Bilateral and Multi-bilateral Agent-Human Negotiations: Two Experiments

Rustam Vahidov
Concordia University
rvahidov@jmsb.concordia.ca

Gregory Kersten
Concordia University
gregory@jmsb.concordia.ca

Dmitry Guimon
Concordia University
dguimon@jmsb.concordia.ca

Abstract

Negotiation is a flexible mechanism for facilitating effective economic exchanges. Electronic negotiations allow participants to negotiate online and use analytical support tools in making their decisions. Software agents can be employed to automate partially or fully the conduct of negotiation process using these tools. This paper aims at investigating the prospects of agent-to-human negotiations in B2C contexts using experiments with human subjects. Two sets of experiments are described: one including bilateral agent-human negotiations, and the other involving multi-bilateral negotiations with agents competing against other participants on the "many" side. The paper discusses the results obtained from both studies and proposes future research directions.

1. Introduction

Electronic marketplaces allow spatially and temporally separated parties to interact over the internet. Negotiation is one important mechanism for facilitating economic transactions in a flexible fashion. In the course of negotiations parties exchange offers in order to jointly explore the possibilities of finding acceptable solutions. Negotiations involving more than a single issue allow for more maneuvering possibilities for the participating parties in search for agreements that would be beneficial to the negotiators due to the potential asymmetry of their preferences.

Electronic negotiation systems (ENS) allow the parties set up negotiation instances, exchange offers over the internet, and reach agreements [1]. In addition to enabling anytime/anywhere mode of interactions, they may organically incorporate analytical facilities for supporting negotiators in their preparation and conduct of negotiations. The extent of such support can range from tools designed for capturing and modeling negotiator's preferences, to providing active advice and critique, and all the way to complete automation of the negotiation conduct.

While electronic auctions have proved to be a popular mechanism in real-life exchanges, despite early optimistic expectations the growth of negotiations as one of the primary mechanisms of conducting online transactions has been slow. In reality only few "negotiating websites" with limited capabilities exist. Such sites include, for example, Priceline, some car dealer sites, and sites that sell unique items (e.g. arts). One possible explanation to the scarcity of negotiating websites is that negotiations imply a relatively high cognitive load, especially if multiple issues are involved (e.g. price, warranty, product attributes, shipment, etc.). This load may translate into a prohibitive cost when day-to-day transactions involving people who are not negotiation experts are concerned. Software agents may circumvent this problem by automating negotiation process while working with customers towards an acceptable deal. Moreover, use of automated tools can also ensure consistency in reaching negotiation outcomes according to the set policies.

In conducting electronic negotiations software agents can be configured to behave in a variety of ways as specified by their tactics and strategies. For example, using time-dependent tactics [2] they can behave competitively or collaboratively, depending on the context and the needs of a business in a given time frame. For example, if demand for company products or services is high, the agents could follow competitive strategies. On the other hand if customer loyalty and retention are the priority, the agents may be configured to behave collaboratively.

While the work on automated negotiations has been extensive in recent times, relatively few experimental studies have been carried out in assessing the potential of human customer vs. software agent negotiations in different negotiation settings.

The purpose of this work is to investigate the prospects of human – software agent negotiations in experimental settings. Two sets of experiments have been conducted to this end. In the first study human and agent negotiators have been paired up in bilateral settings. The second one included multi-bilateral

978-1-4799-2504-9/14 $31.00 © 2014 IEEE
DOI 10.1109/HICSS.2014.39

negotiations where in each instance humans and agents were negotiating with a single human counterpart, thus competing against each other. The paper describes the systems used, experimental setups and the results observed.

2. Related Work

While past research on automated negotiations involving software agents has been extensive [3, 4], relatively little attention has been paid to software agent-to-human negotiations. In light of the scarcity of theoretical work in this direction, this paper aims at testing proposed agent concession-making tactics proposed within the context of automated negotiations. According to [2] negotiation tactics are the "set of functions that determine how to compute the value of an issue… by considering a *single* criterion". In the above work the authors had proposed the following families of tactics: time-dependent, resource dependent, and behavior dependent [2]. The first family guides agent's offer-making behavior based on the time frame allocated for negotiation session. Resource-based tactics determine value of offers based on resource depletion. Behavior-based tactics take into consideration the opponent's behavior in generating offers.

