Analysis of Negotiation Dynamics

Koen Hindriks, Catholijn M. Jonker, and Dmytro Tykhonov

Delft University of Technology, Man-Machine Interaction, Mekelweg 4,
2628CD Delft, The Netherlands
{K.Hindriks,C.Jonker,D.Tykhonov}@tudelft.nl

Abstract. The process of reaching an agreement in a bilateral negotiation to a large extent determines that agreement. The tactics of proposing an offer and the perception of offers made by the other party determine how both parties engage each other and, as a consequence, the kind of agreement they will establish. It thus is important to gain a better understanding of the tactics and potential other factors that play a role in shaping that process. A negotiation, however, is typically judged by the efficiency of the outcome. The process of reaching an outcome has received less attention in literature and the analysis of the negotiation process is typically not as rigorous nor is it based on formal tools. Here we present an outline of a formal toolbox to analyze and study the dynamics of negotiation based on an analysis of the types of moves parties to a negotiation can make while exchanging offers. This toolbox can be used to study both the performance of human negotiators as well as automated negotiation systems.

1 Introduction

Negotiation is an interpersonal decision-making process necessary whenever we cannot achieve our objectives single-handedly [10]. Parties to a negotiation need each other to obtain an outcome which is beneficial to both and is an improvement over the current state of affairs for either party. Both parties need to believe this is the case before they will engage in a negotiation. Although by engaging in a negotiation one party signals to the other party that there is potential for such gain on its side, it may still leave the other party with little more knowledge than that this is so. Research shows that the more one knows about the other party the more effective the exchange of information and offers [9]. Furthermore, humans usually do have some understanding of the domain of negotiation to guide their actions, and, as has been argued, a machine provided with domain knowledge may also benefit from such domain knowledge [3].

It is well-known that many factors influence the performance and outcome of humans in a negotiation, ranging from the general mindset towards negotiation to particular emotions and perception of fairness. As emphasized in socio-psychological and business management literature on negotiation, viewing negotiation as a *joint problem-solving task* is a more productive mindset than viewing negotiation as a *competition* in which one party wins and the other looses [4, 9, 10]. Whereas the latter

M. Klusch et al. (Eds.): CIA 2007, LNAI 4676, pp. 27–35, 2007.

mindset typically induces hard-bargaining tactics and rules out disclosure of relevant information to an opponent, the former leads to joint exploration of possible agreements and induces both parties to team up and search for trade-offs to find a win-win outcome. Different mindsets lead to different negotiation strategies. A similar distinction between hard- and soft-bargaining tactics has also been discussed in the automated negotiation system literature where the distinction has been referred to as either a *boulware* or a *conceder* tactics [2].

Emotions and perception of fairness may also determine the outcome of a negotiation. People may have strong feelings about the "rightness" of a proposed agreement. Such feelings may not always be productive to reach a jointly beneficial and efficient agreement. It has been suggested in the literature to take such emotions into account but at the same time to try to control them during negotiation and rationally assess the benefits of any proposals on the table [4, 10].

Apart from the factors mentioned above that influence the dynamics of negotiation, many other psychological biases have been identified in the literature that influence the outcome of a negotiation, including among others partisan perceptions, overconfidence, endowment effects, and reactive devaluation [8, 10].

In order to gain a better understanding of the negotiation dynamics and the factors that influence the negotiation process it is crucial to not only mathematically evaluate the efficiency of negotiation outcomes but also to look at the pattern of offer exchanges, what Raiffa [8] calls the *negotiation dance*. In the remainder we present part of a formal toolbox to analyze patterns in offer exchanges and present some initial findings in the literature.

2 Towards a Formal Toolbox for Negotiation Dynamics Analysis

The insights of which factors influence the negotiation process as well as outcome as described in the previous section were gained by means of experiments performed e.g. by psychologists, and social scientists. More recently, the development of automated negotiation software has provided a basis to experiment and collect data about the negotiation process through human-computer interaction [1, 7]. Here we introduce part of a toolbox that allows formal analysis of the negotiation dynamics in experiments with humans as well as with machines.

Our interest is in analyzing, classifying and in precisely characterizing aspects of the negotiation dynamics that influence the final agreement of a negotiation. The main interest thus is in proposing concepts and metrics that relate these factors to specific aspects of the negotiation dynamics and to thus also gain a better understanding of the final outcome of a negotiation.

The key concept in the analysis toolbox that we propose is that of various categories or classes of negotiation actions, including in particular the offers made by each party. A proposed offer can be classified based on the utility it provides to the proposing party ("Self") as well as to the other party ("Other"). The possible classifications are visualized in Figure 1.

