

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/330246801>

Bottom-up approaches to achieve Pareto optimal agreements in group decision making

Article in Knowledge and Information Systems · January 2019

DOI: 10.1007/s10115-018-01325-y

CITATIONS

0

READS

26

4 authors:



Victor Sanchez-Anguix

Florida Universitaria

40 PUBLICATIONS 217 CITATIONS

SEE PROFILE



Reyhan Aydogan

Ozyegin University

47 PUBLICATIONS 213 CITATIONS

SEE PROFILE



Tim Baarslag

Centrum Wiskunde & Informatica

69 PUBLICATIONS 658 CITATIONS

SEE PROFILE



Catholijn M. Jonker

Delft University of Technology

541 PUBLICATIONS 6,263 CITATIONS

SEE PROFILE

Some of the authors of this publication are also working on these related projects:



The Eighth International Automated Negotiating Agent Competition (ANAC) [View project](#)



Accounting for Values in Design [View project](#)

Bottom-up approaches to achieve Pareto optimal agreements in group decision making

Victor Sanchez-Anguix^{1,2,3}, Reyhan Aydođan^{4,5}, Tim Baarslag^{6,7} and Catholijn Jonker^{5,8}

¹Florida Universitaria, Carrer del Rei en Jaume I, 2, 46470, Catarroja, Spain;

²Universidad Isabel I, Calle de Fernan Gonzalez, 76, Burgos, Spain;

³Coventry University, Gulson Rd, CV1 2JH, Coventry, United Kingdom

⁴Department of Computer Science, Ozyegin University, Istanbul, Turkey;

⁵ Department of Intelligent Systems, Delft University of Technology, 2628XE Delft, The Netherlands;

⁶ Centrum Wiskunde & Informatica, Science Park 123, 1098 XG Amsterdam, The Netherlands;

⁷ Department of Information and Computing Sciences, Utrecht University, Princetonplein 5, De Uithof, 3584 CC Utrecht, The Netherlands

⁸ Leiden University, Snellius building, 2333CA Leiden, The Netherlands

Abstract.

In this article we introduce a new paradigm to achieve Pareto optimality in group decision making processes: *bottom-up approaches to Pareto optimality*. It is based on the idea that, while resolving a conflict in a group, individuals may trust some members more than others; thus, they may be willing to cooperate and share more information with those members. Therefore, one can divide the group into subgroups where more cooperative mechanisms can be formed to reach Pareto optimal outcomes. This is the first work that studies such use of a bottom-up approach to achieve Pareto optimality in conflict resolution in groups. First, we prove that an outcome that is Pareto optimal for subgroups is also Pareto optimal for the group as a whole. Then, we empirically analyze the appropriate conditions and achievable performance when applying bottom-up approaches under a wide variety of scenarios based on real-life datasets. The results show that bottom-up approaches are a viable mechanism to achieve Pareto optimality with applications to group decision-making, negotiation teams, and decision making in open environments.

Keywords: Agreement technologies; automated negotiation; Pareto optimality; group decision making; multi-agent systems

Received xxx

Revised xxx

Accepted xxx

1. Introduction

Group decision making, in which a group of agents with conflicting preferences aim to reach mutually acceptable decisions, is probably one of the most challenging areas for the decision sciences and related fields. The complexity arises from the preferential conflict among group members as well as the interactions among them during the underlying decision process. This is perhaps one of the reasons why group decision making has received so much attention in the scholarly world. For instance, the negotiation literature has been prolific in studying how different multi-party mechanisms can be used to find unanimous agreements that accommodate the interests of group members from a set of configurable negotiation attributes; multi-objective optimization methods, distributed or not, have also been developed to find optimal solutions for group decision making settings. Regardless of the setting, one of the desired properties of such a solution/agreement is Pareto optimality, proposed by the Italian economist Vilfredo Pareto. Its desirability comes from the fact that, concerning non-Pareto optimal solutions, at least one of the objectives can be improved without worsening the performance of the rest of objectives. Hence, rational decision makers should see no objection in moving from a non-Pareto optimal solution to a Pareto optimal solution.