The current work aims at evaluating the effectiveness of time- and behavior- based tactics when employed by agents in negotiating with agents. The reason for omitting resource-based tactics is that we study single independent negotiation instances whereas resource depletion would be more of an issue if one negotiation outcome affects resource valuation in subsequent negotiation instances.

While thorough coverage of the past work in agent-based negotiation is well beyond the scope of this paper, we will review the representative publications in the context of business exchanges. One could categorize these in accordance with the context of interactions (i.e. C2C, B2B, B2C), and the extent of automation.

One well-known early work in this direction was the construction of the Kasbah electronic marketplace [5, 6]. Targeting primarily the C2C domain the marketplace allowed human users to configure agents, which would then be sent to the marketplace to negotiate with each other using time-based tactics. Three types of agents ranging from competitive to the conceding ones were provided. Negotiations included a single issue, i.e. price.

In B2B applications software agents have been proposed for automating various aspects of supply chain management. For example, in [7] an agent-based architecture has been proposed for dynamic supply chain formation. The agents acting as brokers representing various entities within supply chain negotiated agreements with each other in building up the chain.

There has also been work targeting the B2C transactions. In [8] the authors proposed an agent-based architecture for automated negotiations between businesses and consumers. The buyer agents incorporated such components as searcher and negotiator, while seller agents featured negotiator module whose strategy was set by the sales department.

In [9] the authors have proposed an intelligent sales agent with the capabilities for negotiation and persuasion. The agent employed reinforcement learning in the process. In their experiments with human subjects they found that the agent using persuasion capability has increased buyer's product valuation and willingness to pay.

It has been argued by many that complete automation of real-life negotiations, in particular in business contexts does not seem to be a viable solution (e.g. [10]). Automation in general is applicable only when tasks concerned are well-structured, which is rarely the case in many business situations. However, since efficient policies can be set for multiple daily interactions with the customers regarding the sales of products and services, it seems that a relatively high level of automation may be feasible.

While the work reviewed above concerns fully automated negotiations, there has been some research into sharing responsibilities between human negotiators and negotiation agents. In [11] a system has been proposed where agents actively supported human decision making in the negotiation process. An agent advised the human user on the acceptability of the received offer, helped with the preparation of the counter-offer, and critiqued offers composed by the users when it deemed necessary to intervene.

In [12] an agent-based architecture was proposed for managing multiple negotiations. In this architecture a fleet of agents negotiated deals with customers. These negotiations were monitored by a coordinating agent, which, based on the analysis of situation instructed the negotiating agents to adjust their strategies and reservation levels within the limits of its authority. The overall process was monitored by a human user who could intervene to make changes if necessary.

There has been some experimental work comparing in assessing human-to-agent negotiations. In [9] the authors have described an agent representing a salesperson that employed persuasion and negotiation techniques while interacting with a customer. Persuasion was based on the customer – agent dialogue with the involvement of pre-defined arguments organized into a tree. First, the agent would try to convince a customer to accept an offer. If this did not work, the agent would go into bargaining mode and determine what concession should be made. Price was the single issue in the negotiations. Using the case of a used car sale, the authors conducted both lab and online experiments. Their findings suggested that persuasion increased buyers' product valuation and willingness to pay. Negotiation increased the seller's surplus.

Another related experimental work looked to investigate the effects of framing on the subjective variables when employed by agents using persuasion/ argumentation tactics [13]. Namely, the impacts of gain vs. loss frames adopted in arguments by an agent were studied. In this study subjects were assigned the role of a buyer who had to negotiate purchase of laptops. The issues included unit price, quantity, service level, and delivery terms. Seller was a software agent, and subjects were unaware of it. The authors did not find significant differences in buyer satisfaction with the settlement or with the counter-part when compared across different frames.

The current work is aimed at investigating how software agents perform in agent-to-human dyads as compared to human-human dyads while in multi-issue negotiations in bilateral and multi-bilateral settings. Various types of agents following time- and behavior-based tactics have been configured for the comparison of their performance.

3. Bilateral Negotiations

3.1. System and Negotiation Case

The purpose of experiments with bilateral agent-human negotiations was to assess the effectiveness of different tactics families proposed by past research when employed by software agents. In order to perform experiments with agent/human negotiations a system called DIANA (system for Deal-making Incorporating Agents in Negotiating Agreements) has been developed. The system allows human parties to negotiate in bilateral settings. The parties can specify their preferences towards negotiated issues in a user-friendly fashion using utility scores. Each candidate offer can be assessed using additive utility function. The system also allows pairing up agents with humans. The agents can be configured to follow predefined tactics. To this end the system allows to specify time-dependent tactics using a concession schedule. These schedules can be set graphically and modeled using Bezier curves of up to 3rd order. Another type of agent strategy included is absolute tit-for-tat whereby agents' moves mirrors those made by an opponent.

