity and on which night they will do it. Given that *I* is the set of issues, let *Activity* $\in I$ and *Night* $\in I$ be the possible issues. There are two possible values for the activity—seeing a movie (*M*) or going to a basketball game (*B*) and two possible values for the chosen night—either they will go on a Friday night (*F*) or on Saturday night (*S*). Bob prefers going to a basketball game than to a movie. He also prefers it to be on a Saturday night, as the match should be more interesting than on Friday night. However, if a movie is chosen, he prefers to go on Friday to allow him to catch the game on ‘Pay Per View’ on Saturday. On the other hand, Alice prefers going out to a movie rather than to a basketball game. The best night for a movie would be Saturday as on Saturday is the opening night for the premieres, while if going to a basketball game, she prefers to stay home on Saturday and not engage in any activity. Table 1 reflects the utility values for Bob and Alice derived from their preferences, described above, for all the possible offers, along with the Luce number values and the ranking of the offers. In addition, Table 2 also compares the ranking of the offers between Bob and Alice. Let the reservation price for both sides be equal to 5. For simplicity and demonstration purposes, we assume that there is only one type of utility for each side and no time effect for both sides.

Assume that our agent plays the role of Bob. It will now choose the maximum value among all the minima between the values in lines 7 and 9 in Table 1. Thus, our agent will offer Alice to go to a basketball game on Friday.

We continue now to describe the process formally. 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

Table 1
Example of calculating *QO*

		$\vec{o}_1 = \{M, S\}$	$\vec{o}_2 = \{M, F\}$	$\vec{o}_3 = \{B, S\}$	$\vec{o}_4 = \{B, F\}$
1	$u_b(\vec{o}_i)$	4	6	10	8
2	$u_a(\vec{o}_i)$	10	9	4	6
3	$lu_b(\vec{o}_i)$	$4/28 = 0.14$	$6/28 = 0.21$	$10/28 = 0.36$	$8/28 = 0.29$
4	$lu_a(\vec{o}_i)$	$10/29 = 0.34$	$9/29 = 0.31$	$4/29 = 0.14$	$6/29 = 0.21$
5	$rank_b(\vec{o}_i)$	$1/4 = 0.25$	$2/4 = 0.50$	$4/4 = 1.00$	$3/4 = 0.75$
6	$rank_a(\vec{o}_i)$	$4/4 = 1.00$	$3/4 = 0.75$	$1/4 = 0.25$	$2/4 = 0.50$
7	$lu_b(\vec{o}_i) \cdot rank_b(\vec{o}_i)$	0.04	0.11	0.36	0.21
8	$lu_a(\vec{o}_i) \cdot rank_a(\vec{o}_i)$	0.34	0.23	0.03	0.10
9	$[lu_b(\vec{o}_i) + lu_a(\vec{o}_i)] \cdot rank_a(\vec{o}_i)$	0.49	0.39	0.12	0.25
10	$[lu_a(\vec{o}_i) + lu_b(\vec{o}_i)] \cdot rank_b(\vec{o}_i)$	0.12	0.26	0.49	0.37

Table 2
Ranking of offers

Ranking	Bob	Alice
$4/4 = 1.00$	$\vec{o}_3 = \{B, S\}$	$\vec{o}_1 = \{M, S\}$
$3/4 = 0.75$	$\vec{o}_4 = \{B, F\}$	$\vec{o}_2 = \{M, F\}$
$2/4 = 0.50$	$\vec{o}_2 = \{M, F\}$	$\vec{o}_4 = \{B, F\}$
$1/4 = 0.25$	$\vec{o}_1 = \{M, S\}$	$\vec{o}_3 = \{B, S\}$

all of the calculations in that time period. In Section 4.2 we describe in detail how the belief is dynamically updated. 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, the agent derives the Luce numbers. Since the Luce number is calculated based on a given utility function, we denote the Luce number of an offer, $\vec{o} \in O$, derived from the opponent's believed utility at time t , $u_{opp}^{BT(t)}$, by $lu_{opp}(\vec{o}|u_{opp}^{BT(t)})$, and the Luce number for an offer derived by the agent's own utility simply as $lu(\vec{o})$. We denote our function by $QO(t)$ (standing for Qualitative Offer), where $t \in \mathbf{Time}$ is the time of the offer.

Thus, if the current agent is j , the strategy selects an offer o in time t such that:

$$QO(t) = \arg \max_{o \in O} \min\{\alpha, \beta\}$$

$$\text{where } \alpha = rank(\vec{o}) \cdot lu(\vec{o}) \text{ and } \beta = [lu_{opp}(\vec{o}|u_{opp}^{BT(t)}) + lu(\vec{o})] \cdot rank_{opp}^{BT(t)}(\vec{o}) \quad (3)$$

In the example specified in Example 1 and Table 1, assuming our agent plays the role of Bob, lines 7 and 9 represent α and β , respectively. Then, to calculate the actual offer to propose to Alice, the agent calculates the minimum values between these two values, and finally chooses the maximum between all the minima calculated. That is, the agent will choose \vec{o}_3 since it has the maximum value of the minimum values associated with the four offers (0.04, 0.11, 0.12, 0.21).

Seemingly, our QO function is a non-classical method for generating offers. However, not only were we able to show its efficacy through empirical experiments, in which people used it in negotiations, as we describe in Section 5, we also showed (Section 4.1.3) that it conforms to some properties from classical negotiation theory, which are mainly used by mediators. In addition, our QO function does not build on a fixed opponent's type. That is, eliciting the believed type of the opponent is external to the computation.

The function can work as well with other software agents, and was tested in another two sets of experiments—the first matched our agent against an automated agent that follows a Bayesian Equilibrium strategy, and the second matched our agent against itself. Detailed results are described in Section 5.3.

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

4.1.2. Accepting offers

The agent needs to decide what to do when it receives an offer from its opponent at the current time t . In the following analysis, the terms a and b are used synonymously as the types of agents a and b , respectively. In the analysis we will refer to the automated agent as agent type a and its opponent as agent type b . Similarly, the terms \vec{o}_a and \vec{o}_b respectively denote an offer received from agent a and agent b . If $u_a(\vec{o}_b) \geq u_a(QO(t+1))$ then the automated agent accepts the offer. Otherwise, it 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 the agent's offer. If the difference is lower than a given threshold¹ T , that is $|u_b(QO(t+1)) - u_b(\vec{o}_b)| \leq T$, then there is a high probability that the opponent will be indifferent to its original offer and the agent's counter-offer. Thus, the automated 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 higher utility value for the agent. If the difference is greater than the threshold, i.e., there is a higher probability that the opponent will not accept the agent's counter-offer, the automated agent will accept the opponent's offer with a given probability, which is attached to each outcome, as described below. While in our original version of the agent [14] we did not impose any restrictions on accepting offers, in the new version, acceptance by the agent depends on whether the value of the offer is greater than the agent's reservation price (r). If this is the case, then the probability is calculated based on the ranking number of the offer. The intuition behind this is to enable the agent to also accept offers based on their relative values compared to the reservation price, on an ordinal scale of $[0..1]$. That is, the agent will accept an offer if $u_a(\vec{o}_b) \geq r_a$ and with a probability of $rank(\vec{o}_b)$ it will reject it and make a counter-offer.

Let's return to Example 1 presented above. Assume Alice suggests to Bob to go to a movie on Friday. The utility value of that offer to Bob equals 6. Bob now checks what would be his utility value from an offer he would make to Alice in the next time period. We showed earlier that Bob will suggest an offer to go to a basketball game on Friday.

¹ In the experiments, T was set to 0.05.

The utility value of that offer to Bob is 8. Thus, since the utility value of the offer Bob would make is higher than the one received, Bob does not automatically accept the received offer. Bob needs to take into consideration whether his offer will be accepted or rejected by the opponent. So Bob checks the difference between Alice's utility value of the offer he made and her utility value of the offer she made. The difference in this case equals 3 (as Alice's utility value for Bob's offer equals 6, while the utility value of the offer she made equals 9). If we assume that the threshold for the difference is 0.05 (like in our experiments) then this condition does not hold as well, and Bob needs to continue to check whether he should reject or accept the offer. Since the utility value of the offer received from Alice is higher than the reservation price (5), Bob now decides whether to accept or reject the offer based on the probability derived from the ranking value of the offer. Since the ranking value for Bob from the offer of seeing a movie on Friday is 0.5, Bob will decide according to this probability whether to accept or reject the offer.

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 required for use of our technique in bilateral negotiation, and for the assumption that offers proposed by the agent will also be considered to be accepted by the opponent.

4.1.3. QO: An alternative to the Nash bargaining solution

From Luce and Raiffa ([17], Chapter 6), we employ the definitions of a bargaining problem and the Nash bargaining solution. We denote by $B = \langle (u_a(\cdot), u_b(\cdot)), \vec{d} \rangle$ the bargaining problem with two utilities, u_a and u_b , and a disagreement point, denoted \vec{d} , which is the worst outcome for both agents (in our model, d is equivalent to opting out). The Nash bargaining solution (which 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 ([20], Chapter 15), as defined in the definitions below.

Definition 4.1 (Symmetric). A bargaining problem is symmetric if a bijection $\phi: O \rightarrow O$, also called a symmetry function, exists such that:

- (1) $\phi^{-1} = \phi$,
- (2) $\phi(\vec{d}) = \vec{d}$,
- (3) $\forall \vec{x}, \vec{y} \in O, u_a(\vec{x}) > u_b(\vec{y}) \Leftrightarrow u_b(\phi(\vec{x})) > u_b(\phi(\vec{y}))$.

For example, suppose that two friends want to split \$100 among themselves. Each friend needs to decide how to split the money, but they both receive nothing if they disagree. This problem is a symmetric one (consider the symmetric function given by $\phi(x, y) = (y, x)$).

Definition 4.2 (Efficiency). An outcome $\vec{x} \in O$ is efficient if there is no outcome $\vec{y} \in O$ with $u_j(\vec{y}) > u_j(\vec{x})$ for both $j = \{a, b\}$.

Let's look again at the two friends trying to split the \$100 between themselves. A solution which will leave some of the money undivided between the two is inefficient, in the sense that both friends are better off if the remaining money is split between them.

Definition 4.3 (Equivalence). Two bargaining problems $B = \langle (u_a(\cdot), u_b(\cdot)), \vec{d} \rangle$ and $B' = \langle (u'_a(\cdot), u'_b(\cdot)), \vec{d} \rangle$ are equivalent if there are $\gamma_j > 0$ and $\gamma_j \in \mathbb{R}^+, \delta_j \in \mathbb{R}$, for $j = \{a, b\}$ such that $u'_j = \gamma_j u_j + \delta_j$.