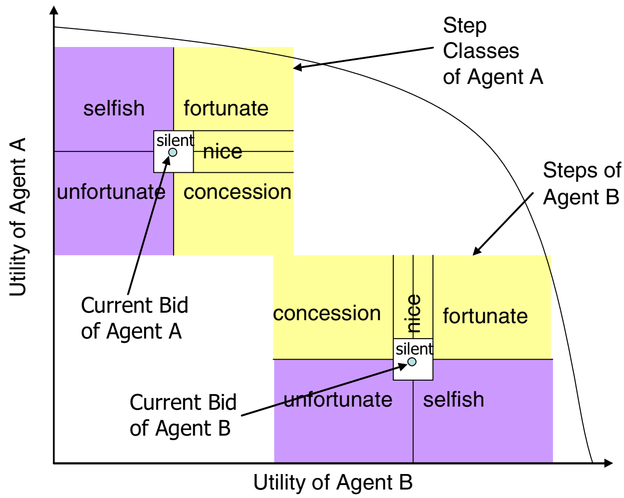

Fig. 1. Visualization of step classes in common outcome space

We distinguish six negotiation step classes, which are formally defined below. Before formally defining the concepts below, some additional notation is introduced. $U_S(b)$ denotes the utility of "Self" with respect to bid *b*. Similarly, $U_O(b)$ denotes the utility of "Other" with respect to *b*. We use $\Delta_a(b, b') = U_a(b')\text{-}U_a(b), a\in\{S,O\}$, to denote the utility difference of two bids *b* and *b'* in the utility space of agent *a*. We also write $\Delta_a(s)$ to denote $\Delta_a(b, b')$ for a step $s = b \rightarrow b'$. Here we present a precise definition of the classes of negotiation steps proposed in [1] extended as discussed above. These step categories define the core of the step-wise analysis method [5].

Definition of Step Classes:

Let s=$b_S \rightarrow b'_S$ be a step in the bidding by Self (the definition for Other is completely symmetric). Then the negotiation step s taken by Self is classified as a:

- *Fortunate Step*, denoted by (S+, O+), iff:
  $\Delta_S(s)>0$, and $\Delta_O(s)>0$.
- *Selfish Step*, denoted by (S+, $O_\leq$), iff:
  $\Delta_S(s)>0$, and $\Delta_O(s)\leq 0$.
- *Concession Step*, denoted by (S-, $O_\geq$), iff:
  $\Delta_S(s)<0$, and $\Delta_O(s)\geq 0$.
- *Unfortunate Step*, denoted by ($S_\leq$, O-), iff:
  $\Delta_S(s)\leq 0$, and $\Delta_O(s)<0$.
- *Nice Step*, denoted by (S=, O+), iff:
  $\Delta_S(s)=0$, and $\Delta_O(s)>0$.
- *Silent Step*, denoted by (S=, O=), iff:
  $\Delta_S(s)=0$, and $\Delta_O(s)=0$.

Observe that the proposed classification is exhaustive, and all step classes are disjoint. These step classes can be used to define additional concepts to analyze the negotiation dance in a particular negotiation. For illustrative purposes, we present just a few additional concepts. For a more extensive overview we refer the reader to [5].

A trace t is a series of negotiation steps as defined above, i.e., transitions $b \to b'$ with b, b' offers. For a given trace the percentage of steps in a particular step class is defined as usual.

Definition. % per Class

The percentage $\%_c(t)$ of class c steps in a trace t is defined by: $\%_c(t) = \#t_c / \#t$.

Negotiation strategies can be designed with specific aims in mind that should be observable as patterns in the negotiation dance. For example, the success of a strategy that is supposed to learn its opponent's preferences can be verified by checking whether the frequency and/or size of unfortunate steps over a negotiation trace decreases. Such patterns can be seen as a measure of *adaptability of a party to its opponent*. Another useful measure of the *sensitivity to the opponent's preferences* can be defined by comparing the percentage of *fortunate, nice and concession steps* that increase the opponent's utility to the percentage of *selfish, unfortunate and silent steps* that decrease it. Intuitively, an agent that only performs steps that increase its opponent's utility can be said to be (very) sensitive to the needs of its opponent.