The negotiation case developed for the experimental study concerned the sale of a desktop computer. This case was chosen since the first set of experiments involved undergraduate students taking an introductory course in information technology. There were five issues including the price, type of monitor, hard drive, service plan, and software loaded. Each option for each issue had a corresponding level of utility (attractiveness), these levels being different for the buyers vs. sellers. In order to calculate the total utility of the offer the issues were assigned different weights. These were then used in an additive utility function to estimate the level of attractiveness of an offer. Agents used this information in order to decide on the acceptability of the received offers and generate offers.

All agents acted on the seller side, and they were not aware of the buyers' preference structures. The weights were slightly different for sellers than buyers to facilitate tradeoffs, which have been considered one of the key integrative negotiation characteristics [14]. Thus, agents would decide on the utility of the next offer first, according to their concession schedules, and then generate the corresponding offer.

We have chosen to use five different concession schedules, three of which were similar to those used in Kasbah experiments, one based on the combination of boulware and conceder tactics, and one representing behavior-based family. Namely, the tactics included: competitive, neutral, collaborative, competitive-then-collaborative, and tit-for-tat. The competitive agents (CM) tend to make smaller concessions in terms of utility of generated offers in the beginning of the negotiation period. However, as they approach the end of the period, they would start making larger concessions in search of an agreement (figure 1).

Neutral strategy (NT) dictates that an agent concedes the constant amount of utility regardless of the time period, i.e. the concession schedule is linear (figure 2). Collaborative schedule (CL) implies making large concessions in the very beginning of the negotiation period in search of a quick agreement. This represents the case where an agent is anxious to sell the product. However, as the agent quickly drops the utility close to the reservation levels, it cannot make large concessions later in the process (figure 3).

Competitive-then-collaborative schedule (CC) models more complex behavior of the agents. In the beginning of the process an agent behaves competitively, however, in the middle of the negotiation period it changes its profile to a collaborative one. Thus, there is an inflexion point in an agent's schedule (figure 4).

