

Unanimously acceptable agreements for negotiation teams in unpredictable domains



Victor Sanchez-Anguix^{a,*}, Reyhan Aydogan^b, Vicente Julian^a, Catholijn Jonker^b

^a Departamento de Sistemas Informáticos y Computación, Universitat Politècnica de València, Camí de Vera s/n, 46022 Valencia, Spain

^b Interactive Intelligence Group, Delft University of Technology, Delft, The Netherlands

ARTICLE INFO

Article history:

Received 28 July 2013

Received in revised form 17 May 2014

Accepted 18 May 2014

Available online 28 May 2014

Keywords:

Automated negotiation

Multi-agent systems

Agreement technologies

ABSTRACT

A negotiation team is a set of agents with common and possibly also conflicting preferences that forms one of the parties of a negotiation. A negotiation team is involved in two decision making processes simultaneously, a negotiation with the opponents, and an intra-team process to decide on the moves to make in the negotiation. This article focuses on negotiation team decision making for circumstances that require unanimity of team decisions. Existing agent-based approaches only guarantee unanimity in teams negotiating in domains exclusively composed of *predictable and compatible issues*. This article presents a model for negotiation teams that guarantees unanimous team decisions in domains consisting of *predictable and compatible*, and also *unpredictable* issues. Moreover, the article explores the influence of using opponent, and team member models in the proposing strategies that team members use. Experimental results show that the team benefits if team members employ Bayesian learning to model their teammates' preferences.

© 2014 Elsevier B.V. All rights reserved.

1. Introduction

In the last decade, there has been an increase in the profit earned by electronic commerce systems. This increase has led to a strong interest of the academic world in researching problems related to e-commerce (Ngai and Wat 2002; Grieger 2003; Wareham et al. 2005). As of today, most e-commerce systems rely on users manually browsing their catalogs and selecting which goods they desire to buy. This task may end up being time consuming and suboptimal in terms of users' preferences, especially as the number of items and services offered on the Web increases. Therefore, it is necessary to propose mechanisms that help costumers take better decisions while saving their time efforts.

Agent-based electronic commerce has been proposed as a solution to such problems (Guttman et al. 1998; Sierra and Dignum 2001; Oliveira and Rocha 2001; He et al. 2003). In an agent-based e-commerce system, autonomous agents act on behalf of their users with the goal of finding and closing satisfactory deals. Automated negotiation is one of the most common approaches when implementing these systems since they allow different electronic parties to reach agreements by exchanging offers and feedback

(Lomuscio et al. 2003; Nguyen and Jennings 2005; Buffett and Spencer 2007; lau 2007; Chan et al. 2008). The benefits of automated negotiation and agent-based e-commerce are many. Being brief, some of the most important include:

- As stated, browsing online catalogs for an optimal deal may be time consuming. The state-of-the-art in automated negotiation can complete complex negotiations for multiple issues in less than a few minutes (Klein et al. 2003a; Williams et al. 2011; Baarslag et al. 2012).
- On the one hand, automated negotiation saves the user from having to browse the entire catalog. Additionally, its personal agent is directed by the preferences of the user in the negotiation, which should result in deals that are adjusted to the personal liking of the individual. Personalization has been reported to increase user satisfaction in many computational systems (Ball et al. 2006; Liang et al. 2007). On the other hand, a dynamic process like automated negotiation allows sellers to adapt their deals to the users' preferences, their current business needs, and their competitor dynamics (He et al. 2003).
- Agreements achieved by human negotiators, suffer from the *leaving money on the negotiation table effect* (Thompson 2003). This means that human negotiators are content with current agreements, which are usually suboptimal, when they could have performed much better. Agents in automated negotiation

* Corresponding author. Tel.: +34 963877350.

E-mail addresses: sanguix@dsic.upv.es (V. Sanchez-Anguix), r.aydogan@tudelft.nl (R. Aydogan), vinglada@dsic.upv.es (V. Julian), C.M.Jonker@tudelft.nl (C. Jonker).

have been reported to provide agreements close to the optimal solution (Lai et al. 2008).

- Compared to centralized and offline approaches (e.g., preference aggregation, recommendation approaches, etc.), automated negotiation is a dynamic and parallel process. For instance, some centralized approaches like preference aggregation are computationally hard especially if the preference space is combinatorial (Chevalyere et al. 2007). On the other hand, recommendation approaches only filter prospective deals, but they do not close specific contracts adapted to business needs. Contrarily, automated negotiation can be adapted to current business needs (e.g., concede to gain customers and close fast deals). Additionally, as stated above, team members are also motivated by their own personal interests. Therefore, it is possible that some team members show opportunistic behavior inside the team. In such cases, preference aggregation may be manipulated by exaggerating preferences. Additionally, each parties' preferences are private, therefore making it difficult for the other parties to exploit and manipulate. This latter factor is important, since nowadays most users in electronic applications care about the information they filtrate in systems (Taylor 2003).

Most negotiation mechanisms proposed for e-commerce settings have focused on solving bilateral or multiparty negotiations where parties are individual agents (Faratin et al. 1998; Zeng and Sycara 1998; Klein et al. 2003b; Nguyen and Jennings 2005; Coehoorn and Jennings 2004; Buffett and Spencer 2007; Lai et al. 2008; Williams et al. 2011; Sanchez-Anguix et al. 2013; Aydođan and Yolum 2012). However, some real life scenarios involve negotiation parties that are not necessarily formed by single individuals. Instead, each party may be formed by more than a single individual. For instance, imagine that a group of travelers wants to go on a holiday together. As a group, they have to negotiate with several travel agencies to get the best travel package for the group. Despite sharing a common goal, each member in the multiplayer party may also be motivated by its own personal interests (Mannix (2005) and Halevy (2008)). Therefore, the group not only faces a possibly difficult negotiation with the travel agency, but it also needs to deal with the conflict present in the group. This type of multi-individual negotiating party has been studied in the social sciences under the name of *negotiation team* (Thompson et al. 1996; Brodt and Thompson 2001).

As far as the authors are concerned, multi-individual parties have been overlooked in automated negotiation research. The use of computational models for negotiation teams opens doors for new types of interesting and novel applications in electronic commerce. The inclusion of agent-based negotiation teams allows for e-commerce systems to deploy dynamic deal mechanisms for groups, making of e-commerce a more social system. Classically, when purchasing for groups in e-commerce systems, one representative takes decisions for the whole group. Either he makes decisions according to his own preferences or the group needs to engage in a human negotiation which is usually a costly process due to different schedules, logistics, lack of communication problems or interpersonal conflict (Behfar et al. 2008). With the inclusion of agent-based negotiation teams these problems are eluded since autonomous agents take decisions jointly while saving time and efforts for their users.

We believe that agent-based negotiation teams could provide potentially interesting new services:

- Electronic markets for groups of travelers: Online travel agencies offer their services by means of online catalogs where users can browse different products like flights, hotels, restaurants, activities, etc. The possibilities for travels are vast, and usually

a single travel operator may offer thousands of possible trip packages/services. Exhaustively looking through this online catalog for an optimal deal becomes an unfeasible task for humans. Additionally, more often than not, travel is a social activity for groups (e.g., friends, family, young people, etc.). Users can benefit from agent-based negotiation teams since they can exhaustively look for deals while taking the preferences of the group into account and saving efforts. Service providers can also benefit from these models since they could adapt their business strategies in a dynamic way and add a level of personalization that may help to retain customers. Moreover, offering the possibility for groups to close travel deals based on their preferences is a value-added service, that as far as we know, is not currently offered by the industry. As an example of its application, users may indicate to their personal agents their desire to go on a travel together. Then, the agents prepare to negotiate with different travel agencies in order to provide a complete and satisfactory travel package for the users. The fact that the negotiation is carried out automatically by electronic agents also gives room to looking for several alternatives in parallel. Once several trip packages have been negotiated, the personal agents may communicate the agreements to users, who can validate them in the last instance.

- Electronic support for agricultural cooperatives: Agricultural cooperatives are supposed to be democratic institutions where groups of farmers join together to save resources for the distribution of their products. One of the main problems of agricultural cooperatives is the principal-agent problem (Ortmann and King 2007). Basically, despite being democratic institutions, agricultural cooperatives are managed by a board of directors who take decisions on behalf of the democratic institution. It has been reported in the literature (Ortmann and King 2007) that dissatisfaction in cooperatives comes from the fact that the goals of members are not aligned with those of the managers. As a novel application for electronic commerce, agent-based negotiation teams may provide support for the processes that are carried out by cooperatives. For instance, the negotiations between agricultural cooperatives and distributors may be supported by an electronic market where the agricultural cooperative is modeled as an agent-based negotiation team. Each member may be represented by an electronic and personal agent that participates in the negotiation team according to the preferences of its owner. This way, if the model is capable of ensuring unanimity with regards to team decisions, it may be possible to avoid the principal-agent problem. Of course, agricultural cooperatives are large institutions and considerable research has still to be done to provide scalable and fair computational models. However, research as the one presented in this article contributes to the obtention of such models in the long term.

- Groups of energy producers in the smart grid: The smart grid is addressed to be the next generation network for electricity distribution (Farhangi 2010). In this network, energy generation may come from geographically distributed small generators (e.g., green energy generators) that have to compete with large energy producers. Decisions at the smart grid have to be taken dynamically since energy production and consumption may vary or face unexpected events (Ramchurn et al. 2012). Recently, agent-based electronic commerce has been proposed as proper paradigm for this scenario due to its dynamic nature and adaptive response (Brazier et al. 2002; Lamparter et al. 2010; Morais et al. 2012; Ramchurn et al. 2012). If small generators want to compete with large generators like power plants, they may need to group together and act together as a single generator. Agent-based negotiation teams can give support for the group decision making of small generators in a dynamic

environment like the smart grid. For instance, an agent-based negotiation team for the smart grid may decide on different contract attributes like energy price for different time slots, contract duration and cancellation fees with different energy consumers.

The applications described above present benefits for electronic commerce systems. However, there are still several issues that need to be solved for deploying real applications based on agent-based negotiation teams due to the novelty of the topic. One of the main issues that should be addressed when designing agent-based negotiation team models is unanimity. The authors argue that, whenever it is possible, it is desirable for the final agreement with the opponent to be unanimously acceptable for all of the team members. When the members of the negotiation team are going to interact in the long term, the intra-team strategy should avoid one or some of the team members being clearly at disadvantage (e.g., unacceptable deal) with respect to the other team members. In the first place, the aforementioned situation may end up in users perceiving unfairness, which may affect commitment to the decision, group attachment, and trust (Korsgaard et al. 1995). And second, but not the least important, users that are not satisfied with agreements found automatically may end up leaving the electronic commerce application.

The existing approaches (Sanchez-Anguix et al. 2011; 2012a,b) have focused on achieving unanimously acceptable agreements for negotiation domains exclusively comprised by *predictable and compatible* issues among the team members. An issue is *predictable and compatible* if the preference order over issue values is the same for team members and this fact is known from the domain (e.g., price in a team of buyers). While some e-commerce domains are exclusively composed by these issues, many domains also contain issues whose preferential ordering over issue values is not known from the domain (i.e., *unpredictable* issues). For instance, it is difficult to predict from a set of cities which ranking represents the preferences of a traveler, which can diverge from the preferences of other travelers.

This article advances the state of the art in agent-based electronic commerce in two different ways. Firstly, it introduces a new model for agent-based negotiation teams, which could support dynamic negotiations for groups of autonomous agents representing their users. Secondly, the present model is capable of assuring that the final agreement is unanimously acceptable for all of the team members in domains that contain both *predictable and compatible* and *unpredictable* issues. We propose an intra-team protocol in which a team mediator helps team members to reach unanimously acceptable decisions. Furthermore, we propose two negotiation strategies for team members: a basic negotiation strategy based on concession tactics and a negotiation strategy using Bayesian learning to model teammates' and opponent's preferences for *unpredictable* issues. The model is capable of outperforming state-of-the-art approaches for agent-based negotiation teams. We describe our general framework in Section 2 and the intra-team protocol that allows team members to reach unanimity in Section 3. After that, we propose two negotiation strategies for team members in Section 4 and we explain why unanimity is guaranteed among team members in Section 5. After analyzing the experiments in Section 6, we relate our work to existing approaches and discuss future lines of work in Section 7.

2. Overview of the negotiation framework

Let A represent a negotiation team consisting of $|A| = M$ different team members and a trusted team mediator med_A , and let $a \in A$ represent a team member in negotiation team A . Let op represent

the opponent party of the negotiation team. The negotiation between team and opponent is carried out in a bilateral fashion, using an alternating-offers protocol (Rubinstein 1982). In this protocol, one of the two parties is the initiating party and sends the first offer to the other party or responding party. The responding party receives the offer and decides whether or not he/she accepts the offer. Accordingly, she or he may accept the current offer or send a counter-offer. If the responding agent sends a counter-offer, the initiating party has to decide whether he/she accepts the counter-offer or not. If the counter-offer is rejected, the process is repeated in a turn-taking fashion until a deal is mutually accepted (successful negotiation) or one of the parties decides to quit the negotiation since its deadline has been reached (failed negotiation). Concerning inter-party communications, the team mediator interacts with the opponent by sending team's proposals and transmitting opponent decisions to team members. The team mediator plays a key role since it coordinates the team members and helps them reach unanimously acceptable deals.

Let X be the object under negotiation, $j \in \{1, \dots, n\}$ be the issues under negotiation, D_j be the negotiation domain or valid values for issue j and $x_j \in D_j$ represent a valid value for issue j . Each agent's preferences are represented by means of a private additive utility function. We assume that there is no preferential interdependency among negotiation issues; that is, the valuation given to a certain issue does not affect preference on the valuation of other negotiation issues. The utility function for an agent in our framework can be formalized as follows:

$$U(X) = w_1 V_1(x_1) + w_2 V_2(x_2) + \dots + w_n V_n(x_n), \quad (1)$$

where w_j represents the importance given to issue j by the agent, and $V_j : x_j \rightarrow [0, 1]$ is a scoring function for issue j that gives the score that the agent assigns to an issue value x_j . It is assumed that $\sum_{j=1}^n w_{aj} = 1$ and $w_{aj} \geq 0$, and then $U(\cdot)$ is a function scaled in $[0, 1]$, where 0 represents the least desirable negotiation deals, and 1 represents the most desirable negotiation deals. For agents, $RU \in [0, 1]$ represents the reservation utility or the minimum level of utility to consider an agreement as acceptable.

In the proposed framework, private information and bounded rationality are assumed. The former has been introduced above: information regarding agents' preferences is private, and so are the strategies and minimum acceptable values of each agent. This is true even among team members, since prior to the negotiation they do not know any information regarding other teammates' preferences. The only information available is obtained via interactions in the intra-team protocol. The latter refers to the fact that given the limited time, information privacy, and limited computational resources, agents cannot calculate the optimal strategy to be carried out during the negotiation. Instead, they employ heuristic strategies that aim to be as good as possible in terms of the achievable utility.

2.1. Unanimously acceptable agreements

Each team member $a \in A$ has a reservation utility $RU_a \in [0, 1]$ that represents the minimum utility that satisfies the team member's need. Each outcome whose utility is lower than the reservation utility is unacceptable for the team member. As stated along this article, we consider that unanimity in a negotiation team is of extreme importance. An offer is unanimously acceptable for a team A if it is acceptable for all of the team members inside the negotiation team:

$$\forall a \in A, U_a(X) \geq RU_a. \quad (2)$$

The proposed intra-team strategy will assure that team members only accept those offers that are unanimously acceptable for

all the team members and that offers proposed to the opponent are over each team members' reservation utilities, thus, making it unanimously acceptable.

2.2. Types of negotiation issues among team members

Among the different negotiation issues that compose the negotiation domain, we consider that there are issues that are *predictable and compatible* among team members and issues that are *unpredictable* among team members.

Formally, we can define an issue j with domain D_j as compatible among team members if for each possible pair of team members $a, b \in A$ and for each pair of issue values $v_1, v_2 \in D_j$, the following expression is true:

$$V_{a,j}(v_2) > V_{a,j}(v_1) \rightarrow V_{b,j}(v_2) \geq V_{b,j}(v_1). \tag{3}$$

Hence, an issue is compatible among team members if one of the team members can increase its utility by selecting a certain issue value with respect to the current assignment, then the rest of team members stay at the same utility or they also increase their utility. Thus, there is no preferential conflict among issue values between the team members, and there is full potential for cooperation among team members with respect to compatible issues. Fig. 1 shows two examples of compatible issues among two agents (top part) and an example of a non compatible issue (bottom part). As it can be observed, in the case of price (top left), both agents obtain a better valuation when choosing a lower price value with respect to a high price value. Thus, Eq. (3) holds and it is a compatible issue for both agents. In the case of the city of destination (top right), the issue is also compatible among the two agents. For any pair of cities, if one of the agents prefers one of the cities with respect to other city, the other agent also holds the same preferential relationship. For instance, both agents prefer Paris to Berlin, Berlin to London, and London to Madrid. However, in the case of the type of room (bottom part), the blue agent prefers an individual room with respect to an apartment, whereas the red

agent prefers exactly the opposite. Thus, there is no full potential for cooperation among team members in that negotiation issue since conflict is present.

