

Let's dans! An analytic framework of negotiation dynamics and strategies

Koen Hindriks, Catholijn M. Jonker* and Dmytro Tykhonov

Man-Machine Interaction Group, Delft University of Technology, Mekelweg 4, 2628 CD, Delft, The Netherlands

Abstract. The “negotiation dance”, as Raiffa calls the dynamic pattern of the bidding, has an important influence on the outcome of the negotiation. The current practice of evaluating a negotiation strategy is to focus on fairness and quality aspects of the agreement. In this article we present the framework DANS (Dynamics Analysis of Negotiation Strategies) for the analysis of the dynamic patterns of the bidding as a means to evaluate the strengths and weaknesses of negotiation strategies for bidding. The method provides the tools to perform a detailed and quantified analysis of a negotiation between two agents in terms of dynamic properties of the negotiation trace. The classification of negotiation steps in the dance plays a central role in the analysis. The method can be applied to tournaments, but can also be used to analyze single 1-on-1 negotiation sessions. The sessions can be played by humans or by software agents. Using DANS we show that some strategies are sensitive to the bidding behaviour of the opponent, and some depend on a correct model of the opponent. DANS helped us in discovering that domain characteristics are important for the analysis of strategies. Some strategies rely heavily on some domain assumptions. Furthermore, the results illustrate that having domain knowledge is not always enough to avoid making unintentional steps. The method is demonstrated in the analysis of three strategies from the literature ABMP, Trade-Off and Bayesian Agent.

Keywords: Automated negotiation, negotiation dynamics, multi-issue negotiation, dynamic properties, analysis

1. Introduction

Negotiation is a type of interaction process between two or more self-interested agents (each with its own aims, needs, and viewpoints) seeking to discover a common ground and to reach an agreement to settle a matter of mutual concern or resolve a conflict (cf. [17]). The negotiation process is also known as the *negotiation dance* [17].

The negotiation dance of exchanging successive offers by negotiating parties affects the negotiation outcome [17]. To gain more insight in the negotiation dynamics, in [1] a classification of negotiation moves was introduced in order to characterize and compare the bidding process of humans and software agents. The results show an overall similarity of the bidding style of humans and the Agent-Based Market Places (ABMP) strategy, a concession-oriented negotiation strategy, see [12]. However, the analysis did not provide insights in why the different kinds of moves were made, nor did it help us understand why and to

what extend these moves affect the outcome of the negotiation. As far as we know, no analytical methods exist that do provide the desired insights.

The analysis method introduced in this paper is a concrete step towards providing such insights. It extends the work presented in [1] by extending and providing a precise characterization of the negotiation move classification and by providing some useful metrics. These metrics in turn are used to define more complex dynamic properties of the negotiation dance to facilitate the analysis of various dynamic properties of the strategies under evaluation. By *negotiation strategy* in this paper we mean the decision mechanism for determining the actions of a negotiator.

Other analytical methods used in the literature typically assess the performance of negotiation strategies in terms of fairness and quality aspects of the agreement (if any) that agents reach. Aspects considered are who wins, the distance of the outcome to the Pareto Efficient Frontier, the Nash Product,

* Corresponding author. E-mail: c.m.jonker@tudelft.nl.

and e.g. the Kalai-Smorodinsky Point (see Section 2.1 for details). Formal definitions of these concepts can be found in e.g., [17]. Such measures of evaluation focus on the *negotiation outcome*.

Instead, the concepts introduced here are intended to facilitate the analysis of typical bidding patterns induced by various negotiation strategies. It is the objective of this paper to propose a method and some metrics that facilitate a *precise characterization* of the negotiation dance. In turn, such a characterization of the dynamics of negotiation may contribute to the identification of explanations for such findings. It is the aim of this paper to at least *partially* identify some of the *reasons* that may explain particular findings, that is, to associate particular aspects of a negotiation problem or strategy with particular extreme values (e.g., minimum or maximum) of the metrics defined below.

We illustrate the use of these concepts for the analysis of *concession tactics*. For example, although it is generally acknowledged that a concession should actually *increase* the utility of the opponent and not just be a move that *decreases* one's own utility, in practice, as we will show, such behaviour is not always achieved by strategies that have been designed to concede towards the opponent. For example, in the well-known case study of the negotiation between Associated Metropolitan Police Officers (AMPO) and New-York city presented by Raiffa in [17] both parties had strongly conflicting preferences on almost all issues. Only with regards to the issue "the fate of the police commissioner, Mr. Daniels" both negotiating parties had the same preferences: they wanted to fire the commissioner. It is clear that if one of the parties would make a concession on that issue both parties would lose rather than gain utility. Such moves that reduce both the agent's own as well as its opponent's utility have been called *unfortunate moves* (cf. [1]). Both humans as well as software agents using the ABMP strategy were observed to make such moves in negotiation experiments reported in [1], but humans made fewer of them.

Another aspect that plays a central role in DANS is the analysis of the domain of negotiation. A negotiation takes place in a certain domain, e.g., real-estate, job negotiation, etc. A *negotiation domain* is a detailed description of a conflict to be resolved by a negotiation. Typically, a negotiation domain is represented by a set of negotiation issues. A *negotiation issue* is a topic of discussion that is of a particular interest in a negotiation, e.g., price, quantity, delivery date, etc. In all informal literature on negotiation it is stressed that the negotiator should prepare with

respect to the domain, the opponent, and the negotiator's own preferences. However, the literature on automated negotiations, the aspect of the domain leaves room for improvement. Some strategies use domain knowledge, but no formal analysis has been made of domains in terms of characterizing properties. In this paper we introduce a number of characterizing properties that proved useful in understanding the strengths and weaknesses of bidding strategies.

The analytical method that we propose in this paper combines the standard analysis described in [18] and the analysis of negotiation dynamics proposed by the authors in [8] to give a better understanding of a negotiation strategy performance. The method takes into account negotiation factors that were introduced in our earlier work [10] that influences the negotiation performance to gain a better understanding of when a negotiation strategy is applicable. In addition, from the analysis presented in this paper we identified a number of design guidelines for developing an efficient negotiation strategy and used them in the design of a negotiation strategy presented in [9].

The paper is organized as follows. The next section discusses related work. Section 3 discusses negotiation factors that are included in the proposed analysis method. In Section 4, we briefly introduce the topic of negotiation dynamics. Section 5 introduces the move-based analysis method and some metrics for analyzing dynamic negotiation properties. Section 6 explains phases of the analysis method. In Section 7, the method is illustrated by analyzing the Trade-Off [5], ABMP [12], and Bayesian learning [9] strategy in various negotiation domains. Finally, the paper concludes with some suggestions for research on automated negotiation derived from the proposed analysis method.

2. Related work

The scope of the current paper concerns the negotiation dynamics as a pattern of offers (cf. [17]). That is, our work concerns bargaining (see e.g., [9,11]), a method for reaching joint agreements by means of exchanging offers according to for example, an alternating offers protocol. With the exception of [1] and [2] all papers in the literature that discuss the quality measures of negotiation strategies, concern outcome analysis, not the negotiation dance. Such an approach has been adopted by negotiation strategies designers to evaluate the quality of proposed strategies, see e.g., [3–5,12–14,19,24]. The evaluation method used in

this literature mainly focuses on the comparison of utility values of the outcomes and number of rounds required to reach an agreement.

This section starts with a summary of measures used for outcome analysis. This work is used within DANS and is complementary to the work presented in later sections of this paper to analyse the dynamics of negotiations.

In order to develop efficient negotiation strategies that are robust against and outperform other strategies, it is important to be able to evaluate the dynamic behaviour induced by negotiation strategies. Therefore, Section 2.2 focuses on the literature regarding aspects useful for the analysis of dynamic patterns of negotiations.

2.1. Optimal solutions and performance metrics

The most common outcome performance metrics (see Fig. 1) used to determine the quality of an agreement with respect to each of the players include the distance to the Pareto Efficient Frontier, the Nash Product, and the Kalai-Smorodinsky outcome, see e.g., [2] and e.g., [17] for formal definitions. Other global measures taken are the (average) number of negotiation rounds R needed to reach an agreement, the number of agreements A reached in a tournament, and the time T taken by each party.

Pareto frontier. A bid is Pareto efficient if given a set of alternatives, no movement from the bid to an alternative exists that can make at least one individual better off without making any other individual worse off. Typically, there exist multiple Pareto optimal solutions. The set of such solutions is called a Pareto frontier.

Nash product. The Nash product is that outcome that maximizes the product of the utilities of the parties. Nash product satisfies certain axioms. It is an invariant to affine transformations, independent to irrelevant alternatives. The Nash solution is always Pareto optimal. In addition, the Nash solution is symmetrical, meaning that if both players have the same utility functions, then symmetry demands that both get equal payoffs.

Kalai-Smorodinsky. The Kalai-Smorodinsky outcome is that point on the Pareto frontier which maintains the ratios of maximal gains. In other words, assuming that the utility functions of the parties are normalized and map into the interval $[0;1]$, the Kalai-Smorodinsky solution is a Pareto efficient outcome that has equal utilities of the two parties.