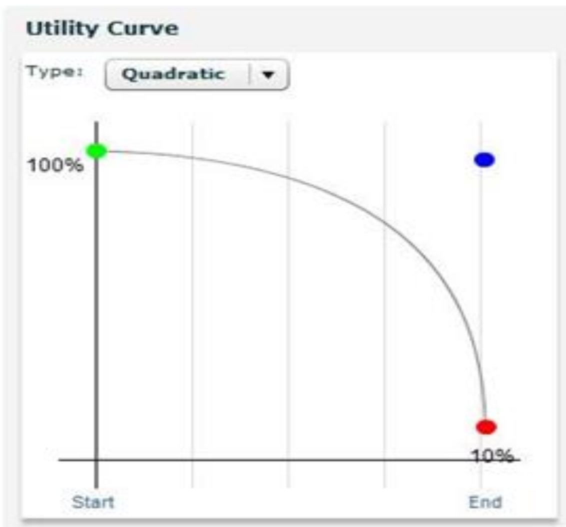

Fig. 1 Competitive schedule

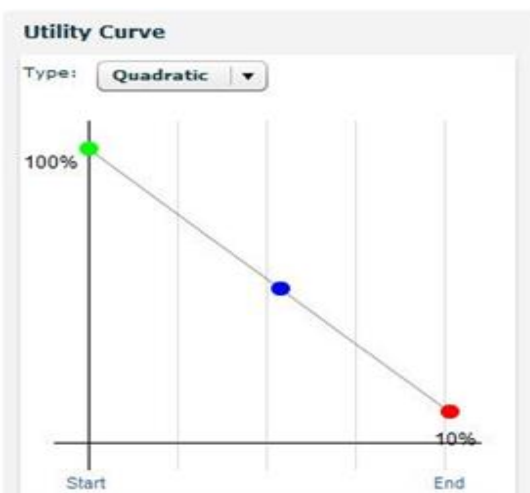

Fig. 2 Neutral schedule

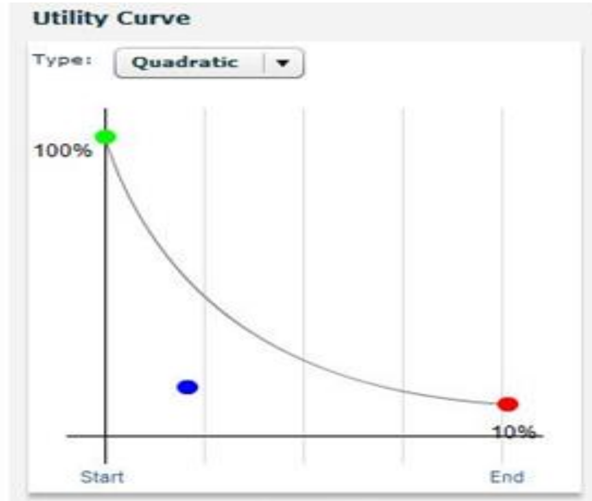

Fig. 3. Collaborative schedule

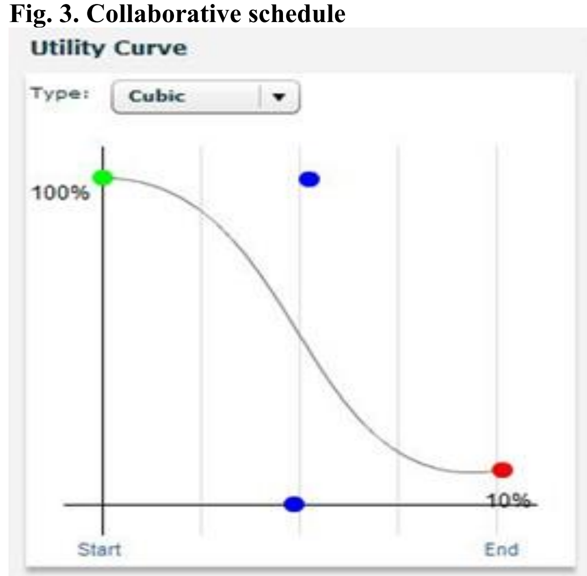

Fig. 4 Competitive-then-collaborative schedule

The reason for introducing this strategy is to imitate the situation when an agent's behavior adjusts due to the overall situation in the market (e.g. the product is not selling well). Moreover, the CC schedule allows introducing less predictable non-obvious behavior, which may be characteristic of human negotiators. (Little circles appearing on the screenshots are used to graphically define the shapes of the curves.)

The final (behavior-based) tactic used is tit-for-tat. These agents do not rely on utility calculations. Rather, they watch the opponent moves and simply mirror them in composing counter-offers. In other words, when an opponent makes a new offer an agent determines the difference between this offer and the previous one made by the opponent, and applies the same difference to its own offer. If, say an opponent made a large change to a price, the agent would do the same.

The agent follows the following algorithm. In the beginning of the process it makes an offer that has highest utility to an agent. It then waits for the opponent to respond. If an opponent agrees, the process terminates. If an opponent makes a counter-offer the agent calculates its acceptable utility level according to the concession schedule employed. If the opponent's offer is equal or higher than the acceptable utility, the agent accepts the offer. Otherwise, the agent generates a new offer according to the acceptable utility level. It takes the opponent's offer as a starting point, and employing hill-climbing algorithm changes it to get close to the set utility level. This heuristic method is used instead of analytical one, since most of the issues are not continuous variables. It then sends this offer to the opponent.

3.2. Experiments

As mentioned earlier, the purpose of the experiments was the assessment of the effectiveness of human-to agent negotiations. Thus, the objective variables measured in the first study included the utility of the agreements, and the proportion of agreements achieved. The subjects in the study were university students within business school enrolled in the introductory course on information technology. Thus, the negotiation case concerning the purchase of a computer was well in line with the learning objectives of the course.

Students were randomly assigned to various treatments, which included pairing up the subjects with various types of agents acting as sellers: collaborative, competitive, neutral, competitive vs. collaborative, and tit-for-tat. We also paired up humans with humans in a control group. Thus, most of the subjects acted as buyers, while some (in a control group) assumed the role of the sellers. Subjects participated only in a single assigned negotiation instance.

The experiment was conducted via the web, whereby subjects could perform their tasks from any location in an asynchronous mode during a two-day period. The subjects were invited to join the negotiations via email containing the link to the system. Subjects were instructed about their tasks, including the case, issues, and their importance, and the use of the system. The system's user interface was designed to emphasize ease of use.

