

Automated Agents that Proficiently Negotiate with People: Can We Keep People out of the Evaluation Loop*

Raz Lin, Yinon Oshrat and Sarit Kraus

Abstract Research on automated negotiators has flourished in recent years. Among the important issues considered is how these automated negotiators can proficiently negotiate with people. To validate this, many experimentations with people are required. Nonetheless, conducting experiments with people is timely and costly, making the evaluation of these automated negotiators a very difficult process. Moreover, each revision of the agent's strategies requires to gather an additional set of people for the experiments. In this paper we investigate the use of Peer Designed Agents (PDAs) – computer agents developed by human subjects – as a method for evaluating automated negotiators. We have examined the negotiation results and its dynamics in extensive simulations with more than 300 human negotiators and more than 50 PDAs in two distinct negotiation environments. Results show that computer agents perform better than PDAs in the same negotiation contexts in which they perform better than people, and that on average, they exhibit the same measure of generosity towards their negotiation partners. Thus, we found that using the method of peer designed negotiators embodies the promise of relieving some of the need for people when evaluating automated negotiators.

Raz Lin

Department of Computer Science, Bar-Ilan University, Ramat-Gan, Israel 52900, e-mail: lin-raz@cs.biu.ac.il

Sarit Kraus

Department of Computer Science, Bar-Ilan University, Ramat-Gan, Israel 52900 and Institute for Advanced Computer Studies, University of Maryland, College Park, MD 20742 USA, e-mail: sarit@cs.biu.ac.il

* This research is based upon work supported in part by the U.S. Army Research Laboratory and the U.S. Army Research Office under grant number W911NF-08-1-0144 and under NSF grant 0705587.

1 Introduction

An important aspect in research on automated negotiation is the design of proficient automated negotiators with people [1, 8, 9, 10, 11, 14, 18, 24]. We refer to these agents as EDNs (Expert Designed Negotiators). EDNs can be used with humans in the loop or without them. EDNs can be used in tandem with people to alleviate some of the efforts required of people during negotiations and also assist people that are less qualified in the negotiation process [4, 9, 24]. Additionally, there may be situations in which EDNs can even replace human negotiators. Another possibility is for people embarking on important negotiation tasks to use these agents as a training tool [3, 16], prior to actually performing the task. Thus, success in developing an automated agent with negotiation capabilities has great advantages and implications.

The design of automated agents that proficiently negotiate with people is a challenging task, as there are different environments and constraints that should be considered (for a recent survey of that describes studies that evaluate automatic agents that negotiate with people, see [13]). While the design issues of EDNs is important, we found little, if any, literature focusing on the important issue of the evaluation of EDNs designed to negotiate with people. Thus, in this paper we only focus on the evaluation process of these EDNs.

The evaluation and validation process of the automated negotiators is a vital part of the design process and allows demonstrating how successful the automated negotiators are. Yet, using people for experimentation purposes is timely and costly, making the evaluation process a very difficult task for researchers. Designing agents that model the human behavior during negotiations only adds to the difficulties, due to the diverse behavior of people which makes it hard to capture it by a monolithic model. For example, people tend to make mistakes, and they are affected by cognitive, social and cultural factors, etc. [12]. Thus, it is commonly assumed that people cannot be substituted in the evaluation process of EDNs designed to negotiate with people.

The question which now arises is, even though people and agents behave differently, whether one can use agents to evaluate EDNs and reflect from the behavior of the EDNs with other agents to their behavior with people. Following this intuition, we turned to the strategy method [15, 20, 21] which is an experimental methodology which requires people to elicit their actions. The assumption behind this method is that people are able to effectively encapsulate their own strategies if they are properly motivated, monetarily or otherwise. This approach is well accepted in experimental economics and has also begun to be used within artificial intelligence research [2, 19]. The application of this methodology within the study of automated negotiator agents implies that peer designed agents (PDAs) that represent the negotiation preferences of diverse people can be developed.

In this paper we present an in-depth investigation into the question of whether PDAs can be used to keep people out of the evaluation loop and thus simplify the evaluation process required by designers of EDNs. As we will demonstrate in the rest of this paper, the use of PDAs has the potential of elevating some of the need for

people in the evaluation of automated negotiators, yet people are still a mandatory factor in the final evaluation loop.

In this paper we provide results of extensive experiments involving more than 300 human subjects and 50 PDAs. The experiments involved negotiations of people that interact with other people, people that interact with PDAs and people that interact with EDNs. The experiments were conducted in two distinct negotiation environments that simulate real-world scenarios that require negotiators to reach agreement about the exchange of resources in order to complete their goals. In each experiment we investigate the behavior and dynamics of people and PDAs with respect to the EDNs in order to find out whether the behavior of PDAs is similar to that of people and whether they can be used as a substitute for people in the evaluation process of automated negotiators. The results shed light as to the prospect of using PDAs to better determine the proficiency of an automated negotiator when matched with people, as well as to compare the behavior of different EDNs.

This paper contributes to research on automated negotiations by tackling the important issue of the evaluation process of automated negotiators. We suggest using PDAs as an unbiased mechanism for evaluating automated negotiators, which can reflect on the behavior of people, as well as allowing fine-tuning and improving the strategy of the automated negotiators. We provide a general methodology for evaluating automated negotiators and different measures to compare the behavior of PDAs and people and show that it is important to understand the different behavioral aspects expressed by them to gain a better perspective on the prospects of the automated negotiator being proficient with people.

The rest of this paper is organized as follows. In Section 2 we review related work in the field of automated negotiators' evaluation. We provide an overview of the problem in Section 3. We continue and describe our methods for evaluating the negotiation behavior and dynamics in Section 4. In Section 5 we present the experiments we conducted, our methodology and our evaluation results. Finally, we provide a summary and discuss the results.

2 Related Work

Important differences exist between designing an automated agent that can successfully negotiate with a human counterpart and designing an automated agent to negotiate with other automated agents. As this paper does not focus on the design of such automated agents, we will not survey related work on this topic. A more detailed review on the design of automated agents capable of negotiating with people can be found, for example, in [13].

To date, in order to replace people in the evaluation loop, one cannot rely on specific automated agents. Instead, we examine the use of peer designed agent as a type of strategy method. Similarly to developing agents, using the strategy method requires subjects to specify their choices for all information sets of the game and not only the ones that occur during the course of a play of a game [15, 21, 20].