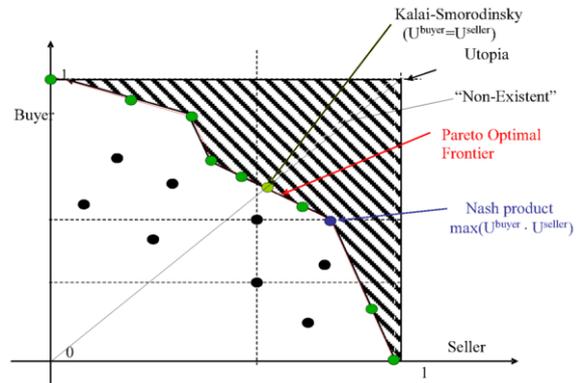


Fig. 1. Standard analysis.

The authors of [14] add the following properties to the usual outcome properties:

- Social welfare: the sum of the utilities of the negotiators for the agreement should be as high as possible.
- Invariance: the solution is invariant under the application of positive affine transformations on the utility functions of both agents.

Nongaillard and co-authors consider not only the here mentioned utilitarian social welfare, but also egalitarian social welfare, see [23].

In [15], a classification scheme is provided that defines some properties that are oriented towards rationality and the use of resources. They identify the following desirable properties for a negotiation protocol and strategies: computational efficiency, communication efficiency, individual rationality, distribution of computation, Pareto efficiency, and symmetry of power between agents. They provide characteristics useful for negotiation system design:

- Cardinality: number of issues, and one-to-one, many-to-one, many-to-many negotiators.
- Agent characteristics: the role it plays (buyer, seller, intermediary), its rationality (perfect or bounded), its knowledge about other agents' preferences, its social behaviour (self-interested vs. altruistic), and its bidding strategy.
- Information parameters: the value of goods (public / private), the nature of goods (discrete / continuous), price quotes, and transaction history.
- Event parameters: validity of bids, visibility of bids (not for one-to-one negotiations), clearing schedule with allocation parameters, timeouts, and a quotes schedule.

2.2. Negotiation dance literature

In [1] and [2], a formalization of the negotiation process is provided together with a set of performance properties that facilitate evaluation of the quality of the agreement reached, based on the work of [18] and [20]. The paper also discusses some dynamic properties of the bidding. The authors used the SAMIN system to analyze the ABMP strategy of [12] playing against itself and playing against human negotiators. The experiments showed that human and ABMP negotiators primarily made concession moves (see for a precise definition Section 5.1). Additionally, it was shown that humans were more diverse, i.e. the types of negotiation moves they performed were more diversified.

In [8] it has been shown that the performance of a negotiation strategy might depend on the negotiation domain and preference profiles of the negotiation parties. The analysis methodology, therefore, should include a mechanism to vary all these factors influencing the negotiation behaviour. The analysis method proposed already includes the following factors: size of the negotiation domain, predictability of the preferences, opposition of preferences, and negotiation strategy of the opponents. The paper shows that these factors can influence learning of opponent's preferences and, as a result, negotiation performance of a strategy.

The initial, informal classification of negotiation moves and the results reported in [1] form the inspiration of the current paper. In combination with the belief expressed by many that the pattern of offers exchanged influences the negotiated outcome, see e.g. [17], this motivated our study of negotiation strategies from the perspective of the negotiation dynamics and the actual moves made.

On the basis of the literature and our own experience we conclude that an analysis of negotiation dynamics requires the use of both theoretical as well as experimental evaluation methods, in which at least the following aspects are attended to:

- competition with other strategies and itself,
- case studies of varying complexity,
- domains with various characteristics, and
- theoretical properties of the dynamics.

For papers that focus on the competition with other strategies the reader is referred to e.g., [5,7,20]. The next sections discuss the value of the other aspects.

3. Negotiation factors

Negotiation always takes place in a setup defined by the negotiation domain and the preference profiles of the parties involved. In [8] it has been shown that performance of the negotiation strategies depends on the negotiation setup. Therefore, we identify a number of factors that can influence a negotiation strategy or one of its components. A number of negotiation factors influencing negotiation behaviour have been reported in [8] and [10]. We reuse these factors in our method.

Size of the negotiation domain. Complexity of the negotiation domain and preference profiles is determined by the size of the negotiation domain. Size of the domain can influence learning performance of the negotiation strategy and, thus, the outcome reached by the strategy [10]. The size of the domain is exponential with respect to the number of issues. Therefore, the experimental setup in the analysis method should have a set of domains ranging from low number of issues to higher number of issues.

Predictability of the preferences. Negotiation strategies can try to exploit the internal structure of the preferences in order to improve one's own efficiency. I.e., the Trade-off strategy proposed in [5] assumes that distance measures can be defined using domain knowledge for the preferences of the opponent. These measures combined with the opponent's offers allow the Trade-off strategy to *predict* opponent preferences and as a result improve efficiency of the bidding. In [8], however, it has been shown that in case of a mismatch of the domain knowledge and the actual structure of the opponent's preferences the performance of a strategy can drastically drop. Therefore, we introduce the notion of the predictability of the preferences into our method.

Issues are called predictable when even though the actual evaluation function for the issue is unknown, it is reasonable to expect some of its global properties. For example, a price issue typically is predictable, where more is better for the seller, and less is better for the buyer, and the normal ordering of the real numbers is maintained; an issue concerning colour, however, is typically less predictable.

Opposition of the preferences. The results of analyzing negotiation dynamics presented in [8] revealed that some negotiation strategies are sensitive to preference profiles with compatible issues. Issues are compatible if the issue preferences of both negotiating parties are such that they both prefer the same alternatives for the given issue. Negotiation strategies

may more or less depend on whether preferences of the negotiating parties are opposed or not on every issue. That is, using some strategies it is harder or even impossible to exploit such common ground and agree on the most preferred option by both parties for compatible issues (humans are reported to have difficulty with this as well; cf. [22]). A selection of preference profiles should therefore take into account that both preference profiles with and without compatible issues are included.

3.1. Negotiation domains and preference profiles

Ideally, negotiation domains used in the analysis should cover all range of the factors sketched above. The selection of the domains presented in this paper is not intended to cover all variations of the domain factors influencing the negotiation performance. The negotiation domains used in this paper are:

The *Second hand car selling* domain, taken from [12], includes 5 issues. Only the buyer's preferences and the price issue are predictable, in the sense that an agent can reliably predict the other agent's preferences associated with an issue.

The *Service-Oriented Negotiation* domain, taken from [5], includes 4 issues. All issues are predictable, i.e. based on available "domain knowledge" preferences can be reliably predicated. The preference profiles have the strongest opposition in our setup.

The *AMPO vs City* domain, taken from [17], includes 10 issues, of which only 8 are predictable. This is the biggest domain in our experimental setup. Information about the opponent's issue priorities is not available, i.e., the weights agents associate with issues.

3.2. Negotiation strategies

The following strategies have been studied. The *ABMP* strategy [12], a concession oriented strategy, which computes bids to offer next without taking domain or opponent knowledge into account. (Experiments were run with a negotiation speed of 0.1 and a concession factor of 1, see [12].)

The *Trade-off* strategy is based on similarity criteria [5], and exploits domain knowledge to stay close to the Pareto Frontier. The "smart" version of this strategy performs nice moves if possible; otherwise it concedes a fixed amount 0.05 (cf. [5]). For the *Service-Oriented Negotiation* domain, we reproduced the results presented in [5].

The *Random Walker* strategy randomly jumps through the negotiation space, and can be run with or without a break-off point (to avoid making offers below that utility). *Random Walker* serves as a "baseline" strategy. This strategy has also been called the *Zero Intelligence strategy* [6].

A negotiation strategy that uses a learning technique based on the Bayesian learning algorithm proposed in [9]. The opponent model in [9] is based on learning a probability distribution over a set of hypotheses about evaluation functions and weights of issues. The probability distribution is defined over the set of hypotheses that represent an agent's beliefs about an opponent's preferences. Structural assumptions about the evaluation functions and weights are made to decrease the number of parameters to be learned and simplify the learning task.

The authors propose two versions of the learning algorithm. In the first version of the algorithm each hypothesis represents a complete utility space as a combination of weights ranking and shapes of the issue evaluation functions. The size of the hypotheses space grows exponentially with respect to the number of issues and thus is intractable for negotiation domains with a high (more than 6) number of issues.

The second version of the algorithm is a scalable variant of the first one. This version of the agent aims to learn a probability distribution over the individual hypotheses about the weight and shape of the issue evaluation function independently of other issues. The computational tractability of the learning is achieved by approximating the conditional distributions of the hypotheses using the expected values of the dependant hypotheses.

4. Negotiation dynamics

In the analysis of negotiation strategies, not only the outcome of a negotiation is relevant, but also the bidding process itself is important. Mistakes made during the bidding can have an enormous impact on both players. Although experienced negotiators confirm this, and it is also recognized by researchers of negotiation strategies for automated negotiation, this hypothesis is difficult to quantify. Examples from human negotiations are of the form: "a wrong offer can upset relationships, even causing the other party to walk away", or "Sometimes an offer that is meant as a concession to the other party confuses the issues. This can only be circumvented if there is enough

trust between the parties to exchange some information on their respective preferences.”