Negotiations had specific start and end dates, of which the subjects were informed. In the course of negotiations subjects had to interact with a given counter-parts and could not switch to a different opponent (to keep the negotiations strictly bilateral). Subjects were also not aware of the outcomes of other negotiations. Furthermore, they were not told whether they were negotiating with a person or a software.

Negotiations began by sellers (agents and humans) making the first offer. The agent sellers then checked for the status of negotiations at fixed intervals of time (every 3 hours). At those points of time, if they have not received new offers, they would wait until the next period of time elapsed. If an offer was received they would evaluate it and would either accept it, or would make a counter-offer. Human sellers were free to check the status of negotiations and make offers at any time during the negotiation period.

Human subjects (buyers or sellers) were free to terminate the negotiation at any time without reaching an agreement with their counter-parts. After either reaching an agreement, or terminating the negotiations the human subjects were asked to complete a questionnaire measuring their perceptions of the outcome, process, and the system. Upon the completion of the experimental task the human negotiators (buyers) were invited to answer the following question: "I was negotiating with: 1) a human; 2) a computer: 3) not sure."

3.3 Results

For the analysis of the results we have selected only those negotiation instances, which featured at least four offers in total. The rationale for this decision was to include only those cases where the subjects took the task seriously. Thus, we ended up having 436 usable negotiation instances. Of these, 65% ended up in an agreement, while in 35% of cases the agreement was not reached.

Figure 5 shows the results of the question related to whether the participants guessed correctly if they were negotiating with humans or computers. The left side shows the results from human-agent dyads, and the right side shows human-human ones. The leftmost bar in each group indicates the number of responses that read "human", the middle one relates to "computer" responses, and the last one shows "not sure" responses.

As one can see, the majority of subjects in the agent-human dyads were not sure if they were interacting with the humans or agents (183 responses). This was followed by the group of subjects who had thought they were negotiating with other humans (114). The smallest group consisted of those who guessed correctly that they were interacting with agents (65). It is interesting to note that some subjects in the human-to-human dyads thought they were interacting with a computer (2 out of 30).

The distribution of answers depended on the type of the agent strategy employed. For example, in competitive-then-collaborative category much larger proportion of subjects thought they were negotiating with a human counter-part as compared to those who had an impression they were dealing with a machine (25 vs. 8). This can be explained by the fact that CC concession schedule results in more complex behavior, less obvious behavior that could be more readily ascribed to humans, rather than machines. Similar, though less prominent results were obtained in competitive agent category (33 vs. 15). On the other hand, the collaborative category was the only one where the number of "human" vs. "machine" responses was equal (21 each). Perhaps, the subjects expected their human counterparts to be more competitive, rather than conceding.

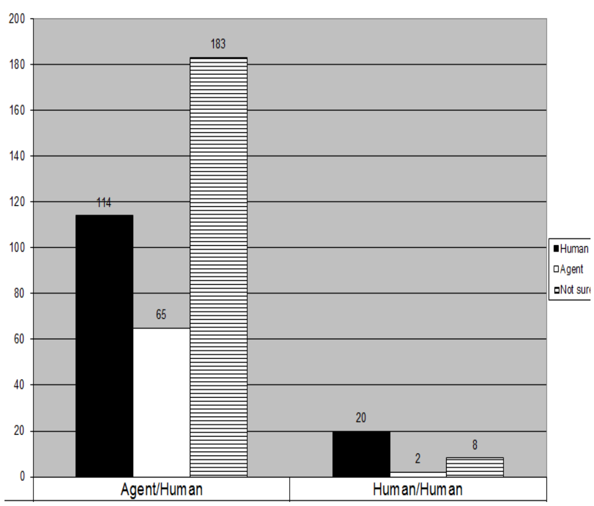

Fig. 5 "I was negotiating with…" agent - human dyads vs. human – human dyads ("I was negotiating with a human": solid fill; "an agent": light fill; "not sure": patterned fill).

Table 1 shows the proportions of agreements for different compositions of dyads. The largest proportion of agreements was reached in the collaborative agent category. This an intuitive result, since collaborative agents make large concessions early in the negotiations process, and thus they have a higher chance of making a deal with the human counterparts. It is interesting to see that human-to-human dyads have a second-lowest record in terms of proportion of agreements made. Thus, the majority of agent-involved dyads have reached more agreements than purely human dyads.