Definition. Sensitivity to Opponent Preferences

The measure for sensitivity of agent a to its opponent's preferences is defined for a given trace t by:

$$\text{sensitivity}_a(t) = \frac{\%_{Fortunate}(t_a) + \%_{Nice}(t_a) + \%_{Concession}(t_a)}{\%_{Selfish}(t_a) + \%_{Unfortunate}(t_a) + \%_{Silent}(t_a)}$$

In case no selfish, unfortunate or silent steps are made we stipulate that $sensitivity(a,t)=\infty$. If $sensitivity_a(t)<1$, then an agent is more or less insensitive to opponent preferences; if $sensitivity_a(t)>1$, then an agent is more or less sensitive to the opponent's preferences, with complete sensitivity for $sensitivity(a,t)=\infty$. Typically, this sensitivity measure varies with different domains and different opponents and averages over more than one trace need to be computed. Note that the notion of sensitivity is *asymmetric*: one agent may be sensitive to the other's preferences, but not vice-versa.

3 Experimental Results

In this section, we present some experimental findings to illustrate the usefulness of our analysis toolbox.

Bosse and Jonker [1] performed two experiments with human subjects. The negotiation dances produced were analyzed with the step analysis method, although silent steps and nice steps were not considered as special cases of the concession step.

In the first experiment eighteen subjects participated and consisted of AI students in The Netherlands (12 males and 6 females). Their age varied between 19 and 27 years. The participants had to negotiate against each other (refered to as HH negotiations) and were motivated by the challenge to obtain the highest utility. Furthermore, they were challenged to outperform the computer in the corresponding Computer-Computer negotiation process (CC) they were also allowed to perform. The computer agent used the ABMP negotiation strategy, see [6]. In the HH process, one person is assigned the role of the buyer, and the other one is assigned the role of the seller. In the CC process, a computer buyer negotiates with a computer seller, both using the profile of the corresponding human negotiator. By keeping the negotiation profile stable over the two processes, it is guaranteed that the utility spaces remains the same, and that the resulting traces are thus comparable.

In the second experiment 76 subjects (43 males and 33 females) participated. The experiment took place during an introductory course for family members of AI students. Most of the participants (about 75%) were parents of the students, their age varying between 45 and 55 years. The other 25% were brothers and sisters of the students, their age varying between 17 and 24 years. Almost all of the participants did not have any background in AI. Education and occupation were a fair representation of the general population in The Netherlands.

The participants formed 38 teams of two persons, and each team was assigned to a computer. Each team was told that they could negotiate as a team against the computer. This deviation was necessary for that occasion, due to a lack of available computers. Each group participated in two negotiation processes: a Human-Computer (HC) process and a CC process. In the HC process, all teams played the buyer role, and use their own personal profile. Computer roles used the ABMP strategy of [6]. In the CC process, a computer buyer uses the profile of the human team. By keeping the negotiation profile stable over the two processes, it is guaranteed that the utility spaces remain the same, and that the resulting traces are thus comparable.

One trend observed in both experiments, is that the Nash distance and the EPP distance (both measures for fairness of the negotiation) were significantly shorter in the CC traces than in the HH traces, see Table 1, and Table 3 shows that they were shorter in the CC traces than in the HC traces. Furthermore, these distances seem to be shorter in the HH traces than in the HC traces. Thus, the CC negotiations turned out to have the "fairest" outcome, followed by the HH traces. The outcomes of the HC traces were the least balanced. This can be seen in the first two cells in Table 3, where the mean (human) buyer utility (0.89) was much higher than the mean (computer) seller utility (0.72). This is an important finding, because when the same negotiation spaces are explored by two computer negotiators, the buyer utility hardly drops (0.87), whilst the seller utility increases significantly (0.83).

Table 1. Performance in Experiment 1

	Buyer Utility	Seller Utility	Pareto Distance	Nash Distance	EPP Distance	Number of rounds
HH traces	0.87	0.80	0.05	0.22	0.16	7.00
CC traces	0.88	0.89	0.03	0.12	0.06	8.00
t-value	0.376	2.807	-0.786	-3.988	-3.463	1.540
p-value	0.717	0.023	0.455	0.004	0.009	0.146

Table 2. Steps made in Experiment 1

	Fortunate (S+ O+)	Concession (S- O+)	Selfish (S+ O-)	Unfortunate (S- O-)
HH, buyer	3 (6.52%)	36 (78.26%)	2 (4.35%)	5 (10.87%)
HH, seller	5 (11.36%)	32 (72.73%)	4 (9.09%)	3 (6.82%)
CC, buyer	0 (0%)	58 (89.23%)	0 (0 %)	7 (10.77 %)
CC, seller	0 (0 %)	48 (82.76%)	0 (0 %)	10 (17.24 %)