From the point of view of automated negotiation, the objective is to stay as close as possible to the Pareto Efficient Frontier. However, in automated negotiations typically no prior information is exchanged about the preferences of the negotiating parties, and none of the players know where the Pareto Efficient Frontier actually is. It thus remains a challenge to stay or end close to that Frontier. To this end, opponent modelling may be used to predict which bids will be appreciated by the other party, see, e.g., [3,9], and [16].

More precisely, five key factors can be identified that shape the outcome of a bilateral negotiation with incomplete information:

- (i) knowledge about the negotiation domain,
- (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.

In this paper, 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 gain a better understanding of the final outcome of a negotiation. The analysis does not take the features of the protocol per se into account but instead focuses on the exchange of offers. In principle, the method allows for generalizations to multi-party negotiations but we do not consider such extensions here. More information on multi-party negotiations can be found in e.g. [19].

5. Step-wise analysis

In bilateral bargaining, the negotiation dynamics is completely represented by the sequence of offers $t = \langle b^1_s, b^2_o, b^3_s \dots \rangle$ exchanged between parties S and O , also called the *negotiation trace*. A negotiation trace is called *closed* if it ends in either an *accept* or *withdraw* move by either party. In this section the basic notions of the step-wise analysis method are defined: classes of negotiation moves, metrics, outcome properties, and patterns over negotiation moves. After that the different phases of the method are defined.

5.1. Negotiation moves

The key unit of negotiation dynamics analysis is a single negotiation step performed by one of the negotiating parties. A negotiation step in bargaining consists of an offer proposed by one party to the other. If this offer is not the first offer proposed by an agent, it typically is computed using at least the previous offer proposed by that agent as input. To record this fact and to facilitate notation below, formally, a negotiation step s by agent a is modeled as a transition from a previous offer b_a to a newly proposed offer b'_a , which is written as $b_a \rightarrow b'_a$. Such moves can be classified based on the associated utility for both parties.

Bosse and Jonker [1] introduced 4 classes of steps: concession, unfortunate, fortunate and selfish steps. Firstly, negotiation strategy should be able to make *concessions*. In a concession, an agent trades in own interest in favour of the opponent to reach an agreement. Unfortunately, in a closed multi-issue negotiation such a move can lead to an alternative that is even worse for the opponent due to lack of information about the opponent's preferences. Such move is called *unfortunate*. In addition, the agent can go up with respect its own and down with respect to the opponents utility function. Such a move is called a *selfish* move. The last class is called the class of fortunate steps, i.e., a move towards the Pareto Efficient Frontier, going up with respect to the utility functions of both parties.

For the step-wise analysis method the classes of negotiation moves of [1] have been extended with two additional classes. In a number of papers, it has been suggested that it is smart not to make concessions too soon, but to move over one's iso-utility lines first [5]. For this reason, a separate category of *nice* moves that move in the direction of the opponent but do not concede with respect to the agent's own utility is introduced. An example of a strategy that is designed to make such moves is the Trade-Off Strategy based on similarity criteria discussed in [5], a variation is proposed in [14]. Additionally, so-called *silent* moves are introduced to represent the fact that parties sometimes repeat their offers, and do not make any concessions at all such as in a Boulware or take it or leave it strategy, see [12].

Each type of move in a negotiation typically has a distinct role or function, though in automated negotiation systems not all of these types of move are taken into account. Fortunate moves happen spontaneously in human negotiations (see [1]). Having a strategy that is able to perform such moves deliber-

ately is beneficial, since such moves can be used to recover from moves away from the Pareto Efficient Frontier, e.g., as the result of *concessions* or *unfortunate moves*. The latter two moves aim at reaching a jointly acceptable outcome. Although in general, it would be best to avoid unfortunate moves when conceding, it is impossible to guarantee this when Opponent's preferences are not completely known. Selfish moves may be performed by an agent to signal to the other party that a previous move is not appreciated. The role of nice and silent moves has been discussed above.

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 "Opponent" with respect to b . We use $\Delta_a(s) = U_a(b') - U_a(b)$, $a \in \{S, O\}$, for a move $s = b \rightarrow b'$ to denote the utility difference of two bids b and b' in the utility space of agent a . Here we present a precise definition of the classes of negotiation moves proposed in [1] extended as discussed above. These move categories define the core of the move-wise analysis method.

Definition of Move Classes. Let $s = b_S \rightarrow b'_S$ be a move in the bidding by Self (the definition for Opponent is completely symmetric). Then the negotiation move s taken by Self is classified as a:

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

The proposed classification is exhaustive, and all move classes are disjoint. (To allow for some marginal errors the areas of the Nice and Silent moves can be stretched somewhat in the analysis. In that way, a move in which only 0.005 of Self's utility is lost would still be classified as e.g., Nice, instead as a Concession.) In a concession move the agent's own utility has to go down but Opponent's utility may stay the same. In such cases, Self can claim that it

made a concession move by arguing that it conceded some of its own resources.

5.2. Step metrics and pattern properties

Having established different types of negotiation moves that are useful in the analysis of negotiation strategies, we now introduce and define metrics in terms of these moves that can be used for the analysis of negotiation traces. First, some additional notation is defined. Given a trace $t = \langle b^1_S, b^2_O, b^3_S \dots \rangle$ of offers, t^i denotes the i^{th} element of this sequence. Let t_S (resp. t_O) denote the sequence of moves from t that are made by agent "Self" (resp. "Opponent") and let class $c \in \{\text{Fortunate, Nice, Concession, Selfish, Unfortunate, Silent}\}$; then t_c denotes the subsequence of moves that belong to class c . Finally, $t_{(a,c)}$, also written t_{ac} , denotes the subsequence of moves by $a \in \{S, O\}$ that belong to class c . The following move metrics are introduced here:

Definition (Number of Moves per Trace). The number of moves $\#t$ in a trace t of length $|t| = n$ is defined as follows: $\#t = |t| - 1$.

Definition (Total Utility per Class). The pair $Total_c(t)$ of sums of utility differences in all moves of class c in a sequence t of moves is defined by:

$$Total_c(t) = (Total_{Sc}(t), Total_{Oc}(t)),$$

where for any agent $a \in \{S, O\}$: $Total_{ac}(t) = \sum_i \Delta_a(t^i_c)$.

Definition (u-Average Utility per Class). The pair $u-Ave_c(t)$ of average differences in utility in all moves in class c in a sequence t of moves is defined by:

$$u-Ave_c(t) = (u-Ave_{Sc}(t), u-Ave_{Oc}(t)),$$

where for any agent $a \in \{S, O\}$:

$$u-Ave_c(t) = \sum_i \Delta_a(t^i_c) / \#t_c.$$

Here $\#t_c$ is the number of moves of class c in trace t . This metric looks at the average utility conceded per negotiation move. A relative measure can be defined in terms of this metric to identify how much utility has been conceded by agent a relative to the other, indicated by agent g : $u-Ave_c(t_a) / u-Ave_c(t_g)$. This figure, if not identical to 1, indicates that one party is a conceiver relative to the other, and that concessions may not have been paced and linked to that of the other party, as is advised by Raiffa [11, p. 128].

Definition (% per Class). The percentage $\%_c(t)$ of class c moves 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 moves over a negotiation trace decreases. Such patterns can be seen as a measure of adaptability of a party to the opponent. Another useful measure of the sensitivity to the opponent's preferences can be defined by comparing the percentage of fortunate, nice and concession moves that increase the opponent's utility to the percentage of selfish, unfortunate and silent moves that decrease it. Intuitively, the more an agent performs moves that increase its opponent's utility the more sensitive to the needs of its opponent, it is said to be.

Definition (Sensitivity to Opponent's Behaviour). The measure for sensitivity of agent a to its opponent's behaviour is defined for a given trace t by:

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

In case no selfish, unfortunate or silent moves are made we stipulate that $\text{BehavSens}_a(a,t) = \infty$. If $\text{BehavSens}_a(t) < 1$, then an agent is more or less insensitive to the Opponent's behaviour; if $\text{BehavSens}_a(t) > 1$, then an agent is more or less sensitive to Opponent's behaviour, with complete sensitivity for $\text{BehavSens}_a(a,t) = \infty$. 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 its opponent's behaviour, but not vice-versa. In Section 7, this metric is used to analyze the sensitivity of two existing negotiation strategies. Furthermore, its relation to the knowledge available to the agent of the opponent's behaviour and the negotiation outcome is discussed.

We expect that the Random Walker strategy to make about 25% of moves in each of the fortunate, unfortunate, selfish, and concession classes due to its randomized selection of bids. This strategy is included in our method as a benchmark strategy and is used to validate our measures (see Table 3).