Despite the similarity, asking subjects to design and program agents is different from the strategy method. Developing agents requires subjects to implement much more complex strategies (e.g., using heuristics and learning algorithms), potentially, to make decisions in situations not originally considered.

The use of PDAs has been extensively studied within the context of the Trading Agent Competition for Supply Chain Management (TAC SCM) [22]. In TAC one needs to design a trading agent that participates in auctions for certain good. The use of PDA's within this domain demonstrates the benefits of a large set of PDAs for evaluation purposes of EDNs. Yet, in this context, both the PDAs and the EDNs were used for interacting with other computer agents, and not for interaction with people.

Grosz *et al.* [5] experimented with people designing agents for a game called Colored Trails. They observed that when people design agents, they do not always follow equilibrium strategies. Moreover, in their analysis they showed that people demonstrated more helpfulness, which led to higher scores, than their designed agents. Chalamish *et al.* [2] report on large-scale experiments in which people programmed agents which were shown to successfully capture their strategy in a set of simple games. They conclude that peer designed agents can be used instead of people in some cases. In another settings, Rosenfeld and Kraus [19] report on experiments done with PDAs designed for optimization problems. Based on the experiments they conclude that theories of bounded rationality can be used to better simulate people's behavior. However, the settings of Chalamish *et al.* [2] were relatively simple, while our settings have richer strategy space and are much more complicated. Hence, it is not straightforward that PDAs can be used to simulate the people's behavior and thus replace them for evaluation purposes.

3 Problem Description

We consider the problem of evaluating the proficiency of EDNs designed to negotiate with people. We consider a general environment of bilateral negotiation in which two agents, either automated negotiators or people, negotiate to reach an agreement on conflicting issues. We consider two distinct bilateral negotiation environments. The first involves a day-to-day scenario in which the parties negotiate to reach an agreement on conflicting goals, while the second involves playing a game. We describe both environments below.

The first negotiation environment involves a multi-attribute multi-issues negotiation environment (see Figure 1). In this environment, 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 (*OPT*), or (c) a predefined deadline, *dl*, is reached, whereby, if a partial agreement is reached it is implemented or, if no agreement is reached, a status quo outcome (*SQ*) is implemented. It is assumed that the agents can take actions during the negotiation process until it terminates. 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 $i \in I$. Therefore, an offer is denoted as a vector $\vec{o} \in O$. It is assumed that the agents can take actions during the negotiation process until it terminates. Let **Time** denote the set of time periods in the negotiation, that is **Time** = $\{0, 1, \dots, dl\}$. Time also has an impact on 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 other agent 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. We use an extension of the model of alternating offers [17, p. 118-121], in which each agent can perform up to $M > 0$ interactions with its counterpart in each time period.

Fig. 1 Bilateral Negotiation: Generating offers screen.

In order to make the settings more realistic, it also involved incomplete information concerning the opponent's preferences. 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 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. The negotiators are aware of the set of possible types of the opponent. However, the exact utility function of the rival is private information.

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 agents (PDAs or EDNs) 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.

In this environment we experimented with two state-of-the-art automated negotiators, *KBAgent* and *QOAgent* which were shown by Oshrat *et al.* [18] and Lin *et al.* [14] to negotiate proficiently with people. Both agents are domain independent and apply non-classical decision making method, rather than focusing on maximizing the expected utility. They also apply different learning mechanism to determine the type of their counterpart. Both agents were shown to reach more agreements and played more effectively than their human counterparts, when the effectiveness is measured by the score of the individual utility. Since they were shown to be proficient negotiators with people they can serve as our baseline for evaluating the PDAs as a strategy method for replacing people in the evaluation loop.

The second negotiation environment involved playing the Colored Trails (CT) game [5] which is a general negotiation test-bed that provides an analogy to task-settings in the real-world². The game is played on a $n \times m$ board of colored squares. Players are issued colored chips and are required to move from their initial square to a designated goal square. To move to an adjacent square, a player must turn in a chip of the same color as the square. Players must negotiate with each other to obtain chips needed to reach the goal square (see Figure 2). 100 points are given for reaching the goal square and 10 points bonus are given for each chip left for each agent at the end of the game. If the player did not reach the goal, 15 points penalty are given for each square from its final position to the goal square. Note that in this game, the performance of the agent does not depend on the outcome of the other player. Agreements are not enforceable, allowing players to promise chips but not transferring them. In addition, each player can see the entire game board.

The simulation environment we used in this setting is adaptable such that different variations of the game can be set. The size of the board, number and color of total chips and chips given to each player can be changed. The automated agents can play both sides in the game, while the human counterpart accesses the game via a web address. The game itself is split into turns, where each turn is divided to a negotiation phase, a commitment phase and a movement phase. In the negotiation

² Colored Trails is Free Software and can be downloaded at <http://www.eecs.harvard.edu/ai/ct>


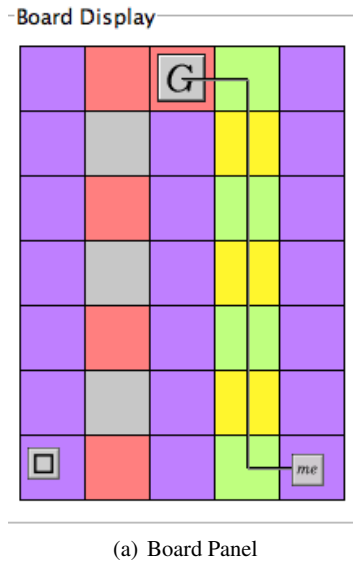


Fig. 2 Snapshots of Colored Trails GUI

phase the players can request or promise to send chips. Then, in the commitment phase, the players can send the actual chips or withdraw from the agreement. This might result in one agent sending the promised chips in return to be given other chips, while the other agent fails to deliver. In the movement phase, the players can choose to move to adjacent squares, given they have the required colored chips. The game terminates when either side reaches the goal square or if no player has moved in three consecutive turns.

An EDN was used in this environment as well. The agent, which is called the Personality Based (*PB*) agent, extends the agent reported by Talman *et al.* [23] to allow it to play proficiently with people. It combines a social utility function that represented the behavioral traits of other participants, as well as a rule-based mechanism that used the utility function to make decisions in the negotiation process. Its behavior is characterized as high reliability and medium generosity. The agent can change its personality level for cooperation and reliability and also model these traits of its opponent. A utility function is used to evaluate each possible action and proposal and randomization is also used to choose between different choices.