To understand the notion of the equivalence problem, assume that the first involves temperature that is measured in Fahrenheit. An equivalent problem would be the same problem with the transformation from Fahrenheit to Celsius ($\gamma_j = 5/9, \delta_j = -160/9$).

Definition 4.4 (Subset Problems). A bargaining problem $B' = \langle (u'_a(\cdot), u'_b(\cdot)), \vec{d} \rangle, u'_j: O' \rightarrow \mathbb{R}^2$ for $j = \{a, b\}$ is a subset of another bargaining problem $B = \langle (u_a(\cdot), u_b(\cdot)), \vec{d} \rangle, u_j: O \rightarrow \mathbb{R}^2$, denoted by $B' \subseteq B$, if $O' \subseteq O$.

For example, the problem of splitting the \$100 between two friends such that both friends receive an equal split is a subset of the problem in which both friends can receive any split of the money.

Given the above definitions we state that a bargaining (or a negotiation) solution should have the following properties:

Property 4.1 (Symmetry). A negotiation solution f satisfies symmetry if for all symmetric problems B with symmetry function ϕ , $f(B) = \phi(f(B))$.

The *symmetry* property above states that if both players have the same bargaining power (since it deals with symmetric negotiation problems), 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 which started the negotiation process. In our example of splitting the \$100 between the two friends, both friends have the same utility function. For the solution to be symmetric, both must receive an equal payoff, that is, an equal distribution of the money.

Property 4.2 (Efficient). A negotiation solution f satisfies efficiency if $f(B)$ is efficient for all B .

Efficiency states that two rational agents will not reach an agreement if its utility is lower for both of them than another possible agreement. This solution is said to be Pareto-optimal. In this case, each friend will not agree on any agreement other than splitting the money equally between themselves in a way that *each* receive exactly half of the money, as every other split will generate a lower payoff for either of them.

Property 4.3 (Invariance). A negotiation solution f satisfies invariance if for all equivalent problems $B = \langle (u_a(\cdot), u_b(\cdot)), \vec{d} \rangle$ and $B' = \langle (u'_a(\cdot), u'_b(\cdot)), \vec{d} \rangle$, $f(B) = f(B')$.

Invariance states that for all equivalent problems B and B' , that is $B' = (\gamma_a \cdot u_a(\cdot) + \delta_a, \gamma_b \cdot u_b(\cdot) + \delta_b)$, $\gamma_a, \gamma_b \in \mathbb{R}^+$, $\delta_a, \delta_b \in \mathbb{R}$, 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. For example, the solution for the problem of splitting the \$100 between the two friends is equivalent to the solution in the case of splitting £500 between themselves instead of dollars. Thus, this solution satisfies the invariance property.

Property 4.4 (Independence of irrelevant alternatives). A negotiation solution f satisfies independence of irrelevant alternatives if $f(B) = f(B')$ whenever $B' \subseteq B$ and $f(B) \subseteq B'$.

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. For example, as we stated above, the problem of splitting the \$100 between two friends in such that both friends receive an equal split is a subset of the problem in which both friends can receive any split of the money. If we assume that also *all* the money has to be distributed between the friends, then the solution of an equal split between the friends satisfies the independence of irrelevant alternatives, as the added alternatives of unequal split, are not part of the solution.

It was shown by Nash [18] that the only solution that satisfies all of these properties is the product maximizing the agents' utilities, described in Eq. (4).

$$\arg \max_{\vec{x} \in O} (u_a(\vec{x}) - u_a(\vec{d})) (u_b(\vec{x}) - u_b(\vec{d})). \quad (4)$$

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 (Eq. (3)), 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 4.5 (Independence of irrelevant alternative solutions). A negotiation solution f satisfies independence of irrelevant alternative 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 $f(B) \subseteq B'$.

In this case, assume that in the problem of splitting the \$100 between two friends any split of the money is legitimate. Also, the problem in which both friends receive an equal split is a subset of the problem in which both friends can receive any split of the money. In this case, the negotiation solution f is a set which consists of every equal split. This solution satisfies the independence of irrelevant alternative solutions, as the added alternatives of unequal split, are not part of the solution.

Proving that the agent's strategy for proposing offers conforms to these properties (Properties 4.1, 4.2, 4.3 and 4.5) is important since although the agent should maximize its own utility, it should also find agreements that would be acceptable to its opponent. The following claims and their proofs lay the theoretical foundation for this.

Theorem 4.1. *The QO function satisfies the properties of symmetry, efficiency and independence of irrelevant alternative solutions.*

The proof of the theorem can be found in Appendix A, Claims A.2, A.3 and A.5. In addition, we show that under some conditions QO also satisfies the property of invariance (see Claim A.4 in Appendix A).

We recall that the Nash bargaining solution should have four properties: symmetry (Property 4.1), efficient (Property 4.2), invariance (Property 4.3) and independent of irrelevant alternatives (Property 4.4). By proving the above theorem, we show that our QO function satisfies only three of the four properties of the Nash bargaining solution. Thus the question is, why not use the Nash solution stated in Eq. (4) instead of our proposed solution? The reasoning for not using the above Nash solution is presented in the next claim.

Claim 4.1. *Let the agreement given by the Nash solution be \vec{x} . If an agreement \vec{y} exists where $lu_a(\vec{y}) \cdot rank_a(\vec{y}) > lu_a(\vec{x}) \cdot rank_a(\vec{x})$ and $lu_a(\vec{y}) \cdot rank_a(\vec{y}) < [lu_a(\vec{y}) + lu_b(\vec{y})] \cdot rank_b(\vec{y})$ then QO's solution will be \vec{y} rather than \vec{x} .*

The proof of the claim is given in Appendix A, Claim A.6. However, to clarify this claim we return to Example 1. Table 1 shows the utility value for Bob and Alice for 4 different offers. Assuming our agent plays the role of Bob, we show that it will suggest to Alice to go to a basketball game on Friday. However, if the agent follows the Nash solution, it will suggest to see a movie on Friday, which is the product maximization of the agents' utilities ($6 \times 9 = 54$), while the QO solution has a product value of $8 \times 6 = 48$. Though, the Nash solution generates a utility value of 6 for Bob, the QO solution generates a value of 8.

We continue to investigate the effects of time on the offers our agent generates. The following definition defines the concept of *time constant discount rate* ([20], Chapter 7):

Definition 4.5 (*Time constant discount rate*). In the case of a time constant discount rate, every agent j has a fixed discount rate $0 < \delta_j < 1$, that is: $u_j(\vec{o}, t) = \delta_j^t u_j(\vec{o}, 0)$.

We show that if both agents have the same time constant discount rate δ , then QO will generate the same solution at each time unit. The proof can be found in Appendix A, Claim A.7.

In the next section we present the component responsible for the belief update regarding the opponent.

4.2. The Bayesian updating rule component

The Bayesian updating rule is based on Bayes' theorem and it provides a practical learning algorithm. Bayes' theorem is generally used for calculating conditional probabilities and basically states how to update or revise beliefs in light of new evidence a posteriori ([12], Chapter 2). The calculation is given in the following formula:

$$P(A|B) = \frac{P(B|A) \cdot P(A)}{P(B)} \tag{5}$$

where $P(A)$ and $P(B)$ are the prior probabilities of A and B , respectively, $P(A|B)$ is the conditional probability of A , given B , and $P(B|A)$ is the conditional probability of B , given A .

We assert that there is a set of different agent types. The bounded rational agent should be matched to one such type. In each time period, the agent consults the component in order to update its belief regarding the opponent's type.

Table 3
Example: Calculating Alice's believed type

		$\vec{o}_1 = \{M, S\}$	$\vec{o}_2 = \{M, F\}$	$\vec{o}_3 = \{B, S\}$	$\vec{o}_4 = \{B, F\}$
1	$u_a(\vec{o}_i), \text{type}^1$	10	9	4	6
2	$u_a(\vec{o}_i), \text{type}^2$	10	7	5	9
3	$lu_a(\vec{o}_i), \text{type}^1$	$10/29 = 0.34$	$9/29 = 0.31$	$4/29 = 0.14$	$6/29 = 0.21$
4	$lu_a(\vec{o}_i), \text{type}^2$	$10/31 = 0.32$	$7/31 = 0.23$	$5/31 = 0.16$	$9/31 = 0.29$

Recall that there are k possible types of agents. At time $t = 0$ the prior probability of each type is equal, that is, $P(\text{type}_{t=0}^i) = \frac{1}{k}, \forall i \in \mathbf{Types}$. Then, for each time period t we calculate the a 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. That is, the believed type is updated every time an offer is received or accepted, thus eventually it is based on the overall total history thus far. Then, this value is assigned to $P(\text{type}_t)$. Using the calculated probabilities, the agent selects the type whose probability is the highest and proposes an offer as if it were the opponent's type. Formally, at each time period $t \in \mathbf{Time}$ and for each type $\in \mathbf{Types}$ and $\vec{o}_t \in O$ (the offer at time period t) we compute:

$$P(\text{type}^i | \vec{o}_t) = \frac{P(\vec{o}_t | \text{type}^i)P(\text{type}^i)}{P(\vec{o}_t)} \tag{6}$$

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

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

$$BT(t) = \arg \max_{i \in \mathbf{Types}} P(\text{type}^i | \vec{o}_t), \quad \forall t \in \mathbf{Time} \tag{7}$$

We will extend Example 1 to demonstrate this. Let's assume that there are two types of possible utilities for Alice ($k = 2$). In the first, given in Table 1, Alice prefers movies over basketball. In the second type, however, let's assume that Alice prefers going to a movie on Friday, and if this is not possible going to a basketball game on Friday. Assume Table 3 reflects the two different types of possible utilities for Alice derived from these preferences.

Initially, a probability of 1/2 is assigned to both types. Let's assume that Alice suggests at time $t = 1$ to go to a basketball on Friday night. Based on this suggestion, our agent needs to update the type it believes Alice to be. Based on Eq. (6), the way $P(\vec{o}_t)$ is calculated and the Luce numbers, we need to update the probability of each type. For simplicity, we omit the time from the calculations given below:

$$P(\text{type}^1 | \vec{o}_4) = \frac{P(\vec{o}_4 | \text{type}^1)P(\text{type}^1)}{P(\vec{o}_4)} = \frac{0.21 \cdot 0.5}{0.21 \cdot 0.5 + 0.29 \cdot 0.5} = 0.42 \tag{8}$$

$$P(\text{type}^2 | \vec{o}_4) = \frac{P(\vec{o}_4 | \text{type}^2)P(\text{type}^2)}{P(\vec{o}_4)} = \frac{0.29 \cdot 0.5}{0.21 \cdot 0.5 + 0.29 \cdot 0.5} = 0.58 \tag{9}$$

Now, based on Eq. (7) we deduce that the believed type for Alice is type^2 .