5.3. Sensitivity to opponent's preferences

A successful negotiation strategy should search for offers that maximize the agent's own utility while

increasing the chance of acceptance by the opponent. We assume that a rational negotiating agent would more easily accept offers with higher utility than those with lower utility. To increase the chance of an acceptance it is, therefore, rational to increase the opponent's utility of an offer without giving in with respect to the agent's own interests by means of trade-offs between the issues. Ideally, such a search procedure would lead to offers on the Pareto Efficient Frontier. Thus, an efficient negotiation strategy should generate offers from the Pareto frontier or at least as close as possible to the frontier depending on the limitation of the strategy.

To generate offers that are close to the Pareto Efficient Frontier a negotiation strategy must use information about the opponent's preferences in addition to its own preferences. Therefore, to assess the performance of a negotiation strategy we need to measure its sensitivity to the opponent's preferences. Given the assumptions of the opponent's rationality we define a measure of sensitivity as follows.

Definition (Sensitivity to Opponent's Preferences). To measure the sensitivity of a strategy to the opponent's preference we calculate the average difference between Opponent's utility of the bids generated by the strategy and the utility of a bid on the Pareto Efficient Frontier:

$$\text{PrefSens}(t_s) = \frac{1}{\#t_s} \sum_{i=1}^{\#t_s} U_o(t'_i) - U_o(t_i), \quad t_i \in t_s$$

where t'_i is defined as

$$t'_i = \arg \max_{U_S(b)=U_S(t_i)} U_o(b), \quad b \in D.$$

Figure 2 visualizes the measure of sensitivity to Opponent's preferences. The figure puts all possible negotiation outcomes in a given negotiation domain on a two-dimensional space. Each dimension represents the utility of an offer with respect to one of the negotiating parties. A negotiation strategy that is perfectly sensitive to Opponent's preferences would propose Pareto efficient offers only and, therefore, would have $\text{PrefSens}(t) = 0$. The higher the measure the worse is the sensitivity of the strategy.

As said before, a strategy that is sensitive to the opponent's preferences requires information on the opponent's preferences. The ABMP strategy belongs to the class of the Time-Dependent Tactics that do not consider the opponent's preferences. Given the definition of the measure of sensitivity to the oppo-

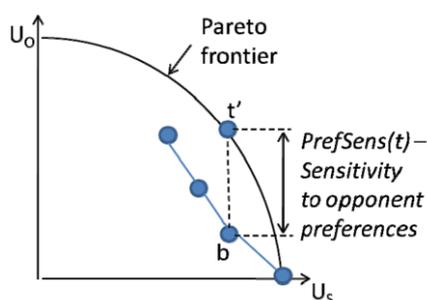


Fig. 2. Sensitivity to opponent's preferences.

nent's preferences, we can expect low sensitivity of the ABMP strategy (see Table 3a).

A typical solution to increase the sensitivity to the opponent's preferences is to use an opponent model. The Bayesian strategy uses a learning algorithm to guess the opponent's preferences and tries to bid on the Pareto Efficient Frontier to increase the chance of acceptance. Thus, the Bayesian strategy should score high on the sensitivity measure in case when learning is successful (see Table 3b). As we saw in the previous section the performance of the Trade-Off strategy strongly depends on the negotiation domain and the opponent's strategy. A similar result is to be expected in the sensitivity of that strategy to the opponent's preferences. The Random Walker does not consider the preferences of the opponent at all and would therefore have low sensitivity.

6. Negotiation analysis methodology

A methodology for analyzing the performance of various negotiation strategies should include the by now standard metrics based on solution concepts such as the Nash product as well as the metrics introduced above to analyze the negotiation moves. The need to include metrics related to the negotiation dance itself is clear from the performance of a Random Walker negotiating agent. As we show in Section 8, even a Random Walker can obtain an outcome close to the Pareto Efficient Frontier and good for the Random Walker provided that its opponent uses a reasonable negotiation strategy. A proper analysis should reveal that such results need to be contributed to the strategy or performance of the opponent rather than to the Random Walker agent. Though this may seem obvious from the performance of a Random Walker in any given negotiation for other negotiation strategies similar conclusions may only be reached by analyzing the performance of that strategy in vari-

ous domains and in combination with different preference profiles. A methodology for analyzing the performance of a negotiating agent thus should involve a careful setup of experiments as well as a range of different metrics to be able to reach sound conclusions about the agent's performance.

The methodology that we propose consists of seven steps to facilitate the setup of a tournament for analyzing negotiation strategies. The first four steps of the proposed method define the tournament whereas the last three steps concern the analysis of the results. Although the method proposed does not provide any guarantees that the proper results will be obtained, it does facilitate and structure the process of setting up negotiation experiments to obtain such results. Moreover, by reiterating the process the results of previously performed experiments can be used to refine the experimental setup and may suggest variations not initially considered.

6.1. Phases of the negotiation analysis method

Based on the above concepts, the analysis method is specified by:

1. Starting point: the strategies to be analysed, and a library of domains and of other strategies that can be used to test the input strategies.

In Section 3 we introduced a number of criteria for selection of negotiation domains, preference profiles and opponent strategies. Ideally, then, one would use an experimental setup based on random sampling of the domains and profiles in order to deal with this problem. However, it is not clear how to setup such a sampling procedure. Therefore, we selected a number of negotiation domains to be used in our experimental setup.

The opponent's negotiation strategy is one of the negotiation factors influencing the agent's negotiation performance [8]. It is important, therefore, to be able to include a wide range of existing negotiation strategies.

2. If necessary, implement the input strategies.

The proposed method is based on the analysis of empirical data and, therefore, requires an implementation of the input negotiation strategies that can be run in an experimental setup. Ideally, the input strategy should be implemented in the same environment with the strategies and negotiation domains and profiles selected in the 1st step. Otherwise, a communication between the strategies must be established and the

negotiation domains and profiles must be translated and made available for the input strategies.

3. Set up a tournament with the selected negotiation strategies and case studies.

A tournament a la [7] is used to experiment with various strategies. In the tournament, strategies play against each other, against themselves, and are applied to varying negotiation domains, with varying preference profiles. Multiple negotiation sessions of a single negotiation setup should be run in the case of non-deterministic negotiation strategies in the tournament.

4. Run the tournament and log every negotiation.

The negotiation log should include information about the tournament and settings of an individual session: the names of the strategies, the domain name, and the preferences of the players. Furthermore, it is important to produce an extensive log with all information needed to calculate the proposed metrics. Therefore, the log should include bids proposed by every negotiating party in a negotiation session and their utilities.

5. Calculate defined metrics.

Given the logs produced in the previous step, the following metrics introduced in this paper should be calculated for every negotiation session:

- utilities of the negotiation outcomes;
- number of negotiation moves;
- number of moves, %, and total utility per class;
- sensitivity to opponent's behaviour;
- sensitivity to opponent's preferences;

6. Apply statistics and produce analytical results.

In the case of non-deterministic negotiation strategies in the tournament as explained in the step 3, average values of the metrics should be calculated and used for analysis.

7. Interpret results, produce graphics.

The goal of the proposed method is to assess the negotiation performance of a negotiation strategy given a number of metrics of the negotiation outcome and dynamics. The assessment is based on the interpretation of the results received from the variation of the values of the negotiation factors introduced in this paper.

Graphics such as the two-dimensional utility plots of the negotiation outcome space (see, for example, Fig. 1) can help to find interpretations of the results.

They highlight the dynamic properties of the negotiation traces proposed here as well as the standard analysis measures.

The proposed analysis method was applied in an experimental setup using the open negotiation environment for heterogeneous negotiating agents, presented in [11]. The environment makes use of advanced software engineering technologies allowing simple and quick integration of existing negotiation strategies. The framework provides a simple application programming interface and a number of auxiliary common services to the agent to simplify the task of strategy implementation. It has a rich library of implemented negotiation strategies.

The analytical module of the environment can be easily extended with new metrics and has advanced logging functionality. Logs of the environment can be exported to widely-used analytical software, such as Excel and Mathematica. The environment has tools to generate a setup for a tournament and run it producing logs with the corresponding information.

7. DANS applied

This section illustrates the DANS analysis method for a combination of strategies and negotiation domains, while focusing on individual moves – in particular *unfortunate moves* – and on the sensitivity of four negotiation strategies with respect to preferences and behaviour of the opponent. Table 1 reminds how the sensitivity values should be interpreted and give extreme values for the high and low sensitivity.

A tournament with the strategies and domains of the previous section was set up and run. A full analysis was made of the type of moves made, which was then used to calculate the average sensitivity ratio's for all tested strategies over multiple runs against all strategies (including itself) in the domains described above.