4 Evaluating the Negotiation Dynamics

In order to analyze whether PDAs can be a suitable replacement for people in evaluating automated negotiators, we focus on the analysis of the negotiation dynamics when people, PDAs and EDNs are involved. The analysis of the negotiation dynamics is a mean of understanding how similar or different the behavior of each population is and its possible effects on the negotiation process. Different parameters investigated with respect to the characteristics of the negotiators and its dynamics.

The first parameter is the final outcome of the negotiation. This includes the final utility score, average end turn and the type of outcome that was reached. The final utility scores can serve as an indication whether people, PDAs and EDNs reach similar outcomes. For the first setting of bilateral negotiation the final outcome can either be a full agreement, a partial agreement, a status quo outcome or an opt out option. For the CT game we compare the percentages of the games ending by reaching the goal square.

The second parameter is the number of proposal exchanged during a negotiation session. This parameter reflects on the strategic behavior of each negotiator. Comparing the different negotiators using this parameter is important, since PDAs and automated negotiators have a larger computation power and they may excess it to send a large number of proposals which can affect the dynamics of the negotiation.

Another parameter that we analyze is the characteristics of the proposals made by the negotiators. For the bilateral negotiation setting we characterize the proposals based on the step-wise analysis method (“negotiation moves”) suggested by Hindriks *et al.* [7] (see Figure 3. In this analysis, each proposal is compared to its preceding proposal (made by the same negotiator) based on the utilities of the offers for both sides. Thus, an offer can be characterized as either:

1. Selfish (better for the proposer, worse for the other side),
2. Fortunate (better for the proposer and for the other side),
3. Unfortunate (worse for both sides),
4. Concession (worse for the proposer, better for the other party),
5. Silent (both utilities are changed only within a given small threshold), or
6. Nice oriented (utility for the proposer is only within a given small threshold, and higher for the other party).

Formally, let $u_{diff_A}^t$ and $u_{diff_B}^t$ denote the difference between the utilities of side A and B , respectively, at time t , that is: $u_{diff_A}^t = u_A(\vec{o}, t+1) - u_A(\vec{o}, t)$ (and similarly for side B). Let T_A and T_B denote a small threshold for sides A and B , respectively. In our settings we chose the threshold to be equal to the time discount value. A proposal made by agent A at time $t+1$ is characterized based on the algorithm listed in Listing 1.

For the CT settings, the characteristics of the proposals is based on the type of the proposal. Three proposal types were available for the players, each represents the ratio between the number of chips the player is willing to send and in return the number of chips it requires:

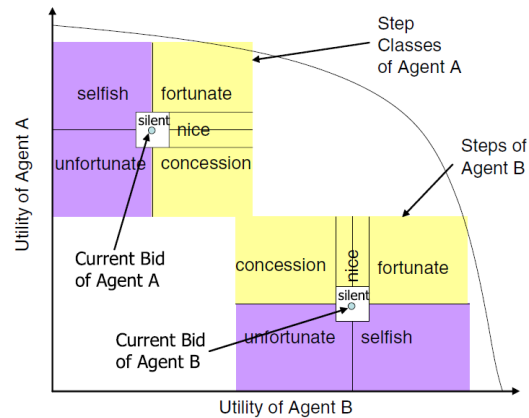


Fig. 3 Classification of negotiation moves. Figure taken from [6].

- A *Give* proposal is a proposal in which the player proposes to send less chips than it will receive (e.g., send 1 chip in exchange of receiving 2 chips from the other party).
- A *Take* proposal is one in which the player proposes to give more than it receives.
- A *Reciprocal* proposal is one in which the two sides send and receive the same number of chips.

The type of the proposal the players choose to propose can reflect on the cooperativeness level of each player. The cooperativeness level or the reliability level can also be reflected by the extent of fulfilling the agreements.

5 Experiments

The experiments were conducted using the simulation environments mentioned in Section 3. We begin by describing the environments which were used in the different experiments and then continue to describe the experimental methodology and results.

5.1 The Negotiation Environments

For the bilateral negotiation environment two domains were used, which were described by Lin *et al.* [14]. The first domain is a Job Candidate domain, which is related to the subjects' experience, and thus they could better identify with it. In this domain, 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 can-

Listing 1 Negotiation Moves Characteristics

```

1: if ( $u_{diff_A}^t > T_A$ ) then
2:   if ( $u_{diff_B}^t > 0$ ) then
3:     Proposal is fortunate
4:   else
5:     Proposal is selfish
6:   end if
7: else
8:   if ( $u_{diff_A}^t \geq (-1) \cdot T_A$ ) then
9:     if ( $u_{diff_B}^t > T_B$ ) then
10:      Proposal is nice
11:    else
12:      if ( $u_{diff_B}^t \geq (-1) \cdot T_B$ ) then
13:        Proposal is silent
14:      else
15:        if ( $u_{diff_A}^t > 0$ ) then
16:          Proposal is selfish
17:        else
18:          Proposal is unfortunate
19:        end if
20:      end if
21:    end if
22:  else
23:    if ( $u_{diff_B}^t \geq (-1) \cdot T_B$ ) then
24:      Proposal is concession
25:    else
26:      Proposal is unfortunate
27:    end if
28:  end if
29: end if

```

didate wish to formalize the hiring terms and conditions of the applicant. In this scenario, 5 different attributes are negotiable with a total of 1,296 possible agreements that exist.

The second domain involved reaching an agreement between England and Zimbabwe evolving from the World Health Organization’s Framework Convention on Tobacco Control, the world’s first public health treaty. The principal goal of the convention is “to protect present and future generations from the devastating health, social, environmental and economic consequences of tobacco consumption and exposure to tobacco smoke”. In this domain, 4 different attributes are under negotiation, resulting with a total of 576 possible agreements.

In both domains if an agreement is not reached by the end of the allocated time a status quo agreement is implemented. In addition, time also has an impact and the sides might lose or gain as time progresses. In the Job candidate domain both sides lose as time progresses, while in the England-Zimbabwe domain, England gains while Zimbabwe loses as time progresses. Also, each side can choose to opt out of the negotiation at any time. As there is also incomplete information in each domain, we assume that there are three possible types of agents for each role. These types are

associated with different additive utility functions. The different types are characterized as ones with short-term orientation regarding the final agreement, long-term and a compromising orientation.