Table 3. Performance in Experiment 2

	Buyer Utility	Seller Utility	Pareto Distance	Nash Distance	EPP Distance	Number of rounds
HC traces	0.89	0.72	0.05	0.30	0.23	8.84
CC traces	0.87	0.83	0.06	0.17	0.10	8.91
t-value	-1.729	3.684	0.309	-5.161	-6.228	0.066
p-value	0.092	0.001	0.759	0.000	0.000	0.948

Table 4. Steps made in Experiment 2

	Fortunate (S+ O+)	Concession (S- O+)	Selfish (S+ O-)	Unfortunate (S- O-)
HC, buyer	23 (7.62%)	232 (76.82%)	17 (5.63%)	30 (9.93%)
HC, seller	2 (0.68%)	251 (85.37%)	0 (0 %)	41 (13.95%)
CC, buyer	0 (0 %)	287 (94.41 %)	0 (0 %)	17 (5.59 %)
CC, seller	0 (0 %)	267 (90.51 %)	0 (0 %)	28 (9.49 %)

Apparently the ABMP strategy used by the computer seller is not robust to being exploited by a human buyer. This observation is supported by the data in Tables 2 and 4. In both situations, the computers made more unfortunate steps than the humans. In addition, the computer sellers made more unfortunate steps than the computer buyers.

A last important finding concerns the diverse bidding behaviour by humans. As shown in Table 2, human negotiators sometimes make steps that improve the utility for both parties. Of course, doing this has the risk of making selfish steps. In its current state, the ABMP agent hardly makes these kinds of steps. Nevertheless, in some cases the unpredictable human behavior actually resulted in better results.

The extended step-wise analysis technique has been applied to a number of negotiation strategies for software agents [5]. A tournament with the strategies was set up and run. The following strategies have been studied: The ABMP strategy [6], a concession oriented strategy, which computes bids to offer next without taking domain or opponent knowledge into account, the Trade-off strategy is based on similarity criteria [3], and exploits domain knowledge to stay close to the Pareto Frontier. The

Random Walker strategy serves as a “baseline” strategy. It randomly jumps through the negotiation space, and can be run with or without a break-off point.

A full analysis was made of the type of steps made on AMPO vs City domain [9]. Here we show some examples of negotiation dances typical for the negotiation strategies we selected for the tournament (see figure 2).

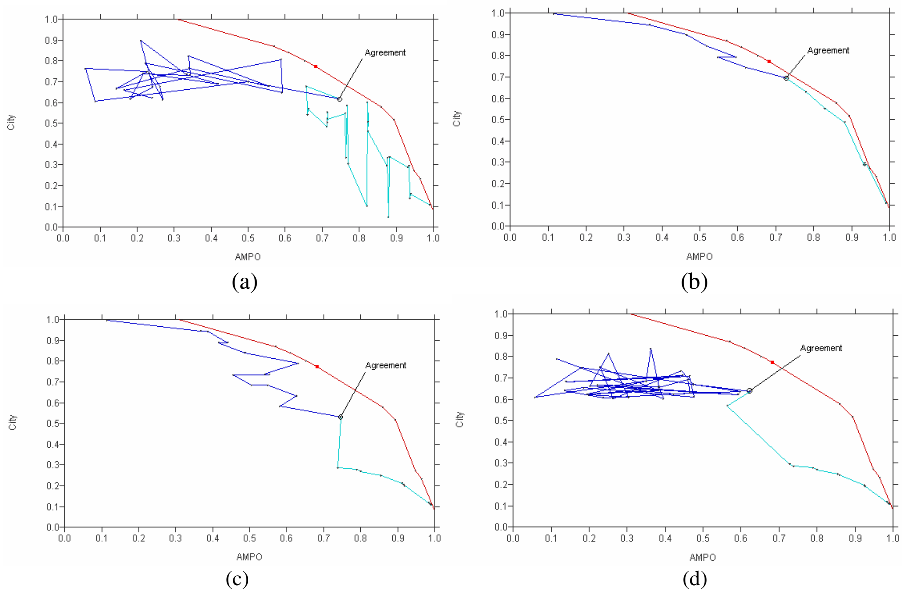

Fig. 2. Dynamics of negotiation process for: a) Random Walker (City) vs Trade-Off strategy (AMPO), b) Trade-Off (City) vs Trade-Off strategy (AMPO), c) Trade-Off (City) vs ABMP strategy (AMPO), d) Random Walker (City) vs ABMP strategy (AMPO)

The Trade-off strategy uses bids sent by its opponent to estimate what is the best possible trade-off between issues using similarity criteria. It assumes rationality of the opponent’s strategy, which means that values of the issues in the opponent’s bid represent its preferences. The Random-Walker strategy (fig. 2a), however, does not match this assumption and, as a result, the Trade-Off is rather strongly affected by it and performs multiple unfortunate steps.