7.1. Outcome analysis

The traditional (and valuable) method for analysis of strategies in terms of the outcome and its properties is applied first to the chosen combination of

Table 1
Extreme values for the sensitivity measures

	High sensitivity	Low sensitivity
BehavSens	∞	0
PrefSens	0	1

Table 2
Result of standard analysis

Negotiation Domain	Strategy of Agent A	Strategy of Agent B																		
		ABMP				Trade-Off				Random Walk (break-off=0.6)				Bayesian				Average		
		Utility		Distance to		Utility		Distance to		Utility		Distance to		Utility		Distance to		Utility	Distance to	
Agent A	Agent B	Nash	Kalai	Agent A	Agent B	Nash	Kalai	Agent A	Agent B	Nash	Kalai	Agent A	Agent B	Nash	Kalai	Agent A	Nash	Kalai		
Car domain	ABMP	0.83	0.93	0.08	0.09	0.90	0.83	0.10	0.09	0.86	0.76	0.18	0.17	0.82	0.96	0.10	0.11	0.85	0.12	0.11
	Trade-Off	0.72	0.94	0.19	0.20	0.84	0.92	0.08	0.08	0.87	0.76	0.19	0.17	0.84	0.94	0.08	0.08	0.82	0.13	0.14
	Random	0.78	0.83	0.17	0.17	0.74	0.89	0.18	0.18	0.73	0.76	0.26	0.25	0.74	0.86	0.19	0.19	0.75	0.20	0.20
	Bayesian	0.78	0.96	0.14	0.15	0.87	0.94	0.04	0.05	0.74	0.85	0.19	0.19	0.89	0.92	0.03	0.03	0.82	0.10	0.11
	Average		0.91	0.15	0.15		0.90	0.10	0.10		0.78	0.21	0.20		0.92	0.10	0.10			
AMPO vs City	ABMP	0.73	0.63	0.16	0.11	0.70	0.59	0.19	0.15	0.58	0.65	0.18	0.18	0.65	0.70	0.10	0.09	0.67	0.16	0.13
	Trade-Off	0.62	0.53	0.27	0.23	0.74	0.73	0.06	0.01	0.67	0.58	0.21	0.17	0.65	0.70	0.10	0.09	0.67	0.16	0.12
	Random	0.69	0.68	0.10	0.07	0.54	0.60	0.25	0.24	0.65	0.65	0.15	0.12	0.53	0.72	0.18	0.20	0.60	0.17	0.16
	Bayesian	0.65	0.70	0.10	0.09	0.74	0.72	0.07	0.01	0.68	0.50	0.28	0.24	0.73	0.70	0.08	0.03	0.70	0.14	0.09
	Average		0.63	0.16	0.12		0.66	0.14	0.10		0.60	0.21	0.18		0.70	0.12	0.10			
Service Oriented Negotiation	ABMP	0.51	0.59	0.36	0.33	0.65	0.49	0.38	0.32	0.73	0.66	0.19	0.13	0.76	0.84	0.02	0.06	0.66	0.23	0.21
	Trade-Off	0.52	0.65	0.31	0.29	0.79	0.77	0.09	0.02	0.55	0.61	0.31	0.28	0.76	0.82	0.03	0.05	0.66	0.18	0.16
	Random	0.63	0.60	0.28	0.23	0.63	0.77	0.14	0.15	0.66	0.70	0.17	0.14	0.78	0.78	0.08	0.00	0.67	0.17	0.13
	Bayesian	0.79	0.75	0.11	0.03	0.81	0.71	0.15	0.08	0.80	0.74	0.12	0.05	0.76	0.82	0.03	0.05	0.79	0.10	0.05
	Average		0.65	0.26	0.22		0.68	0.19	0.14		0.68	0.20	0.15		0.82	0.04	0.04			

strategies and negotiation domains as explained in the previous section.

presents the results of the outcome analysis. The average values are calculated over all opponents' strategies per negotiation domain and role in the domain. The ABMP strategy shows excellent performance (average utility of 0.85 for Role A and 0.91 for Role B) on the Car domain for which it was designed. The Trade-Off strategy underperforms the ABMP strategy on this domain with average utilities of 0.82 and 0.90 for the roles A and B respectively. The Bayesian agent performs somewhat better than the Trade-Off strategy (average utility of 0.82 for Role A and 0.92 for Role B). Random Walker is definitely a layman in this domain.

In the SON domain the Bayesian agent clearly outperforms the other strategies (average utility of 0.79 for Role A and 0.82 for Role B). Interestingly, the Trade-Off strategy performance varies from one opponent to another. It does well when playing against itself (utility of 0.79 for Role A and 0.77 for Role B) or against Bayesian agent (utility of 0.76 for Role A and 0.71 for Role B) but fails against ABMP (utility of 0.52 for Role A and 0.49 for Role B) and the Random Walker (utility of 0.55 for Role A and 0.67 for Role B). Unfortunately, the ABMP strategy is not able to take advantage of the Trade-Off strategy but leaves better alternatives on the table. The Random Walker shows a performance that is on average better than that of the Trade-Off and ABMP strategies (average utility of 0.67 for Role A and 0.68 for Role B). However, this might be due to the selected reservation value of Random Walker (every bid will have a utility of at least 0.6). Setting the reservation value even higher would lead to even better results.

The results in the AMPOvsCity domain are similar to that of the SON domain. The distance of the nego-

tiation outcome reached by the Bayesian strategy and the Nash and Kalai-Smorodinsky outcomes is bigger than in the SON domain. This can be explained by the significantly larger size of the AMPOvsCity domain that causes some degradation in the learning performance of the Bayesian strategy, see [10]. Due to the large size of the AMPOvsCity domain the Random Walker and the ABMP strategy show the worst performance among the other strategies.

The standard analysis method that focuses on the final outcome does not explain why such a performance is observed, e.g., the Trade-Off and the ABMP strategy underperform against the Random Walker strategy. Therefore, we propose the negotiation dynamics analysis method to get more information from the experimental results and get insights into the negotiation process.

7.2. Sensitivity to opponent's behaviour

Theoretically, over all domains and against all strategies, Random Walker would have a sensitivity value of 1. As expected, the Random-Walker shows approximately equal percentage of moves for the unfortunate, fortunate, selfish, and concession classes. It makes almost no nice or silent moves. Small deviation of the results can be explained by structural features of a negotiation domain.

ABMP shows an overall BehavSens of 4.15, Trade-off 4.13, Bayesian 6.21 and Random Walker 1.08. Note, that the efficiency of the agreement (see previous section) does not always correlate strongly with the BehavSens values of the strategies. The sensitivity scores of ABMP and Trade-Off can be better understood by considering the domains in which they played. Figure 3 shows typical runs in the AMPO vs City domain and the Pareto Efficient Frontier. The

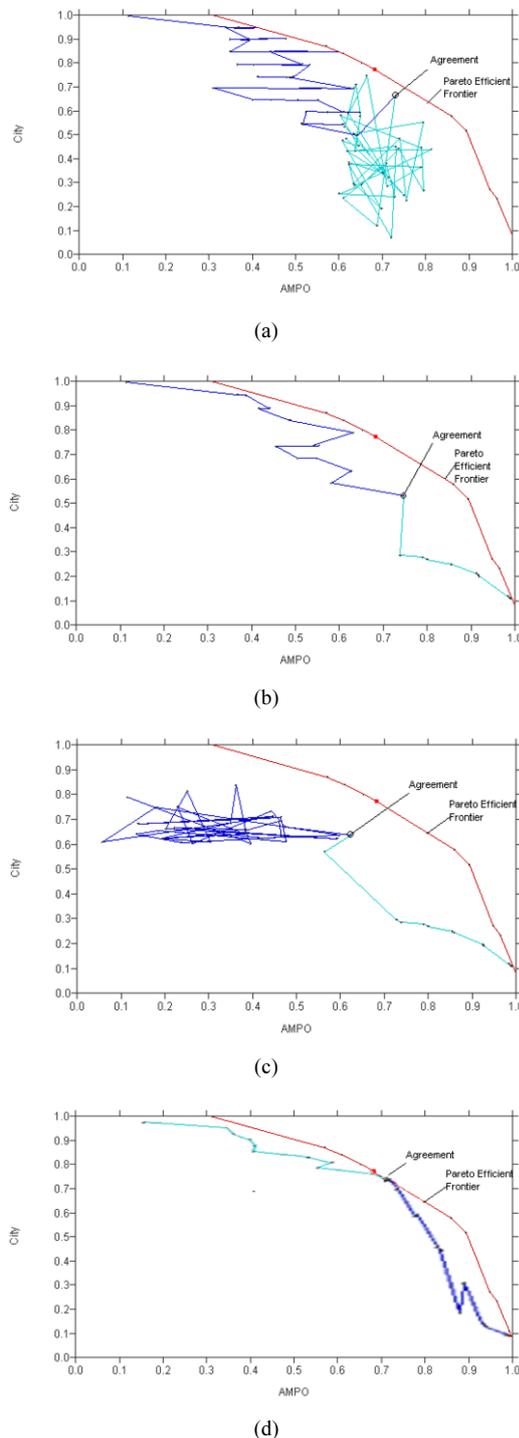


Fig. 3. Dynamics of negotiation process for: a) Trade-Off (City) vs Random Walker strategy (AMPO), b) Trade-Off (City) vs ABMP strategy (AMPO), c) Random Walker (City) vs ABMP strategy (AMPO), d) Trade-Off (City) vs Bayesian strategy (AMPO).

Pareto Efficient Frontier is built according to its definition (see [17], pp. 227) using exhaustive search.