In the CT game, a 7×5 board was used. Two types of games were used, where each differed by the ability of each player to reach the goal square with or without the assistance of the other player:

1. *Asymmetric* game, which was characterized by one of the player having 15 chips and being *dependant* of the other player and needed to exchange chips in order to reach the goal square, while the other player had 24 chips and was *independent* of its counterpart, and thus could reach the goal square without doing any exchange.
2. *Symmetric* game, which was characterized by the two players having 24 chips and being dependent and needing each other's chips to reach the goal square.

5.2 Experimental Methodology

We ran an extensive set of simulations, consisting of more than 300 human negotiators and more than 50 PDAs. The human negotiators were mostly computer science undergraduate and graduate students, while a few were former students who are currently working in the Hi-Tech industry. Each subject served only one specific role in the negotiations (e.g., in the bilateral negotiations either the employer role or the job candidate one, and in the CT game environment either the dependent player or the independent player). 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. The score function was private information and unknown for the other side.

The PDAs were automated negotiators designed by people. The students were given a task to implement an efficient automated agent for a given negotiation environment (different students designed PDAs for the different negotiation environments). The implementation was done in the same simulation environment as the negotiation itself. The students were provided skeleton classes, having all the necessary server-communication functionality, to help them implement their agents. This also allowed them to focus on the strategy and the behavior of the agent, and eliminate the need to implement the communication protocol or the negotiation protocol. In addition, it provided them with a simulation environment in which they could test their agents and their strategies. The students were able to first negotiate or play the CT game before submitting their PDAs.

5.3 Experimental Results

The main goal of the experiments was to analyze whether the strategy method of PDAs can be used to replace people in the evaluation process of EDNs designed

to negotiate proficiently with people. In addition, we wanted to find whether this method can also be used to evaluate and compare different automated negotiators and obtain from it which will be a more proficient negotiator with people.

5.3.1 Evaluating EDNs when Matched with PDAs versus when Matched with People

In this section we analyze the final outcomes of the negotiation, mainly the final utility, the negotiation's duration and the success of it, to understand whether these parameters can be used, when matching EDNs with PDAs, as indicators to the results if the EDNs are later matched with people.

5.3.1.1 Analyzing Utility Outcomes

Table 1 summarizes the final utilities achieved by each side in each experiment for the Job candidate and England-Zimbabwe domains, while Table 2 summarizes the final utilities in the CT game environment. All results are statistically significant within the $p < 0.05$ range. To try and understand whether PDAs can be used in replacement of people to predict the performance of the EDNs we compare the results of the final utility values when the PDAs were involved, that is PDAs versus EDNs and PDAs versus PDAs, and when people were involved, that is, people versus EDNs and people versus people.

When the *KBAgent* was matched with PDAs it was able to achieve higher utility values than the average of the PDAs matched against themselves (lines (1),(2) in Table 1). This is also consistent with the *KBAgent*'s achievements when matched with people (lines (3),(4) in Table 1). Similar results are attained by the second EDN in this environment, the *QOAgent* (lines (2),(5) as compared to lines (4), (6) in Table 1).

A similar phenomenon was observed in the CT game. When the *PB* agent played in the symmetric settings and in the asymmetric game as the independent role, the final utilities achieved by it were higher than the average utilities of the PDAs. When it played the dependent role in the asymmetric game, its final utility was lower than the average utility of the PDAs (lines (1),(2) in Table 2). The same relation is revealed when comparing the *PB*'s utility when playing with people and the average utilities of people playing with one another (lines (3),(4) in Table 2).

In the bilateral negotiation domain we had two distinct EDNs (the *KBAgent* and the *QOAgent*). Thus, it is interesting to see whether the performance of them when matched with PDAs can be used as a prediction of whom will perform better when matched with people. The *KBAgent* was shown to perform better when matched with people than the *QOAgent* (lines (3),(6) in Table 1). In three out of the four sides in the two domains, this is also reflected when they are matched with the PDAs, with the *KBAgent* achieving higher utility scores than the *QOAgent* (lines (1),(5) in Table 1).

	Job Can. Domain		Eng-Zim Domain	
	u_{employer}	$u_{\text{job can}}$	u_{eng}	u_{zim}
(1) <i>KBAgent vs. PDAs</i>	437.7	415.8	720.0	-14.5
(2) <i>PDAs vs. PDAs</i>	368.2	355.1	251.8	-83.7
(3) <i>KBAgent vs. People</i>	472.8	482.7	620.5	181.8
(4) <i>People vs. People</i>	423.8	328.9	314.4	-160.0
(5) <i>QOAgent vs. PDAs</i>	466.1	396.8	663.4	-36.5
(6) <i>QOAgent vs. People</i>	417.4	397.8	384.9	35.3

Table 1 Final utility results in the bilateral negotiation environment.

Thus, the results support the hypothesis that the final utility values can serve as a good indication for evaluating the proficiency of the automated negotiator. Moreover, they can also be used to compare between different EDNs and reflect on their proficiency when matched with people.

	Asymmetric game		Symmetric game
	$u_{\text{independent}}$	$u_{\text{dependent}}$	$u_{\text{dependent}}$
(1) <i>PB vs. PDAs</i>	180.53	35.00	131.36
(2) <i>PDAs vs. PDAs</i>	178.38	45.25	111.48
(3) <i>PB vs. People</i>	187.08	81.94	157.83
(4) <i>People vs. People</i>	181.45	97.26	130.67

Table 2 Final utility results in the CT game environment.

5.3.1.2 Analyzing the Characteristics of the Negotiation Ending

Two more relevant questions are whether the end period of the negotiation and the type of agreements reached or whether the goal was reached when EDNs are matched with PDAs can also serve as an indication to the proficiency of the EDNs when matched with people. Tables 3 and 4 summarize the average end turn and the percentages of the negotiations terminated with full agreements, status quo (SQ) agreements, opting out (OPT) or partial agreements in the Job candidate and the England-Zimbabwe domains, respectively, while Table 5 summarizes the average end turn and the percentages of reaching the goal square in the CT game settings.