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 ([20], Chapter 12):

- (1) *Consistency.* Agent i 's belief should be consistent with its initial belief and with the possible actions of its opponents. Whenever possible, an agent should update its beliefs. If, after any history, all the actions of agent j 's in the given sequence of actions, regardless of its type, indicate that it has to take the same action, and this action is indeed taken by agent j , then agent i 's beliefs remain as they were before the action was taken. On the other hand, if an action is taken by j and this action can only be attributed to a single type of agent j , type l , then i believes with a probability of 1 that j 's type is indeed of type l . The agent uses the same reasoning about its opponent j 's beliefs based on the given sequence of actions and updates its model of j 's beliefs in a similar manner.

- (2) *Never dissuaded once convinced.* 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, that is, the probability of this type is 0, it is never dissuaded from its viewpoint.

The results of the experiments indeed show that in the majority of the experiments the agent believed that its opponent is of the correct type with the highest probability amongst all possible opponent's types.

5. Experiments

We developed a simulation environment which is adaptable such that any scenario and utility function, 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. The 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.

We conducted experiments on two distinct domains to test the efficiency of the proposed agent.² In the experiments, human subjects were matched both against the automated agent and against other human counterparts. These experiments show that the agent is capable of negotiating in various domains. That is, only the utility functions play a role, and not the scenario nor the domain. In addition, the experiments show the benefits achieved for both sides in the agreements (in terms of utility and time) when using an automated agent as compared to human only negotiations. In the following subsections we describe the two domains and the experimental methodology and we review the results.

5.1. Experimental domain

The experimental domains match the problem definitions described in Section 2. In the first domain one agent gains as time advances, while the other loses; the status quo value for one of the agents is much higher than for the opponent, and there is an option to reach partial agreements. In the second domain, both agents lose as time advances, and the status quo value for both players is quite similar. In both domains we modeled three possible agent types, and thus a set of six different utility functions was created for each domain. These sets describe the different types or approaches towards the negotiation process and the other party. For example, the different approaches can describe the importance each agent attaches to the effects of the agreement over time. One agent might have a long term orientation regarding the final agreement. This type of agent would favor agreements which are concerned more with future outcomes of the negotiations, than those focusing only on solving the present problem. On the other hand, another agent might have a short term orientation which focuses on solving only the burning issues under negotiation without dealing with future aspects that might arise from the negotiation or its solutions. Finally, there can also be agents with a compromise orientation. These agents try to find the middle grounds between the possible agreements.

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.

One of the domains was based on an international crisis, and the subjects had to play a role that was outside of their normal experience. On the other hand, the second domain was more related to the subjects' experience, so they could identify with it better. We describe the two domains in the following subsections. Detailed score functions for both domains can be found in Appendix B. A snapshot of one of the negotiation experiments in the second domain is given in Appendix C.

5.1.1. The World Health Organization's Framework Convention on Tobacco Control domain

In this scenario England and Zimbabwe negotiate in order to reach an agreement evolving from the World Health Organization's Framework Convention on Tobacco Control, the world's first public health treaty. The principal goal

² Preliminary results on the first domain were presented in [14].

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”.

The leaders of both countries are about to meet at a long scheduled summit. They must reach an agreement on the following issues:

- (1) *The total amount to be deposited into the Global Tobacco Fund to aid countries seeking to rid themselves of economic dependence on tobacco production.* This issue has an impact on the budget of England and on the effectiveness of short-range and long-range economic benefits for Zimbabwe. The possible values are (a) \$10 billion, (b) \$50 billion, (c) \$100 billion, or (d) no agreement. Thus, a total of 4 possible values are allowed for this issue.
- (2) *Impact on other aid programs.* This issue affects the net cost to England and the overall benefit for Zimbabwe. If other aid programs are reduced, the economic difficulties for Zimbabwe will increase. The possible values are (a) no reduction, (b) reduction equal to half of the Global Tobacco Fund, (c) reduction equal to the size of the Global Tobacco Fund, or (d) no agreement. Thus, a total of 4 possible values are allowed for this issue.
- (3) *Trade issues.* Both countries can use trade policy to extract concessions or provide incentives to the other party. They can use restrictive trade barriers such as tariffs (taxes on imports from the other country) or they can liberalize their trade policy by increasing imports from the other party. There are both benefits and costs involved in these policies: tariffs may increase revenue in the short run but lead to higher prices for consumers and possible retaliation by affected countries over the long run. Increasing imports can cause problems for domestic industries. But it can also lead to lower consumer costs and improved welfare. Thus, the possible values are divided between Zimbabwe's (a) reducing tariffs or (b) increasing tariffs on imports, and England's (a) reducing or (b) increasing imports. Both can also choose not to agree on this. Thus, a total of 9 possible values are allowed for this issue.
- (4) *Creation of a forum to explore comparable arrangements for other long-term health issues.* This issue relates to the precedent that may be set by the Global Tobacco Fund. If the fund is established, Zimbabwe will be highly motivated to apply the same approach to other global health agreements. This would be very costly to England. The possible values are (a) creation of a fund, (b) creation of a committee to discuss the creation of a fund, (c) creation of a committee to develop an agenda for future discussions, or (d) no agreement. Thus, a total of 4 possible values are allowed for this issue.

Consequently, a total of 576 possible agreements exist ($4 \times 4 \times 3 \times 3 \times 4 = 576$). While for the first two issues there are contradicting preferences for England and Zimbabwe, for the last two issues there are options which might be jointly preferred by both sides.

Each turn in the scenario is equivalent to a week of the summit, while the summit is limited to 14 weeks. If no agreement is reached within the specified time limit, the Framework Convention will be seen as an empty document, devoid of any political significance. This will be a blow to England, which has invested political capital to reach an agreement, in the hope of gaining support for other, perhaps more important, international agreements in the future. It will also, however, save England money in the near term. For Zimbabwe, failure to reach an agreement will create a major financial hardship and deprive it of a precedent that can be used for future negotiations. Consequently, England is better able to accept a failure than is Zimbabwe. This outcome is modeled for both agents as the status quo outcome.

Opting out of the negotiation is also an option. Opting out by England means trade sanctions imposed by England on Zimbabwe (including a ban on the import of tobacco from Zimbabwe), while if Zimbabwe opts out then it will boycott all British imports. However, if England opts out it also saves the funds that would have been spent on the Tobacco Fund, and if Zimbabwe opts out it loses the opportunity for financial gain and for assistance in reducing the health problems that arise from tobacco use. Consequently, England will likely be more willing to opt out if the negotiations are not going its way, and Zimbabwe will be more willing to continue negotiations until agreement is reached.

Time also has an impact on the negotiations. Creation of the fund is more urgent for Zimbabwe than for England. Consequently, Zimbabwe has an incentive to reach an agreement earlier rather than later; thus as time advances Zimbabwe's utility reduces. On the other hand, England gains as time advances, as it postpones the time at which it must transfer money to the fund.

Taking into account the different types of agents, we can say, for example, that an agent representing Zimbabwe with a short term orientation, will focus on a short term redistribution of resources, insist on the largest possible

current assistance and help with long-term health problems, as well as trade concessions. On the other hand, an agent representing England with the same short term orientation, for example, will aim to minimize current cost, limit impact on trade, and maintain its economic and political position in the near term.

5.1.2. The Job Candidate domain

In this scenario, a negotiation takes place after a successful job interview between an employer and a job candidate. In the negotiation both the employer and the job candidate wish to formalize the hiring terms and conditions of the applicant. In contrast to the England–Zimbabwe scenario, some issues *must* be agreed upon to achieve even a partial agreement. Below are the issues under negotiation:

- (1) *Salary*. This issue dictates the total net salary the applicant will receive per month. The possible values are (a) \$7000, (b) \$12,000, or (c) \$20,000. Thus, a total of 3 possible values are allowed for this issue.
- (2) *Job description*. This issue describes the job description and responsibilities given to the job applicant. The job description has an effect on the advancement of the candidate in his/her work place and his/her prestige. The possible values are (a) QA, (b) programmer, (c) team manager, or (d) project manager. Thus, a total of 4 possible values are allowed for this issue.
- (3) *Social benefits*. The social benefits are an addition to the salary and thus impose an extra expense on the employer, yet they can be viewed as an incentive for the applicant. The social benefits are divided into two categories: company car and the percentage of the salary allocated, by the employer, to the candidate's pension funds. The possible values for a company car are (a) providing a leased company car, (b) no leased car, or (c) no agreement. The possible value for the percentage of the salary deposited in pension funds are (a) 0%, (b) 10%, (c) 20%, or (d) no agreement. Thus, a total of 12 possible values ($3 \times 4 = 12$) are allowed for this issue.
- (4) *Promotion possibilities*. This issue describes the commitment by the employer regarding the fast track for promotion for the job candidate. The possible values are (a) fast promotion track (2 years), (b) slow promotion track (4 years), or (c) no agreement. Thus, a total of 3 possible values are allowed for this issue.
- (5) *Working hours*. This issue describes the number of working hours required by the employee per day (not including over-time). This is an integral part of the contract. The possible values are (a) 8 hours, (b) 9 hours, or (c) 10 hours. Thus, a total of 3 possible values are allowed for this issue.

In this scenario, a total of 1296 possible agreements exist ($3 \times 4 \times 12 \times 3 \times 3 = 1296$).

Each turn in the scenario equates to two minutes of the negotiation, and the negotiation is limited to 28 minutes. If the sides do not reach an agreement by the end of the allocated time, the job interview ends with the candidate being hired with a standard contract, which cannot be renegotiated during the first year. This outcome is modeled for both agents as the status quo outcome.

Each side can also opt-out of the negotiation if it feels that the prospects of reaching an agreement with the opponent are slim and it is impossible to negotiate anymore. Opting out by the employer entails the postponement of the project for which the candidate was interviewing, with the possible prospect of its cancellation and a considerable amount of expenses.

Opting-out by the job candidate will make it very difficult for him to find another job, as the employer will spread his/her negative impression of the candidate to other CEOs of large companies.

Time also has an impact on the negotiation. As time advances the candidate's utility decreases, as the employer's good impression has of the job candidate decreases. The employer's utility also decreases as the candidate becomes less motivated to work for the company.