Figure 3a shows a run of Trade-Off, representing the City, versus Random Walker (with a break-off set to 0.6), playing AMPO. The Random Walker strategy is insensitive with respect to the opponent's behaviour (BehavSens is always approximating 1) and strategy selects an offer in a random way. The Trade-off strategy uses the opponent's offers to select counteroffers and expects some rational behaviour of the opponent's strategy. This fact causes the unfortunate moves produced by the Trade-Off strategy (36% of all moves for Role A and 37% for Role B). This results in a low BehavSens of the Trade-Off strategy in negotiation against Random Walker (1.04 for Role A and 1.33 for Role B), see the columns for Trade-Off in Table 3a. However, the Trade-Off strategy in negotiations against itself and the Bayesian strategy makes significantly less unfortunate moves due to more rational and efficient behavior of the opponent and, as a result achieves much higher BehavSens scores (against itself: 5.25 for Role A and 6.14 for Role B; against Bayesian: 8.09 for Role A and 6.83 for Role B). In general, Trade-Off has high BehavSens in this domain (average value 4.07 for Role A and 3.89 for Role B).

Figure 3b shows Trade-Off (as City) vs ABMP (as AMPO) in which ABMP makes mostly concessions and silent moves by following its concessions tactic regardless of the moves of the opponent. It has rather low BehavSens (average value 1.56 for Role A and 2.93 for Role B) in this domain. ABMP shows more rational behaviour than Random Walker and, therefore, in this domain Trade-Off really exploits the available domain knowledge. The percentage of unfortunate steps in this negotiation drops to 27% for Role A and 32% for Role B.

Figure 3c shows Random Walker (City) vs ABMP (AMPO). ABMP in principle concedes on all issues, determining the size of the concession on the difference between the utilities of its own bid and that of its opponent. Unlike the Trade-Off strategy it does not use previous bids of the opponent to get insight into the opponent's preferences and as a result does not adapt much to the strategy of the opponent. Such a strategy will make unfortunate steps in case there are issues with compatible preferences (concession on such an issue would decrease utility of the opponent). AMPOvsCity does not have any issues with compatible preferences and, thus, ABMP does not make unfortunate steps. Given opposed preferences for every issue such a strategy typically produces

Table 3

(a) Results of the step wise analysis

Negotiation Domain	Strategy of Agent A	Metric	Strategy of Agent B															
			ABMP								Trade-Off							
			Number of steps				Sensitivity				Number of steps				Sensitivity			
Fort	Nice	Cons	Silent	Selfish	Unfort.	Behav.	Opp.Prefs	Fort	Nice	Cons	Silent	Selfish	Unfort.	Behav.	Opp.Prefs			
Car domain	ABMP	Agent B	7%	6%	67%	6%	0%	14%	3.98	0.26	0%	23%	47%	12%	0%	18%	2.34	0.28
		Agent A	1%	3%	54%	20%	0%	21%	1.41	0.24	1%	12%	42%	25%	0%	21%	1.19	0.20
	Trade-Off	Agent B	5%	16%	60%	6%	0%	13%	4.23	0.25	3%	14%	53%	18%	0%	13%	2.24	0.05
		Agent A	0%	10%	45%	20%	0%	25%	1.23	0.12	2%	4%	63%	23%	0%	7%	2.27	0.04
	Random Walk	Agent B	1%	1%	81%	8%	2%	8%	4.63	0.23	1%	29%	22%	18%	1%	30%	1.07	0.31
		Agent A	29%	3%	22%	8%	17%	22%	1.18	0.45	28%	9%	23%	2%	17%	21%	1.47	0.49
	Bayesian	Agent B	4%	12%	66%	6%	0%	12%	4.32	0.25	4%	10%	58%	19%	0%	9%	2.57	0.02
		Agent A	10%	43%	39%	3%	0%	5%	11.50	0.03	3%	12%	62%	20%	0%	3%	3.35	0.01
	Average	Agent B	4%	9%	68%	6%	1%	12%	4.29	0.25	2%	19%	45%	17%	0%	17%	2.06	0.17
		Agent A	1%	6%	49%	21%	2%	21%	1.32	0.22	2%	13%	46%	21%	0%	18%	1.90	0.15
Service Oriented Negotiation	ABMP	Agent B	0%	0%	91%	9%	0%	0%	10.55	0.39	0%	6%	62%	5%	0%	26%	2.19	0.16
		Agent A	0%	2%	83%	16%	0%	0%	5.40	0.39	0%	4%	67%	28%	0%	0%	2.55	0.48
	Trade-Off	Agent B	0%	3%	67%	29%	0%	0%	2.42	0.35	0%	0%	81%	1%	0%	18%	4.31	0.09
		Agent A	0%	0%	77%	6%	0%	17%	3.30	0.20	0%	5%	85%	3%	0%	7%	9.02	0.01
	Random Walk	Agent B	6%	1%	86%	6%	1%	0%	11.63	0.38	2%	39%	20%	3%	1%	35%	1.57	0.13
		Agent A	20%	2%	26%	8%	24%	21%	0.91	0.53	17%	4%	31%	3%	26%	18%	1.08	0.42
	Bayesian	Agent B	0%	1%	87%	4%	0%	8%	7.33	0.38	0%	8%	80%	3%	0%	9%	7.33	0.01
		Agent A	0%	3%	84%	4%	0%	9%	6.69	0.01	0%	9%	82%	3%	0%	3%	15.17	0.01
	Average	Agent B	1%	1%	83%	12%	0%	2%	7.98	0.38	0%	13%	61%	3%	0%	22%	3.85	0.10
		Agent A	1%	3%	81%	15%	1%	2%	6.14	0.42	1%	9%	68%	6%	0%	16%	4.90	0.13
AMPO vs City	ABMP	Agent B	0%	0%	65%	35%	0%	0%	1.89	0.44	0%	0%	56%	12%	0%	32%	1.27	0.36
		Agent A	0%	12%	38%	49%	0%	0%	1.02	0.43	0%	4%	64%	29%	0%	3%	2.13	0.40
	Trade-Off	Agent B	0%	4%	75%	21%	0%	0%	3.78	0.40	0%	6%	80%	5%	0%	9%	6.14	0.03
		Agent A	0%	0%	68%	5%	0%	27%	2.17	0.35	0%	8%	76%	6%	0%	10%	5.25	0.01
	Random Walk	Agent B	0%	0%	59%	37%	1%	3%	1.47	0.46	1%	26%	30%	6%	0%	37%	1.33	0.45
		Agent A	21%	4%	25%	3%	25%	23%	0.98	0.82	22%	3%	26%	4%	22%	23%	1.04	0.82
	Bayesian	Agent B	0%	3%	79%	18%	0%	0%	4.56	0.39	0%	0%	82%	6%	0%	6%	6.83	0.04
		Agent A	12%	45%	34%	5%	0%	5%	9.10	0.04	10%	48%	35%	4%	0%	3%	13.29	0.03
	Average	Agent B	0%	2%	70%	28%	0%	1%	2.93	0.42	0%	8%	62%	7%	0%	21%	3.89	0.22
		Agent A	1%	7%	52%	38%	0%	2%	1.56	0.42	1%	8%	63%	10%	0%	19%	4.07	0.21

(b) Results of the step wise analysis (continued)