Figure 1b shows that the Trade-Off strategy performs well against itself because both opponent’s use the same assumption of rationality of the negotiation strategies. In addition, good predictability[1] of the issues produces only few unfortunate steps. Here unfortunate step is a result of a mismatch in the weights of the similarity functions and actual weights of the opponent’s issues cause wrong trade-offs between issues and results in unfortunate steps.

ABMP strategy has very similar negotiation paths in all experiment due to its concession strategy and high dimensionality (10 issues) of the negotiation space. ABMP

[1] Here predictability of issue means that a distance function on the values of this issue can be easily defined using common sense knowledge. I.e., price issue has good predictability ($10 is most likely to be closer to $20 than to $30) and color has poor predictability.

strategy always concedes on every issue that prevents it from making unfortunate steps in case of strong mismatch in issues weights between opponents. From the other side, such concession algorithm does not allow trade-offs between issues and thus brings the bids further from the Pareto-frontier (see fig. 1c). This effect has an impact on the Trade-off strategy causing multiple unfortunate steps.

ABMP is not affected by the opponent's strategy because of its concession algorithm. However, it is essential for a negotiation strategy efficiency to be able to make trade-offs between issues that are unequally important for the opponents. Thus, such concession strategy pushes bids away from the Pareto-frontier.

4 Conclusion

It is important to gain a better understanding of the negotiation dance, the exchange of offers between parties, in a more formal way. In order to do so a toolbox for analyzing this exchange needs to be developed. Such a toolbox, elements of which were outlined in this extended abstract, may provide a basis for relating and explaining the moves of negotiating parties to five key factors that shape the outcome of a bilateral negotiation with incomplete information: (i) knowledge about the negotiation domain (e.g. the market value of a product or service), (ii) one's own and one's opponent's preferences, (iii) process attributes (e.g. deadlines), (iv) the negotiation strategies, and (v) the negotiation protocol.

Many challenges remain for developing the toolbox envisaged. One important extension of the toolbox is to introduce *benchmark problems* for bilateral negotiation that can be used to evaluate automated negotiation systems. Additional experimental data is required to refine the concepts and to develop new concepts that need to be included in the toolbox. Also such data may provide insights into relating experimental results to the key factors identified above.

References

[1] Bosse, T., Jonker, C.M.: Human vs. Computer Behaviour in Multi-Issue Negotiation. In: Proceedings of the 1st International Workshop on Rational, Robust, and Secure Negotiations in Multi-Agent Systems, pp. 11–24. IEEE Computer Society Press, Los Alamitos (2005)

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

[3] Faratin, P., Sierra, C., Jennings, N.: Using Similarity Criteria to Make Negotiation Trade-Offs. Journal of Artificial Intelligence 142(2), 205–237 (2003)

[4] Fisher, R.R., Ury, W.L.: (and for the latest edition B. Patton), Getting to Yes: Negotiating Agreement Without Giving In. Penguin Books, 1981, 1992 (2003)

[5] Hindriks, K.V., Jonker, C.M., Tykhonov, D.: Negotiation Dynamics: Analysis, Concession Tactics, and Outcomes, Submitted to: 2007 IEEE/WIC/ACM International Conference on Intelligent Agent Technology IAT (2007)

[6] Jonker, C.M., Treur, J.: An Agent Architecture for Multi-Attribute Negotiation. In: Proc. of the 17th Int. Joint Conference on AI (IJCAI'01), pp. 1195–1201 (2001)

[7] Lin, R., Kraus, S., Wilkenfeld, J., Barry, J.: An Automated Agent for Bilateral Negotiation with Bounded Rational Agents with Incomplete Information. In: Proceedings of the 17th European Conference on Artificial Intelligence (ECAI'06), pp. 270–274 (2006)
[8] Mnookin, R.H., Peppet, S.R., Tulumello, A.S.: Beyond Winning: Negotiating to Create Value in Deals and Disputes. Harvard University Press, Cambridge (2000)
[9] Raiffa, H., Richardson, J., Metcalfe, D.: Negotiation Analysis: The Science and Art of Collaborative Decision Making. Harvard University Press, Cambridge (2002)
[10] Thompson, L.L.: The Heart and Mind of the Negotiator. Pearson Prentice Hall (2005)