The concept of predictability and unpredictability (Hindriks and Tykhonov 2010, Marsa-Maestre et al. 2013) is related to vertical and horizontal issues found in economics literature (Stole 1995). The definition of predictable issues matches with vertical issues, while the definition of unpredictable issues matches with horizontal issues. From this point on, we will use the concepts of unpredictable/predictable and we will briefly introduce them. An issue is *predictable* for an agent if the preference ordering of issue values is known in the negotiation domain. Therefore, an issue is *compatible and predictable* among team members if the preferences regarding issue values are known in the negotiation domain and increasing the utility of one of the team members by selecting one specific issue value results in other team members staying at the same utility or also increasing their respective utilities. For instance, from the examples in Fig. 1, one can consider that inside a team of buyers the price is a compatible and *predictable* issue among team members since it is known that all of the buyers prefer low prices to high prices, and reducing the price results in all of the buyers increasing their utility or staying at the same utility. On the other hand, an issue is *unpredictable* among team members if the preference ordering of the issue values cannot be accurately predicted and Eq. (3) may not hold for that issue. In the case of Fig. 1, the city of destination is a compatible issue among the two agents. Nevertheless, in a travel negotiation domain it is not true that all of the travelers will hold the same preference ranking over the issue values and without additional knowledge, the preference ordering may not be predicted accurately. Hence, it is an unpredictable issue. With respect to the type of room, the preference ordering over issue values may vary for the travelers. Moreover, we cannot predict their preference ordering directly, thus making the issue unpredictable.

In this framework, *PR* denotes the set of *predictable and compatible* issues among team members, while *UN* denotes the set of *unpredictable* issues.

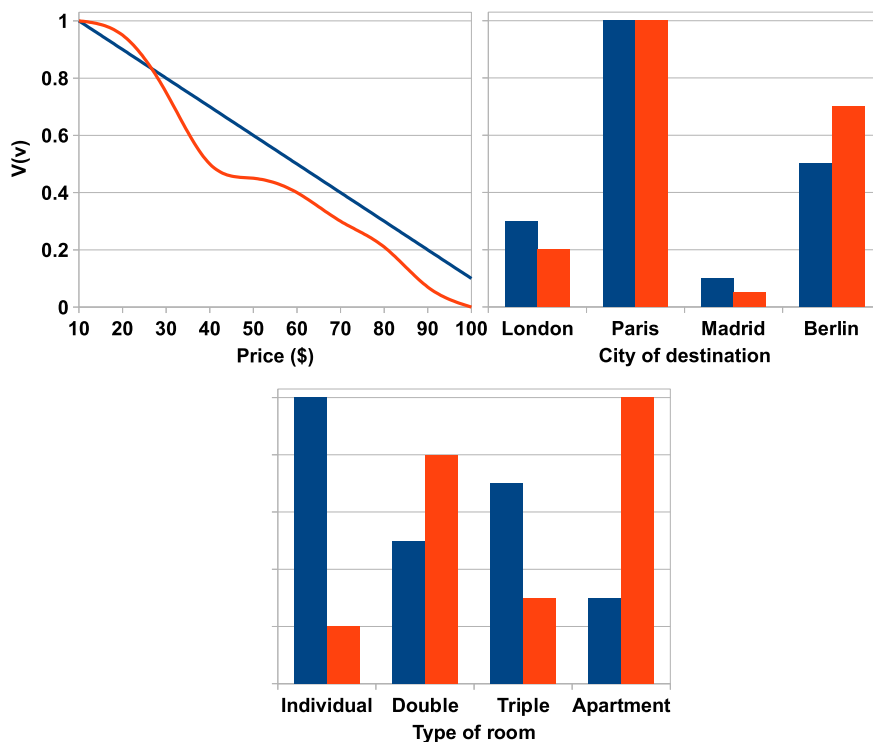


Fig. 1. Two compatible issues among two agents (top) and a non compatible issue among two agents (bottom).

2.3. Forbidden unpredictable partial offers among team members

We define an unpredictable partial offer X' as a partial offer that has a concrete instantiation of all the unpredictable issues in UN . The utility of an unpredictable partial offer is calculated as $U_a(X') = \sum_{j \in UN} w_{a,j} V_{a,j}(X_j)$.

For a team member $a \in A$, an unpredictable partial offer X' will never be part of an acceptable offer (i.e., it will never be an unambiguously acceptable offer for the team) when the sum of the utility of X' and the maximum utility that can be obtained from predictable issues $\max PR_a = \sum_{j \in PR} w_{a,j}$ is less than its reservation value RU_a , since any full offer that completes X' is below the reservation utility. For a team member a , we refer to the set of unpredictable partial offers that will never be part of an acceptable offer as *forbidden unpredictable partial offers*, F_a (see Eq. (4)).

$$F_a = \{X' \mid U_a(X') + \max PR_a < RU_a\} \quad (4)$$

It is worth noting that F_a does not represent the whole negotiation space that is unacceptable for a , but just a portion of it. In fact, some unpredictable partial offers that are not contained in F_a , can become unacceptable when the agent does not get the value needed from predictable issues. The size of F_a may grow as the reservation utility increases. Thus, agents with high reservation utilities are expected to have larger sets of F_a than agents with low reservation utilities.

2.4. Case of study

In this article we have employed a case of study (i.e., a negotiation domain) that is extracted from a possible tourism electronic market. The case of study is used to illustrate and test the proposed negotiation framework.

A group of travelers wants to go on a holiday together and arrange their accommodation. The group negotiates with a hotel on the following issues.

- *Price (p)*: It represents the price per night that each traveler pays to the hotel for the booking service. The value goes from 200\$, which is the minimum rate applicable by the hotel, to 400\$, which is the maximum rate found in the hotel. This negotiation issue is considered to be predictable and compatible among team members since all of the travelers obviously prefer low prices to high prices. Contrarily, the hotel prefers high prices to low prices.
- *Cancellation fee (cf)*: This issue represents the amount of the final price that each friend pays if the reservation is canceled. Possible values for this negotiation issue go from 0% to 50%. This is a predictable and compatible issue among team members since all of the travelers prefer low cancellation fees to high cancellation fees. On the contrary, the opponent prefers high cancellation fees to low cancellation fees.
- *Arranged foods included (af)*: The hotel may also offer some meals included in deal with the travelers. The type of meal plans included are *none*, *breakfast*, *breakfast + lunch*, *breakfast + dinner*, *lunch + dinner*, and *all*. In our negotiation scenario, we have considered that this negotiation issue is unpredictable among team members since preferences of team members on this issue may vary and it cannot be assumed to be same for each member.
- *Type of room (tr)*: The four travelers can be accommodated in different types of room depending on their preferences. More specifically, the hotel offers *4 individual rooms*, *2 twin rooms*, *1 triple and 1 individual room*, or *1 apartment*. The type of room is an unpredictable negotiation issue among team members.
- *Payment method (pm)*: The amount of money paid by the travelers may be paid by different methods. The hotel allows for the payment to be made in *cash*, via *credit card*, by *bank transfer*,

in a *3 months deferred payment* through the bank, and in a *6 months deferred payment*. This negotiation issue is unpredictable since team members may prefer to choose different payment methods and we cannot predict their preference ordering directly.

- *Room orientation (ro)*: If possible, the team members can decide upon an orientation for the balcony of their rooms. The different options are *inner garden*, *main street*, *pool*, *sea*, and *outer garden*. This issue is also considered an unpredictable issue among team members.
- *Free amenity (fa)*: As a token of generosity for booking as a group, the hotel offers one free service to all of the team members. More specifically, the team members can choose between *gym service*, *free wi-fi*, *1 free drink per day*, *1 free spa session*, *pool service*, *cable tv service*, and one *free guided tour*. Since the preferences of team members vary for this issue and no assumption about their preferences can be made, this issue is also considered as unpredictable.

To sum up, for this case study we have that $PR = \{p, cf\}$ and $UN = \{af, tr, pm, ro, fa\}$ with a total of 4200 different combinations of discrete issue values (af, tr, pm, ro, fa) and two real issues (p, cf). We assume that the team mediator knows which issues are predictable and can apply an operator that determines the best value for team members from a given set. For unpredictable issues, team members can have different types of valuation functions and the mediator does not know which issue values are better for team members. Each team member may assign different weights (i.e., priorities) to negotiation issues and the opponent's valuation functions and issue weights may be different from those of team members. The team mediator does not know the weights given by agents to the different issues.

3. Intra-team protocol

In a negotiation involving a negotiation team, the intra-team protocol defines how and when decisions are taken regarding the negotiation. In this framework, we propose an intra-team protocol that is governed by the trusted team mediator med_A . Basically, the team mediator regulates the interactions that can be carried out among team members and, accordingly, helps team members reaching unanimous acceptable decisions inside the team during the negotiation. The proposed protocol is clearly differentiated into two different phases: *Pre-negotiation* and *Negotiation*. On the one hand, during the pre-negotiation, the mediator helps team members identifying potential offers that are not unanimously acceptable for every teammate. On the other hand, during the negotiation the mediator coordinates the offer proposal mechanism, which is composed of a voting process for unpredictable issues and an iterated building process for predictable issues, and the offer acceptance mechanism for offers that come from the opponent. We describe those phases in a detailed way in the following Sections 3.1 and 3.2. An overview of all of the communications carried out in the negotiation model are depicted in Fig. 2. It specifies the protocols carried out within the team and the communications carried out with the opponent by means of Agent UML (Bauer et al. 2001) sequence diagrams. More detailed views of the intra-team protocols for the pre-negotiation, evaluation opponent's proposals and proposing offers can be observed in Figs. 3–5 respectively.

3.1. Pre-negotiation phase

In the pre-negotiation phase, the mediator coordinates the following intra-team protocol to discover the set of forbidden

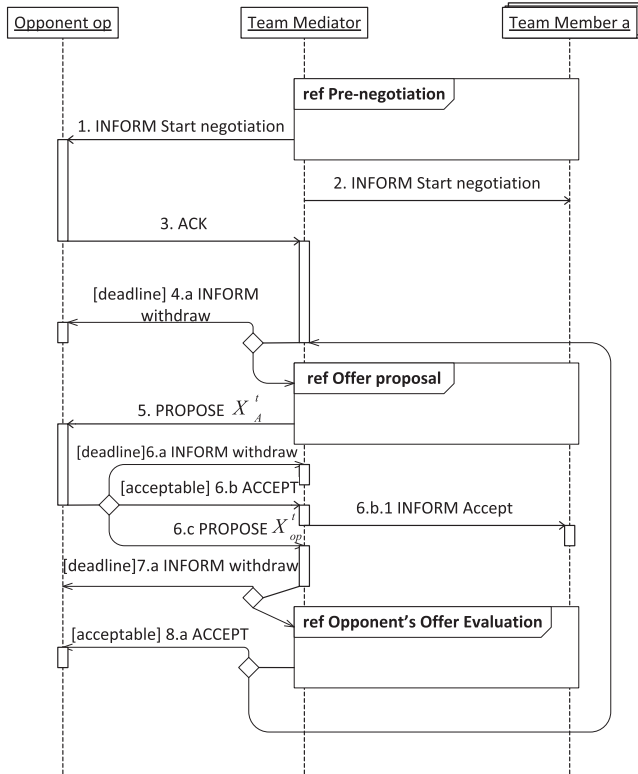


Fig. 2. Overview of the communications carried out by the team mediator.

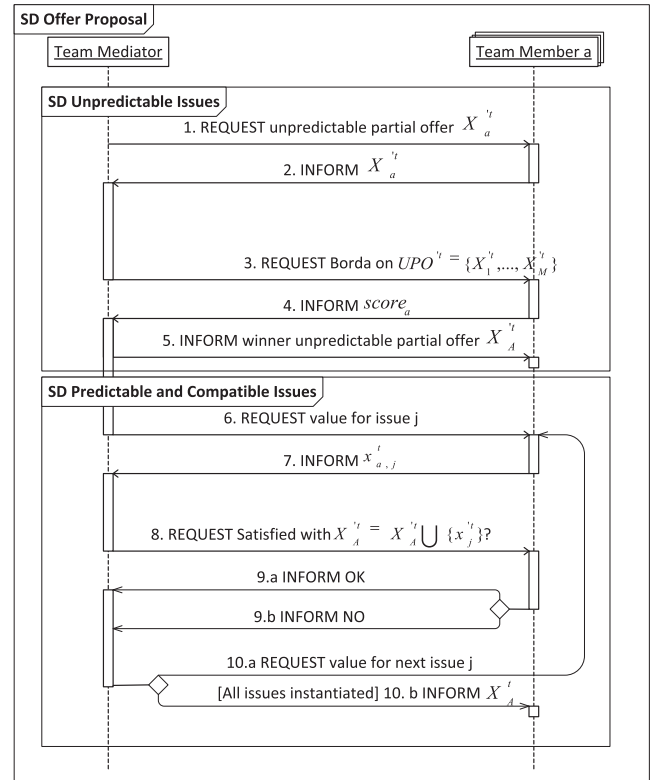


Fig. 5. Overview of the intra-team protocol carried out to propose offers to the other party.




Fig. 3. Overview of the intra-team protocol carried out during the pre-negotiation.

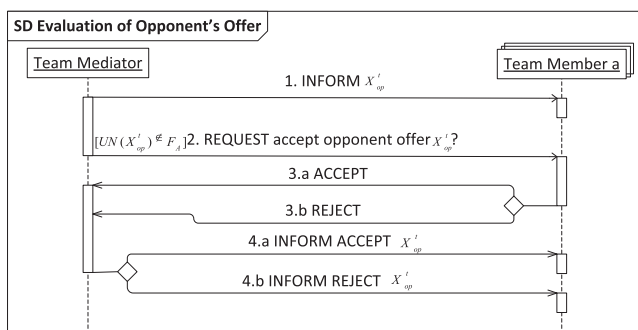


Fig. 4. Overview of the intra-team protocol employed to evaluate opponent's offers.

unpredictable partial offers F_A for the team. The set of forbidden unpredictable partial offers for the team, F_A , is defined as $F_A = \{X' \mid \exists a \in A, X' \in F_a\}$. This means that any unpredictable partial offer in F_A is never part of an acceptable offer for at least one team member. Thus, these unpredictable partial offers should be avoided for the team since the goal of the negotiation model is reaching unanimously acceptable agreements.

A formal description of the pre-negotiation protocol is presented in Fig. 3. The picture describes the protocol by means of Agent UML sequence diagrams. According to the proposed protocol, the team mediator initiates the pre-negotiation phase by asking each team member a to calculate its own set of forbidden unpredictable partial offers F_a (message 1 in Fig. 3). Each team member builds its own (forbidden) set as requested, and it is communicated to the mediator privately (message 2 in Fig. 3). When the mediator receives the sets from the team members, it aggregates them in order to construct the set of forbidden unpredictable partial offers for team A, $F_A = \bigcup_{a \in A} F_a$. Then, the team mediator makes public the list of forbidden unpredictable partial offers of the team F_A (message 3 in Fig. 3). It should be stated that, since any unpredictable partial offer in this set will prevent one of the team members from reaching its reservation utility, the team is not allowed to generate an offer involving any of these partial offers in F_A . After the team mediator has shared F_A with team members, the negotiation phase starts.

The reader may realize that it is possible that during this phase, most of the *unpredictable* partial offers are pruned. In that case, it means that there is little potential for cooperation among team members. This issue can be observed by the team mediator prior to starting any negotiation process. In that case, the team mediator may suggest the team not to negotiate and save the computational resources used in the negotiation. If the team is not static and can be dynamically formed, it may suggest team members to disband the team and look for other potential partners. However, this team/coalition formation (Gaston and desJardins 2005; Rahwan et al. 2009) is outside of the scope of this work since we focus on studying the performance of the negotiation model. We consider the use of the information provided with forbidden *unpredictable* partial offers for negotiation team formation as a future line of

work. In general, combinations of team members that prune a small portion of the space should be more similar among them, and it should be more easy to achieve cooperation.

3.2. Negotiation phase

In the negotiation, two mechanisms are carried out at each round: a mechanism for deciding to accept/reject the opponent's offer (Evaluation of Opponent's Offer), and a mechanism for proposing an offer to the opponent (offer proposal). For the former, a unanimity voting process is employed, while for the latter an offer building process is governed by the team mediator.