5.2. Experimental methodology

We evaluated the performance of the agent against human subjects, all of whom were computer science undergraduates at Bar-Ilan University in Israel. The experiment involved 88 simulations with human subjects, divided into 44 pairs, such that 44 simulations were run for each domain. Each simulation was divided into two parts: (i) negotiating against another human subject, and (ii) negotiating against the automated agent. While the subjects knew that they will negotiate against both an automated agent and against another human, they did not know in advance against whom they played. Also, in order not to bias the results as a consequence of the subjects' familiarity 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 reaching a full agreement, opting out, or reaching the deadline without an agreement. Prior to the experiments, the subjects were given oral instructions regarding the experiment and the domain. The subjects were instructed to play based on their score functions and to achieve the best possible agreement for them.

5.3. Experimental results

The main goal of the experiments was to verify that the automated agent is capable of achieving better agreements than a human playing the same role, and to facilitate an earlier end to the negotiation as compared to negotiations without the agent. A secondary goal was to check whether indeed the agent facilitated an increase in the social welfare of the outcome, that is, improved the utility scores for both parties, as compared to negotiations without an automated agent. When analyzing the results we use three types of statistical tests:

- *t-test*: A statistical hypothesis test in which the test statistic has a *t-distribution* if the null hypothesis is true. This test requires a normal distribution of the measurements ([2], Chapter 3). Thus, it is used in our analysis for comparing utility values, which have continuous values.
- *Wilcoxon signed-rank test*: A non-parametric alternative to the paired *t-test* for the case of two related samples or repeated measurements on a single sample. This test does not require any assumptions regarding the distribution of the measurements ([22], Chapter 5). This test is used in our analysis for comparing discrete samples.
- *Fisher's Exact test*: Fisher's exact test is a statistical significance test used in the analysis of categorical data where sample sizes are small. This test is used to examine the significance of the association between two variables in a 2×2 contingency table ([8], Chapter 14.4). We use this test in order to show the correlation between the type of agreements reached (full agreement or partial) and the type of negotiators who reach them (two humans or a human and an automated agent).

As we mentioned earlier, we experimented in two distinct domains. Tables 4 and 5 summarize the average utility values of all the negotiations, the average ranking of the agreements reached, and the average of the sums of utility values and ranking of the agreements in all the experiments in the England–Zimbabwe domain and the Job Candidate domain, respectively. H_{Zim} , H_{Eng} , H_{Can} and H_{Emp} denote the utility value gained by people playing the role of Zim-

Table 4
Utility values, ranking values, sums of utility values and sums of ranking values of final negotiations in the England–Zimbabwe domain

Parameter	Avg	Stdev
Q_{Eng} vs. H_{Zim}	565.1	283.30
$rank(Q_{Eng})$ vs. $rank(H_{Zim})$	0.82	0.17
H_{Eng} vs. H_{Zim}	331.8	210.6
$rank(H_{Eng})$ vs. $rank(H_{Zim})$	0.58	0.19
Q_{Zim} vs. H_{Eng}	18.45	223.1
$rank(Q_{Zim})$ vs. $rank(H_{Eng})$	0.64	0.13
H_{Zim} vs. H_{Eng}	−92.6	247.90
$rank(H_{Zim})$ vs. $rank(H_{Eng})$	0.60	0.15
H_{Zim} vs. Q_{Eng}	−322.55	265.94
$rank(H_{Zim})$ vs. $rank(Q_{Eng})$	0.41	0.16
H_{Eng} vs. Q_{Zim}	311.50	204.79
$rank(H_{Eng})$ vs. $rank(Q_{Zim})$	0.57	0.18
Sum- H_{Eng} vs. Q_{Zim}	330	222.8
Sum- $rank(H_{Eng})$ vs. $rank(Q_{Zim})$	1.21	0.07
Sum- H_{Zim} vs. Q_{Eng}	242.5	409.4
Sum- $rank(H_{Zim})$ vs. $rank(Q_{Eng})$	1.25	0.03
Sum- H_{Eng} vs. H_{Zim}	239.2	298.8
Sum- $rank(H_{Eng})$ vs. $rank(H_{Zim})$	1.17	0.07

Table 5
Utility values, ranking values, sums of utility values and sums of ranking values of final negotiations in the Job Candidate domain

Parameter	Avg	Stdev
Q_{Can} vs. H_{Emp}	409	93.95
$rank(Q_{Can})$ vs. $rank(H_{Emp})$	0.75	0.19
H_{Can} vs. H_{Emp}	309.7	140.2
$rank(H_{Can})$ vs. $rank(H_{Emp})$	0.56	0.29
Q_{Emp} vs. H_{Can}	437.3	121.7
$rank(Q_{Emp})$ vs. $rank(H_{Can})$	0.77	0.19
H_{Emp} vs. H_{Can}	410.6	114.0
$rank(H_{Emp})$ vs. $rank(H_{Can})$	0.75	0.20
H_{Can} vs. Q_{Emp}	342.45	114.40
$rank(H_{Can})$ vs. $rank(Q_{Emp})$	0.58	0.24
H_{Emp} vs. Q_{Can}	448.82	82.41
$rank(H_{Emp})$ vs. $rank(Q_{Can})$	0.74	0.21
Sum- H_{Emp} vs. Q_{Can}	852.8	132.2
Sum- $rank(H_{Emp})$ vs. $rank(Q_{Can})$	1.49	0.23
Sum- H_{Can} vs. Q_{Emp}	779.7	199.0
Sum- $rank(H_{Can})$ vs. $rank(Q_{Emp})$	1.35	0.24
Sum- H_{Emp} vs. H_{Can}	720.3	212.5
Sum- $rank(H_{Emp})$ vs. $rank(H_{Can})$	1.30	0.27

babwe or England and the role of the job candidate or the employer, respectively, and Q_{Zim} , Q_{Eng} , Q_{Can} and Q_{Emp} denote the utility value gained by the QO agent playing either role in either domain.

The utility values range from -575 to 895 for the England role and from -680 to 830 for the Zimbabwe role, and in the Job Candidate domain from 170 to 620 for the employer role and from 60 to 635 for the job candidate role. The status quo value in the beginning of the negotiation was 150 for England and -610 for Zimbabwe, and in the second domain it was 240 for the employer and -160 for the job candidate. England had a fixed gain of 12 points per time period, while Zimbabwe had a fixed loss of -16 points. In the Job Candidate domain both players had a fixed loss per time period—the employer of -6 points and the job candidate of -8 points per period.

In both domains similar results were achieved. Thus, in this section we elaborate mainly on the results of the first domain. Later, we discuss the results in both domains.

5.3.1. Results of negotiations against people

First, we examine the final utility values of all the negotiations for each player, and the sums of the final utility values. When the automated agent played the role of England the average utility value achieved by the automated agent was 565.1 , while the average for the human playing the role of England was 331.8 . The results show that our agent achieves significantly higher utility values as opposed to a human agent playing the same role (using the 2-sample t -test: $t(22) = 3.10$, $p < 0.004$). (This was also the case when the automated agent played the role of the job candidate in the second domain (using the 2-sample t -test: $t(22) = 2.76$, $p < 0.008$.) On the other hand, when the agent played the role of Zimbabwe, there was no significant difference between the utility values of the agent and the human player, though the average utility value for the automated agent was higher (18.45) than that of the humans (-92.6). 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 the Zimbabwe side, which has a negative time cost (when QO played Zimbabwe the average end turn of the negotiation was 5 , while when the humans played Zimbabwe the average end turn was 7). Thus, accepting agreements sooner rather than later allows the agent to gain higher utility values than the human playing the same side. On the other hand, when the agent played the role of England, the average end turn for the negotiation was 7 and the same average was achieved when the humans played the role of England.

The above results are also supported by the results received from ranking the agreements. When the automated agent played the role of Zimbabwe, the average ranking it achieved was similar to the ranking the human players attained playing the same role (0.64 and 0.60). On the other hand, when the automated agent played the role of

England it achieved significantly higher ranking values than the human playing the same role, with an average of 0.82 as compared to only 0.58 (using the 2-sample Wilcoxon test, $p < 0.002$).

Comparing the sum of utility values of both negotiators, based on the role the agent played, we show that this sum is higher when the agent is involved in the negotiations. When the automated agent played the role of Zimbabwe, the sum of utility values was 330 as opposed to only 239.2 when two humans were involved. When the automated agent played the role of England, the sum of utility values was 242.5, which is only marginally higher than the score of 239.2 reached by the human subjects. (In the second domain the sum was also higher. Furthermore, when the automated agent played the role of the job candidate the sum was even significantly higher, using the 2-sample *t*-test: $t(22) = 2.48$, $p < 0.002$, when compared to negotiations in which no automated agent was involved.) When comparing the sum of the rankings, we note that when the automated agent was involved the sum of rankings was higher than when only humans were involved (an average of 1.21 and 1.25 when the automated agent played the role of Zimbabwe and England, respectively, and an average of 1.17 when the human players played against each other). However, this is only significant when the automated agent played the role of England (using the 2-sample Wilcoxon test, $p < 0.001$).

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 64% of the negotiations involving only people ended with a full agreement, more than 72% of the negotiations involving the automated agent ended with a full agreement (and in the Job Candidate domain, respectively 72% and 86%). Using the *Fisher's Exact test* we determined that a correlation exists between the kind of opponent agent (be it an automated agent or a 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 an automated agent ($p < 0.006$ for both domains).

In the next section we discuss the results of the experiments in both domains, when negotiating against people.

5.3.2. Discussion: Results against people

The results of the experiments, described above, show that the automated agent achieved higher utility values than the human counterpart. This can be explained by the nature of our agent both in reference to accepting offers and generating offers. Using the decision making mechanism we allow the agent to propose agreements that are good for it, but also reasonable for its opponent. In addition, the automated agent makes straightforward calculations. It evaluates the offer based on its attributes, and not based on its *content*. In addition, it also places more weight on the fact that it loses or gains as time advances. This is not the case, however, when analyzing the logs of the people. It seems that people put more weight on the content of the offer than on its value. This was more evident in the Job Candidate domain with which the human subjects could more easily identify.

Yet, this does not explain why, in both domains, these results are significant only for one of the sides. In the England–Zimbabwe domain, the results are significant when the agent played the role of England, while in the Job Candidate domain these results are significant when it played the role of the job candidate. It is interesting to note that our results, which show that the automated agents play significantly better when playing one of the sides, are not unique. Kraus et al. [11] also experimented with an automated agent playing against humans. While they experiment with a single-issue negotiation in one domain only (i.e. a fishing dispute domain—which is different from ours) they also showed that their agent design, which has a different design and logic than the one implemented by our agent, played significantly better only when playing one of the sides.