With regard to the duration of the negotiation, while in the Job candidate domain the negotiation lasted longer when the *KBAgent* was matched with people than when it was matched with PDAs, it lasted shorter time in the England-Zimbabwe domain (lines (5),(8) in Tables 3 and 4). The CT game also lasted shorter when the *PB* was

	End Turn	Full	SQ	OPT	Partial
People vs. People					
(1)	6.68	91%	2%	4%	2%
PDA vs. PDAs					
(2)	5.56	80%	4%	7%	9%
KBAgent vs. PDAs					
(3) <i>KBAgent_{Employer}</i>	6.14	86%	14%		
(4) <i>KBAgent_{Job can.}</i>	5.85	77%	8%	8%	8%
(5) Average	6	81%	11%	4%	4%
KBAgent vs. People					
(6) <i>KBAgent_{Employer}</i>	7.17	100%			
(7) <i>KBAgent_{Job can.}</i>	6.13	100%			
(8) Average	6.68	100%			

Table 3 Average end turn and percentages of reaching full agreements, status quo, opting out or partial agreements in the Job Candidate bilateral negotiation domain.

	End Turn	Full	SQ	OPT	Partial
People vs. People					
(1)	9.08	70%	8%	11%	11%
PDA vs. PDAs					
(2)	8.30	69%	8%	12%	12%
KBAgent vs. PDAs					
(3) <i>KBAgent_{England}</i>	11.50	60%			40%
(4) <i>KBAgent_{Zimbabwe}</i>	10.00	50%	10%		40%
(5) Average	10.75	55%	5%		30%
KBAgent vs. People					
(6) <i>KBAgent_{England}</i>	10.06	94%			6%
(7) <i>KBAgent_{Zimbabwe}</i>	6.13	100%			
(8) Average	8.15	97%			3%

Table 4 Average end turn and percentages of reaching full agreements, status quo, opting out or partial agreements in the England-Zimbabwe bilateral negotiation domain.

matched with people than when it was matched with the PDAs (lines (3),(4) in Table 5). Thus the negotiation duration when EDNs are matched with PDAs cannot be used as a good indication to the duration when the EDNs are matched with people.

With regard to the way the negotiation ended, when the *KBAgent* is matched with people more full agreements are reached than when it is matched with PDAs (as shown in lines (5),(8) in Tables 3 and 4). Similarly, the *PB* agent reaches the goal square in more cases when matched with people than when matched with other PDAs (as shown in lines (3),(4) in Table 5).

	Asymmetric game			Symmetric game	
	End	Goal		End	Goal
	Turn	Reached		Turn	Reached
	<i>Ind Dep</i>	<i>Ind</i>	<i>Dep</i>	<i>Dep</i>	
(1) People vs. People	4.1	97%	74%	5.8	60%
(2) PDA vs. PDA	4.5	98%	32%	4.8	47%
(3) PB vs. PDA	6.1	5.9	100%	4.6	73%
(4) PB vs. People	3.8	4.4	100%	4.3	87%

Table 5 Average end turn and percentages of reaching the goal in the CT game settings. *Ind* stands for the independent player and *Dep* for the dependent player.

5.3.2 Evaluating the Performance and Behavior of People versus PDAs

It was demonstrated that PDAs behave differently than people (e.g., [5, 19]), yet our experiments indicated that playing against PDAs can reflect on the results when the EDNs are matched with people. Thus, it is important to understand whether the behavior of people is similar to that of PDAs when either matched with the same population or with the EDNs.

Our results show that most of the differences in the behavior of PDAs as compared to people lie in the negotiation dynamics. In the following subsections we investigate different parameters with respect to the negotiation dynamics.

5.3.2.1 Investigating the Final Outcome

With respect to the final outcome, when people were matched with people they were able to achieve higher utilities than when the PDAs were matched with other PDAs in one of the roles (the employer role in the Job candidate domain and the England role in the England-Zimbabwe domain) and worse in the other roles (lines (2),(4) in Table 1), while in the CT game they were better in both roles (lines (2),(4) in Table 2). Moreover, in most occasions the PDAs demonstrated lower percentages in reaching full agreements or reaching the goal square when matched against themselves than when people were matched against people (lines (1),(2) in Tables 3, 4 and 5).

5.3.2.2 Investigating the Negotiation Dynamics

To bolster our confidence from these results we examined the patterns of behavior demonstrated by people and the PDAs when matched with the EDNs. To make it concise, we only present the results on one of the domains and negotiation's sides, though the results are similar in the other domains and sides. Figure 4 compares the final utilities achieved by PDAs and people when matched with the EDNs in the job

candidate domain when playing the role of the employer, while Figure 5 compares the times in which the negotiations terminated. Note, that we compare between the behavior of people and PDAs and not the EDNs behavior. The results demonstrate the similarity and trend between people and PDAs when matched with EDNs. For example, in Figure 4 we can observe that PDAs achieve somewhat higher utilities when matched with the *QOAgent* as compared to the *KBAgent*. The same trend is then observed when people are matched with both agents.

Another parameter we investigated is the number of proposals exchanged during the negotiation. Tables 6 and 7 summarize the total number of exchanges for the bilateral negotiation domains and the CT game settings, respectively. In most of the times, more exchanges are made when PDAs are involved than when only people are involved. This can be explained by the fact the automated agents have more computing power, making it easier for them to evaluate messages and send more messages. When the *KBAgent* was matched with people in three out of the four settings the average number of proposals it received from the other party was lower when received from people than when received from the PDAs (lines (3),(4) in Table 6). Similarly, in two out three cases the average number of proposals received from people was lower when matched with the *PB* agent than when the PDAs played against the *PB* agent (lines (3),(4) in Table 7). It is interesting to observe the exception when people exchange more proposals than agents. This happens in the bilateral negotiation domain when people play the job candidate role or the Zimbabwe role and in the CT game settings in the symmetrical game when both players are dependent of each other. It seems that the cause lies in the fact that in these roles the people play the “underdog” roles and have a greater incentive to propose more messages in the hope of one of them being accepted by the other side. The other side, in return, has the incentive to make the negotiation lasts longer or toughen its stands.

When comparing the PDAs’ only negotiations to people’s only negotiations slightly more exchanges are made by people in the Job candidate domain, yet higher exchanges are made by PDAs in the England-Zimbabwe domain (lines (1),(2) in Table 6). As for the CT game settings, while in the symmetric settings more proposals are made by people, in the asymmetric game settings more proposals are made by




Fig. 4 Comparing overall performance between people and PDAs when matched with EDNs when playing the role of the employer in the job candidate domain.