Negotiation Domain	Strategy of Agent A	Metric	Strategy of Agent B															
			Random Walk (break-off=0.6)								Bayesian							
			Number of steps				Sensitivity				Number of steps				Sensitivity			
Fort	Nice	Cons	Silent	Selfish	Unfort.	Behav.	Opp.Prefs	Fort	Nice	Cons	Silent	Selfish	Unfort.	Behav.	Opp.Prefs			
Car domain	ABMP	Agent B	26%	3%	23%	9%	21%	18%	1.08	0.45	0%	5%	82%	3%	0%	10%	6.69	0.02
		Agent A	1%	6%	55%	16%	6%	16%	1.60	0.21	1%	4%	47%	23%	0%	25%	1.09	0.22
	Trade-Off	Agent B	23%	6%	32%	7%	18%	14%	1.53	0.40	4%	11%	61%	18%	0%	6%	3.17	0.02
		Agent A	2%	23%	19%	19%	1%	36%	0.77	0.32	5%	14%	58%	20%	0%	3%	3.35	0.12
	Random Walk	Agent B	23%	1%	23%	7%	22%	24%	0.90	0.39	2%	28%	37%	8%	0%	15%	2.91	0.29
		Agent A	26%	4%	22%	4%	18%	27%	1.06	0.44	22%	3%	24%	8%	22%	21%	0.96	0.48
	Bayesian	Agent B	23%	4%	22%	7%	23%	19%	1.01	0.45	0%	4%	78%	8%	0%	10%	4.56	0.02
		Agent A	15%	25%	33%	10%	0%	17%	2.70	0.12	0%	6%	86%	4%	0%	4%	11.50	0.01
	Average	Agent B	24%	3%	25%	8%	21%	19%	1.13	0.42	2%	12%	65%	9%	0%	10%	4.33	0.09
		Agent A	26%	5%	23%	5%	18%	23%	1.17	0.47	7%	22%	55%	9%	0%	7%	7.26	0.04
Service Oriented Negotiation	ABMP	Agent B	19%	2%	28%	7%	2%	21%	1.62	0.33	1%	12%	68%	13%	0%	6%	4.26	0.01
		Agent A	5%	1%	86%	6%	2%	0%	12.02	0.40	0%	3%	86%	10%	0%	9%	4.60	0.41
	Trade-Off	Agent B	17%	3%	30%	3%	26%	21%	0.99	0.38	0%	14%	70%	14%	0%	2%	5.25	0.01
		Agent A	2%	27%	33%	2%	0%	36%	1.63	0.22	2%	5%	78%	12%	0%	3%	5.67	0.08
	Random Walk	Agent B	20%	2%	30%	4%	26%	19%	1.05	0.38	0%	20%	56%	12%	0%	12%	3.17	0.01
		Agent A	20%	3%	27%	2%	27%	21%	1.00	0.44	23%	3%	25%	3%	23%	23%	1.04	0.45
	Bayesian	Agent B	21%	3%	28%	4%	26%	18%	1.08	0.31	0%	9%	79%	11%	0%	1%	7.33	0.01
		Agent A	0%	6%	81%	5%	0%	8%	6.69	0.03	0%	1%	87%	10%	0%	2%	7.33	0.01
	Average	Agent B	19%	3%	29%	4%	20%	20%	1.19	0.35	0%	14%	68%	13%	0%	5%	5.00	0.01
		Agent A	20%	3%	27%	4%	25%	21%	1.01	0.46	0%	5%	84%	6%	0%	6%	8.97	0.02
AMPO vs City	ABMP	Agent B	16%	7%	28%	4%	23%	22%	1.03	0.39	3%	13%	78%	5%	0%	1%	15.67	0.05
		Agent A	2%	6%	50%	41%	1%	0%	1.38	0.42	2%	4%	57%	34%	0%	3%	1.70	0.42
	Trade-Off	Agent B	18%	5%	22%	12%	21%	20%	0.85	0.58	0%	10%	77%	10%	0%	3%	6.69	0.03
		Agent A	0%	14%	30%	20%	0%	36%	0.79	0.46	2%	8%	79%	9%	0%	2%	8.09	0.02
	Random Walk	Agent B	22%	5%	23%	1%	23%	26%	1.00	0.76	0%	18%	60%	12%	0%	10%	3.55	0.08
		Agent A	19%	4%	28%	2%	26%	22%	1.04	0.87	21%	5%	22%	5%	24%	23%	0.92	0.49
	Bayesian	Agent B	23%	2%	25%	5%	22%	23%	1.00	0.55	1%	15%	69%	12%	0%	3%	5.67	0.01
		Agent A	5%	29%	27%	20%	0%	6%	2.35	0.07	0%	18%	63%	14%	0%	5%	4.26	0.02
	Average	Agent B	20%	5%	24%	5%	22%	23%	0.97	0.57	1%	14%	71%	10%	0%	4%	7.89	0.04
		Agent A	21%	4%	25%	3%	24%	23%	1.00	0.75	7%	35%	40%	11%	0%	5%	7.25	0.04

concession moves (see Fig. 3b,c). This results in a relatively high BehavSens in negotiations against Random Walker compared to that of the Trade-Off strategy (average value 1.47 for Role A and 1.89 for Role B). BehavSens for the ABMP strategy increases in negotiations against Trade-Off and Bayesian but has relatively smaller values than that of the Trade-Off strategy (against Trade-Off: 2.13 for Role A and 3.78 for Role B; against Bayesian: 1.70 for Role A and 4.56 for Role B).

Figure 3d shows Trade-Off (City) vs Bayesian (AMPO). The negotiation behaviour of the Bayesian strategy is very similar to that of the Trade-Off strategy: stays very close to the Pareto frontier and makes few unfortunate moves. The BehavSens of the Bayesian strategy is very high in this negotiation (13.29 for Role A and 6.69 for Role B). Unlike the Trade-Off strategy the Bayesian strategy does not rely on domain knowledge to generate offers that are close to the Pareto frontier. Instead it tries to learn the opponent's preference profile from the opponent's offers. The learned preference profile allows the Bayesian strategy to respond to the opponent with more nice and concession moves and keep the number of unfortunate moves at minimum. As a result, it remains rather sensitive to the opponent's behaviour even when it negotiates against Random Walker (value 2.35 for Role A and 3.55 for Role B).

7.3. Sensitivity to opponent's preferences

The sensitivity to the opponent's preferences is measured in the same experimental setup as for the outcome analysis method and the step-wise method. Table 3 also presents the sensitivity results with respect to the opponent's preferences. As expected, the Bayesian strategy is the most sensitive to the opponent's preferences with average values per domain of $PrefSens \leq 0.09$. The Bayesian strategy is less sensitive when negotiating against the ABMP and Random Walker strategies than against the other strategies. This can be explained from the fact that the assumptions used in the learning algorithm do not hold for those strategies and the learned opponent's preference profile in these negotiations has lower quality.

The Trade-Off strategy is somewhat less sensitive to the preferences of the opponent (average $PrefSens \leq 0.22$ in the various domains) than the Bayesian strategy. The Trade-Off strategy can be efficient but has difficulties with domains with low predictability of the preferences. For example, its sensitivity to the opponent's preferences is not as good in the 2nd Hand Car Selling domain ($PrefSens \leq 0.17$ as in the SON

domain ($PrefSens \leq 0.13$) despite the fact that the later domain is much bigger than the earlier. The similarity functions for the 2nd Hand car domain often do not match the preferences of the opponent. In addition, the weights of the similarity function do not match the opponent's importance factors of the negotiation issues. The SON domain does not have information about the weights of the similarity functions but the issues preferences perfectly match the similarity functions and thus the sensitivity to the opponent's preferences of the Trade-Off strategy increases.

The sensitivity to the opponent's preferences for the Trade-Off strategy depends on the strategy used by the opponent. In negotiations against ABMP and Random Walker the Trade-Off strategy is less sensitive to the opponent's preferences (e.g., in SON domain it has $PrefSens \leq 0.09$ in negotiations against Trade-Off and Bayesian whilst it has $PrefSens \leq 0.22$ in negotiations against ABMP and Random Walker). This variation can be explained by the fact that the Trade-Off strategy tries to match the opponent's preferences by maximizing the similarity of the its offers with those of the opponent. If the opponent's offers are far from the Pareto frontier (such as in case of ABMP and RandomWalker) the Trade-Off strategy would not be able to match the opponent's preferences.

The ABMP strategy shows rather robust performance but it is often outperformed by the smarter strategies. It has overall a rather low sensitivity to the opponent's preferences (average values over the various domains are $PrefSens \leq 0.42$). Still it is more sensitive than the Random Walker that, as expected, has the lowest sensitivity to the opponent's strategy (average values over the various domains are $PrefSens \leq 0.75$). Note, that the $PrefSens$ value of the ABMP strategy does not vary much over the opponent's strategy because it uses the opponent's offers to determine the size of a concession but does not try to match the opponent's preferences.

7.4. DANS analysis results

In summary the DANS analysis results are as follows. The opponent sensitivity analysis shows a direct link between the correctness and/or completeness of the domain knowledge and opponent sensitivity. The Trade-Off strategy is very sensitive to an opponent given complete information. In that case, the similarity functions exactly match the opponent's preferences and the weights exactly represent the issue importance factors of the opponent and the sen-

sitivity metric for the behaviour of the opponent is converging to zero. Intuitively, the Trade-Off strategy would be more efficient in a smaller domain due to the smaller search space. However, the incomplete domain knowledge in the 2nd Hand Car Selling domain does not allow the Trade-Off strategy to fully use its potential to search for efficient outcomes. The experiments show that if less domain knowledge is available, Trade-Off makes more unfortunate moves.

In general, when issues are predictable, the chance of making an unfortunate step becomes small. This aspect becomes clear in the car domain, where the seller's preferences (Role A) are rather predictable, but the buyer's preferences (Role B) vary a lot.

We conclude that it is impossible to avoid unfortunate moves without sufficient domain knowledge or opponent knowledge. Indeed, the similarity criteria functions used in the Trade-Off Strategy provide general information about the negotiation problem, but do not take into account the specific attributes of the negotiating parties. In any particular case, a negotiator may deviate from the generalized domain model in various ways.

On the other hand the Bayesian strategy does not use domain knowledge and tries to learn the opponent's preference during negotiation. The learning algorithm of the Bayesian strategy allows it to remain sensitive to the opponent's behaviour and preferences regardless of the completeness and correctness of the available domain knowledge.

Sensitivity to the opponent's preference of the ABMP strategy does not seem to be influenced by the opponent's strategy unlike the Trade-Off strategy. The ABMP strategy shows rather robust performance but it is often outperformed by the smarter strategies, such as the Trade-Off and the Bayesian strategies, in terms of outcome utilities. The Bayesian strategy similar to the ABMP strategy show somewhat more robust behaviour than the Trade-Off strategy. This can be explained by the way it learns the opponent's preferences. It does not require that the opponent's offers stay close to the Pareto frontier. A better model of the opponent's preferences allows the Bayesian strategy to be more sensitive to the opponent's preferences and, finally, reach better negotiation outcomes.