In the original version of the agent [14] we believed this to be attributed to the fact that the agent is more eager to accept agreements than people. To this end, we updated the agent and made it less eager (and more conservative) when it comes to accepting agreements. While in the original version of the agent we did not impose any restrictions on accepting agreements (agreements were accepted purely on probability based on the ranking of the offer), in the current version, acceptance by the agent depends on whether the value of the offer is greater than the agent's reservation price (r). While still many negotiations ended by the agent accepting the offer (and not by the agent proposing the winning offer), it allowed us to improve the scores of the agent. However, the fact still remained that these results are only significant for one of the roles in each domain.

Another possible explanation for this phenomenon can be found by examining the logs of the negotiations and the values of the agreements. In both domains we can see that the British side and the job candidate sides are the more dominant sides and have more leverage than the other side. For example, for England the fact that it gains as time

advances could place more pressure on the other side to accept agreements. For the job candidate side, a psychological interpretation could serve as an explanation. It seems that the job candidate's side has less to lose in the negotiation. While both the employer and the job candidate lose as time passes, the status quo agreement ensures the hiring of the candidate.

5.3.3. Results of an automated agent playing against another automated agent

In this set of experiments, we matched our automated agent against another automated agent. We conducted two sets of experiments. In both experiments the agents negotiated in both domains—the England–Zimbabwe domain and the Job candidate domain. In the first we matched our automated agent against itself. In the second set of experiments we matched it against another automated agent which followed a Bayesian equilibrium approach. In a Bayesian game, agents face uncertainties about the characteristics (types) of other agents. This imperfect information can be modeled by letting *Nature* select the agents' types. Agents have initial beliefs about the type of each agent and can update their beliefs according to Bayes' rule. Since the agents only know their own types, they must seek to maximize their expected payoff, given their beliefs about the other players. Note that our domains and experimental settings can be viewed as Bayesian games. The Bayesian Nash equilibrium is then defined as a strategy profile and beliefs specified for each agent about the types of the others agents that maximizes the expected payoff for each agent given their beliefs about the other agents' types and given the strategies played by the other agents ([20], pp. 24–29). Recall that there are k possible types of agents. Both our automated agent and the agent which followed a Bayesian equilibrium approach assume that the initial prior probability of each type is equal, that is, $P(\text{type}_{i=0}^i) = \frac{1}{k}, \forall i \in \mathbf{Types}$.

In the first set of experiments, when the automated agent was matched against itself, most of the agreements were reached by the earlier turns (by the third round in the England–Zimbabwe domain and by the second round in the Job Candidate domain). The average utility values for the *QO* agent playing England was 325.13 and for the *QO* agent playing Zimbabwe 79.93 and 499.58 and 423.06 when playing the role of the employer and job candidate, respectively. When looking at the utility values gained by the automated agent itself in both domains, the automated agent's results are higher than the results obtained by humans when playing the same role against either a human or an automated agent (the results are also significant when the automated agent played the role of Zimbabwe with $p < 0.004$ and both the employer and job candidate with $p < 0.002$ for both roles).

In addition, the average sum of utility values, 405.07 in the England–Zimbabwe domain and 922.65 in the Job Candidate domain, is also higher (significantly higher in the Job Candidate domain with $p < 0.002$ when it played the role of the employer and $p < 0.02$ when it played the role of the job candidate) than the sum obtained when either the automated agent played against people (330 when it played the role of Zimbabwe and 242.5 when it played the role of England; 779.7 when it played the role of the employer and 852.8 when it played the role of job candidate) and it was significantly higher than negotiations in which only people were involved (an average of 239.2 and $p < 0.017$ in the England–Zimbabwe domain and an average of 720.3 and $p < 0.001$ in the Job Candidate domain).

We can see that in both domains when the automated agent is matched against itself, it reaches better agreements for both sides. This can be attributed to the decision making component of the agent, which, as we described above, allows the agent both to generate agreements that are good for it, but also reasonable for the other side, and also to accept such agreements when they are proposed by their rival.

In the next set of experiments, we matched our automated agent against an automated agent that followed the Bayesian equilibrium strategy.

In the first domain, when the equilibrium agent played the role of Zimbabwe and the *QO* agent played the role of England, most of the negotiations ended by the early time periods, while in the second domain, for both roles the negotiation ended early (third time period in the England–Zimbabwe domain and second time period in the Job Candidate domain). The average final utility values of the negotiations were 398.38 for the *QO* agent playing the role of England and 61.38 for the equilibrium agent. In the Job candidate domains the values were 488.28 for the *QO* agent playing the role of the job candidate and 426.89 for the equilibrium agent and 459.8 for the *QO* agent playing the role of the employer and 488.4 for the equilibrium agent. In all these cases, the final utility values for the *QO* agent were higher than the average utility values achieved by the humans playing either against our automated agent or against themselves. On the other hand, when the equilibrium agent played the role of England, all of the negotiations lasted until the last time period and eventually ended with a status quo agreement, giving England a very high score of 981 and Zimbabwe a very low score of –548. The differences in the results between the two domains is that in the England–Zimbabwe domain, an agent playing Zimbabwe loses as time advances, so the equilibrium agent playing the

role of Zimbabwe is highly motivated to propose an attractive offer to the opponent to facilitate the termination of the negotiation sooner rather than later. In the job-candidate domain, however, both sides lose as time advances, and thus when the equilibrium agent played either side the negotiation ended quickly and did not drag on until the last turn.

Though we did not run simulations of the equilibrium agent against human agents, the utility values of the opponent from the offers suggested by the equilibrium agent are much lower than the final utility values of the human negotiations. By also analyzing the simulation process of the human negotiations, we can deduce that without incorporating any heuristics into the equilibrium agent, the human players would not have accepted the offers proposed by it. Thus, when the equilibrium agent would play the role of England the negotiation might be dragged out until the last turn with the implementation of a status quo agreement. However, perhaps the human playing the role of Zimbabwe would have given up and preferred opting-out, resulting in an outcome which would have been worse (utility-wise) for the equilibrium agent than the status quo outcome.

6. Conclusions

This paper presents an automated agent design for bilateral negotiation with bounded rational agents where there is incomplete information regarding the opponent's utility preferences. The results show that the agent is indeed capable of negotiating successfully with human counterparts and reaching efficient agreements. In addition, the results demonstrate that the agent plays at least as well as, and in the case of one of the two roles, achieved significantly higher utility values, than the human player. By running the experiments on two distinct domains we have shown that it is quite straightforward to adapt the simulation environment and the agent to any given scenario.

We have developed an automated negotiation environment. However, we do not intend to replace humans in negotiation, but rather to use the model as an efficient decision support tool or as a training tool for negotiations with people. Thus, this model can be used to support training in real life negotiations, such as: e-commerce, and it can also be used as the main tool in conventional lectures or online courses, aimed at turning the trainee into a better negotiator.

We have shown the importance of designing an automated negotiator that can negotiate efficiently with humans and we have shown that indeed it is possible to design such a negotiator. We believe that the results of our research can be particularly useful for constructing agents in open environments where uncertainty prevails. By pursuing non-classical methods of decision making it could be possible to achieve greater flexibility and effective outcomes. As we have shown, this can also be done without constraining the model to the domain. Thus these agents could be extremely useful in e-commerce environments and e-negotiations.

Most negotiation tools today are domain-dependent and focus on a single negotiation issue (e.g., see [23]). These tools do not provide an efficient training and learning experience for the trainee. Instead of providing the trainee with a wide range of test cases, they constrain him/her to a predefined scenario, which is only a fragment of the variety of scenarios he/she might encounter in the real world. We have demonstrated that our automated negotiation environment is adaptable such that any scenario and utility function, expressed as a single issue or multi-issue attributes, can be used, with no additional changes in the configuration of the interface of the simulations or the automated agents. The automated agents can play either role in the negotiation. In addition, our environment embodies an automated agent that plays against the trainee. This allows the trainee to use it *anytime* in order to test his/her capabilities and note improvements.

Although much time was spent on designing the mechanism for generating an offer, the results show that most of the agreements reached were offered by the human counterpart. This, indeed, allowed for more agreements to be reached when the automated agent was involved (as compared to negotiations in which only humans were involved). Nonetheless, a careful investigation should be made to examine how the offer generation mechanism can be improved. Another direction for future would be to improve this mechanism in order to allow the agent to make more than one offer per turn and to add more 'personality' to the agent, by allowing it to interact more with the opponent and adapt its approach based on this interaction and the pressure of time.

Appendix A. Theoretical proofs

To prove Theorem 4.1 we will prove the following claims. Combining those proofs depicts the correctness of our theorem. In the following analysis, the terms a and b are used synonymously as the types of agents a and b , respectively. In the analysis we will refer to the automated agent as the agent of type a and its opponent as the agent

of type b . We also assume that there is a unique negotiation solution. Recall also that in our model, the disagreement point d is equivalent to opting out.

Claim A.1. *QO always generates an agreement which is not worse than the disagreement point,³ \vec{d} , for both agents.*

Proof. Since \vec{d} is the worst outcome for both agents, then:

$$\forall \vec{x} \in O \quad \text{rank}_j(\vec{x}) \geq \text{rank}_j(\vec{d}) \tag{A.1}$$

Assume, by contradiction, that *QO* generated \vec{d} as the offer. Following the *QO* function (Eq. (3)), we obtain:

$$\vec{d} = \arg \min\{\alpha_d, \beta_d\} \quad \text{where } \alpha_d = \text{rank}_a(\vec{d}) \cdot lu_a(\vec{d}) \text{ and } \beta_d = [lu_a(\vec{d}) + lu_b(\vec{d})] \cdot \text{rank}_b(\vec{d}) \tag{A.2}$$

Also assume that there is another agreement $\vec{x} \in O$, $\vec{x} \neq \vec{d}$ such that:

$$\begin{aligned} \alpha_x &= \text{rank}_a(\vec{x}) \cdot lu_a(\vec{x}) \\ \beta_x &= [lu_a(\vec{x}) + lu_b(\vec{x})] \cdot \text{rank}_b(\vec{x}) \end{aligned} \tag{A.3}$$

Thus, we can distinguish between two options. In the first, $\alpha_d < \beta_d$. In this case, following our assumption, α_d must now be selected as the maximum of all the minima. Now, if the agreement \vec{x} satisfies $\alpha_x < \beta_x$, then α_x will be chosen as the minima and must be compared to α_d when choosing the maximum. Since \vec{d} is the worst outcome, and following Eq. (A.1), $\alpha_x > \alpha_d$, which contradicts our assumption that \vec{d} was chosen by *QO*. Thus, we find that $\alpha_x > \beta_x$. In this case, β_x is compared to α_d . Following Eq. (A.1) and Property 3.1 we obtain for every $\vec{x} \neq \vec{d}$ also $lu_j(\vec{x}) \geq lu_j(\vec{d})$. Thus, $\beta_x > \alpha_d$, which means that \vec{d} cannot have been chosen by *QO*.