3.2.1. Evaluation of Opponent's Offer

This mechanism is carried out each time the team mediator receives an offer from the opponent. Since the main goal of the proposed intra-team strategy is achieving unanimously acceptable agreements for the team, a unanimity voting is carried out to decide whether or not the opponent's offer is acceptable for the team. With this mechanism, as long as one of the team members is not satisfied with the opponent's offer, the offer is not accepted by the team, precluding the team from reaching agreements that are not unanimously acceptable. A formalization of the protocol followed in this mechanism can be observed in Fig. 4. The picture shows the formalization employing sequence diagrams from Agent UML. The intra-team protocol used for this mechanism goes as follows. First, the team mediator receives the offer X^t from the opponent at time t . If X^t involves any forbidden unpredictable partial offer in F_A , the opponent offer is automatically rejected. However, the opponent's offer is also informed to team members in order to allow each team member to process the new information leaked by the opponent if they see it necessary (message 1 in Fig. 4). Otherwise, if the combination of unpredictable issue values is not in F_A , in order to see whether the offer is unanimously acceptable for team members, the mediator makes the opponent's offer public among team members and starts an anonymous voting process (message 2 in Fig. 4). Each team member $a \in A$ states to the mediator whether he is willing to accept X^t (positive vote) or to reject it (negative vote) at that specific instant (messages 3.a or 3.b in Fig. 4). Since our aim is to guarantee unanimity, the offer is only accepted if all of the team members emit a positive vote (message 4.a in Fig. 4). Otherwise, the offer is rejected and a counter-offer is proposed as explained in Section 3.2.2.

3.2.2. offer proposal

Proposing an offer to the opponent is a complex task, since the space of offers may be huge and the preferences of the team members should be reflected in the offer sent to the opponent, and, in our case, the offer sent should be unanimously acceptable for team members. The process is divided into two sub-phases: constructing an unpredictable partial offer, and setting up predictable issues. In both phases, the team mediator acts according to Algorithm 3.1. We include another formal description of the interactions between the mediator and a team member during the offer proposal. This information can be found in Fig. 5, which depicts the intra-team protocol specified in and Agent UML sequence diagram.

- *Constructing an unpredictable partial offer:* The first step is to propose an unpredictable partial offer, a partial offer which has all of the unpredictable issues instantiated. Since team members know from F_A the list of unpredictable partial offers that will not result in unanimously acceptable offers under any circumstance, any offer proposed by the team should avoid being constructed from unpredictable partial offers found in F_A . The method used to propose offers to the opponent relies on the fact that unpredictable issues are those where intra-team

conflict may be present, whereas there is full potential for cooperation in predictable and compatible issues. Hence, in order to build an offer to be sent to the opponent, it seems more appropriate to jointly set unpredictable issue values first and then, depending on the remaining needs of team members, allow team members to set compatible and predictable issues as they require for reaching their demands. The proposed mechanism for the first part, proposing an unpredictable partial offer, is based on voting and social choice. The voting process goes as follows.

1. The mediator asks each team member to anonymously propose one unpredictable partial offer X_a^t (message 1 in Fig. 5).
 2. Each team member privately sends its proposal to the mediator, who gathers all of the proposals in a list that will be later sent to team members. If any unpredictable partial offer proposed by a is contained in F_A , the mediator automatically ignores this proposal (message 2 in Fig. 5).
 3. Once all of the proposals have been gathered, the mediator makes public the list of proposal UPO^t among team members and opens a Borda scoring process (Nurmi 2010) on proposed candidates (message 3 in Fig. 5).
 4. Each team member anonymously scores candidates and sends the scores to the team mediator (message 4 in Fig. 5).
 5. The team mediator sums up scores and selects the winner candidate with the highest score X_A^t , making it public among team members (message 5 in Fig. 5). This candidate, an unpredictable partial offer, will be the base for the full offer that will be sent to the opponent.
- *Setting up predictable and compatible issue values:* Once unpredictable issues have been set, it is necessary to set predictable and compatible issues to construct a complete offer. As it has been stated along this article, there is full potential for cooperation among team members in these issues since increasing the utility of one of the team members by selecting one issue value will result in the other team members staying at the same utility or increasing their utility. Obviously, the selected unpredictable partial offer will not satisfy equally the needs of all the team members. Nevertheless, team members can make use of predictable and compatible issues to satisfy their remaining needs while not generating conflict inside the team. To complete the partial offer X_A^t , an iterative mechanism that we proposed in (Sanchez-Anguix et al. 2012a) is used to build the final offer issue per issue. The mechanism follows an order for predictable and compatible issues that is constructed by the mediator at each round according to the history of the opponent's concessions. The rationale used to build this order is that the opponent would concede less on those predictable issues more important for him in the first negotiation rounds, whereas it would concede more on those predictable issues that are less important. Thus, the order established by the team mediator attempts to order predictable and compatible issues in ascending order of importance for the opponent. The general idea behind this ordering is attempting to satisfy team members' demands with those predictable issues less important for the opponent first. The order is updated as new information becomes available from the offers sent by the opponent. Based on this order, the iterative mechanism goes as follows.
6. The mediator selects the first predictable issue j and asks team members, given the current partial offer X_A^t , the necessary value x_j for j to get as close as possible to their current demands (message 6 in Fig. 5).
 7. Accordingly, each team member a informs the mediator privately about the most convenient value $x_{a,j}$ for that issue (message 7 in Fig. 5). To decide the final value x_j for the issue j , the trusted mediator aggregates agents' opinions (since the issue is predictable) by means of a function that, for

team members, returns the most preferred issue value from a given set ($best(\cdot)$). After deciding the value x_j , X_A^t is updated with x_j .

8. The mediator asks the team whether or not the new partial offer is already satisfactory at round t (message 8 in Fig. 5).
9. Each team member emits an affirmative response if the current partial offer covers its current demands and a negative response if it still has not covered its demands (message 9.a or 9.b in Fig. 5). Those agents that agree with the current state of X_A^t leave the iterative mechanism for this offer since they already are satisfied with the current partial offer. The process steps back to the selection of the next issue.
10. The process continues until all of the predictable issues have been set or until all of the team members have left the iterative mechanism. In the latter case, the remaining issues are set attempting to maximize the opponent's preferences. Once the offer is complete, it is announced among team members and sent to the opponent (message 10.b in Fig. 5).

Algorithm 3.1. Pseudo-algorithm for the offer construction from the point of view of the mediator. Send (*message* \rightarrow *condition*) means that *message* is sent to every agent that fulfills *condition*

```

1:
2: /*Proposing an unpredictable offer*/
3: Send (REQUEST  $X_A^t \rightarrow \forall a \in A$ )
4: Receive (INFORM  $X_A^t \leftarrow \forall a \in A$ )
5:  $UPO^t = (\bigcup_{a \in A} X_A^t) - F_A$ 
6: Send (REQUEST Borda on  $UPO^t \rightarrow \forall a \in A$ )
7: Receive (INFORM  $score_a \leftarrow \forall a \in A$ )
8:  $X_A^t = \operatorname{argmax}_{X' \in UPO^t} \sum_{a \in A} score(a, X')$ 
9: Send (INFORM  $X_A^t \leftarrow \forall a \in A$ )
10:  $order = build\_predictable\_order()$ ;  $A' = A$ 
11:
12: /*Setting predictable issues*/
13: for all  $j \in order$  do
14:   Send (REQUEST value for  $j \rightarrow \forall a \mid a \in A'$ )
15:   Receive (INFORM  $x_{a,j} \leftarrow \forall a \mid a \in A'$ )
16:    $x_j = best(\{x_{a,j} \mid a \in A'\})$ 
17:    $X_A^t = X_A^t \cup \{x_j\}$ 
18:   Send (REQUEST Satisfied with  $X_A^t? \rightarrow \forall a \in A'$ )
19:   for all  $a \in A'$  do
20:     Receive (INFORM  $ac'_a(X_A^t) \leftarrow a$ )
21:     if  $ac'_a(X_A^t) = true$  then
22:        $A' = A' - \{a\}$ 
23:     end if
24:   end for
25:   if  $A' = \emptyset$  then
26:     break;
27:   end if
28: end for
29: for all  $j \in order \wedge issue\_not\_set(j)$  do
30:    $x_j = maximize\_for\_opponent(j)$ 
31:    $X_A^t = X_A^t \cup \{x_j\}$ 
32: end for
33:  $X_A^t = X_A^t$ 

```

4. Team members' strategies

The team mediator defines the coordination mechanisms inside the team. However, each team member's internal strategy has a great effect on team dynamics. In this article, we propose two types

of strategy for team members. According to the first strategy (i.e., our basic team member), the team member only proposes unpredictable partial offers based on its own utility. In the second strategy, team members model the preferences of the team and the opponent on unpredictable partial offers. Then, in the mechanism employed to set unpredictable partial offers, each team member selects the candidate that guarantees that it can reach its current aspirations at time t , and maximizes the probability of being acceptable for the opponent and the team. The learning mechanism employed by these team members is Bayesian learning (i.e., Bayesian team member).

4.1. Basic strategy for team members

Since negotiations are time-bounded in our framework, we consider that team members have to perform some kind of concession if an agreement is to be found. For this purpose we have designed basic team members as agents whose demands are controlled by an individual and private concession strategy. More specifically, the concession strategy for a team member $a \in A$ is based on time-based tactics $s_a(t)$ (Faratin et al. 1998, Lai et al. 2008). It estimates the utility demanded by a at time t by using the formula in Eq. (5), where RU_a is its reservation utility, T is the negotiation deadline, and β_a is the concession speed, which determines how fast the agent's demands are lowered towards RU_a .

$$s_a(t) = 1 - (1 - RU_a) \times \left(\frac{t}{T}\right)^{\frac{1}{\beta_a}} \quad (5)$$

Based on this concession tactic, each team member participates in the intra-team protocol with their demands regulated by his private concession tactic. Next, we define how team members take their decisions: evaluating the opponent's offer, and proposing an offer for the opponent.

4.1.1. Evaluation of opponent's offer

Given an offer X^t proposed by the opponent at instant t , the team member emits a positive vote in the unanimity voting process if it reports a utility greater than or equal to its current demands $s_a(t)$. Otherwise, a negative vote is emitted.

$$ac_a(X^t) = \begin{cases} \text{true} & \text{if } s_a(t) \leq U_a(X^t) \\ \text{false} & \text{otherwise} \end{cases} \quad (6)$$

4.1.2. Offer proposal

As documented in Section 3.2, team members interact at three points during the offer proposal. First, they propose an unpredictable partial offer to the team mediator. Since each team member a has its demands regulated by a time-based tactic, when proposing an unpredictable partial offer to the mediator at instant t , the proposed unpredictable partial offer X_a^t fulfills:

$$X_a^t \notin F_A \wedge (U_a(X_a^t) + \max PR_a \geq s_a(t)) \quad (7)$$

Hence, agent a selects an unpredictable partial offer which is not forbidden inside the team (since it will be ignored by the team mediator) and whose utility allows him to achieve or surpass its current demands at time t . This way, the team member assures that if its proposed unpredictable partial offer is the winner of the Borda voting process, it can reach its current demands. However, one should be aware that many unpredictable partial offers may fulfill Eq. (7). Therefore, it is necessary to select one of them as the proposed candidate. Being our basic team member, from the set of partial offers that fulfill Eq. (7), a team member selects one of the candidates randomly.

The second time that a team member interacts with the team mediator is for scoring unpredictable partial offers that have been

proposed by team members. For scoring candidate partial offers in the Borda voting process, a basic team member orders the candidates according to the partial utility reported by each of the candidates. That is, the team member assigns the highest score to the partial offer whose utility is the highest for itself, and the second highest score to the partial offer whose utility is the second best one, and so forth.

Finally, team members also interact with the mediator during the mechanism used to set *predictable and compatible* issues. When team members are asked about a value for issue j , each team member communicates anonymously the value $x_{a,j}$. The value is the one that, given the current partial offer X_A^t , gets the utility of the new partial offer as close as possible to its current aspiration $s_a(t)$. Taking normalized additive utility functions, it can be calculated as:

$$x_{a,j} = \begin{cases} \operatorname{argmax}_{x \in D_j} V_{a,j}(x) & \text{if } U_a(X_A^t) + w_{a,j} \leq s_a(t) \\ \operatorname{argmin}_{x \in D_j} V_{a,j}(x) & \text{otherwise} \\ x \in D_j \\ \wedge \\ w_{a,j} V_{a,j}(x) \geq (s_a(t) - U_a(X_A^t)) \end{cases} \quad (8)$$

where $s_a(t)$ is the utility demanded by the agent a at round t , $U_a(X_A^t)$ is the utility reported by the current partial offer, and $w_{a,j} V_{a,j}(x)$ is the weighted utility reported by the value demanded by the agent. The value asked for issue j is the closest one to the current demands of the agent. On the one hand, if the agent cannot reach its current demands by just setting issue j , it asks for the value that reports the highest utility. On the other hand, if the agent can reach or surpass its current demands by setting j , it asks for the value that makes the new partial offer the closest to the current demands. In the same iterative process, team members still have to declare whether or not they are satisfied with the different partial offers that are constructed. Team members follow a similar criterion to the method proposed to determine if an opponent offer is acceptable at t . Basically, a partial offer is acceptable for an agent a at t if it reports a utility greater than or equal to the aspiration level marked by its concession strategy:

$$ac'_a(X_A^t) = \begin{cases} true & \text{if } U_a(X_A^t) \geq s_a(t) \\ false & \text{otherwise} \end{cases} \quad (9)$$

where *true* indicates that the partial offer is acceptable at its current state for agent a , and *false* indicates the opposite.

4.2. Bayesian-based strategy for team members

The Bayesian-based negotiation strategy for a team member is based on modeling the team's (as a whole) and its opponent's preferences on unpredictable issues, and acting accordingly. For this purpose, two Bayesian models are employed to predict if unpredictable partial offers are acceptable for both teammates and the opponent. One of the Bayesian models is employed to capture the preferences of the team on unpredictable issues, whereas the other is used for capturing the preferences of the opponent on unpredictable issues. The strategy used to evaluate the opponent's offer is the same as the one described in the basic strategy.

4.2.1. Bayesian learning

Bayesian learning is a probabilistic learning method based on Bayes' theorem (Russell and Norvig 2003). Given a certain set of hypothesis H and some observation e , Bayesian learning attempts to compute the probability $p(h | e)$ that a certain hypothesis h is true after observing e . In our case, we want to determine whether or not the proposed offer will be acceptable for the opponent (or the team) ($H = \{acc, \neq gacc\}$) given a certain unpredictable partial

offer ($e = X^t$) where *acc* stands for "acceptable" and $\neq gacc$ stands for "unacceptable".

Since we assume that there is no interdependence among negotiation issues, we can consider that each negotiation issue contributes individually to the acceptability of an offer/unpredictable partial offer. Thus, applying Bayes' theorem under independence assumption we have:

$$p(acc | X^t) = \frac{p(acc) \prod_{j \in UN} p(x_j | acc)}{\sum_{H \in \{acc, \neq gacc\}} p(H) \prod_{j \in UN} p(x_j | H)} \quad (10)$$

where $p(acc)$ is the prior probability for an unpredictable partial offer to be acceptable, $p(\neq gacc)$ is the prior probability for an unpredictable partial offer to be non-acceptable, and $p(x_j | acc)$ is the conditional probability that assuming X^t as acceptable, it has the value of the j issues instantiated to x_j .

We consider positive examples S_{acc} as those examples that correspond to the acceptable hypothesis (*acc*) and negative examples $S_{\neq gacc}$ as those examples that correspond to the not acceptable hypothesis ($\neq gacc$). For the opponent's model, we employ unpredictable partial offers that have appeared in opponent's offers as positive examples, and unpredictable partial offers that appear in offers rejected by the opponent as negative samples. For the team's model, we use F_A and those opponent's offers rejected by team members as the set of negative examples. Winners in the Borda votings (i.e., unpredictable partial offers contained in offers sent to the opponent) are considered as positive examples. For computing $p(x_j | h)$, we use the proportion between the number of times that x_j appears in hypothesis h (*acc* or $\neq gacc$) and the total number of examples for h :

$$p(x_j | h) = \frac{|\{x_j \in S_h\}|}{|S_h|} \quad (11)$$

The reasons for employing Bayesian learning are varied. The most important one is that it allows online updating of the model as new samples become available. This is important in a process like negotiation, where at each interaction new information becomes available regarding the opponent's/teammates' preferences. If a computationally expensive learning mechanism was used, it would not be possible to include the new information in the model as it becomes available. Furthermore, the learning mechanism is computationally cheap since it mainly involves counting. This is also important in a real application since it allows for simultaneous negotiation threads to exist, which should be maintained with different opponents to look for the best alternatives in an electronic marketplace.

4.2.2. Offer proposal

Bayesian models are employed to help in the selection of the unpredictable partial offer that is proposed to the other team members. Bayesian team members propose at t unpredictable partial offers in the set defined in Eq. (7). Bayesian models help to select a candidate from that set.

However, it is reasonable to think that in the first interactions Bayesian model do not accurately represent other agents' preferences. For that purpose, a team member invests part of the negotiation time t_{exp} in exploring the negotiation space and collecting information regarding the opponent's and the team's preferences. As long as the negotiation process has not surpassed t_{exp} , the team member just selects randomly one of candidate unpredictable partial offers as basic team members do. Meanwhile, the Bayesian models are continuously updated with the new information that becomes available during the negotiation. After reaching the time threshold, the team member starts to use Bayesian models in order to select the unpredictable partial offer to be proposed to the mediator during the offer proposal phase.