Fig. 5 Comparing end-turn between people and PDAs when matched with EDNs when playing the role of the employer in the job candidate domain.

Average Proposals Number				
	Job Can. Domain	Eng-Zim Domain		
	Employer	Job can.	England	Zimbabwe
People vs. People				
(1)	7.06	6.87	8.00	9.13
PDAs vs. PDAs				
(2)	6.04	6.13	11.66	11.73
KBAgent vs. PDAs				
(3) PDAs's proposals	10.85	7.14	12.80	27.90
KBAgent vs. People				
(4) People's proposals	5.44	8.94	7.25	11.76

Table 6 Average number of proposals sent in the bilateral negotiation settings.

the PDAs when matched with other PDAs, as compared to proposals made by people (lines (1),(2) in Table 7). In both settings, it seems that more exchanges are made by the PDAs when the domain involves two sides, in which one has significantly more leverage than the other (in the England-Zimbabwe domain, England gains more as time progresses while Zimbabwe loses, and in the asymmetric settings the independent role can reach the goal without needing the dependent role). Thus, it seems to be the case that each side stands firmly and is inclined less to concede.

This leads us to the question whether the average behavior of people and the average behavior of PDAs differ. The results indeed demonstrate that there are differences in the type of exchanges made by people when negotiating with EDNs and those made by the PDAs. Recall that we characterize the negotiation moves in the bilateral negotiation domains as selfish, fortunate, concession, unfortunate, concession, silent or nice (cf. Section 4), where cooperation is characterized by moves of type fortunate, concession, silent or nice. In the CT game settings we defined three types of proposals (*Give*, *Take* and *Reciprocal*, from which cooperation is characterized by *Take* exchanges.

Average Proposals Number			
	Asymmetric game		Symmetric game
	Independent	Dependent	Dependent
People vs. People			
(1)	3.16	2.87	4.15
PDA's vs. PDA's			
(2)	3.54	3.31	3.28
PB vs. PDA's			
(3) PDA's's proposals	2.68	2.95	2.09
PB vs. People			
(4) People's proposals	2.39	2.08	2.30

Table 7 Average number of proposals sent in the CT game settings.

The cooperation levels differ between people and PDAs in all settings we tested. Tables 8 and 9 summarize the negotiation moves for the Job candidate and the England-Zimbabwe domains, respectively. When matched with the EDNs, while the majority of the proposals made by people are concession oriented in both domains (lines (7),(8) in Tables 8 and 9), the majority of the proposals made by the PDAs are nice oriented and silent oriented in the Job candidate domain (lines (5),(6) in Table 8) and concession oriented and silent oriented in the England-Zimbabwe domain (lines (5),(6) in Table 9). In the Job candidate domain, people are also succeeding in making more fortunate offers and less silent offers than the PDAs when playing both roles against the same population (lines (1),(3) and (2),(4) in Table 8). It is also interesting to note that people are less selfish when playing the role of England and more selfish when playing the role of Zimbabwe, while the opposite is expressed in the behavior of the PDAs (lines (1),(3) and (2),(4) in Table 9). This might be derived from the leverage England role has over Zimbabwe and the fact that it gains as time progresses allows people to play more generously when playing England, while playing more rigidly when playing as Zimbabwe. A somewhat similar observation was also reported in experiments with the CT game run by Grosz *et al.* [5].

Figures 6 and 7 display the cooperation levels of people and PDAs in the Job candidate domain and England-Zimbabwe domain, respectively. We can observe similarity in the cooperation levels of both people and PDAs in the Job candidate domain (75.35% and 69.57% when people are matched with people, 77.78% and 65.03% when people are matched with the *KBAgent*, 72.01% and 75.51% when PDAs are matched with PDAs and 70.77% and 71.59% when PDAs are matched with the *KBAgent*), however the cooperation level is again inverted in the England-Zimbabwe domain. People are more cooperative when playing the England role (71.62% when matched with people and 70% when matched with the *KBAgent*) and less cooperative when playing the Zimbabwe role (59.77% and 61.2% when playing against people and against the *KBAgent*, respectively). Moreover, a similar trend is

	Se	F	U	C	Si	N
People vs. People						
(1) People _A	10.92%	19.37%	13.73%	24.65%	29.23%	2.11%
(2) People _B	11.96%	17.03%	18.48%	19.93%	31.88%	0.72%
PDA vs. PDA						
(3) PDA _A	6.70%	14.20%	21.30%	20.35%	32.13%	5.32%
(4) PDA _B	9.04%	12.69%	15.45%	18.98%	37.85%	5.99%
KBAgent vs. PDA						
(5) PDA _A	16.15%	5.38%	13.08%	15.38%	13.08%	36.92%
(6) PDA _B	19.32%	10.23%	9.09%	12.50%	27.27%	21.59%
KBAgent vs. People						
(7) People _A	12.50%	15.28%	9.72%	47.22%	9.72%	5.56%
(8) People _B	20.98%	16.78%	13.99%	34.27%	7.69%	6.29%

Table 8 Negotiation moves in the Job Candidate bilateral negotiation domain. Se = Selfish, F = Fortunate, U = Unfortunate, C = Concession, Si = Silent, N = Nice. Side A is the Employer and side B is the job candidate.

	Se	F	U	C	Si	N
People vs. People						
(1) People _A	14.86%	10.54%	13.51%	43.51%	13.51%	4.05%
(2) People _B	30.47%	10.93%	9.77%	25.58%	17.67%	5.58%
PDA vs. PDA						
(3) PDA _A	35.93%	14.10%	10.40%	26.19%	9.75%	3.64%
(4) PDA _B	19.11%	3.36%	7.36%	37.77%	9.94%	22.47%
KBAgent vs. PDA						
(5) PDA _A	19.49%	3.39%	11.86%	39.83%	18.64%	6.78%
(6) PDA _B	23.60%	2.25%	7.49%	16.48%	31.46%	18.73%
KBAgent vs. People						
(7) People _A	23.00%	12.00%	7.00%	49.00%	4.00%	5.00%
(8) People _B	31.69%	7.10%	7.10%	40.98%	11.48%	1.64%

Table 9 Negotiation moves in the England-Zimbabwe bilateral negotiation domain. Se = Selfish, F = Fortunate, U = Unfortunate, C = Concession, Si = Silent, N = Nice. Side A is England and side B is Zimbabwe.

also noticed by the behavior of the PDAs. The cooperation levels when they are matched with other PDAs are 53.67% and 73.53% when playing the role of England and Zimbabwe, respectively, while playing against the *KBAgent* the cooperation levels are high and similar when playing both roles (68.64% and 68.91% when playing the roles of England and Zimbabwe, respectively).