Competitive agents were able to reach an agreement in 53% of cases. Competitive-then-collaborative agents have made agreements in 75% of cases, falling between the CL and CM categories, but higher than neutral category. The lowest number of agreements was achieved in tit-for-tat category. This is the only agent strategy that does not employ utility function, and, thus it does not necessarily drop its utility level to the minimum towards the end of the period. Overall, agent-human pairs achieved agreements in 66% of cases vs. 50% exhibited by HH dyads.

Table 1. Proportions of agreements

Category	Agreements, %
All agent categories	66
Competitive	53
Neutral	70
Collaborative	82
Competitive-collaborative	75
Tit-for-tat	43
Human-human	50

Table 2 compares the utilities of reached agreements for sellers and buyers across different categories. In human-human dyads the sellers achieved much lower utility levels than buyers. This could be explained by the adopted reference frames. Since both sellers and buyers in this category were undergraduate student subjects, they tended to shift the price levels downwards to what they consider to be acceptable regions. Nonetheless, as it can be seen from the table, the human sellers had reached the lowest levels of utility.

The highest average utility was achieved by tit-for-tat agents (72.4). However, as already mentioned, they performed worst in terms of proportion of agreements reached. In terms of proportion of agreements the competitive agents have performed slightly better than human sellers. However, utility-wise these agents have considerably outperformed their human "colleagues" (63.2 vs. 35.9). Collaborative agents did only slightly better than humans, reaching 36.5 utility. However, they had much higher proportion of agreements. Competitive-then-collaborative agents have reached the average utility level of 40.4, and the neutral ones had a slightly higher value of 43.8. Overall, agents did better than human negotiators (46.8 vs. 35.9).

Table 2. Utilities of agreements

Category	Seller utility	Buyer utility
All agent categories	46.8	65.6
Competitive	63.2	44.9
Neutral	43.8	69.7
Collaborative	36.5	79.0
Competitive-collaborative	40.4	71.9
Tit-for-tat	72.4	36
Human-human	35.9	73.0

4. Multi-bilateral Negotiation Experiments

The first set of experiments involving bilateral negotiations showed that, overall, the agents have outperformed human sellers. However, in these settings the human counterparts had limited choice in regards with selecting the counterpart. The choice was restricted to negotiating with a given agent, or quitting. In reality, negotiators would normally have alternatives in choosing other sources or counter-parts.

In order to evaluate agent performance in more realistic and competitive environment one has to provide human subjects with choices. In this respect giving subjects the opportunity to consider offers from a variety of sources would render more realistic results concerning the performance of agent tactics.

To this end we have also conducted experiments in multi-bilateral settings. These experiments have differed from the first set in two important aspects. First, the negotiation case has been changed to better adapt to multi-bilateral scenario. While purchase of a computer is a good starting point for studying bilateral agent-human interactions, we felt that procurement scenarios are more suitable to business context for multi-bilateral case, whereby a buyer gets to consider alternative proposals from a number of sellers. Second, multi-bilateral scenario implies increased competition on the many side, we were mostly interested on the effectiveness of competitive vs. collaborative tactics employed by the agents. Thus, our second set of experiments had employed a procurement scenario and two types of tactics: collaborative vs. competitive.

4.1 Systems and Negotiation Case

The second set of experiments was a part of a larger study aimed at comparing auction and negotiation mechanisms. The case involved negotiation over the provision of transportation services to a milk producing company. The experiments were conducted with human participants using the Invite negotiation platform that allows flexible configuration of a range of different exchange mechanisms.

For the human negotiators a mechanism implementation called Imbins that allows for multi-bilateral multi-issue negotiations was provided using the Invite platform. The agents in the DIANA system could negotiate with counterparts by receiving and sending offers from/to the Invite system. The communication between the agents and the Invite platform was enabled via limited set of XML messages. Thus, the agents could negotiate through the Invite against and alongside the human negotiators (figure 6).

One important difference between the second and first study concerned the inclusion of the messaging capability. Since free-style messaging was allowed between human negotiators, we had to provide basic messaging mechanism for the agents as well. The mechanism consisted of rules with trigger conditions which included such criteria as time passed from last offer, time remaining to deadline, and others.

In the case scenario three negotiation issues were involved: standard rate (i.e. the normal rate for transporting a unit of product); rush rate (i.e. the rate at which rush orders are handled); and penalty for delay (i.e. penalty for not delivering in time, as the product is time-sensitive). Each of these issues had fifteen different option levels, for a total of 3,375 possible agreement alternatives.