The heuristic used in the selection of the candidate is proposing an unpredictable partial offer that is both acceptable for the team and the opponent. The model has an additional parameter named p_{esc} . It represents the probability of avoiding the Bayesian proposal model and using the random proposal model as described in the basic team member model when the negotiation time has gone beyond t_{exp} . This parameter is included in the model in order to: (i) further explore the negotiation space; (ii) escape from local optima induced by inaccurate Bayesian models (e.g., wrong samples, limited number of samples, etc.). We can formalize the selection as follows:

$$X_a^t = \begin{cases} \operatorname{argmax}_{X \in B} \sum w_b p_b(\operatorname{acc} | X)_{b \in \{A, op\}} & \text{if } \begin{cases} \text{rand} \leq p_{esc} \\ \wedge \\ t \geq t_{exp} \end{cases} \\ \text{random_partial_offer}(B) & \text{otherwise} \end{cases} \quad (12)$$

where B is the set of candidate unpredictable partial offers that fulfill Eq. (7), rand is a random number, $p_A(\operatorname{acc} | X)$ is the probability for a candidate unpredictable partial offer to be acceptable for the team, $p_{opp}(\operatorname{acc} | X)$ is the probability for the candidate unpredictable partial offer to be acceptable for the opponent, and w_A and w_{op} ¹ represent the weights given to the acceptability of the unpredictable partial offer for the team and the opponent, respectively (i.e., we will refer to them as Bayesian weights). Varying these Bayesian weights allow team members to show different behaviors depending on their inclination to satisfy either the team or the opponent with the unpredictable partial offer.

5. Provably Unanimously Acceptable Decisions

As stated in the introduction, one of our research goals is proposing a negotiation team model that is able to guarantee unanimously acceptable team decisions. Next, we show that under the assumption of rationality,² team members are able to achieve unanimously acceptable final agreements, if an agreement is found. For that matter, let us employ *reductio ad absurdum* (reduction to absurdity).

If X is the final agreement, let us suppose that Eq. (2) (unanimously acceptable) is violated in a negotiation: unanimity is not reached because a obtained a utility below its reservation utility.

$$\exists a \in A, U_a(X^t) + \sum_{j \in PR} w_{aj} V_{aj}(x_j) < RU_a \quad (13)$$

The final agreement is found when (1) team members accept an opponent's offer or (2) the opponent accepts a team's offer. Next, we show that in both cases, Eq. (13) is never true.

1. When the team members accept an opponent's offer, a unanimity voting process has been carried to decide whether or not to accept the final offer. The offer is only accepted if all of the team members have emitted a positive vote. Since a rational agent a would never have incentive to emit a positive vote of the offer reported a utility below its reservation utility, this scenario is never true due to the intra-team mechanism.
2. When the opponent accepts a team's offer X , this offer has been necessarily proposed by the intra-team mechanism mentioned in Section 3.2.2. The offer can be decomposed into an unpredictable partial offer X^t and the instantiation of predictable issues. The team member a is not able to get over its reservation utility if and only if $X^t \in F_a$ or when $X^t \notin F_a$ and a could not get what

it demanded in predictable issues. A rational agent has no incentive to exclude a forbidden unpredictable partial offer X^t when declaring F_a . Since $F_a = \bigcup_{a \in A} F_a$ and the mediator ignores unpredictable partial offers in F_a , an unpredictable partial offer X^t that forms a team offer is never in F_a . If $X^t \notin F_a$ then the agent can accomplish to satisfy the following expression $U_a(X^t) + \max PR_a \geq RU_a$. Agent a could not get over its reservation value because he could not demand the most of predictable issues. However, when the team mediator aggregates predictable issues inside the team, the team mediator selects the highest value for team members in the list of values proposed by them. This makes possible for a to obtain the maximum utility from predictable issues. Hence, Eq. (13) never holds if the negotiation ends with agreement.

Since both possible scenarios are never true under our initial assumption, we have shown by *reduction ad absurdum* that, if a final agreement is found, it is unanimously acceptable among team members. Another issue is the presence of exaggerating agents (i.e., agents that exaggerate to get the most from the negotiation). In our setting, even if team members exaggerate and decide to include in F_a unpredictable partial offers that are acceptable but report low utility, or demand more than they need from predictable issues, if a final agreement is found it will be unanimous among team members. However, by doing so, they may be pruning negotiation space and lowering the probability of finding agreements. This is an interesting situation that we plan to study in the future.

6. Experiments

In this section, we explore the behavior of the proposed negotiation model in different scenarios. The proposed framework has been implemented in GENIUS (Lin et al. 2014), a simulation framework for automated negotiation that allows researchers to test their frameworks and strategies against state-of-the-art agents designed by other researchers. Recently, GENIUS has become a widespread tool that increases its repository of negotiating agents with the annual negotiation competition (Baarslag et al. 2012).

In order to assess the performance of the proposed negotiation approach, we have performed different experiments. All of the experiments have been carried out in the negotiation domain (or case study) introduced in Section 2.4. The first experiment (Section 6.1) studies the performance of the proposed model when facing single opponent agent. The comparison is carried out in scenarios with different degrees of team's preference dissimilarity. In the second experiment, we study the performance of our negotiation team model when facing another negotiation team in bilateral negotiations. In the third experiment (Section 6.3) we study how the Bayesian weights w_A and w_{op} , which control the importance given to the preferences of the team and the opponent in the unpredictable partial offer proposed to teammates, impact the performance of the proposed model when team members employ the Bayesian strategy. Finally, we conduct an experiment to study the effect of team members' reservation utility on the performance of the proposed negotiation model (Section 6.4).

6.1. Performance Against a Single Opponent

In this first set of experiments we study the performance of the proposed negotiation team model when facing a single opponent. The study is carried out with an emphasis on observing if the performance of the team is higher by employing Bayesian team members rather than basic team members. Due to the fact that we are interested in open environments, we study how the team

¹ $w_A + w_{op} = 1$.

² Rational agents seek to improve their current welfare. Thus, they would not take actions that lead to utilities below their reservation utilities.

performance varies with team configurations ranging from scenarios where no team members plays the Bayesian strategy (i.e., all Basic team members) to situations where all the team members play the Bayesian strategy. The performance of the team is measured using the average team joint utility.³ As an additional measure of optimality, we also measure the distance to the closest Pareto optimal point. In this case, the Pareto frontier is computed taking into account the team joint utility and the utility of the opponent. Our initial hypotheses were:

- **H1** As more Bayesian team members form the team, the team is able to obtain average team joint utilities that are higher than or equal to those configurations with less Bayesian team members.
- **H2** As more Bayesian team members form the team, the team is able to obtain average opponent utilities that are higher than or equal to those configurations with less Bayesian team members.

Since Bayesian team members take the preferences of the team and the opponent into account when deciding which offer is sent to the other part, they should be able guarantee equal or higher average team joint utility and opponent utility than basic team members. But in no case, they should not obtain lower team joint utility and opponent utility. As a consequence of both hypothesis, the distance to the closest Pareto optimal point should be also equal or lower than that obtained by configurations solely composed by basic team members.

In order to compare the proposed model with other models in the literature, we also included the Similarity Borda Voting model (i.e., SBV) (Sanchez-Anguix et al. 2011) in our experiment. SBV is a mediated intra-team strategy that is able to guarantee semi-unanimity regarding team decisions. The mediator imposes a unanimity voting process to decide on whether or not to accept the opponent's offer, while team members propose offers to be sent to the opponent by means of a similarity heuristic that takes into account the last offer proposed by the opponent, and the last offer proposed by the team. A Borda voting process is used in order to decide on which offer is sent to the opponent. The reason to include this intra-team strategy in our study is due to the fact that it has been documented to achieve similar results to intra-team strategies that guarantee unanimity under certain circumstances for domains solely composed by predictable issues (Sanchez-Anguix et al. 2011). In order to adapt this approach for domains with unpredictable issues, we use a similarity heuristic that uses Euclidean distance for real/integer issues and string matching for other types of issues. Due to the fact that our proposed model guarantees unanimously acceptable agreements and SBV does not, we formulated the following hypothesis:

- **H3** Teams exclusively formed by basic team members and teams exclusively formed by Bayesian team members obtain equal or higher average team joint utility than teams following the Similarity Borda Voting model.

The performance of this first experiment is analyzed in three scenarios with different degrees of preference dissimilarity among team members: very similar, average similarity, and very dissimilar preferences. For this reason, we introduce a measure for measuring team members' preference similarity in different scenarios. The authors proposed a method for calculating preference dissimilarity in teams based on the utility difference of offers among teammates (Sanchez-Anguix et al. 2011). The dissimilarity between two teammates $a, b \in A$ can be measured as:

$$D(U_a(\cdot), U_b(\cdot)) = \frac{\sum_{\forall X} |U_a(X) - U_b(X)|}{\# \text{ possible offers}} \quad (14)$$

Due to the fact that a team may be composed of more than two members, it is necessary to provide a team dissimilarity measure. The team dissimilarity measure is calculated as the average of the dissimilarity between all of the possible pairs of teammates. For this experiment, we decided to explore teams whose preferences are dissimilar, teams whose preferences are similar, and teams with an average degree of similarity/dissimilarity (i.e., average similarity). For the scenario of dissimilar preferences, 9 negotiation cases were randomly generated (i.e., a combination of 3 different negotiation teams consisting of four team members with 3 different opponents), while 9 negotiation cases were randomly generated for the similar preferences scenario (i.e., a combination of 3 different negotiation teams consisting of four team members with 3 different opponents) and 12 negotiation cases were randomly generated for the average similarity scenario (i.e., a combination of 4 different negotiation teams consisting of four team members with 3 different opponents). Since we consider that in practice it is less likely to meet extreme cases such as dissimilar or similar teams, we decided to increase the number of negotiation cases in the average similarity scenario.

As for the single opponents, we decided to test the negotiation team models against different families of opponents. More specifically, we followed the categorization of negotiation strategies proposed by Baarslag et al. (Baarslag et al. 2011), which divides negotiation strategies into four categories:

- **Competitors:** They hardly concede, independently of opponent behavior. **Agent K** is a competitor agent (Kawaguchi et al. 2011, Baarslag et al. 2011) from the 2010 ANAC competition (Baarslag et al. 2012) that adjusts its aspirations (i.e., target utility) in the negotiation process considering to an estimation of the maximum utility that will be offered by the other party. More specifically, the agent gradually reduces its target utility based on the average utility offered by the opponent and its standard deviation.
- **Matchers:** They concede when they perceive that the opponent concedes, and they do not concede if they perceive that the other party does concede. **Nice Tit-for-Tat** is a matcher agent (Baarslag et al. 2011, 2013) from the 2011 ANAC competition that reciprocates the other party's moves by means of a Bayesian model of the other party's preferences. According to the Bayesian model, the Nice Tit-for-Tat agent attempts to calculate the Nash point and it reciprocates moves by calculating the distance of the last opponent offer to the aforementioned point. When the negotiation time is reaching its deadline, the Nice TFT agent will wait for an offer that is not expected to improve in the remaining time and accept it in order to secure an agreement.
- **Conceders:** They yield independently of the opponent behavior. **Conceder** is an implementation of the time-based concession tactics proposed by Faratin et al. (Faratin et al. 1998) categorized by Baarslag et al. as conceder. For the Conceder agent $\beta_{op} = 2$, which leads to large concessions towards the reservation utility in the first rounds.
- **Inverter:** They respond by implementing the opposite behavior shown by the other party. **Boulware** is an implementation of the time-based concession tactics proposed by Faratin et al. (Faratin et al. 1998) categorized by Baarslag et al. as inverter. In the case of the Boulware agent, the concession speed is set to $\beta_{op} = 0.2$. Hence, the agent concedes very insignificantly during most of the negotiation and it concedes very quickly as the deadline approaches.

The reservation utility of each team member was set to $RU_a = 0.5$ to represent scenarios where team members have average aspirations. Additionally, for each team member (i.e., basic,

³ We consider the joint utility of the team to be the product of the utilities of the team members. Since the utility of an agent is between 0 and 1, the team joint utility tends to be lower as more team members are present.

Bayesian and SBV) the concession speed was randomly selected from a uniform distribution $\beta_a = U[0.5, 1]$. In the case of Bayesian members, the time of exploration was set to $t_{exp} = 70\%$ and the probability of escape after the exploration phase was set to $p_{esc} = 30\%$.⁴ Therefore, Bayesian models are not used unless a 70% of the negotiation time (126 s) has passed. Initially, we set Bayesian members to care equally about the acceptability of unpredictable partial for the team and the opponent $w_A = w_{op} = 0.5$.

Following the type of setting used in the annual agent competition, the negotiation time was set to $T = 180$ seconds. Each opponent strategy was faced against each negotiation team model in every possible negotiation case. A total of 20 repetitions were done per negotiation case in order to capture stochastic variations in negotiation strategies. Therefore, $3 \times 3 \times 3 \times 4 \times 20 = 2160$ (team preference profiles \times opponent preference profiles \times team negotiation models \times opponent strategies \times repetitions) negotiations were simulated in the similar scenario, 2160 negotiations were simulated in the dissimilar scenario, and 2880 negotiations were simulated in the average similarity scenario.

Table 1 shows the average team joint utility and opponent utility for the cases where all of the team members either play the Basic strategy, the Bayesian model, or the team employs the SBV team negotiation model. It also shows the Euclidean distance to the closest point in the Pareto frontier. An ANOVA test ($\alpha = 0.05$) with a Bonferroni post hoc analysis was carried out to assess statistical differences among the different measures gathered. Those measures that are statistically the best configurations for each column are highlighted in bold style. All of the claims in this experimental section are supported by the ANOVA test with the Bonferroni post hoc analysis. The average negotiation time taken by each method is included in Table 2. Fig. 7 shows the evolution of the average team joint utility and the average opponent utility as more team members play the Bayesian strategy. We have also included some examples of agreements obtained in the different negotiation scenarios and how they relate to the Pareto frontier.⁵ These results can be observed in Fig. 6. Next, we analyze the results.

6.1.1. Results for the first hypothesis

- **H1** As more Bayesian team members form the team, the team is able to obtain average team joint utilities that are higher than or equal to those configurations with less Bayesian team members.

First, we focus on the situations when all of the team members either play the basic strategy or the Bayesian strategy (Tables 1,2). It can be observed that when team members' preferences are similar, both types of teams perform equally in terms of the average team joint utility. This result is consistent with H1, since both prove to be statistically equivalent with the ANOVA test with Bonferroni post hoc analysis. The reason why Bayesian models do not give an advantage over the basic model in the similar scenario can be explained since team members are similar and there is no necessity to carry out team modeling. The distance to the closest Pareto optimal point is also very similar for both team configurations, which can be also explained due to the fact that team modeling is not necessary due to team members' similar preferences.

As conflict is introduced inside the team by making team members' preferences more dissimilar (i.e., *dissimilar and average similarity scenarios*, middle and bottom part of Table 1), it can be observed that usually the team formed exclusively by Bayesian

⁴ These values were found to be the best ones after carrying out a grid search over values of t_{exp} and p_{esc} in a subset of test negotiation scenarios.

⁵ The quadratic root of the team joint utility is taken to convert the results to the same scale (remember that the team joint utility is the product of for team members' utilities).

team members gets the statistically highest average team joint utility, which is also coherent with our hypothesis H1 and refines our hypothesis for these scenarios. In this case, the teammates' preferences are no longer similar and some sort of modeling mechanism is needed in order to guide the intra-team negotiation towards agreements that are good for all of the team members. The only exception is found in the conceder case, where the performance in terms of the team joint utility was found to be statistically equivalent among the team exclusively formed by basic team members and the team exclusively formed by Bayesian members. Taking a closer look at the negotiation traces, we observed that, in all of the negotiations, the exploration time t_{exp} was never surpassed. Since the Conceder agent concedes rapidly in the negotiation, the team's demands are also met early. Therefore, Bayesian models do not get to be used. In fact, the average negotiation time against Conceder agents was 61.7, 76.3, and 88 s respectively in the similar scenario (see Table 2), the average similarity scenario (see Table 2), and the dissimilar scenario (see Table 2). All of them are below the threshold of 126 s delimited by t_{exp} . As a result, the team members have not used their Bayesian model while generating their proposals and they are equivalent to the team formed by basic members. This is also consistent with H1, since in no case the team formed by Bayesian members gets statistically lower results than the team formed by basic members.

If we observe the evolution of the average team joint utility in Fig. 7,⁶ there is a tendency to increase the average team joint utility as more Bayesian members are included (triangle shaped data series on Fig. 7) in situations where team members' preferences have an average similarity or they are very dissimilar. This tendency is more pronounced against Agent K (left plot in Fig. 7) and and Boulware agents (right plot in Fig. 7). However, when team members' preferences are very similar the team performance remains at statistically equivalent values (ANOVA test with Bonferroni analysis) for the average team joint utility. The results of these graphics are coherent with our findings in Table 1 and H1.

In conclusion, we have found that as more Bayesian team members form the team, the team is able to obtain average team joint utilities that are higher than or equal to those configurations with less Bayesian team members. Being more specific, we have been able to detect that, as long as there is preferential conflict among team members (i.e., average similarity among team members, and very dissimilar team members), and the opponent does not concede early in the negotiation, more Bayesian team members result in higher team joint utility.