In the second option we have $\alpha_d > \beta_d$. Following similar considerations as the ones stated in the first case we also obtain that β_d cannot be chosen as the maximum, and thus \vec{d} cannot have been chosen by *QO*. \square

Claim A.2. *QO satisfies symmetry (Property 4.1).*

Proof. Let $B = \langle (u_a(\cdot), u_b(\cdot)), \vec{d} \rangle$ be a symmetric negotiation problem with a symmetric function ϕ . Let \vec{x}^* be the negotiation solution generated by *QO*. We need to show that $\phi(\vec{x}^*)$ is also a negotiation solution, i.e., $(u_a(\phi(\vec{x}^*)), u_b(\phi(\vec{x}^*))) = (u_a(\vec{x}^*), u_b(\vec{x}^*))$. From the definition of the bargaining problem and Claim A.1 we know that

$$u_a(\vec{x}^*) \geq u_a(\vec{d}) \tag{A.4}$$

$$u_b(\vec{x}^*) \geq u_b(\vec{d}) \tag{A.5}$$

Since ϕ is a symmetric function we find that (Definition 4.1)

$$u_a(\phi(\vec{x}^*)) \geq u_a(\phi(\vec{d})) = u_a(\vec{d}) \tag{A.6}$$

$$u_b(\phi(\vec{x}^*)) \geq u_b(\phi(\vec{d})) = u_b(\vec{d}) \tag{A.7}$$

Assume, by contradiction, that $\phi(\vec{x}^*)$ is not a negotiation solution. Then, there is $\vec{y} \in O$ such that

$$u_a(\vec{y}) \geq u_a(\phi(\vec{x}^*)) \tag{A.8}$$

$$u_b(\vec{y}) \geq u_b(\phi(\vec{x}^*)) \tag{A.9}$$

But since ϕ is a symmetric function we attain

$$\begin{aligned} u_a(\phi(\vec{y})) &\geq u_a(\phi \circ \phi(\vec{x}^*)) = u_a(\vec{x}^*) \\ u_b(\phi(\vec{y})) &\geq u_b(\phi \circ \phi(\vec{x}^*)) = u_b(\vec{x}^*) \end{aligned} \tag{A.10}$$

From Eq. (A.10) we reveal that there is another solution, $\phi(\vec{y})$, in the original negotiation problem which is not worse than \vec{x}^* , which is the negotiation solution *QO* generated. However, this contradicts the fact that \vec{x}^* is the unique negotiation solution in the original problem. \square

³ If the disagreement point is unique, then *QO* always generates an agreement which has a higher utility value than the disagreement point.

Claim A.3. *QO satisfies efficiency (Property 4.2).*

Proof. We will prove this by contradiction. Let \vec{x}^* be the solution. That is:

$$\vec{x}^* = \arg \min\{\alpha_x, \beta_x\} \quad \text{where } \alpha_x = \text{rank}_a(\vec{x}^*) \cdot lu_a(\vec{x}^*) \text{ and } \beta_x = [lu_a(\vec{x}^*) + lu_b(\vec{x}^*)] \cdot \text{rank}_b(\vec{x}^*) \quad (\text{A.11})$$

Assume that \vec{x}^* is not efficient, that is, there is $\vec{y} \in O$ such that $u_a(\vec{y}) > u_a(\vec{x}^*)$ and $u_b(\vec{y}) > u_b(\vec{x}^*)$. For \vec{y} to be chosen by *QO*:

$$\vec{y} = \arg \min\{\alpha_y, \beta_y\} \quad \text{where } \alpha_y = \text{rank}_a(\vec{y}) \cdot lu_a(\vec{y}) \text{ and } \beta_y = [lu_a(\vec{y}) + lu_b(\vec{y})] \cdot \text{rank}_b(\vec{y}) \quad (\text{A.12})$$

QO generates a set of minima from which it selects the maximum as \vec{x}^* . We will distinguish between four possible selections:

(1) $\alpha_x < \beta_x$.

Here there are two possible options:

(a) $\alpha_y < \beta_y$.

Since *QO* selected agreement \vec{x}^* we know that $\alpha_x > \alpha_y$. However, this contradicts the assumption that $u_a(\vec{y}) > u_a(\vec{x}^*)$ (since $\text{rank}_a(\vec{y}) > \text{rank}_a(\vec{x}^*)$ and $lu_a(\vec{y}) > lu_a(\vec{x}^*)$).

(b) $\alpha_y > \beta_y$.

That is, β_y is chosen as the minimum. Since *QO* selects α_x as the maximum element in the set we obtain $\beta_x > \alpha_x > \beta_y$. However, since $u_a(\vec{y}) > u_a(\vec{x}^*)$ and $u_b(\vec{y}) > u_b(\vec{x}^*)$ we reveal that $lu_a(\vec{y}) > lu_a(\vec{x}^*)$, $\text{rank}_a(\vec{y}) > \text{rank}_a(\vec{x}^*)$, $lu_b(\vec{y}) > lu_b(\vec{x}^*)$, $\text{rank}_b(\vec{y}) > \text{rank}_b(\vec{x}^*)$. However, this requires that $\beta_y > \beta_x$. Thus, the contradiction assumption is wrong.

(2) $\alpha_x > \beta_x$.

As in the previous case, there are two possible options:

(a) $\alpha_y < \beta_y$.

Since *QO* selected β_x then $\beta_x > \alpha_y$. By the assumption, $u_a(\vec{y}) > u_a(\vec{x}^*)$. Thus, we obtain $\beta_x > \alpha_y > \alpha_x$ which is in contradiction with our base case (Case A.3). Thus, the contradiction assumption is wrong.

(b) $\alpha_y > \beta_y$.

Since *QO* selects β_x as the maximum element in the set of all minima we find that $\beta_x > \beta_y$. However, since $u_a(\vec{y}) > u_a(\vec{x}^*)$ and $u_b(\vec{y}) > u_b(\vec{x}^*)$ we reveal that $lu_a(\vec{y}) > lu_a(\vec{x}^*)$, $\text{rank}_a(\vec{y}) > \text{rank}_a(\vec{x}^*)$, $lu_b(\vec{y}) > lu_b(\vec{x}^*)$, $\text{rank}_b(\vec{y}) > \text{rank}_b(\vec{x}^*)$. However, this requires that $\beta_y > \beta_x$. Thus, the contradiction assumption is wrong. \square

Claim A.4. *QO satisfies invariance (Property 4.3) under the following conditions:*

$$(1) u_i(\vec{o}) < u_j(\vec{o}) \Leftrightarrow u'_i(\vec{o}) < u'_j(\vec{o});^4$$

$$(2) \forall \vec{o}, \vec{p} \in O \quad u_i(\vec{o}) > u_j(\vec{p}) \Rightarrow u'_i(\vec{o}) > u'_j(\vec{p}).$$

Proof. Let $B = \langle (u_a(\cdot), u_b(\cdot)), \vec{d} \rangle$ and $B' = \langle (u'_a(\cdot), u'_b(\cdot)), \vec{d} \rangle$ be equivalent problems. Let \vec{x}^* be the negotiation solution generated by *QO*. We need to show that $f(B) = f(B')$, that is, the same negotiation solution is obtained by $f(B')$. From Eq. (3), we observe that we perform a maximum over the set of minima which is generated from u_a and u_b . Condition 1 above guarantees us that if $u_a(\vec{x}) \leq u_b(\vec{x})$ then also $u'_a(\vec{x}) \leq u'_b(\vec{x})$. Thus, the set of minima remains identical using u_a, u_b and u'_a, u'_b .

Note that since we perform a linear transformation where $\alpha_i > 0$ and $\alpha_j \in \mathbb{R}^+, \beta_i \in \mathbb{R}$ then

$$u_i(\vec{x}) > u_i(\vec{y}) \Rightarrow u'_i(\vec{x}) > u'_i(\vec{y}). \quad (\text{A.13})$$

Condition 2 and Eq. (A.13) above guarantee that the preference relation among this set will also remain the same. Thus, we will attain the same solution \vec{x}^* . The value of this solution for B' can be obtained using the same linear transformation, that is, $QO_{B'}(t) = \alpha_a x^* + \beta_i$. \square

⁴ This condition is applicable in domains such as ours, in which there are contradictory preferences.

Claim A.5. *QO satisfies independence of irrelevant alternative solutions (Property 4.5).*

Proof. Let $B = \langle (u_a(\cdot), u_b(\cdot)), \vec{d} \rangle$ and $B' = \langle (u'_a(\cdot), u'_b(\cdot)), \vec{d} \rangle$ be negotiation problems, where $B' \subseteq B$. Let X be the set of all possible negotiation solutions generated by QO_B ($|X| = 1$ if there is only one maximum, and $|X| > 1$ if there are several maxima) such that $X \subseteq B'$. We need to show that X is also the set of all possible negotiation solutions generated by $QO_{B'}$. Let $\vec{x}^* \in X$ be an arbitrary solution. That is, \vec{x}^* is chosen as the maximum element of all the values in the minima set. Since the agreement associated with \vec{x}^* is also in B' and since B' contains only a subset of the agreements in B , \vec{x}^* must also be the maximum element in the minima set for B' . That is, \vec{x}^* is also chosen as the maximum. \square

Claim A.6. *Let the agreement given by the Nash solution be \vec{x} . If an agreement \vec{y} exists where $lu_a(\vec{y}) \cdot rank_a(\vec{y}) > lu_a(\vec{x}) \cdot rank_a(\vec{x})$ and $lu_a(\vec{y}) \cdot rank_a(\vec{y}) < [lu_a(\vec{y}) + lu_b(\vec{y})] \cdot rank_b(\vec{y})$ then QO's solution will be \vec{y} rather than \vec{x} .*

Proof. From the assumption we know the $lu_a(\vec{y}) \cdot rank_a(\vec{y})$ is chosen as the minimum. We distinguish between the two possible cases:

(1) $lu_a(\vec{x}) \cdot rank_a(\vec{x}) < [lu_a(\vec{x}) + lu_b(\vec{x})] \cdot rank_b(\vec{x})$.