There was a single buyer in the case, who needed transportation services for milk deliveries. There were three transportation companies who engaged in negotiations with the buyer. Thus, the case describes multi-bilateral multi-issue negotiation settings. The service sellers had somewhat different preference structures and reservation levels than the buyers.

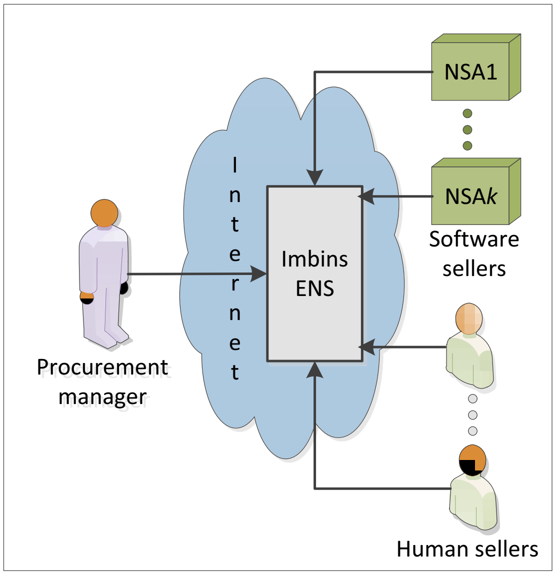

Figure 6. Mixed multi-bilateral e-negotiations via Imbins

The subjects were recruited among students from different countries, including Canada, Taiwan, Poland, Brazil, and Ukraine. In some experimental treatments all of the participants were humans. In others, some of the sellers were software agents equipped with the messaging capability. Two types of agents were employed: competitive and collaborative. They were configured in a way similar to one described in the previous section. Figure 7 shows the actual concession-making by the two types of agents recorded in the experiments along with the aggregated human behavior shown in terms of buyer utility. As one can see, they do follow competitive and collaborative schedules.

The sellers (both human and agent) were negotiating with one of the two type of buyers: collaborative vs. competitive. Tables 3 and 4 show the proportion of agreements reached by the agents and humans.

In total there were 116 multi-bilateral negotiations between two and four human sellers and between 0 and 2 NSAs. There were 363 human sellers competing against each other and 78 software agents. The negotiations took between 62.5 and 100.6 hrs., on average. The most effective and shortest negotiations involved four human sellers.

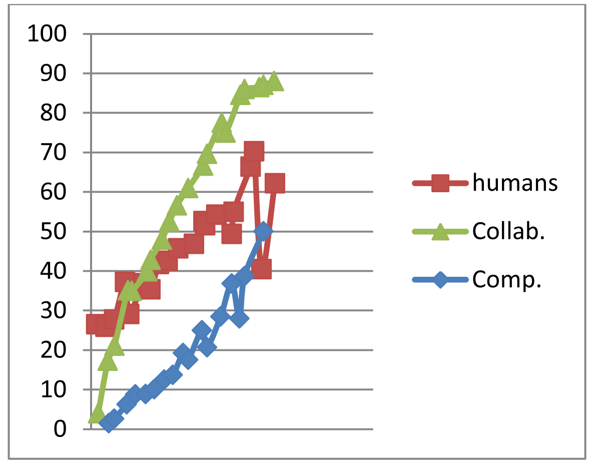

Figure 7. Concession-making by agents and humans

In all settings, only cooperative agents reached agreements, both in the 3+1 and 2+2 negotiations; in the latter there was one competitive and one cooperative NSA. The overall agreement rate for the agents (number of agreements by agents divided by the total number of agreements was 42%). However, competitive agents made no agreements. In the negotiations where cooperative agents participated they were able to reach 69% agreement rate, which is an impressive result for this type of agents, given that they were in competition with two or three human sellers. Furthermore, the cooperative agents were able to achieve the average seller profit of 23 compared to 13.3 reached by human negotiators. Also on average winner agents made more offers (6.9) as compared to humans (4.4).

Table 3. Proportion of agreements in negotiation with cooperative buyers

	Collab. agent	Comp. agent	Two agents
Human	4	10	3
%	36.4%	100.0%	27.3%
Agent	7	0	8
%	63.6%	0.0%	72.7%

Table 4. Proportion of agreements in negotiation with competitive buyers

	Collaborative agent	Competitive agent
Human	3	11
%	30.0%	100.0%
Agent	7	0
%	70.0%	0.0%

5. Conclusions

The purpose of this paper was to experimentally assess the prospects of agent/human negotiations. In the first study we designed negotiating software agents capable of participating in bilateral negotiations. We also conducted mixed negotiation experiments: agents were negotiating with humans in bilateral settings. In the second study we conducted experiments involving multi-bilateral negotiations with agents competing against human sellers to win a contract.