6.1.2. Results for the second hypothesis

- **H2** As more Bayesian team members form the team, the team is able to obtain average opponent utilities that are higher than or equal to those configurations with less Bayesian team members.

For the average opponent utility, the Bayesian team obtained significantly better results than the basic team only in the scenario where the team faces Agent K and team members are dissimilar with regards to their preferences (see Table 1, bottom part). The same pattern is found in Fig. 7, where we can observe that the average opponent utility increases as more Bayesian team members are present in the team (first plot in the second row of Fig. 7). In other cases, the Bayesian and the basic team obtain statistically equivalent results to each other, and, in some situations, the basic team obtained significantly better results. More specifically, when the team faces Nice Tit-for-Tat, the basic team obtains significantly better results than the Bayesian team. We can also

⁶ Results for the Conceder agent are omitted since the Bayesian models are not employed.

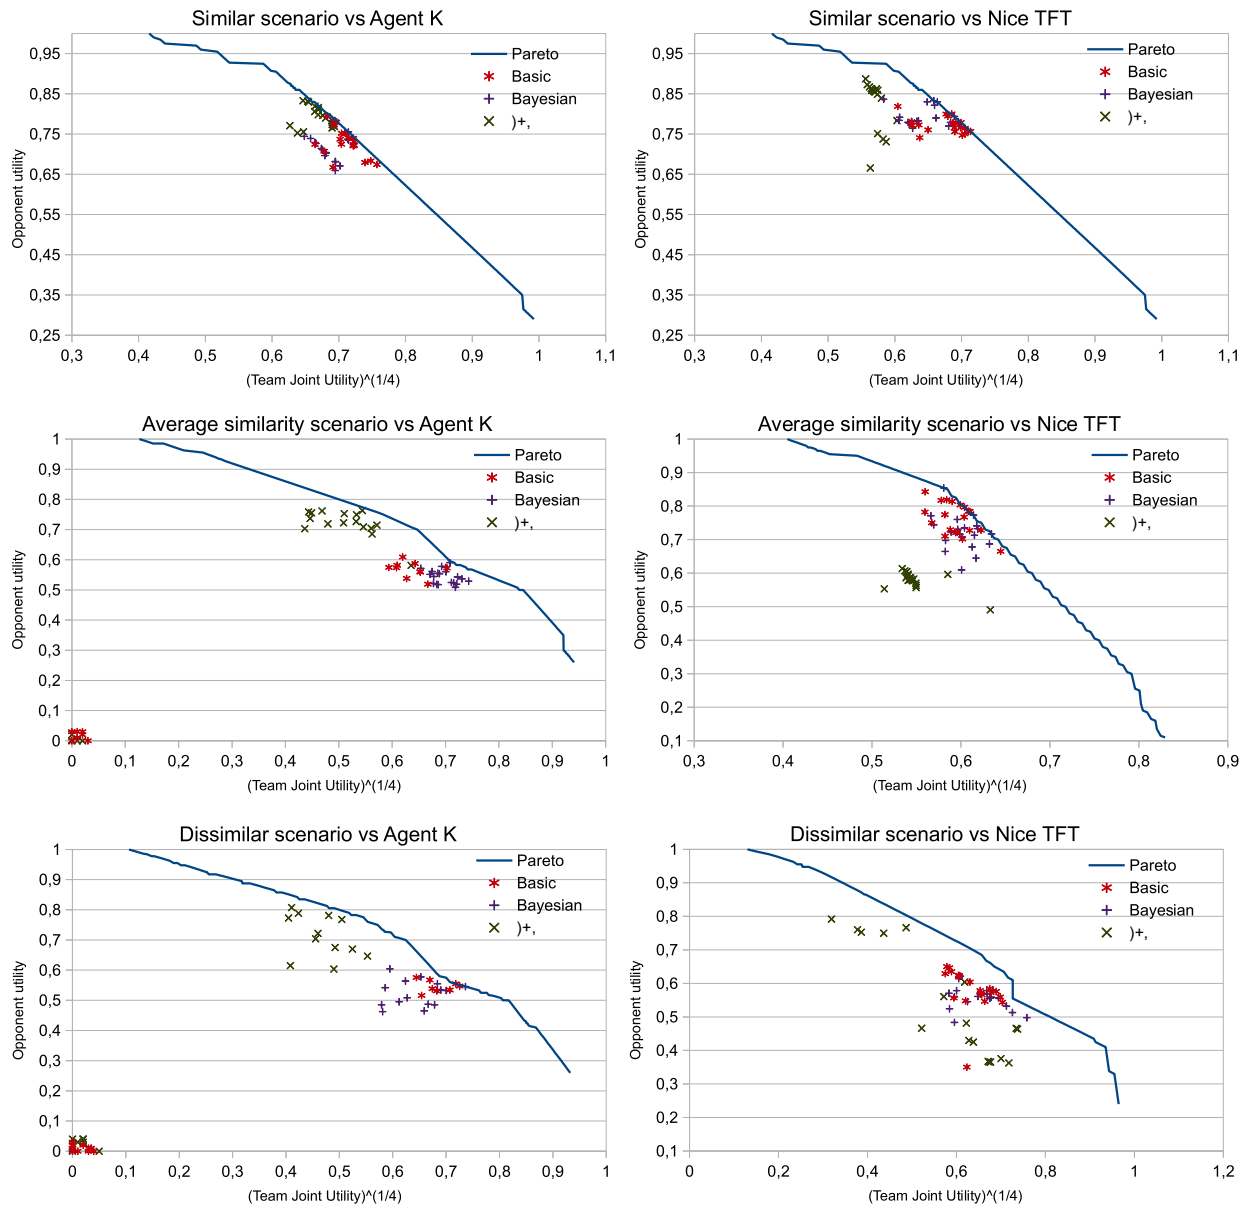


Fig. 6. Examples of agreements and their distance to the Pareto frontier.

Table 1

Average team joint utility (T.), the average opponent utility (Op.), and the average Euclidean distance to the closest Pareto optimal point (D.) for the first set of experiments. Bold values are 95% confidence calculated by ANOVA with post-hoc analysis.

	Agent K			Nice Tit-for-Tat			Boulware			Conceder		
	T.	Op.	D.	T.	Op.	D.	T.	Op.	D.	T.	Op.	D.
<i>Similar</i>												
SBV	0.181	0.743	0.070	0.150	0.694	0.130	0.184	0.755	0.064	0.552	0.482	0.037
Basic	0.259	0.683	0.065	0.173	0.760	0.067	0.223	0.696	0.078	0.561	0.468	0.045
Bayesian	0.263	0.690	0.058	0.164	0.746	0.080	0.224	0.695	0.080	0.557	0.472	0.043
	T. J	Op.	D.	T. J	Op.	D.	T. J	Op.	D.	T. J	Op.	D.
<i>Average similarity</i>												
SBV	0.168	0.629	0.065	0.137	0.562	0.116	0.170	0.598	0.070	0.324	0.428	0.074
Basic	0.211	0.574	0.070	0.141	0.691	0.050	0.210	0.585	0.060	0.386	0.414	0.052
Bayesian	0.248	0.583	0.034	0.158	0.669	0.047	0.224	0.574	0.045	0.390	0.414	0.050
	T. J	Op.	D.	T. J	Op.	D.	T. J	Op.	D.	T. J	Op.	D.
<i>Dissimilar</i>												
SBV	0.07	0.522	0.168	0.160	0.457	0.157	0.128	0.547	0.110	0.257	0.430	0.110
Basic	0.174	0.397	0.180	0.184	0.572	0.055	0.254	0.505	0.053	0.472	0.367	0.046
Bayesian	0.209	0.457	0.118	0.196	0.559	0.60	0.271	0.489	0.058	0.475	0.367	0.044

Table 2
Average time (seconds) for negotiations in the first set of experiments. K (Agent K), N. TFT (Nice Tit-for-Tat), B. (Boulware), C. (Conceder).

	K	N. TFT	B.	C.
<i>Similar</i>				
SBV	148.7	164.3	139.4	61.0
Basic	141.5	162.5	145.4	62.8
Bayesian	142.3	165.2	144.0	61.7
<i>Average similarity</i>				
SBV	155.1	177.6	153.1	74.1
Basic	153.1	174.8	154.9	77.4
Bayesian	150.1	175.0	154.7	76.3
<i>Dissimilar</i>				
SBV	163.8	175.7	156.6	73.4
Basic	162.2	176.6	160.1	87.7
Bayesian	163.7	177.2	160.5	88.0

observe this pattern in Fig. 7. As more Bayesian members are introduced, the average opponent utility slightly decreases (middle plot in the second row of Fig. 7).

These findings only support partially our hypothesis H2 since we found a set of scenarios where the basic team provides better utility to the opponent (i.e., when facing Nice Tit-for-Tat). We analyzed the trace of different negotiations against Nice Tit-for-Tat and Boulware opponents. In the former case, we could observe that close to the end of the negotiation the Nice Tit-for-Tat opponent had only sent on average 5 different unpredictable partial offers in a domain that has 4200 different unpredictable partial offers. This behavior results in scarce information for any learning mechanism. In the case of negotiations against Boulware agents, one should consider that the Boulware strategies concede only towards the end of the negotiation and, most of the time, the aspirations are high. Thus, most of the samples gathered by the Bayesian classifier when facing Boulware agents correspond to offers with high demands where usually only the best issue values appear. Other issue values do not appear in the samples or they have their frequency misinterpreted with respect to the utility that they actually report. Therefore, the learning mechanism misinterprets the preferences of the opponent and the team formed by Bayesian members is not able to obtain statistically better utility for the opponent than the team formed by basic members.

The behavior of Nice Tit-for-Tat and Boulware opponents also has a direct consequence on the distance to the closest Pareto optimal point obtained by both team configurations (Bayesian and basic). Despite the fact that the Bayesian configuration is capable of obtaining statistically better results for the team joint utility by learning the preferences of the team, the utility reported to the opponent is usually lower than the one reported by the basic configuration. The only exception to this case are scenarios against Agent K, where the Bayesian configuration obtains a statistically higher utility for the opponent. The inability to model the opponents' preferences in the case of the Boulware and Nice Tit-for-Tat opponents, results in a higher team joint utility (due to team modeling) at the cost of reducing the utility received by the opponent. Hence, there is not an improvement in the distance to the closest Pareto optimal point. These findings can also be observed in some of the examples included in Fig. 6. The agreements found by the basic team configuration and the Bayesian configuration tend to be found at the same distance to the Pareto frontier (the Bayesian configuration tending to populate regions with higher team joint utility). The exception to this rule are negotiations against Agent K, where the basic configuration tends to populate regions of no agreement (close to the axis origin).

6.1.3. Results for the third hypothesis

- **H3** Teams exclusively formed by basic team members and teams exclusively formed by Bayesian team members obtain equal or higher average team joint utility than teams following the Similarity Borda Voting model.

It can be observed that when *team members' preferences are similar* (top part of Table 1), both basic and Bayesian models are statistically equivalent to each other and better than SBV with respect to the average team joint utility. As conflict is introduced inside the team by making team members' preferences more dissimilar (i.e., *dissimilar and average similarity scenarios*, middle and bottom part of Table 1), the team get statistically lower average team joint utility by employing the SBV model. Basic and Bayesian models outperform SBV with respect to the average team joint utility since they are able to guarantee unanimously acceptable agreements,

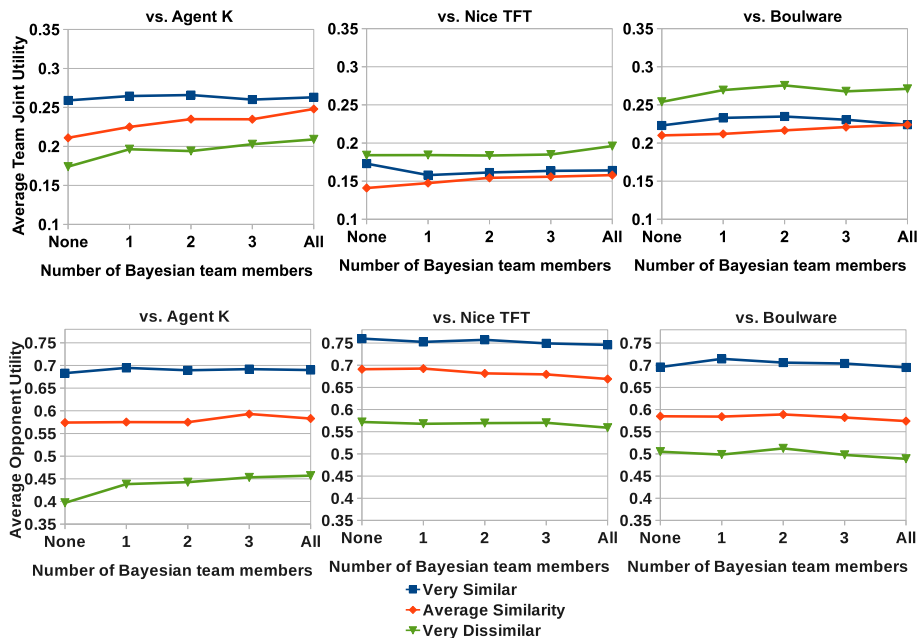


Fig. 7. Evolution of the average team joint utility and opponent utility as more Bayesian team members are introduced in the team.

while SBV does not guarantee such condition. This result supports and refines our hypothesis H3, since, in general, the team formed exclusively by Bayesian team members and the team formed by basic team members obtains statistically higher team joint utilities than the SBV model. There is only one exception to this refinement. The basic team, the Bayesian team and SBV perform statistically equal in terms of the average team joint utility only when the opponent is a conceder and team members' preferences are similar. Since the opponent concedes rapidly in the first rounds, the three models obtain equivalent team joint utility due to the opponent concessions and the fact that conflict is almost nonexistent among team members. However, this finding is still consistent with our initial guess H3.

Regarding the optimality of the solutions found by the model proposed in this article with respect to SBV, it is possible to observe that both Bayesian and basic configurations obtain statistically lower distance to the Pareto frontier than SBV as long as preferential conflict between team members is present (i.e., average similarity and dissimilar scenarios). Only when the team is very similar, SBV gets statistically equal distances to the model proposed in this article. However, despite the fact that the distance to the Pareto frontier may be statistically equal, we can observe in Fig. 6 that agreements found by SBV tend to populate areas that are closer to the lowest team joint utility. This result can be explained due to the fact that unanimity cannot be guaranteed in the team, and some team members end up with low utility agreements.

In conclusion, we found that teams exclusively formed by basic team members and teams exclusively formed by Bayesian team members obtain equal or higher average team joint utility than teams following the Similarity Borda Voting model. More specifically, we found that, in general, the results for the team formed by Bayesian team members and the team formed by basic team members obtains statistically higher results than the SBV model. The distance to the Pareto frontier also shows higher quality and more optimal agreements for teams employing the model proposed in this article. This result is important, since it shows that the present model, not only guarantees unanimously acceptable agreements, but it also ensures that better results are obtained with respect to other state-of-the-art team negotiation models.

6.2. Performance against another team

In this experiment we analyze the performance of the proposed model when two teams face each other. More specifically, we simulate negotiations where one negotiation team represents a group of four travelers, and the other negotiation team represents the board of managers for a hotel, which consists also of four managers. The negotiation has a common deadline of $T = 180$ seconds. Both parties negotiate with each other by means of the alternating bilateral protocol, but they employ the negotiation team model proposed in this article to coordinate and take team decisions. The goal of the experiment is determining if Bayesian team member improve the performance of the team when it faces another team.

In this setting, we use 4 different team preference profiles to represent the group of travelers and 2 different team preference profiles to represent the board of managers. Since we are interested in open environments, we consider different team configurations from the perspective of strategy profiles: No team member plays the Bayesian strategy (0–0), half of the team members play the Bayesian strategy in one team (2–0), all of the team members play the Bayesian strategy in one team (4–0), half of the team members are Bayesian in both teams (2–2), all of the team members are Bayesian in one team and half of them are Bayesian in the other team (4–2), and all of the team members are Bayesian

(4–4). For each strategy profile, each group of travelers' is faced against each board of managers 20 times to capture stochastic variations. Therefore, a total of $4 \times 2 \times 20 = 160$ negotiations is carried out per team strategy profile, giving a total of 960 negotiations for this experiment.

Our initial guess is that more Bayesian team members will help to obtain higher team joint utilities due to the learning and proposal mechanism used to take into account the preferences of the team and the preferences of the opponent. More specifically, we formulated the following hypotheses:

- **H4** As long as only one team includes Bayesian team members (configurations 2–0 and 4–0), the average team joint utility for both teams will be higher than the average team joint utility obtained by negotiations where no Bayesian team member participates (configuration 0–0).
- **H5** Those configurations where both teams include Bayesian team members (configurations 2–2, 4–2, and 4–4) will obtain higher average team joint utilities for both teams than configurations where only one team includes Bayesian team members (configurations 2–0 and 4–0).