In this case, $lu_a(\vec{x}) \cdot rank_a(\vec{x})$ is also chosen as the minimum and since $u_a(\vec{y}) > u_a(\vec{x})$, QO will prefer \vec{y} over \vec{x} .

(2) $lu_a(\vec{x}) \cdot rank_a(\vec{x}) > [lu_a(\vec{x}) + lu_b(\vec{x})] \cdot rank_b(\vec{x})$.

In this case, $[lu_a(\vec{x}) + lu_b(\vec{x})] \cdot rank_b(\vec{x})$ is chosen as the minimum but since $u_a(\vec{y}) > u_a(\vec{x})$, which requires $lu_a(\vec{y}) \cdot rank_a(\vec{y}) > lu_a(\vec{x}) \cdot rank_a(\vec{x})$, QO will prefer \vec{y} over \vec{x} . \square

Claim A.7 (Identical discount rate). *If both agents have the same time constant discount rate, QO will generate the same solution at each time unit.*

Proof. Let \vec{x}_0^* be the solution generated by QO at time $t = 0$. We need to show that $\forall t > 0, \vec{x}_t^* = \vec{x}_0^*$. Since both agents multiply their utilities by the same factor, $0 < \delta_j < 1$, the preference relations among their own utility and between their utilities remain the same. Thus, the probability of accepting an offer and the ranking of the offers also remain the same. As a result, the agreement generated by QO at time t will be the same as the one generated at the previous time unit $t - 1$, and thus will be the same as the one generated at time $t = 0$. \square

Appendix B. Score functions

Tables B.1 and B.2 present the score functions for both negotiators, in both domains. While the human subject is given his own score function at the beginning of the negotiation, he is also given three additional score functions which model the different possible types of his opponent.

Appendix C. Negotiation example

The following is a snapshot of a negotiation log between the automated agent and a human counterpart in the Job-Candidate domain. The automated agent played the role of the job candidate (denoted QO_{Can}) and the human played the role of the employer (denoted H_{Emp}). The following is the internal log of our program. The people themselves communicated and received messages via GUI windows. Examples of the main screen, an offer generation screen and a receiving offer screen are given in Figs. C.1, C.2 and C.3, respectively.

The log below depicts a successful negotiation which enabled the sides to reach an agreement after 5 turns. Alternating offers between the automated agent and the human can be observed until the offer is accepted by the automated agent in the third turn (offer #6). Yet, the negotiation is not over since this is a partial agreement. This agreement is enforced and the agents cannot withdraw from it, unless both sides agree to a new agreement (or one side decides to opt out). We can see that the sides still need to resolve the issue of the leased car. However, since the negotiation is not over, both sides are free to propose additional offers, not necessarily ones which refer to the unresolved issue. Indeed this is what the automated agent does. Eventually, the human player proposes to resolve the issue of the leased car (offer #12). The automated agent agrees and thus the negotiation terminates with a full agreement.

Table B.1

The England–Zimbabwe domain (i) short-term, (ii) long-term and (iii) compromise orientation score functions

Outcomes	Zimbabwe Outcome Weight/Importance			England Outcome Weight/Importance		
	(i)	(ii)	(iii)	(i)	(ii)	(iii)
Size of fund	50%	10%	20%	50%	10%	30%
\$100 Billion	9	5	6	−5	1	2
\$50 Billion	2	2	4	2	3	4
\$10 Billion	−5	−3	2	10	6	6
No agreement	−8	−6	−2	7	−1	−2
Impact on other aid	30%	10%	20%	30%	10%	30%
No reduction	8	6	3	−4	1	0
Reduction is equal to half of the fund size	0	0	0	4	2	3
Reduction is equal to the fund size	−3	−3	−2	10	3	5
No agreement	−5	−4	−4	−7	0	−2
Trade policy	10%	30%	30%	10%	30%	10%
Zimbabwe will reduce tariffs on imports	−6	−3	−4	3	4	5
Zimbabwe will increase tariffs on imports	3	6	4	−3	−6	−6
England will increase imports	7	8	10	−4	−8	−5
England will reduce imports	−8	−9	−8	4	6	4
No agreement	0	0	0	0	0	0
Forum on other health issues	10%	50%	30%	10%	50%	30%
Creation of fund	9	8	7	−8	7	4
Creation of committee to discuss creation of fund	3	5	5	2	4	7
Creation of committee to develop agenda	−5	−6	3	6	−2	1
No agreement	−6	−8	−3	1	−4	−2
Time effect	−16	−16	−16	12	12	12
Status quo	−610	−500	−210	150	−210	−180
Opting out	−530	−520	−240	−105	−240	−75

To: H_{Emp} From: QO_{Can}

Message Type: Offer (id = 1)

Turn: 1

Message: <Salary: 20,000 NIS, Job Description: Programmer, Leased Car: With leased car, Pension Fund: 20%, Promotion Possibilities: Fast promotion track, Working Hours: 8 hours>

To: QO_{Can} From: H_{Emp}

Message Type: Response

Turn: 1

Message: Offer (id = 1) refused

To: QO_{Can} From: H_{Emp}

Message Type: Offer (id = 2)

Turn: 1

Message: <Salary: 12,000 NIS, Job Description: Programmer, Leased Car: Without leased car, Pension Fund: 10%, Promotion Possibilities: Fast promotion track, Working Hours: 10 hours>

To: H_{Emp} From: QO_{Can}

Message Type: Response

Table B.2

The Job Candidate domain (i) short-term, (ii) long-term and (iii) compromise orientation score functions

Outcomes	Job Candidate Outcome Weight/Importance			Employer Outcome Weight/Importance		
	(i)	(ii)	(iii)	(i)	(ii)	(iii)
Salary	20%	30%	15%	20%	15%	10%
7000 NIS	3	2	3	8	7	7
12,000 NIS	6	6	5	6	6	6
20,000 NIS	8	9	6	3	3	4
Job description	15%	25%	20%	20%	30%	20%
QA	2	−2	2	4	2	3
Programmer	4	3	4	6	6	6
Team Manager	5	6	6	4	3	4
Project Manager	6	8	8	2	1	3
Leased car	20%	5%	10%	10%	10%	10%
Without leased car	−5	−5	−2	3	4	5
With leased car	5	5	2	−2	2	4
No agreement	0	0	0	0	0	0
Pension fund	10%	5%	10%	10%	10%	10%
0% pension fund	−2	−2	−2	3	6	6
10% pension fund	3	4	3	4	4	4
20% pension fund	5	6	5	3	3	3
No agreement	0	0	0	0	0	0
Promotion possibilities	5%	25%	35%	10%	20%	20%
Slow promotion track	4	1	−2	3	8	6
Fast promotion track	5	5	5	3	5	4
No agreement	0	0	0	0	0	0
Working hours	30%	10%	10%	30%	15%	30%
10 hours	3	3	4	8	8	9
9 hours	5	4	5	6	6	6
8 hours	7	5	6	3	4	3
Time effect	−8	−8	−8	−6	−6	−6
Status quo	160	135	70	240	306	306
Opting out	150	75	80	210	150	215

Turn: 1

Message: Offer (id = 2) refused

To: QO_{Can} From: H_{Emp}

Message Type: Offer (id = 3)

Turn: 2

Message: <Salary: 12,000 NIS, Job Description: Programmer, Leased Car: Not discussed, Pension Fund: 10%, Promotion Possibilities: Fast promotion track, Working Hours: 9 hours>

To: H_{Emp} From: QO_{Can}

Message Type: Offer (id = 4)

Turn: 2

Message: <Salary: 20,000 NIS, Job Description: Programmer, Leased Car: With leased car, Pension Fund: 20%, Promotion Possibilities: Fast promotion track, Working Hours: 9 hours>

To: H_{Emp} From: QO_{Can}

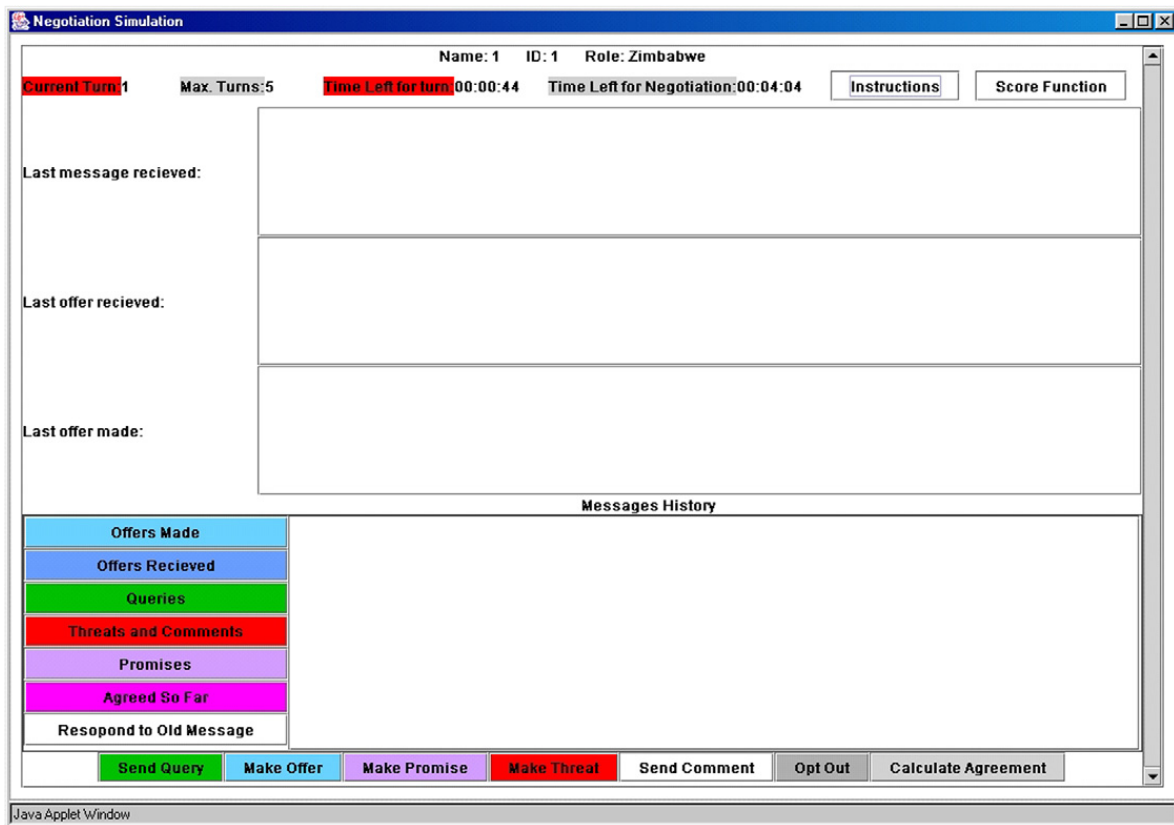


Fig. C.1. Main negotiation screen.