We also measured the cooperative levels, reflected by the types of exchanges, in the CT game settings, as summarized in Tables 10 and 11. Here as well, there is an apparent difference in the behavior of the PDAs and people. When playing

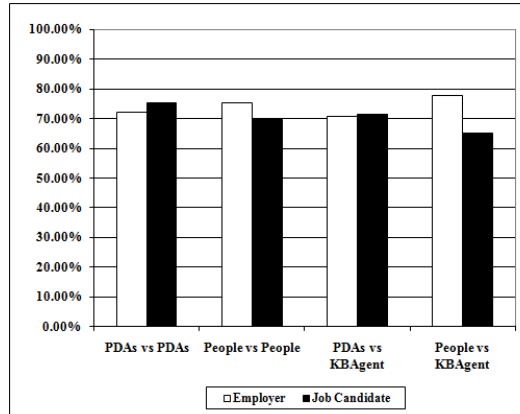


Fig. 6 Cooperation levels in the Job candidate domain.

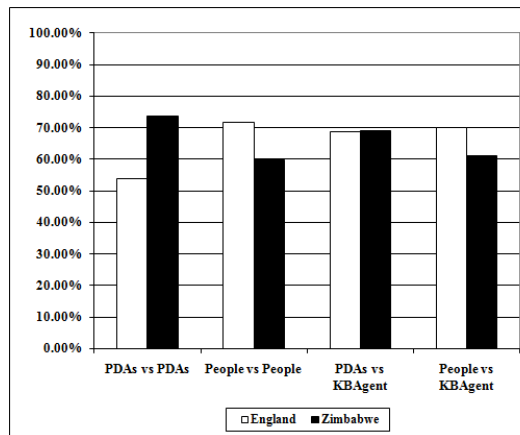


Fig. 7 Cooperation levels in the England-Zimbabwe domain.

with the *PB* agent in the asymmetric settings, the majority of proposals were *reciprocal* (46.51%) when people played the independent role and 48% of type *Give* when playing the dependent role. Yet, the majority of offers when the PDAs were matched with the *PB* agent were 39.22% of type *Give* when they played the role of the independent player and 51.79% *reciprocal* when playing the dependent role (lines (5),(7) and (6),(8) in Table 10). In addition, in both cases the PDAs fully kept 50% of the agreements while people had fully kept 58% and 19% when playing the independent and dependent roles, respectively. More full agreements were also fully kept in the symmetric settings when the PDAs were involved as compared to when people were involved (lines (1),(2) and (3),(4) in Table 11).

	<i>Give</i>	<i>Reciprocal</i>	<i>Take</i>	Fully Kept	Partially Kept	Unkept
People vs. People						
(1) <i>Ind</i>	14.3%	28.6%	57.1%	61%	10%	29%
(2) <i>Dep</i>	55.1%	19.1%	25.8%	32%	39%	29%
PDAs vs. PDAs						
(3) PDAs _{<i>Ind</i>}	15.4%	38.3%	46.4%	58%	2%	40%
(4) PDAs _{<i>Dep</i>}	21.5%	51.5%	67%	53%	2%	30%
PB vs. PDAs						
(5) PDAs _{<i>Ind</i>}	33.3%	39.2%	27.4%	50%	0%	50%
(6) PDAs _{<i>Dep</i>}	25.0%	51.8%	23.2%	50%	10%	41%
PB vs. People						
(7) People _{<i>Ind</i>}	20.9%	46.5%	32.6%	58%	11%	31%
(8) People _{<i>Dep</i>}	48.0%	36.0%	16.0%	19%	38%	44%

Table 10 Negotiation moves in the CT game asymmetric settings. *Ind* stands for the independent player and *Dep* for the dependent player.

	<i>Give</i>	<i>Reciprocal</i>	<i>Take</i>	Fully Kept	Partially Kept	Unkept
People vs. People						
(1) <i>Dep</i>	26.8%	37.5%	35.6%	41%	21%	38%
PDAs vs. PDAs						
(2) PDAs _{<i>Dep</i>}	11.1%	62.9%	26.0%	58%	7%	35%
PB vs. PDAs						
(3) PDAs _{<i>Dep</i>}	15.2%	63.0%	21.7%	78%	0%	22%
PB vs. People						
(4) People _{<i>Dep</i>}	24.6%	47.8%	13.3%	54%	21%	24%

Table 11 Negotiation moves in the CT game symmetric settings. *Dep* stands for the dependent player.

6 Conclusions

The importance of designing proficient automated negotiators to negotiate with people and evaluating them cannot be overstated. Yet, evaluating agents against people is a tiresome task, due to the cost and time required. In this paper we presented an extensive systematic experimentation to answer the question whether people can be kept out of the evaluation loop when evaluating automated negotiators designed specifically to negotiate proficiently with people. To do so, we evaluated several negotiation behavioral parameters in an extensive set of experiments with people and with peer designed agents. In the bottom line, our results reveal that playing well against peer designed agents can reflect on the proficiency of the automated negotia-

tor when matched with people. Moreover, we showed that while PDAs results with different negotiation outcomes than people, there is a common behavioral pattern when they are both matched with EDNs.

Another interesting observation is that the EDNs strategy was mainly similar when matched with people and when matched with PDAs. The *KBAgent* mostly sent concession offers when playing both with PDAs and with people in the bilateral negotiation domains. The *PB* agent, when playing in the symmetric game settings sent mainly *reciprocal* offers, when playing the independent role sent *Take* offers and offered *Give* exchanges when it played the dependent role in the asymmetric game. This is regardless of whether it played with people or PDAs; only the specific offers had been adapted to the specific players. With respect to reliability, the *PB* agent was highly reliable, yet its exact reliability level varied between 0.72 when playing the dependent role in the asymmetric game with people to 0.96 when playing with PDAs in the symmetric case. The consistent strategies with some adaptation the agent made for the specific player it was matched with probably contributed to the similar outcomes with people and PDAs.