We configure the parameters as we did in our previous experiment. The results of the experiment can be observed in Fig. 8. The blue points represent the average team joint utility for the group of travelers, whereas the red points represent the average team joint utility for the board of managers. The graphic shows an increasing average team joint utility as the total of Bayesian team members in both sides increases. The worst results for both teams are obtained when all of the team members act as the basic team member (0–0). This result is explainable due to the large number of negotiations that finished with no agreement (80 out of 160 negotiations, a 50% of failure). As long as one of the sides implements the Bayesian strategy (2–0, 4–0), both teams benefit by obtaining higher average team joint utilities. An ANOVA test with Bonferroni post hoc analysis ($\alpha = 0.05$) confirmed that both configurations 2–0 and 4–0 obtain statistically different and higher team joint utilities for both teams than the configuration 0–0. This result confirmed our initial hypothesis H4.

It can be observed that the next relevant increase in the average team joint utility of both teams is present as long as both sides apply the Bayesian strategy (2–2, 4–2, 4–4). An ANOVA test with Bonferroni post hoc analysis ($\alpha = 0.05$) revealed that the averages obtained by configurations 2–2, 4–2, and 4–4 are statistically different and higher than the averages obtained by 2–0 and 4–0. Hence, both teams obtain higher average team joint utilities as long as both teams include Bayesian team members, supporting our hypothesis H5.

In conclusion, we have been able to determine that when two teams face each other by means of the proposed model, both teams benefit by obtaining higher team joint utilities when Bayesian team members participate in the negotiation. This is especially




Fig. 8. Evolution of the average team joint utility with the different strategy profiles (Blue = group of travelers, Red = board of managers). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

true when Bayesian team members are distributed between both teams, which obtains the highest team joint utilities for both sides.

6.3. Analyzing the impact of Bayesian weights for the proposal of unpredictable partial offers

Recalling from Section 4.2, there are two weight parameters that control how important the opponent's and team's preferences are while generating the unpredictable partial offer (respectively w_{op} , w_A). w_A represents how important it is for the team members to make an unpredictable partial offer that is acceptable for the team, whereas w_{op} represents how important it is for us to make an unpredictable partial offer that is acceptable for the opponent. The use of these weights is not trivial, since one should consider that, it only refers to the acceptability of the unpredictable partial offer by one of the two parties. A complete offer is composed by the predictable and unpredictable issues. Therefore, for instance, using a high value of w_{op} may not have the desired effect on the opponent unless unpredictable issues are important for the opponent. Additionally, one should also consider that the more acceptable an unpredictable is for the team/opponent, the more utility it should report.

In this experiment we explore the impact of these weights in a wide variety of situations. More specifically, we study how different values for w_{op} and w_A affect situations where the team gives more importance to unpredictable issues than the opponent, situations where the opponent gives more importance to unpredictable issues than the team, and situations where both team and the opponent give the same importance to unpredictable issues.

To assess the importance given by an agent to unpredictable partial issues, we consider the sum of unpredictable issue weights in its utility function.

$$I_a = \sum_{j \in UN} w_{i,j} \quad (15)$$

We consider that when $I_a \in [0.0, 0.33]$ the agent a gives low importance to unpredictable issues, when $I_a \in [0.33, 0.66]$ it gives average importance to unpredictable issues, and when $I_a \in [0.66, 1.0]$ the agent gives high importance to unpredictable issues. We generated 8 random negotiation cases where team members give a high importance to unpredictable issues and the opponent gives low (4 cases) and average (4 cases) importance to unpredictable issues, 8 different randomly generated negotiation cases where team members give a low importance to unpredictable issues and the opponent gives average (4 cases) and high (4 cases) importance to unpredictable issues, and 12 negotiation cases where the opponent and the team give the same importance to unpredictable issues (4 cases where both give low importance, 4 cases where both give average importance, and 4 cases where both give high importance).

We tested three configurations for Bayesian teams: standard Bayesian members that give the same importance to the acceptability of the unpredictable partial offer by the opponent and the team $w_A = w_{op} = 0.5$ (Normal), Bayesian members that give more importance to the acceptability of the unpredictable partial offer by the opponent $w_A = 0.25$ $w_{op} = 0.75$ (Opponent Oriented), and Bayesian members that give more importance to the acceptability of the unpredictable partial offer by the team $w_A = 0.75$ $w_{op} = 0.25$ (Team Oriented). As for the opponent's strategies, we selected Agent K, Nice TFT and Boulware.

In situations where the opponent gives more importance to unpredictable issues than the team, the team should be able to obtain higher average team joint utility by playing high values for w_{op} , satisfying opponent needs with unpredictable issues, and demanding more on predictable issues. If the team gives more importance to unpredictable issues, the team should be able to

obtain higher average team joint utility by playing high values for w_A , satisfying opponent needs with predictable issues, and demanding more on unpredictable issues. If both teams give the same importance to unpredictable issues, the team should be able to guarantee higher average team joint utility by giving the same weight to w_{op} and w_A since no party gives more importance to unpredictable issues. Attending to these initial guesses, we formulated the following hypotheses:

- **H6** When the opponent gives more importance to unpredictable issues, the team obtains higher team joint utility by using high values for w_{op} and demanding more on predictable issues ($w_A = 0.25$ $w_{op} = 0.75$).
- **H7** When the team gives more importance to unpredictable issues, the team obtains higher team joint utility by using high values for w_A and letting the opponent demand on predictable issues ($w_A = 0.75$ $w_{op} = 0.25$).
- **H8** When both parties give the same importance to unpredictable issues, the team obtains higher team joint utility by giving the same importance to the opponent's and team's preferences on unpredictable issues ($w_A = w_{op} = 0.5$).

For each negotiation case, we repeated the negotiation 20 times in order to capture stochastic variations in strategies. Therefore, a total of 1680 negotiations were carried out in this experiment. The results of this experiment can be observed in Table 3. It shows the average of the joint team utility and proportion of negotiations that finished with success (agreement) in the scenario. An ANOVA test ($\alpha = 0.05$) with Bonferroni post hoc analysis was carried out to detect statistically different averages. The best configurations for each of the three scenarios are highlighted in bold font style. All of the claims of this experimental section are supported with the aforementioned ANOVA test with post hoc analysis.

Focusing on the case of negotiations against Agent K, it can be appreciated that when unpredictable issues are more important for the opponent (bottom part of Table 3), the best results are obtained by taking an opponent oriented approach: proposing unpredictable partial offers that are likely to be acceptable for the opponent and satisfy remaining members' aspirations by demanding on predictable and compatible issues, which are less important for the opponent. This finding supports H6. As for the scenario where unpredictable issues are more important for the team (middle part of Table 3), it is clearly observed that the best choice for team joint utility is to give a high weight to w_A , thus employing a team oriented approach. Since unpredictable issues are more important for the team, they should satisfy their needs as much as possible with proposed unpredictable partial offers and demand less on predictable issues, which are more important for the opponent, supporting our initial hypothesis H7. Finally, the last scenario corresponds to the case where unpredictable issues have the same importance for both parties (top part of Table 3). In this case, there may be more conflict between the team and its opponent since the parties do not have a clear trade-off opportunity such as increasing the demand on unpredictable issue while decreasing the demand on predictable issues as appeared in two previous cases. One can observe that the best team joint utility is obtained when using the standard team members ($w_A = w_{op} = 0.5$). Since both parties give the same importance to unpredictable issues, it seems natural to give the same importance to the acceptability of the unpredictable partial offer by the team and the opponent, supporting our hypothesis H8. The team oriented approach is clearly worse than the rest of approaches since many negotiations (only a 58.7% were successful. See top part of Table 3) ended in failure due to the team being too demanding and not satisfying the opponent's preferences.

Table 3

Impact of w_A and w_{op} on the average team joint utility in different scenarios and proportion of negotiations that finished with success (% Ag.). Bold values are 95% confidence calculated by ANOVA with post-hoc analysis.

	vs. Agent K		vs. Nice TFT		vs. Boulware	
	T. Joint	% Ag.	T. Joint	% Ag.	T. Joint	% Ag.
<i>Equal importance on unpredictable issues</i>						
Normal	0.168	84.2	0.137	100	0.202	100
Opponent Oriented	0.155	91.1	0.126	100	0.188	100
Team Oriented	0.116	58.7	0.188	100	0.206	100
<i>Unpredictable issues more important for the team</i>						
Normal	0.213	100	0.135	100	0.185	100
Opponent Oriented	0.200	100	0.154	100	0.189	100
Team Oriented	0.248	100	0.196	100	0.213	100
<i>Unpredictable issues more important for the opponent</i>						
Normal	0.280	100	0.186	100	0.326	100
Opponent Oriented	0.296	100	0.192	100	0.300	100
Team Oriented	0.271	92.0	0.259	100	0.340	100

When negotiating against Nice TFT and Boulware, the results are different. It can be appreciated that, generally, the team oriented approach always reports statistically better results from the point of view of the average team joint utility. This means that the team should indistinctly select highly acceptable unpredictable partial offers for the team when facing opponents like Nice TFT and Boulware agents. These findings do not support our hypotheses H6, H7, and H8, and drove a more in-depth analysis and study. There are two factors that should be taken into account. First, as we detected in Section 6.1, the Bayesian models that are learnt from both agents are not adequate due to the lack of learning samples detected in the case of Nice TFT and the misinterpretation of the importance of issue values in the case of the Boulware agent. This factor precludes the agents from finding good agreements for both parties when using an opponent oriented approach in scenarios where unpredictable issues are more important for the opponent. Since the Bayesian models misinterpret the preferences of Boulware and Nice TFT agents, it is not possible to create *win-win* situations. Second, by selecting a team oriented approach, the Bayesian team members are selecting the more acceptable (i.e., the best) unpredictable partial offers for the team. These unpredictable partial offers report high utility for the team. Thus, once completed with predictable issues, these offers are expected to report high utility for the team. Differently to Agent K, Nice TFT and Boulware agents are not competitor agents since they will eventually accept the team's offer when the deadline is approaching. In the case of the Boulware agent, it concedes quickly with respect to its aspirations as the deadline approaches, eventually meeting the requirements of the demanding team offer. As stated in (Baarslag et al. 2013), when the negotiation time is reaching its deadline, the Nice TFT agent will wait for an offer that is not expected to improve in the remaining time in order to secure an agreement. Hence, the team is able to exploit the other party by selecting the team oriented approach. Due to these circumstances, hypotheses H6, H7, and H8 were not supported for agents Boulware and Nice Tit-for-Tat.

In conclusion, hypotheses H6, H7 and H7 are only partially supported. In general, we have discovered that depending on the type of opponent agent, the values for weights w_{op} , w_A have different effects on the negotiation. When facing exploitable agents like Nice TFT and Boulware, the team benefits if the team members take the team oriented approach, selecting those unpredictable partial offers that are more acceptable for the team. If the team faces competitors like Agent K, the team should match the values for w_{op} and w_A depending on the importance given by each party to unpredictable issues. It is acknowledged that, depending on the opponent

and the desired behavior, the team should select different values for the Bayesian proposal weights. A mechanism for adjusting Bayesian weights based on the type of opponent is considered as future work.

6.4. Analyzing the impact of the reservation utility

In this experiment, we investigated the impact of the reservation utility of team members on the team joint utility. As explained in Section 3.1, team members jointly prune a part of the negotiation space (i.e., a set of unpredictable partial offers) which does not contain, with absolute certainty, any unanimously acceptable offer. This pruning is related with the reservation utility of team members, which represents the minimum acceptable utility by team members. Any offer with a utility lower than the reservation utility is not acceptable for the agent.

Lower reservation utilities make it easier to obtain the needed utility by just setting compatible and predictable issues. Thus, each team member needs to prune less negotiation space with the unpredictable partial offers sent to the team mediator. Presumably, a joint list of forbidden unpredictable partial offers (i.e., the negotiation space that is pruned) with lower reservation utilities is smaller than lists constructed with higher reservation utilities. This leaves more room for finding an agreement with the opponent. However, if team members have low reservation utilities, despite having more room for finding an agreement, they may end up with low utility agreements in the end. On the contrary, with higher reservation utilities, it is harder to obtain the needed utility with compatible and predictable issues. Therefore, each team member may need to prune more negotiation space and the joint list of forbidden unpredictable partial offers will be larger than the list constructed with lower reservation utilities. In fact, if team members set high aspirations with their reservation utility, it may end up with all the negotiation space being pruned. If an agreement is found under these conditions, it may lead to team members achieving high levels of utility.

In this experiment, we test the impact of different levels of reservation utility on the team joint utility. More specifically, as we did in Section 6.1, we tested teams employing the Bayesian model against different families of strategies: competitor (i.e., Agent K), matcher (i.e., Nice Tit-for-Tat), inverter (i.e., Boulware), and conceiver (i.e., conceiver). As an additional dimension to our study, we also introduced preference similarity among team members. Therefore, teams are tested in the scenario where team members' preferences are dissimilar, the scenario where team members' preferences have an average degree of similarity, and the scenario where team members' preferences are similar.

We configured three types of Bayesian teams with different levels of reservation utilities: a team with a relatively low reservation utility $RU_a = 0.35$, a team with a moderate reservation utility $RU_a = 0.5$, and a team with a high reservation utility $RU_a = 0.65$. We expect that playing higher reservation utilities against competitor like Agent K will result in lower average team joint utility due to many negotiations ending with no agreement. On the other hand, we expected that playing high reservation utilities against Conceder agents, inverter agents like Boulware, and the special case of Nice Tit-for-Tat, would result in higher average team joint utility due to the fact that both agents should eventually concede towards the other parties' demands. In the case of Nice Tit-for-Tat, we have observed that when the deadline is approaching it attempts to secure a deal, making this agent candidate to be exploited by setting a high reservation utility. Therefore, we formulated the following hypotheses:

- **H9** Playing a high reservation value for team members ($RU_a = 0.65$) will result in the lowest average team joint utility against Agent K.
- **H10** Playing a high reservation value for team members ($RU_a = 0.65$) will result in the highest average team joint utility against Conceder, Boulware, and Nice Tit-for-Tat.

These types of teams ($RU_a = 0.35, RU_a = 0.5, RU_a = 0.65$) were faced in every scenario and negotiation case against every type of opponent for 20 repetitions. We gathered information on the team joint utility and the utility obtained by the opponent, and an ANOVA ($\alpha = 0.05$) with Bonferroni post hoc analysis was carried out to determine results that are statistically better than the rest. All of the claims of this experimental section are supported by the ANOVA test with the Bonferroni post hoc analysis.

Table 5 shows the results of this experiment in terms of the average joint utility. A bold font style is used to highlight those Bayesian team configurations that are statistically the best option against each opponent. Additionally, the same table shows the proportion of negotiations that finished with success (agreement) in this experiment. Table 4 shows the average percentage of unpredictable partial offers that were pruned in the pre-negotiation

depending on the team configuration and team preference similarity.

With respect to H10, it can be observed that independently of the degree of dissimilarity among team members' preferences, team members obtained statistically better team joint utility by setting high reservation utilities ($RU_a = 0.65$) against Boulware, Conceder, and Nice Tit-for-Tat opponents. These results support our initial hypothesis H10.

Nevertheless, H9 is only partially supported. When facing Agent K, playing the highest reservation utility configuration lead to the highest team joint utilities when team members' preferences are similar or they have an average similarity (top and middle part of Table 5), which is opposite to our initial hypothesis H9. In this case, the demands of the team are still not high enough to preclude agreements with the competitor agent. However, when team members' preferences are dissimilar (bottom part of Table 5), we can observe how setting a high reservation utility (i.e., $RU_a = 0.65$) gradually becomes the worst possible course of action when facing Agent K, making H9 true for this situation. The reason for this behavior is mainly explained due to the decrease in the number of successful negotiations. If we observe the similar scenario, the number of successful negotiations when facing Agent K and $RU_a = 0.65$ is 99.98%. If we change to average similarity scenarios, the number of successful negotiations is 74.2% when facing Agent K with $RU_a = 0.65$. The same measure is decreased to 18.3% in the dissimilar scenario. If we observe Table 4, as team dissimilarity increases, the number of unpredictable partial offers to be pruned is larger. This leaves less negotiation space to be played with Agent K. Agent K is a competitor agent that attempts to concede as less as possible by estimating the maximum utility that can be obtained from the opponent and employing a limit of compromise when the opponent takes a hard stance. First of all, if reservation utilities are high, it can be considered that team members play a hard stance. Second, if too much negotiation space is pruned, it may be feasible that the set of remaining unpredictable partial offers precludes Agent K from reaching its limit of compromise. Thus, employing high reservation utilities against a competitor agent like Agent K may result, as we have observed in this case, in an increase in the number of failed negotiations and lower team joint utilities.

In conclusion, H10 is supported by our findings, but H9 is only partially supported. We have observed that team members may benefit from playing high reservation utilities against Conceder, Boulware, and Nice Tit-for-Tat. If faced against competitors like Agent K, setting high reservation utilities may prune too much negotiation space, especially when team members are dissimilar. This results in negotiation spaces that may not contain the minimum limits established by competitor agents, thus, ending negotiations with failure. For other scenarios against agent K, the

Table 4
Average percentage of unpredictable partial offers pruned in the pre-negotiation.