8. Conclusion

This paper shows that an analysis of the negotiation dance [17] is important for the understanding

and improvement of negotiation strategies. The DANS analysis method introduced in this paper focuses on the classification of negotiation moves and metrics over this classification. The classification enables us to relate the *intent* of a strategy in making a negotiation step with the actuality of the *perception* of that step by the opponent. For example, a strategy might be concession oriented, i.e., moves are intended to be concessions, but in reality some of these moves might be *unfortunate*, meaning that although the proposer of the bid is giving in, from the perception of the receiver, the bid is actually worse than the previous bid.

By testing strategies over various domains and against various opponents patterns emerge of when such unfortunate moves occur. These patterns are related to dynamic properties such as the sensitivity of strategies. Experiments with DANS show, for example, that the Trade-Off strategy is rather *responsive* to the behaviour of the opponent, in that it follows the behaviour of the opponent. If that is rather wild, such as the random behaviour of the Random Walker, Trade-Off shows a high percentage of unfortunate moves.

Experiments further show that the occurrence of unfortunate moves is related to features of the negotiation domain and the extent to which such features are incorporated in the strategy. The same holds for knowledge about the preferences of the opponent. To better understand the relative importance of each relation, we have emphasized the distinction between *domain knowledge* and *opponent knowledge*.

We think it is impossible to avoid unfortunate moves without sufficient domain knowledge or opponent knowledge. Domain knowledge provides generalized information about the negotiation problem, but does not necessarily match with individual preferences of negotiating parties. Opponent knowledge concerns individual information and as such is not transferable to other opponents. Therefore, we advocate a combination of domain and opponent knowledge.

The DANS analysis method focuses explicitly on properties of interest to the researcher. The combination of statistical methods and graphical representation is strong: Humans process graphs faster than tables with numbers, however, the number of experiments typically done make it impossible for the human to view every graph produced by the experiment. A more general aspect is that graphs of long negotiation dances become hard to grasp; what may look like a neat series of nice moves, might actually be a mixture of unfortunate and nice moves.

The examples in the paper show how the focus on the percentage of unfortunate moves makes it possible for the DANS method provide particularly insightful graphs, such as the graph of the Trade-Off vs Random Walker that provides insight into the sensitivity of the Trade-Off strategy. Furthermore, the sensitivity analyses were used to show that the Bayesian strategy is able to overcome the inefficiency of the Trade-Off strategy in domains for which the available domain knowledge is incomplete and/or incorrect.

9. Future work

We believe that our results also show the need for *benchmark problems* for bilateral negotiation. An interesting direction for future research in this area would be to propose a measure for exploitability of a negotiation strategy. A good negotiation strategy must be able to withstand an inefficient opponent strategy, such as Random Walker, and a strategy that tries to exploit its opponent.

In this paper we showed that the negotiation domain can have strong influence on performance of a negotiation strategy. Thus, another interesting direction for the future work is to develop a design method to generate varying negotiation domains and preference profiles and add this domain- and profile generation method to DANS. While it is impossible to test a negotiation strategy in all possible negotiation scenarios the method should be able to give a good spread of negotiation domains, preference profiles and negotiation strategies such that it covers factors such as the size of the negotiation domain, predictability of the preferences, opposition of preferences, and opponent behaviour (i.e., strategy).

References

- [1] T. Bosse and C.M. Jonker, Human vs. computer behaviour in multi-issue negotiation, in: *Proc. of the 1st International Workshop on Rational, Robust, and Secure Negotiations in Multi-Agent Systems*, IEEE Computer Society Press, 2005, pp. 11–24.
- [2] T. Bosse, C.M. Jonker, and J. Treur, Experiments in human multi-issue negotiation: Analysis and support, in: *Proc. of the 3rd Int. Joint Conf. on Autonomous Agents and Multi-Agent Systems*, IEEE Computer Society Press, 2004, pp. 672–679.
- [3] R.M. Coehoorn and N.R. Jennings, Learning an opponent's preferences to make effective multi-issue negotiation trade-offs, in: *Proc. of the 6th International Conference on E-Commerce*, Delft, The Netherlands, 2004, pp. 59–68.
- [4] P. Faratin, C. Sierra, and N.R. Jennings, Negotiation decision functions for autonomous agents, *Int. Journal of Robotics and Autonomous Systems* **24**(3–4) (1998), 159–182.
- [5] P. Faratin, C. Sierra, and N. Jennings, Using similarity criteria to make negotiation trade-offs, *Journal of Artificial Intelligence* **142**(2) (2003), 205–237.
- [6] D.K. Gode and S. Sunder, Allocative efficiency in markets with zero intelligence (zi) traders: Market as a partial substitute for individual rationality, *Journal of Political Economy* **101**(1) (1993), 119–137.
- [7] P. Henderson, S. Crouch, R.J. Walters, and Q. Ni, Comparison of some negotiation algorithms using a tournament-based approach, *Agent Technologies, Infrastructure, Tools and Applications for E-Services*, Springer-Verlag, 2003, pp. 137–150.
- [8] K. Hindriks, C. Jonker, and D. Tykhonov, Negotiation dynamics: Analysis, concession tactics, and outcomes, in: *International Conference on Intelligent Agent Technology IAT'07*, IEEE/WIC/ACM, 2007, pp. 427–433, ISBN 0-7695-3027-3.
- [9] K. Hindriks and D. Tykhonov, Opponent modelling in automated multi-issue negotiation, in: *Proc. of 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS'08)*, Padgham, Parkes, Müller, and Parsons, eds, 2008, pp. 331–338.
- [10] K.V. Hindriks and D. Tykhonov, Towards a quality assessment method for learning preference profiles in negotiation, in: *Proc. of the 10th Int. Workshop on Agent Mediated Electronic Commerce (AMEC'08)*, 2008, pp. 5–17.
- [11] K. Hindriks, C. Jonker, and D. Tykhonov, Towards an open negotiation architecture for heterogeneous agents, in: *Proc. of 12th Int. Workshop on Cooperative Information Agents*, Springer-Verlag, 2008, pp. 264–279.
- [12] C.M. Jonker and J. Treur, An agent architecture for multi-attribute negotiation, in: *Proc. of the 17th Int. Joint Conference on AI (IJCAI'01)*, 2001, pp. 1195–1201.
- [13] L.-K. Soh and X. Li, Investigating adaptive, confidence-based strategic negotiations in complex multiagent environments, *Journal of Web Intelligence and Agent Systems* **6**(3) (2008), 313–326.
- [14] R. Lin, S. Kraus, J. Wilkenfeld, and J. Barry, An automated agent for bilateral negotiation with bounded rational agents with incomplete information, in: *Proc. of the 17th European Conference on Artificial Intelligence (ECAI'06)*, 2006, pp. 270–274.
- [15] A.R. Lomuscio, M. Wooldridge, and N.R. Jennings, A classification scheme for negotiation in electronic commerce, *Int. Journal of Group Decision and Negotiation* **12**(1) (2000), 31–56.
- [16] C. Mudgal and J. Vassileva, Bilateral negotiation with incomplete and uncertain information, in: *Proc. of the Workshop on Cooperative Information Agents*, Springer, 2000, pp. 107–118.
- [17] H. Raiffa, *The Art and Science of Negotiation*, Harvard University Press, 1982.
- [18] H. Raiffa, J. Richardson, and D. Metcalfe, *Negotiation Analysis: The Science and Art of Collaborative Decision Making*, Harvard University Press, 2002.
- [19] R. Ros and C. Sierra, A negotiation meta strategy combining trade-off and concession moves, *Journal of Autonomous Agent and Multiagent Systems* **12**(2) (2006), 163–181.
- [20] J.S. Rosenschein and G. Zlotkin, *Rules of Encounter: Designing Conventions for Automated Negotiation Among Computers*, MIT Press, 1994.

- [21] D. Zeng and K. Sycara, Bayesian learning in negotiation, in: *Proc. AAAI Symposium on Adaptation, Co-Evolution, and Learning in Multiagent Systems*, AAAI Press, 1996.
- [22] L.L. Thompson, *The Heart and Mind of the Negotiator*, Pearson Prentice Hall, 2005.
- [23] A. Nongaillard, Ph. Mathieu, and B. Jaumard, A multi-agent resource negotiation for the utilitarian social welfare, in: *Proc. of the 9th International Workshop Engineering Societies in the Agents World (ESAW2008)*, A. Artikis, G. Picard, and L. Vercoeur, eds, LNAI, Vol. 5485, 2008.
- [24] Y. Yasumura, T. Kamiryo, S. Yoshikawa, and K. Uehara, Acquisition of a concession strategy in multi-issue negotiation, *Journal of Web Intelligence and Agent Systems* 7(2) (2009), 161–171.