Message Type: Response

Turn: 2

Message: Offer (id = 3) refused

To: QO_{Can}

From: H_{Emp}

Message Type: Response

Turn: 2

Message: Offer (id = 4) refused

To: H_{Emp}

From: QO_{Can}

Message Type: Offer (id = 5)

Turn: 3

Message: <Salary: 12,000 NIS, Job Description: Team Manager, Leased Car: With leased car, Pension Fund: 20%, Promotion Possibilities: Fast promotion track, Working Hours: 9 hours>

To: QO_{Can}

From: H_{Emp}

Message Type: Response

Turn: 3

Message: Offer (id = 5) refused

To: QO_{Can}

From: H_{Emp}

Message Type: Offer (id = 6)

Fig. C.2. Generating offers screen.

Turn: 3

Message: <Salary: 12,000 NIS, Job Description: Programmer, Leased Car: Not discussed, Pension Fund: 20%, Promotion Possibilities: Slow promotion track, Working Hours: 9 hours>

To: H_{Emp}

From: QO_{Can}

Message Type: Response

Turn: 3

Message: Offer (id = 6) *agreed*

To: QO_{Can}

From: H_{Emp}

Message Type: Offer (id = 7)

Turn: 4

Message: <Salary: Not discussed, Job Description: Not discussed, Leased Car: Without leased car, Pension Fund: Not discussed, Promotion Possibilities: Not discussed, Working Hours: Not discussed >

To: H_{Emp}

From: QO_{Can}

Message Type: Offer (id = 8)

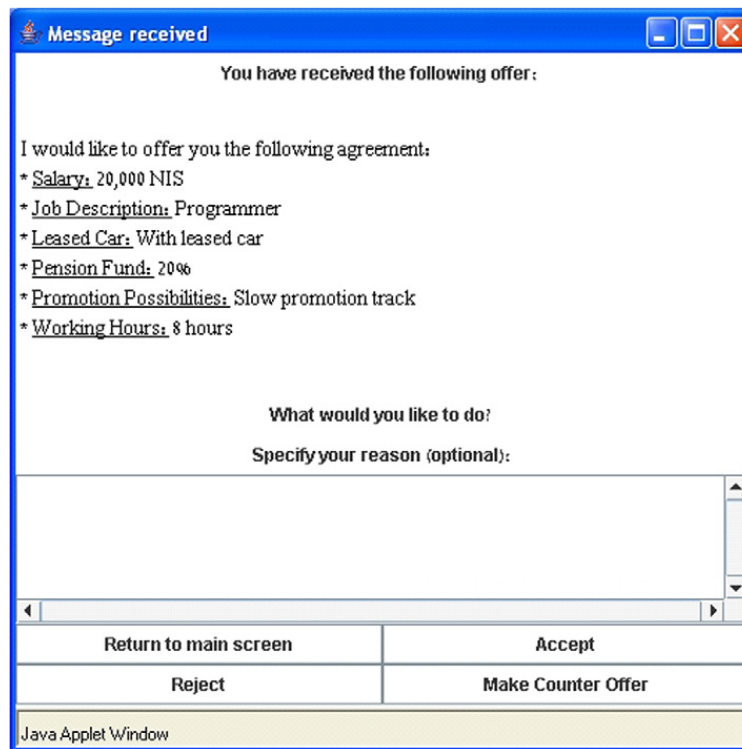


Fig. C.3. Receiving offers screen.

Turn: 4

Message: <Salary: 12,000 NIS, Job Description: Programmer, Leased Car: With leased car, Pension Fund: 20%, Promotion Possibilities: Fast promotion track, Working Hours: 9 hours>

To: QO_{Can}

From: $HEmp$

Message Type: Response

Turn: 4

Message: Offer (id = 8) refused

To: $HEmp$

From: QO_{Can}

Message Type: Response

Turn: 4

Message: Offer (id = 7) refused

To: QO_{Can}

From: $HEmp$

Message Type: Offer (id = 9)

Turn: 4

Message: <Salary: Not discussed, Job Description: Not discussed, Leased Car: Without leased car, Pension Fund: Not discussed, Promotion Possibilities: Fast promotion track, Working Hours: Not discussed>

To: $HEmp$

From: QO_{Can}

Message Type: Response

Turn: 5

Message: Offer (id = 9) refused

To: QO_{Can}

From: H_{Emp}

Message Type: Offer (id = 10)

Turn: 5

Message: <Salary: Not discussed, Job Description: Not discussed, Leased Car: Without leased car, Pension Fund: Not discussed, Promotion Possibilities: Fast promotion track, Working Hours: Not discussed>

To: H_{Emp}

From: QO_{Can}

Message Type: Response

Turn: 5

Message: Offer (id = 10) refused

To: H_{Emp}

From: QO_{Can}

Message Type: Offer (id = 11)

Turn: 5

Message: <Salary: 12,000 NIS, Job Description: Team Manager, Leased Car: With leased car, Pension Fund: 20%, Promotion Possibilities: Slow promotion track, Working Hours: 9 hours>

To: QO_{Can}

From: H_{Emp}

Message Type: Response

Turn: 5

Message: Offer (id = 11) refused

To: QO_{Can}

From: H_{Emp}

Message Type: Offer (id = 12)

Turn: 5

Message: <Salary: Not discussed, Job Description: Not discussed, Leased Car: With leased car, Pension Fund: Not discussed, Promotion Possibilities: Not discussed, Working Hours: Not discussed>

To: H_{Emp}

From: QO_{Can}

Message Type: Response

Turn: 5

Message: Offer (id = 12) *agreed*

End Negotiation: a full agreement was reached:

<Salary: 12,000 NIS, Job Description: Programmer, Leased Car: With leased car, Pension Fund: 20%, Promotion Possibilities: Slow promotion track Working Hours: 9 hours>

QO_{Can} Score: 468.0

H_{Emp} Score: 436.0

References

- [1] R.I. Brafman, M. Tennenholtz, An axiomatic treatment of three qualitative decision criteria, *Journal of the ACM* 47 (3) (2000) 452–482.
- [2] S.R. Brown, L.E. Melamed, *Experimental Design and Analysis*, Sage Publications, Inc., CA, 1990.
- [3] C.M. Capra, J.K. Goeree, R. Gomez, C.A. Holt, Anomalous behavior in a traveler's dilemma? *American Economic Review* 89 (3) (1999) 678–690 (June 1999).

- [4] D. Dubois, H. Prade, R. Sabbadin, Decision-theoretic foundations of qualitative possibility theory, *European Journal of Operational Research* 128 (2001) 459–478.
- [5] P. Faratin, C. Sierra, N.R. Jennings, Using similarity criteria to make issue trade-offs in automated negotiations, *AIJ* 142 (2) (2002) 205–237.
- [6] S. Fatima, M. Wooldridge, An agenda based framework for multi-issue negotiation, *AIJ* 152 (2004) 1–45.
- [7] S. Fatima, M. Wooldridge, N.R. Jennings, Bargaining with incomplete information, *AMAI* 44 (3) (2005) 207–232.
- [8] J.D. Gibbons, S. Chakraborty, *Nonparametric Statistical Inference*, Marcel Dekker, New York, 1992.
- [9] P.T. Hoppman, *The Negotiation Process and the Resolution of International Conflicts*, University of South Carolina Press, Columbia, SC, 1996.
- [10] R. Keeney, H. Raiffa, *Decisions with Multiple Objective: Preferences and Value Tradeoffs*, John Wiley & Sons, New York, 1976.
- [11] S. Kraus, P. Hoz-Weiss, J. Wilkenfeld, D.R. Andersen, A. Pate, Resolving crises through automated bilateral negotiations, *Artificial Intelligence* 172 (1) (2008) 1–18.
- [12] T. Leonard, J.S.J. Hsu, *Bayesian Methods—An Analysis for Statisticians and Interdisciplinary Researchers*, Cambridge University Press, Cambridge, 1999.
- [13] R.J. Lin, S.T. Chou, Bilateral multi-issue negotiations in a dynamic environment, in: *Proc. of the AAMAS Workshop on Agent Mediated Electronic Commerce (AMEC V)*, Melbourne, Australia, 2003.
- [14] R. Lin, S. Kraus, J. Wilkenfeld, J. Barry, An automated agent for bilateral negotiation with bounded rational agents with incomplete information, in: *Proc. of ECAI-06*, 2006, pp. 270–274.
- [15] R.D. Luce, *Individual Choice Behavior: A Theoretical Analysis*, John Wiley & Sons, New York, 1959.
- [16] R.D. Luce, The choice axiom after twenty years, *Journal of Mathematical Psychology* 15 (1977) 215–233.
- [17] R.D. Luce, H. Raiffa, *Games and Decisions—Introduction and Critical Survey*, John Wiley & Sons, New York, 1957.
- [18] J.F. Nash, The bargaining problem, *Econ.* 18 (1950) 155–162.
- [19] M. Oprea, An adaptive negotiation model for agent-based electronic commerce, *Studies in Informatics and Control* 11 (3) (2002) 271–279.
- [20] M.J. Osborne, A. Rubinstein, *A Course In Game Theory*, MIT Press, Cambridge, MA, 1994.
- [21] E. Rasmusen, *Games and Information: An Introduction to Game Theory*, third ed., Blackwell Publishers, February 2001.
- [22] S. Siegel, *Non-Parametric Statistics for the Behavioral Sciences*, McGraw–Hill, New York, 1956.
- [23] Stitt Feld Handy Group Online Negotiation Course, <http://www.negotiate.tv/>, 2007.
- [24] M. Tennenholtz, On stable social laws and qualitative equilibrium for risk-averse agents, in: *KR*, 1996, pp. 553–561.
- [25] J. Wilkenfeld, K. Young, D. Quinn, V. Asal, *Mediating International Crises*, Routledge, London, 2005.
- [26] D. Zeng, K. Sycara, Bayesian learning in negotiation, in: S. Sen (Ed.), *Working Notes for the AAAI Symposium on Adaptation, Co-evolution and Learning in Multiagent Systems*, Stanford University, CA, 1996, pp. 99–104.