	Similar (%)	Average similarity (%)	Dissimilar (%)
$RU_a = 0.35$	0.4	11.6	35.3
$RU_a = 0.50$	23.8	34.2	72.6
$RU_a = 0.65$	73.7	81.8	90.8

Table 5
Average joint Utility (T. Joint) for teams composed by Bayesian team members with different reservation utilities and proportion of negotiations that finished with success (% Ag.). Bold values are 95% confidence calculated by ANOVA with post-hoc analysis.

	Agent K		Nice Tit-for-Tat		Boulware		Conceder	
	T. Joint	% Ag.	T. Joint	% Ag.	T. Joint	% Ag.	T. Joint	% Ag.
<i>Similar</i>								
$RU_a = 0.35$	0.195	100	0.117	100	0.160	100	0.526	100
$RU_a = 0.50$	0.263	100	0.164	100	0.224	100	0.557	100
$RU_a = 0.65$	0.350	99.98	0.286	100	0.354	100	0.635	100
<i>Average similarity</i>								
$RU_a = 0.35$	0.167	100	0.090	100	0.136	100	0.342	100
$RU_a = 0.50$	0.248	100	0.158	100	0.224	100	0.390	100
$RU_a = 0.65$	0.242	74.2	0.268	100	0.313	100	0.470	100
<i>Dissimilar</i>								
$RU_a = 0.35$	0.193	100	0.115	100	0.173	100	0.373	100
$RU_a = 0.50$	0.209	86.6	0.196	100	0.271	100	0.475	100
$RU_a = 0.65$	0.068	18.3	0.346	100	0.409	100	0.580	100

team can benefit from setting high reservation utilities since the resulting negotiation space still has room to accommodate an agreement with the competitor agent. In general, team members should be cautious when setting the reservation utility since it may end up in more failures.

6.5. Team performance with risk attitudes

One scenario that should be considered in multi-agent systems is agent's attitude towards risk. Some team members may be more willing to choose actions that guarantee a safer agreement, while other may prefer to go for more profitable but less probable agreements. Classically, agents can show a risk seeking, a risk averse, and a risk neutral attitude. The goal of this experiment is determining how risk attitudes affect the performance of the proposed negotiation model. In this experiment, we decided to test three different patterns of behavior:

- *Risk averse team member*: This team member selects from the list of available unpredictable partial offers (see Eq. (7)), those that guarantee the current aspirations of the team member $s_a(t)$ the best unpredictable partial offer according to the acceptance probability for the opponent by using the Bayesian mechanism proposed in this article. When the team sets predictable issue, the maximum utility obtainable with other unpredictable partial offers may be higher since they may provide a higher utility of oneself. However, in order to secure a deal, the team member selects the unpredictable partial offer that is supposed to be more acceptable to the opponent even if the maximum achievable utility is lower. Therefore, the team member bases its choices on the acceptability of the offer by the opponent party instead of the maximum achievable utility.
- *Risk seeking team member*: This team member selects for the list of available unpredictable partial offers (see Eq. (7)) the best unpredictable partial offer according to the utility reported by one's own utility function. Hence, this unpredictable partial offer represents the choice that enables the team member to get the maximum achievable utility in the negotiation domain. By filling predictable issues by the team, the team member is more likely to get closer to its maximum achievable utility, even if this action reduces acceptability of the proposed offer for the opponent.
- *Risk neutral team member*: This behavior is represented by the Bayesian team member presented in this article.

Our initial hypothesis is that the proposed model is robust to risk attitudes. By robustness, we mean that the proposed negotiation model will be able to obtain a team joint utility higher than or comparable to the team joint utility obtained by other state-of-the-art models like SBV:

- **H11** Different configurations of team members' risk attitudes will yield a team joint utility that is statistically higher or equal to the results obtained by SBV.

The experimental parameters of the previous experiments were repeated ($RU_a = 0.5$, $\beta_a = U[0.5, 1]$), and we selected four different team configurations. The first one is composed by four neutral team members (Neutral), the second is composed by four risk seeking team members (Seeker), the third is composed by three risk averse team members (Averse), and the last team configuration is composed by two neutral team members, one risk seeker, and one risk averse (Mix). As in the previous experiments, the team members faced different opponent profiles in scenarios where team members' preference profiles are similar, scenarios with average similarity of team members' preferences, and scenarios

Table 6

Team joint utility obtained by the different team risk configurations. K (Agent K), N, TFT (Nice Tit-for-Tat), B. (Boulware), C. (Conceder). Bold values are 95% confidence calculated by ANOVA with post-hoc analysis.

	K	N, TFT	B.	C.
<i>Similar</i>				
Averse	0.241	0.160	0.238	0.372
Neutral	0.259	0.164	0.224	0.557
Mix	0.250	0.216	0.267	0.520
Seeker	0.169	0.248	0.283	0.615
SBV	0.181	0.150	0.184	0.552
<i>Average similarity</i>				
Averse	0.188	0.147	0.186	0.308
Neutral	0.248	0.158	0.224	0.390
Mix	0.231	0.171	0.222	0.383
Seeker	0.149	0.162	0.224	0.477
SBV	0.168	0.137	0.170	0.324
<i>Dissimilar</i>				
Averse	0.203	0.167	0.245	0.373
Neutral	0.209	0.196	0.271	0.475
Mix	0.228	0.183	0.265	0.452
Seeker	0.060	0.247	0.294	0.566
SBV	0.07	0.16	0.128	0.257

where team members have dissimilar preferences. The team faces the same opponent strategies presented in the previous experiments. We gathered information on the team joint utility.

The results of this experiment can be found on Table 6. Results highlighted in bold represent the statistically higher team joint utility configurations for each negotiation scenario (ANOVA, $\alpha = 0.05$). As it can be observed, the risk seeking configuration usually gets a higher team joint utility as long as the opponent faced is not a competitor. This configuration is able to obtain one of the best team joint in all negotiation scenarios with Nice Tit-for-Tat, Boulware, and Conceder. When negotiating against Agent K, many negotiations end with failure since the opponent also has high aspirations. Other opponents like Nice Tit-for-Tat, Boulware, or Conceder match the high aspirations of the risk seeking team at one point or another of the negotiation, resulting in a higher team joint utility.

The mixed configuration usually gets the next highest team joint utility, being able to obtain statistically equivalent results in some scenarios like negotiating against Nice Tit-for-Tat in the average similarity scenario, and statistically higher results like negotiations against Agent K in dissimilar scenarios. This configuration obtains the best team joint utility in all of the scenarios involving against Agent K.

The neutral configuration, obtains statistically good results in some scenarios where team members have similar preferences and the Bayesian mechanism is able to learn the opponent preferences properly (e.g., negotiating against Agent K, some scenarios against Boulware and Nice Tit-for-Tat). The risk averse configuration is only able to obtain some of the top results in very specific scenarios like negotiations against in similar settings against Agent K.

However, in no case any of the team configurations proposed in this article is worse than SBV. All of the results are statistically higher than those obtained by SBV, except for those obtained by the risk seeking configuration against Agent K, which are statistically equivalent to those obtained by SBV. Therefore, despite being affected by team heterogeneity, the proposed model is robust and it is able to obtain results that are at least equal to the state-of-the-art (i.e., SBV), outperforming it in many situations. This result supports our initial hypothesis H11.

7. Related work

The contributions of this article to the automated negotiation community can be divided into two different categories: general contributions to the field of automated negotiation, and

contributions to the specific field of agent-based negotiation teams. Next, we analyze the contributions to each of these fields.

7.1. Automated negotiation with single individual parties

The artificial intelligence community has focused on bilateral or multi-party negotiations where parties are composed of single individuals. The main contribution of our work with respect to the general field of automated negotiation is that we support negotiation parties composed by multiple individuals. Apart from that, in the next paragraphs we discuss other contributions of our present negotiation framework with respect to works in automated negotiation.

Faratin et al. (Faratin et al. 1998) introduced some of the most widely used families of concession tactics in negotiation. The authors proposed concession strategies for negotiation issues that are a mix of different families of concession tactics. The authors divide these concession tactics into three different families: (i) time-dependent concession tactics; (ii) behavior-dependent concession tactics; and (iii) resource-dependent tactics. Our negotiation framework also considers time as crucial element in negotiation. Therefore, team members employ time tactics inspired in those introduced by Faratin et al. However, the authors do not propose any explicit preference learning mechanism.

In Zeng and Sycara (Zeng and Sycara 1998), the authors argue about the benefits of using Bayesian models in negotiation and they study a bilateral negotiation case where the buyer attempts to learn the reservation price of the seller by updating its beliefs with Bayesian learning. Despite the fact that it introduces the use of Bayesian learning in negotiation, the article only focuses on single issue models. One of our team member models also uses Bayesian learning as a method for learning other agents' preferences. The main different resides in the fact that our Bayesian approach attempts to model which instantiations of unpredictable issues are acceptable for the opponent and the team in multi-issue negotiations.

Ehtamo et al. (Ehtamo et al. 2001) propose a *mediated* multi-party negotiation protocol which looks for joint gains in an iterated way. The algorithm starts from a tentative agreement and moves in a direction according to what the agents prefer regarding some offers' comparison. Results showed that the algorithm converges quickly to Pareto optimal points. However, the work proposed by Ehtamo et al. does not support unpredictable issues and multiple individual parties.

Klein et al. (Klein et al. 2003a) propose a *mediated* negotiation model which can be extended to multiple parties. Their main goal is to provide solutions for negotiation processes that use complex utility functions to model agents' preferences. The negotiation attributes are binary and no longer independent. Our work supports multiple individual parties and negotiation issues with unrestricted domain (e.g., real, integer, discrete, binary, etc.).

In Coehoorn et al. (Coehoorn and Jennings 2004), the authors propose the use of kernel density estimation for the estimation of the importance weights of the linear additive utility function. The agent calculates tuples composed of the difference between pairs of consecutive offers, the estimated weight for the issue, and the probability density of the weight. These tuples form a three dimensional kernel that is used along the other kernels to calculate an estimation of the real issue weight. Our proposed model is capable of learning in negotiations where domains are also composed by discrete issues, which is not supported by Coehoorn et al. Moreover, the learning mechanism proposed for our team members deals with the information in one single negotiation, whereas the aforementioned mechanism learns over several negotiations.

Later, Narayanan et al. (Narayanan and Jennings 2006) present a negotiation framework where pairs of agents negotiate over a

single issue. The authors assume that agents' strategies may change over time. Non-stationary Markov chains and Bayesian learning are employed to tackle the uncertainty in this domain, and eventually converge towards the optimal negotiation strategy. In our case, we focus on one single negotiation process, and our team members learn over the information provided by the current negotiation. Additionally, we consider negotiations where multiple issues are involved.

Another example of the use of Bayesian learning in negotiation is presented by Buffett et al. (Buffett and Spencer 2007). In the aforementioned article, a bilateral framework is presented in a domain where agents negotiate over a set of binary issues. A Bayesian classifier is employed to classify opponent's preferences into classes of preference relations. Groups of similar preference relations are grouped according to the k-means algorithm prior to the negotiation process. Our model does not require learning prior to the negotiation, which may require a costly learning process every negotiation to avoid domain dependent classifiers. Moreover, we consider any kind of issue type in the negotiation domain.

Carbonneau et al. (Carbonneau et al. 2008) propose a neural network that takes as input the negotiation history of a bilateral negotiation with continuous issues and an offer to make an estimation of the opponent's counter-offer. This approach requires that an artificial neural network is trained per negotiation case. Similarly, the same authors propose an improvement over their previous work in (Carbonneau et al. 2011). It aims to make a predictive model that does not depend on the negotiation case. The model takes pairs of negotiation issues as inputs of the neural network, where one of the issues is considered the primary issue (i.e., independent variable) and the other issue is considered the secondary issue (i.e., dependent variable). Differently to these works, our proposed model does not rely on information from past negotiations. It only employs information gathered in the present process.

Robu et al. (Robu and La Poutré 2008) introduce a bilateral negotiation model where agents represent their preferences by means of utility graphs. The negotiation domain is formed of bundles of items that can be either included or excluded in a final deal. Utility graphs are graphical models that relate negotiation issues that are dependent. Nodes represent negotiation issues whereas arcs connect issues that have some joint effect on the utility function (i.e., positive for complementary issues, and negative for substitutable issues). Hence, utility graphs represent binary dependencies between issues. The authors propose a negotiation scenario where the buyer's preferences and the seller's preferences are modeled through utility graphs. The seller is the agent that carries out a more thorough exploration of the negotiation space in order to search for agreements where both parties are satisfied. With this purpose, the seller builds a model of the buyer's preferences based on historic information of past deals and expert knowledge about the negotiation domain. Our work does not consider dependencies between issues, however it is capable of supporting every type of domain for negotiation issues (i.e., real, binary, discrete, etc.).

In (Aydoğan and Yolum 2012), a concept-based learning method is proposed for modeling opponent preferences and generating well-targeted offers. In that method, the preferences of the opponent are represented via disjunctive and conjunctive constraints. In this article, our aim is also to find the agreement earlier by means of learning the other participants' preferences but the preferences are represented by means of additive utility functions. In our case, an offer that is rejected by the opponent may become acceptable over time because the opponent may concede; therefore, we choose a probabilistic learning method.

Williams et al. (Williams et al. 2012) present a negotiation framework for coordinating multiple bilateral negotiations with different opponents. The agent simultaneously negotiates with

different opponent in order to acquire a desired good at the best possible condition. The framework makes use of optimization techniques and probabilistic information in order to carry out this coordination. Similarly, Mansour et al. (Mansour and Kowalczyk 2012) present a meta-strategy for coordination multiple negotiations with different sellers. The meta-strategy adjusts the concession speed according to the current state of the multiple negotiation threads. The problem of agent-based negotiation teams is different since multiple agents collaborate in the same party to get a deal from an opponent, instead of competing between each other to get an individual deal.

7.2. Agent-based negotiation teams

As far as we are concerned, only our previous works (Sanchez-Anguix et al. 2011, 2012a,b) have considered negotiation teams in computational models. In (Sanchez-Anguix et al. 2011, 2012a) four different computational models for a negotiation team negotiating with a single opponent are presented. These four models attempt to gather four different minimum levels of unanimity regarding team decisions: representative approach (RE, no unanimity), similarity simple voting (SSV, majority/plurality), similarity Borda voting (SBV, semi-unanimity), and full unanimity mediated (FUM, unanimity).

The RE model is based on the selection of team members as representative of the team. The representative acts on behalf of the group by taking decisions according to its own negotiation strategy and utility function.

SSV and SBV are models based on the presence of a mediator that coordinates voting processes. In the case of SSV, a majority voting is employed to determine whether or not the opponent's offer is accepted, and a majority voting is used to select which offer is sent from those offers proposed by team members to be sent to the opponent. In the case of SBV, a unanimity voting is designated as the mechanism to decide if opponent's offers are accepted and Borda count is used to decide which offer proposed by team members is sent to the opponent. In both cases, SSV and SBV team members decide which offers are sent to the opponent using similarity heuristics that consider the opponent's and team's last offers.

FUM is a mediated model where the opponent's offers are evaluated by means of a unanimity voting process, and team's offers are built issue per issue by aggregation rules. These models were introduced at AAMAS 2011 (Sanchez-Anguix et al. 2011) as the first approach to tackle problems involving negotiation teams.

Later, we studied the special properties of FUM, given that it was the intra-team negotiation model capable of guaranteeing unanimity regarding team decisions at each negotiation round (Sanchez-Anguix et al. 2012b). We proved how unanimity is guaranteed in FUM, how the intra-team strategy is robust against certain types of manipulation attacks, and how team members did not have incentive to strongly deviate from the proposed model.

We provided a full fledged experimental analysis of the four intra-team strategies in different negotiation environments (Sanchez-Anguix et al. 2012a). The results showed that FUM was able to guarantee better results for the negotiation team in most negotiation scenarios. Only SBV is able to guarantee similar results in a limited number of negotiation cases.

Even though these four models cover different levels of unanimity with regards to team decisions, they were initially designed to provide solutions for negotiation domains that are exclusively composed by *predictable and compatible* issues among team members (e.g., price, quality, and dispatch time in a team of buyers).

Domains exclusively composed by *compatible and predictable* issues among team members cover a range of feasible negotiation domains. However, a relatively large number of negotiation

domains also include *unpredictable* issues among team members. RE, SSV, and SBV can be easily adapted to domains that include *unpredictable* issues among team members by using a different similarity heuristic. FUM, the model that guarantees unanimity with regards to team decisions, cannot be used in these domains since, in the offer proposal mechanism, it aggregates issue values based on the fact that all of the negotiation issues are *compatible and predictable* among team members.

As stated, whenever it is possible, it is desirable for the final agreement to be unanimously acceptable for all of the team members since the opposite situation may end up in users perceiving unfairness, which may affect commitment to the decision, group attachment, and trust (Korsgaard et al. 1995), and users that are not satisfied with agreements found automatically may end up leaving the electronic application.