Our findings indicate that in strictly bilateral negotiations software agents, overall, have outperformed human subjects in terms of agreement utilities and proportion of agreements. Moreover, competitive agents performed better than collaborative ones. However, as our second experiment revealed, when agents have to compete in multi-bilateral settings, competitive agents failed to win agreements, while the cooperative agents were able to secure a good portion of favourable contracts.

These findings suggest that, in general more adaptive, context-aware and knowledge-based agent solutions may be preferable to simple time-dependent agent tactics. While competitive agents did well in the bilateral settings, their counterparts had no other choices in terms of negotiation parties. In more competitive settings, though, this time-based robotic behavior by agents failed to produce desirable results, as agents were not responsive to the less certain dynamics of negotiations. The direction for future research would include designing adaptive negotiation agent solutions, whereby agents would be able to recognize current situation and choose their tactics accordingly.

10. References

[1] Kersten, G., and Noronha, S.J., "Www-Based Negotiation Support: Design, Implementation, and Use", Decision Support Systems, 25(1999, pp. 135-154.

[2] Faratin, P., Sierra, C., and Jennings, N.R., "Negotiation Decision Functions for Autonomous Agents", Robotics and Autonomous Systems, 24(3-4), 1998, pp. 159-182.

[3] Beam, C., and Segev, A., "Automated Negotiations: A Survey of the State of the Art", Wirtschaftsinformatik, 39(3), 1997, pp. 263-268.

[4] Jennings, N.R., Faratin, P., Lomuscio, A.R., Parsons, S., Wooldridge, M.J., and Sierra, C., "Automated Negotiation: Prospects, Methods and Challenges", Group Decision and Negotiation, 10(2), 2001, pp. 199-215.

[5] Chavez, A., Dreilinger, D., Guttman, R., and Maes, P., "A Real-Life Experiment in Creating an Agent Marketplace", in (Nwana, H.S., and Azarmi, N., 'eds.'): Software Agents and Soft Computing, Springer-Verlag, 1997, pp. 160-179.

[6] Maes, P., Guttman, R.H., and Moukas, A.G., "Agents That Buy and Sell", Communications of the ACM, 42(3), 1999, pp. 81-87.

[7] Wang, M., Wang, H., Vogel, D., Kumar, K., and Chiu, D.K.W., "Agent-Based Negotiation and Decision Making for Dynamic Supply Chain Formation", Engineering Applications of Artificial Intelligence, 22(7), 2009, pp. 1046-1055.

[8] Huang, C.-C., Liang, W.-Y., Lai, Y.-H., and Lin, Y.-C., "The Agent-Based Negotiation Process for B2c E-Commerce", Expert Systems with Applications, 37(1), 2010, pp. 348-359.

[9] Huang, S.-L., and Lin, F.-R., "The Design and Evaluation of an Intelligent Sales Agent for Online Persuasion and Negotiation", Electronic Commerce Research and Applications, 6(3), 2007, pp. 285-296.

[10] Lin, R., and Kraus, S., "Can Automated Agents Proficiently Negotiate with Humans?", Communications of the ACM, 53(1), 2010, pp. 78-88.

[11] Chen, E., Vahidov, R., and Kersten, G.E., "Agent-Supported Negotiations in the E-Marketplace", International Journal of Electronic Business, 3(1), 2005, pp. 28-49.

[12] Vahidov, R., "Situated Decision Support Approach for Managing Multiple Negotiations", in (Gimpel, H., Jennings, N.R., Kersten, G.E., Ockenfels, A., and Weinhardt, C., 'eds.'): Negotiation, Auctions, and Market Engineering, Springer Berlin Heidelberg, 2008, pp. 179-189.

[13] Yang, Y., See, Y., Ortony, A., and Tan, J., "Subjective Effectiveness in Agent-to-Human Negotiation: A Frame X Personality Account": Argumentation in Multi-Agent Systems, Springer Berlin / Heidelberg, 2010, pp. 134-149.

[14] Raiffa, H., Richardson, J., and Metcalfe, D., Negotiation Analysis. The Science and Art of Collaborative Decision Making,, Harvard University Press, Cambridge, 2003.