Surprisingly, the results also demonstrated that there is a difference in the behavior of PDAs and people, even though the PDAs were supposed to be designed in a way that encapsulates people's strategy. Our results provide analysis of these differences by understanding the negotiation dynamics that occur when people and PDAs are involved, and demonstrate the importance of an efficient opponent modeling mechanism and the need for adapting to different attitudes towards the negotiation.

There are fundamental benefits of using PDAs instead of people. First, PDAs are accessible 24/7 and can be used whenever needed. In addition, PDAs are not biased and thus can be used several times to assess the EDN's behavior. Thus, they allow the agent designer to revise and change her agent with the ability to evaluate each design and compare it to previous designs. Lastly, it allows different EDNs to be matched on the same set of PDAs and obtain an objective evaluation of the results.

While people cannot be kept completely out of the evaluation loop, we demonstrated the promise embodied in peer designed agents for evaluation purposes of automated negotiators. Thus, evaluating on peer designed agents could and should serve as a first extensive attempt to validate the agent's proficiency and strategy design before continuing on to evaluation with people.

As noted before, people behave differently than PDAs. Future work warrants careful investigation on the differences between the behavior of PDAs and people. This investigation might allow for a better understanding of people and the better design of automated agents specifically designed to negotiate with people.

7 Acknowledgments

We thank Ya'akov (Kobi) Gal for his helpful remarks and suggestions.

References

1. A. Byde, M. Yearworth, K.-Y. Chen, and C. Bartolini. AutONA: A system for automated multiple 1-1 negotiation. In *Proceedings of the 2003 IEEE International Conference on Electronic Commerce (CEC)*, pages 59–67, 2003.
2. M. Chalamish, D. Sarne, and S. Kraus. Programming agents as a means of capturing self-strategy. In *Proceedings of the Seventh International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1161–1168, 2008.
3. M. Fleming, D. Olsen, H. Stathes, L. Boteler, P. Grossberg, J. Pfeifer, S. Schiro, J. Banning, and S. Skochelak. Virtual reality skills training for health care professionals in alcohol screening and brief intervention. *Journal of the American Board of Family Medicine*, 22(4):387–398, 2009.
4. Y. Gal, A. Pfeffer, F. Marzo, and B. J. Grosz. Learning social preferences in games. In *Proceedings of the Nineteenth AAAI Conference on Artificial Intelligence*, pages 226–231, 2004.
5. B. Grosz, S. Kraus, S. Talman, and B. Stossel. The influence of social dependencies on decision-making: Initial investigations with a new game. In *Proceedings of the Third International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 782–789, 2004.
6. K. Hindriks, C. Jonker, and D. Tykhonov. Analysis of negotiation dynamics. In M. Klusch, K. Hindriks, M. Papazoglou, and L. Sterling, editors, *Cooperative Information Agents XI*, volume 4676 of *Lecture Notes in Computer Science*, pages 27–35. Springer Berlin / Heidelberg, 2007.
7. K. Hindriks, C. Jonker, and D. Tykhonov. Negotiation dynamics: Analysis, concession tactics, and outcomes. In *Proceedings of the 2007 IEEE/WIC/ACM International Conference on Intelligent Agent Technology (IAT)*, pages 427–433, 2007.
8. C. M. Jonker, V. Robu, and J. Treur. An agent architecture for multi-attribute negotiation using incomplete preference information. *Autonomous Agents and Multi-Agent Systems*, 15(2):221–252, 2007.
9. R. Katz and S. Kraus. Efficient agents for cliff edge environments with a large set of decision options. In *Proceedings of the Fifth International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 697–704, 2006.
10. S. Kraus, P. Hoz-Weiss, J. Wilkenfeld, D. R. Andersen, and A. Pate. Resolving crises through automated bilateral negotiations. *Artificial Intelligence*, 172(1):1–18, 2008.
11. S. Kraus and D. Lehmann. Designing and building a negotiating automated agent. *Computational Intelligence*, 11(1):132–171, 1995.
12. D. A. Lax and J. K. Sebenius. Thinking coalitionally: party arithmetic, process opportunism, and strategic sequencing. In H. P. Young, editor, *Negotiation Analysis*, pages 153–193. The University of Michigan Press, 1992.
13. R. Lin and S. Kraus. Can automated agents proficiently negotiate with humans? *Communications of the ACM*, 53(1):78–88, January 2010.
14. R. Lin, S. Kraus, J. Wilkenfeld, and J. Barry. Negotiating with bounded rational agents in environments with incomplete information using an automated agent. *Artificial Intelligence*, 172(6-7):823–851, 2008.
15. T. Offerman, J. Potters, and H. A. A. Verbon. Cooperation in an overlapping generations experiment. *Games and Economic Behavior*, 36(2):264–275, 2001.
16. D. E. Olsen. Interview and interrogation training using a computer-simulated subject. In *Interservice/Industry Training, Simulation and Education Conference*, 1997.
17. M. J. Osborne and A. Rubinstein. *A Course In Game Theory*. MIT Press, Cambridge MA, 1994.
18. Y. Oshrat, R. Lin, and S. Kraus. Facing the challenge of human-agent negotiations via effective general opponent modeling. In *Proceedings of the Eighth International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 377–384, 2009.

19. A. Rosenfeld and S. Kraus. Modeling agents through bounded rationality theories. In *Proceedings of the Twenty-First International Joint Conference on Artificial Intelligence (IJCAI)*, pages 264–271, 2009.
20. R. Selten, K. Abbink, J. Buchta, and A. Sadrieh. How to play (3x3)-games: A strategy method experiment. *Games and Economic Behavior*, 45(1):19–37, 2003.
21. R. Selten, M. Mitzkewitz, and G. R. Uhlich. Duopoly strategies programmed by experienced players. *Econometrica*, 65(3):517–556, 1997.
22. TAC Team. A trading agent competition. *IEEE Internet Computing*, 5(2):43–51, 2001.
23. S. Talman, M. Hadad, Y. Gal, and S. Kraus. Adapting to agents’ personalities in negotiation. In *Proceedings of the Fourth International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 383–389, 2005.
24. D. Traum, S. Marsella, J. Gratch, J. Lee, and A. Hartholt. Multi-party, multi-issue, multi-strategy negotiation for multi-modal virtual agents. In *Proceedings of the Eighth International Conference on Intelligent Virtual Agents (IVA)*, pages 117 – 130, 2008.