The model proposed in this article advances the state of the art in agent-based negotiation teams by solving both problems: it guarantees unanimity with regards to team decisions, and it supports *unpredictable* negotiation issues, which were not specifically supported in previous models. For that matter, the negotiation domains is split between those issues that are *compatible and predictable* among team members, and those issues that are *unpredictable*. In the former case, part of the mechanism employed in FUM is employed. By doing so, the model is capable of guaranteeing that team members are able to get as much as they need from *predictable and compatible* issues. In the latter case, the team discards those combinations of unpredictable partial offers that preclude the team from reaching unanimously acceptable agreements, even if the most is obtained from *predictable and compatible* issues. Then, in the offer proposal mechanism team members select one of remaining unpredictable partial offers, and complete the *predictable and compatible* issue with the values that they need to make it at least unanimously acceptable. The idea behind this splitting, is cooperating as much as possible on those issues that are *predictable and compatible* among team members to create less intra-team conflict in *unpredictable* issues.

8. Conclusions and future work

In this article we have presented a new mediated team negotiation model for a team as a multi-individual party negotiating with an opponent in the alternating offers protocol. The present model is capable of assuring unanimously acceptable agreements for all of the team members. It takes advantage of the categorization of negotiation issues as predictable and compatible, and unpredictable. The former are those issues whose preferential order over issue values is known from the negotiation domain and it is common among team members (e.g., price in a team of buyers), whereas the latter are those issues whose preferential order over issue values is not known in the negotiation domain. In the case of predictable and compatible issues, there is full potential for cooperation among team members since if one of the team members demands more from the issue, the other team members are also benefited. Our negotiation model takes advantage of this property. During the pre-negotiation, each team member shares with a team mediator those unpredictable partial offers (i.e., partial offers that have all of the unpredictable issues instantiated) that, even if the team demands the most from predictable issues, preclude the agent from achieving its reservation value. A joint list forbidden unpredictable partial offers is constructed by the team mediator from the lists received from the team members. In the negotiation, the team mediator coordinates a unanimity voting process to decide whether or not to accept offers received from the opponent. As for the mechanism employed to decide on which offer should be sent to the opponent, the team mediator

coordinates two processes: a proposing and voting process where each team member suggests an unpredictable partial offer not included in the forbidden list followed by a Borda voting on candidates received, and an iterated process where predictable issues are set issue per issue attending to the demands of the team members.

We have proposed two different types of team members for the current model: a basic team member that proposes unpredictable partial offers during the negotiation solely guided by its own utility function, and a Bayesian team member that suggests unpredictable partial offers based on the preferences of the team and the preferences of the opponent. Results have shown that, as long as preferential conflict is present in the team, team members have an incentive to employ the Bayesian strategy over the basic strategy. In any case, we have shown that both approaches outperform other existing models for negotiation teams. Moreover, we have determined that when two negotiation teams face each other, both teams benefit from including Bayesian team members in the negotiation. Finally, we have shown that team members may benefit from playing higher reservation utilities against conceders, matchers and inverters. Nevertheless, setting high reservation utilities may become the worst option as team members' preferences are more dissimilar and the opponent plays a competitor strategy.

The topic of agent-based negotiation teams remains largely unexplored compared to negotiation involving single individual parties. The present work has focused on agent-based negotiation teams where team members have different preferences but they share the same knowledge regarding the negotiation domain. One potential area of research is modeling negotiation teams where team members differ in their knowledge of the negotiation domain and their skills related to the negotiation process. As another line of future work, one could consider the problem of forming negotiation teams based on the individual list of unpredictable partial offers. Lists that are more similar may suggest team members that are more similar, which, if put together in the same team, may result in more cooperative teams. Related to team formation, dynamic negotiation teams where team members may join and leave the team during the negotiation can be considered an alternative line of research.

Acknowledgements

This research is partially supported by TIN2012-36586-C03-01 of the Spanish government and PROMETEOII/2013/019 of Generalitat Valenciana. Other part of this research is supported by the Dutch Technology Foundation STW, applied science division of NWO and the Technology Program of the Ministry of Economic Affairs; the Pocket Negotiator Project with Grant No. VICI-Project 08075.

References

- Aydoğan, R., Yolum, P., 2012. Learning opponent's preferences for effective negotiation: an approach based on concept learning. *Autonomous Agents Multi-Agent Systems* 24, 104–140.
- Baarslag, T., Hindriks, K.V., and Jonker, C.M. Towards a quantitative concession-based classification method of negotiation strategies. In *Agents in Principle, Agents in Practice. Lecture Notes of The 14th International Conference on Principles and Practice of Multi-Agent Systems*, 2011, 143–158.
- Baarslag, T., Hindriks, K., Jonker, C., Kraus, S., Lin, R., 2012. The first automated negotiating agents competition (ANAC 2010). In: *New Trends in Agent-Based Complex Automated Negotiations. Studies in Computational Intelligence*, Vol. 383. Springer, Berlin/Heidelberg, pp. 113–135, ISBN 978-3-642-24695-1.
- Baarslag, T., Hindriks, K.V., Jonker, C.M., 2013. A Tit for tat negotiation strategy for real-time bilateral negotiations. *Studies in Computational Intelligence*, Vol. 435. Springer, Berlin Heidelberg, pp. 229–233.
- Ball, D., Coelho, P.S., Vilares, M.J., 2006. Service personalization and loyalty. *Journal of Services Marketing* 20 (6), 391–403.
- Bauer, B., Müller, J.P., Odell, J., 2001. *Agent UML: A formalism for specifying multiagent interaction*. Agent-Oriented Software Engineering, Vol. 1957. Springer, Berlin, pp. 91–103.
- Behfar, K., Friedman, R.A., and Brett, J.M. The team negotiation challenge: defining and managing the internal challenges of negotiating teams. In *Proceedings of the 21st Annual Conference for the International Association for Conflict Management (IACM-2008)*, 2008.
- Brazier, F.M., Cornelissen, F., Gustavsson, R., Jonker, C.M., Lindeberg, O., Polak, B., Treur, J., 2002. A multi-agent system performing one-to-many negotiation for load balancing of electricity use. *Electronic Commerce Research and Applications* 1 (2), 208–224.
- Brodth, S., Thompson, L., 2001. Negotiating teams: a levels of analysis. *Group Dynamics* 5 (3), 208–219.
- Buffett, S., Spencer, B., 2007. A bayesian classifier for learning opponents' preferences in multi-object automated negotiation. *Electronic Commerce Research and Applications* 6 (3), 274–284.
- Carbonneau, R., Kersten, G., Vahidov, R., 2008. Predicting opponent's moves in electronic negotiations using neural networks. *Expert Systems with Applications* 34 (2), 1266–1273.
- Carbonneau, R.A., Kersten, G.E., Vahidov, R.M., 2011. Pairwise issue modeling for negotiation counteroffer prediction using neural networks. *Decision Support Systems* 50 (2), 449–459.
- Chan, C.-C.H., Cheng, C.-B., Hsu, C.-H., 2008. Bargaining strategy formulation with CRM for an e-commerce agent. *Electronic Commerce Research and Applications* 6 (4), 490–498.
- Chevalerey, Y., Endriss, U., Lang, J., Maudet, N., 2007. A short introduction to computational social choice. In: *SOFSEM 2007: Theory and Practice of Computer Science. Lecture Notes in Computer Science*, Vol. 4362. Springer, Berlin/ Heidelberg, pp. 51–69.
- Coehoorn, R.M., and Jennings, N.R. Learning on opponent's preferences to make effective multi-issue negotiation trade-offs, in: *The 6th International Conference on Electronic Commerce (ICEC'04)*, 2004, 59–68.
- Ehtamo, H., Kettunen, E., Hamalainen, R.P., 2001. Searching for joint gains in multi-party negotiations. *European Journal of Operational Research* 130 (1), 54–69.
- Faratin, P., Sierra, C., Jennings, N.R., 1998. Negotiation decision functions for autonomous agents. *International Journal of Robotics and Autonomous Systems* 24 (3–4), 159–182.
- Farhangi, H., 2010. The path of the smart grid, power and energy magazine. *IEEE* 8 (1), 18–28.
- Gaston, M.E., desjardins, M., 2005. Agent-organized networks for dynamic team formation. *Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems (AAMAS'05)*. ACM, New York, NY, USA, pp. 230–237, ISBN 1-59593-093-0.
- Grieger, M., 2003. Electronic marketplaces: a literature review and a call for supply chain management research. *European Journal of Operational Research* 144 (2), 280–294.
- Guttman, R.H., Moukas, A.G., Maes, P., 1998. Agent-mediated electronic commerce: a survey. *Knowledge Engineering Review* 13 (2), 147–159, ISSN 0269-8889, doi:<http://dx.doi.org/10.1017/S0269888998002082>.
- Halevy, N., 2008. Team negotiation: social, epistemic, economic, and psychological consequences of subgroup conflict. *Personality and Social Psychology Bulletin* 34, 1687–1702.
- He, M., Jennings, N.R., Leung, H.-F., 2003. On agent-mediated electronic commerce. *IEEE Transactions on Knowledge and Data Engineering* 15 (4), 985–1003.
- Hindriks, K.V., Tykhonov, D., 2010. Towards a quality assessment method for learning preference profiles in negotiation. In: *Agent-Mediated Electronic Commerce and Trading Agent Design and Analysis*. Springer, pp. 46–59.
- Kawaguchi, S., Fujita, K., Ito, T., 2011. Compromising strategy based on estimated maximum utility for automated negotiation agents competition (ANAC-10). *Modern Approaches in Applied Intelligence*, Vol. 6704. Springer, Berlin/ Heidelberg, pp. 501–510.
- Klein, M., Faratin, P., Sayama, H., Bar-Yam, Y., 2003a. Protocols for negotiating complex contracts. *IEEE Intelligent Systems* 18 (6), 32–38, ISSN 1541-1672.
- Klein, M., Faratin, P., Sayama, H., Bar-Yam, Y., 2003b. Negotiating complex contracts. *Group Decision and Negotiation* 12, 111–125.
- Korsgaard, M., Schweiger, D., Sapienza, H., 1995. Building commitment, attachment, and trust in strategic decision-making teams: the role of procedural justice. *Academy of Management Journal*, 60–84.
- Lai, G., Sycara, K., Li, C., 2008. A decentralized model for automated multi-attribute negotiations with incomplete information and general utility functions. *Multiagent and Grid Systems* 4 (1), 45–65, ISSN 1574-1702.
- Lamparter, S., Becher, S., and Fischer, J.-G. An agent-based market platform for smart grids. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: Industry track*, International Foundation for Autonomous Agents and Multiagent Systems, 2010, 1689–1696.
- Lau, Raymond Y.K., 2007. Towards a web services and intelligent agents-based negotiation system for {B2B} eCommerce. *Electronic Commerce Research and Applications* 6 (3), 260–273, ISSN 1567-4223.
- Liang, T.-P., Lai, H.-J., Ku, Y.-C., 2007. Personalized content recommendation and user satisfaction: theoretical synthesis and empirical findings. *Journal of Management Information Systems* 23 (3), 45–70.
- Lin, R., Kraus, S., Baarslag, T., Tykhonov, D., Hindriks, K., Jonker, C.M., 2014. Genius: an integrated environment for supporting the design of generic automated negotiators. *Computational Intelligence* 30, 48–70.
- Lomuscio, A.R., Wooldridge, M., Jennings, N.R., 2003. A classification scheme for negotiation in electronic commerce. *Group Decision and Negotiation* 12, 31–56, ISSN 0926-2644.
- Mannix, E., 2005. Strength in numbers: negotiating as a team. *Negotiation* 8 (5), 3–5.

- Mansour, K., Kowalczyk, R., 2012. *A Meta-Strategy for Coordinating of One-to-Many Negotiation over Multiple Issues*. Foundations of Intelligent Systems, Vol. 122. Springer, Berlin/Heidelberg, pp. 343–353, ISBN 978-3-642-25663-9.
- Marsa-Maestre, I., Klein, M., Jonker, C.M., Aydogan, R. From problems to protocols: towards a negotiation handbook, decision support systems, 2013, doi:http://dx.doi.org/10.1016/j.dss.2013.05.019.
- Morais, H., Pinto, T., Vale, Z., Praça, I., 2012. Multilevel negotiation in smart grids for VPP management of distributed resources. *IEEE Intelligent Systems* 27 (6), 8–16.
- Narayanan, V., Jennings, N.R., 2006. Learning to Negotiate Optimally in Non-stationary Environments. In: *Cooperative Information Agents X*, 10th International Workshop, CIA 2006, Edinburgh, UK, September 11–13, 2006, Proceedings. Lecture Notes in Computer Science, Vol. 4149. Springer, pp. 288–300.
- Ngai, E.W., Wat, F., 2002. A literature review and classification of electronic commerce research. *Information & Management* 39 (5), 415–429.
- Nguyen, T.D., Jennings, N.R., 2005. Managing commitments in multiple concurrent negotiations. *Electronic Commerce Research and Applications* 4 (4), 362–376, ISSN 1567-4223, developments in intelligent support for e-Commerce negotiation applications The 6th International Conference on Electronic Commerce.
- Nurmi, H., 2010. Voting systems for social choice. *Handbook of Group Decision and Negotiation*, 167–182.
- Oliveira, E., Rocha, A., 2001. Agents advanced features for negotiation in electronic commerce and virtual organisations formation process. In: *Agent Mediated Electronic Commerce*. Springer, pp. 78–97.
- Ortmann, G., King, R., 2007. Agricultural cooperatives I: history, theory and problems. *Agrekon* 46 (1), 18–46.
- Rahwan, T., Ramchurn, S.D., Jennings, N.R., Giovannucci, A., 2009. An anytime algorithm for optimal coalition structure generation. *Journal of Artificial Intelligence Research* 34 (1), 521–567, ISSN 1076-9757.
- Ramchurn, S.D., Vytelingum, P., Rogers, A., Jennings, N.R., 2012. Putting the 'smarts' into the smart grid: a grand challenge for artificial intelligence. *Communications of the ACM* 55 (4), 86–97.
- Robu, V., La Poutré, J.A., 2008. Retrieving the Structure of utility graphs used in multi-item negotiation through collaborative filtering of aggregate buyer preferences. In: *Rational, Robust and Secure Negotiations*. *Computational Intelligence*, Vol. 89. Springer.
- Rubinstein, A., 1982. Perfect equilibrium in a bargaining model. *Econometrica* 50, 155–162.
- Russell, S., Norvig, P., 2003. *Artificial Intelligence: A Modern Approach*. Pearson Education.
- Sanchez-Anguix, V., Julian, V., Botti, V., and Garcia-Fornes, A. Analyzing intra-team strategies for agent-based negotiation teams. In 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'11), 2011, 929–936.
- Sanchez-Anguix, V., Julian, V., Botti, V., Garcia-Fornes, A., 2012a. Reaching unanimous agreements within agent-based negotiation teams with linear and monotonic utility functions. *IEEE Transactions on Systems, Man and Cybernetics, Part B* 42 (3), 778–792.
- Sanchez-Anguix, V., Dai, T., Semnani-Azad, Z., Sycara, K., and Botti, V. Modeling power distance and individualism/collectivism in negotiation team dynamics. In 45 Hawaii International Conference on System Sciences (HICSS-45), 2012, 628–637.
- Sanchez-Anguix, V., Valero, S., Julian, V., Botti, V., Garcia-Fornes, A., 2013. Evolutionary-aided negotiation model for bilateral bargaining in Ambient Intelligence domains with complex utility functions. *Information Sciences* 222, 25–46.
- Sierra, C., Dignum, F., 2001. Agent-mediated electronic commerce: scientific and technological roadmap. In: *Agent Mediated Electronic Commerce*. Springer, pp. 1–18.
- Stole, L.A., 1995. Nonlinear pricing and oligopoly. *Journal of Economics & Management Strategy* 4 (4), 529–562.
- Taylor, H., 2003. Most people are privacy pragmatists who, while concerned about privacy, will sometimes trade it off for other benefits. *The Harris Poll* 17, 19.
- Thompson, L., 2003. *The Mind and Heart of the Negotiator*. Prentice Hall Press, Upper Saddle River, NJ, USA.
- Thompson, L., Peterson, E., Brodt, S., 1996. Team negotiation: an examination of integrative and distributive bargaining. *Journal of Personality and Social Psychology* 70 (1), 66–78.
- Wareham, J., Zheng, J.G., Straub, D., 2005. Critical themes in electronic commerce research: a meta-analysis. *Journal of Information Technology* 20 (1), 1–19.
- Williams, C., Robu, V., Gerding, E., and Jennings, N. Using gaussian processes to optimise concession in complex negotiations against unknown opponents. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI'11)*, 2011, 432–438.
- Williams, C.R., Robu, V., Gerding, E.H., and Jennings, N.R. Negotiating concurrently with unknown opponents in complex, real-time domains. In *20th European Conference on Artificial Intelligence*, Vol. 242, 2012, 834–839.
- Zeng, D., Sycara, K., 1998. Bayesian learning in negotiation. *International Journal of Human-Computer Studies* 48 (1), 125–141, ISSN 1071-5819, doi:http://dx.doi.org/10.1006/ijhc.1997.0164.