Despite the benefits arising from reaching a decision that is Pareto optimal, practitioners know that reaching a Pareto optimal agreement is not a straightforward practice. The reality is that, in open and dynamic environments, agents rarely know each others' preferences and, to make it worse, the risk of exploitation by manipulation precludes agents from sharing their complete preference profiles with all of the other group members. As a consequence, agents must seek alternative mechanisms to obtain Pareto optimal agreements. To overcome this lack of information, existing mechanisms invite a significant number of interactions among agents and still cannot guarantee Pareto optimality.

A number of works in the field focus on finding a global Pareto optimal solution by involving all agents at the same time (Heiskanen, Ehtamo and Hämäläinen, 2001; Lai, Li and Sycara, 2006; Hara and Ito, 2013; Zhenh, Chakraborty, Nilanjan, Dai and Sycara, 2015), which may lead to complex interactions and lengthy decision making processes. However, we believe that, in many situations, agents can benefit from taking a bottom-up approach: calculating Pareto optimal outcomes in subgroups. In other words, we pursue the question of whether or not it is possible to estimate some Pareto optimal outcomes without explicitly interacting with all of the agents in a group. In essence, solving the Pareto optimal set problem in a smaller group may be less complex than in larger groups. For instance, there may be less exploitation risks in subgroups of trusted agents, or agents may trust other subgroup members more and, therefore, be willing to share more preferential information or cooperate to a greater degree.

To the best of the authors' knowledge, this work constitutes the first attempt at introducing bottom-up approaches to achieve Pareto optimal agreements. This article expands our initial study carried out in (Sanchez-Anguix, Aydoğan, Baarslag and Jonker, 2017), where we introduced the proof that shows that an outcome that is Pareto optimal in a subgroup is also Pareto optimal in a larger group containing the subgroup as long as agent preferences are strict, and in which we reported about some preliminary and small scale experiments to test applicability of bottom-up approaches to Pareto optimality. In this present article, we reintroduce the proof presented in (Sanchez-Anguix et al., 2017) and

provide extra lemmas and corollaries that help to further understand how Pareto optimality in groups behaves in theory. In addition to this, we report about new experiments to understand how bottom-up approaches to Pareto optimality are affected by the degree of conflict of the domain, how the group size affects the applicability of bottom-up approaches to Pareto optimality, and we report about large scale experiments to test not only the applicability of bottom-up approaches to Pareto optimality, but also the performance of these approaches in practice for a broad variety of synthetic and real domains. Furthermore, we analyze in what scenarios bottom-up approaches to Pareto optimality are effective. Where possible, we also provide recommendations for the application of bottom-up approaches to Pareto optimality.

In order to make a point for bottom-up approaches to Pareto optimality, in Section 2, we first prove that any Pareto optimal outcome in a subgroup is also Pareto optimal in a larger group that contains the subgroup, as long as agents' preferences follow a strict order. Then, in Section 3 we report the experimental setting that was designed to study both the applicability and the prospective empirical performance of bottom-up approaches in a wide variety of scenarios. Finally, we discuss relevant and related work in Section 4 and provide conclusions and future lines of work in Section 5.

2. Bottom-up approaches to Pareto optimality

In this section we provide the theoretical foundation that makes it possible to apply a bottom-up approach to achieve Pareto optimal outcomes in a group of agents seeking an agreement. More specifically, we show that, under agents' strict preferences on outcomes, an outcome that is Pareto optimal in a subgroup is also Pareto optimal in a larger group containing the subgroup. Hence, an outcome that is identified as Pareto optimal in a subgroup by using any mechanism, can be later used as a Pareto optimal outcome in a posterior decision making phase involving all of the group. For the sake of simplicity, we will call this proof the bottom-up Pareto optimality proof from this point onward. In addition to this proof, we identify an extra lemma and corollary that throws some light on how Pareto optimality behaves in a group, thus affecting the application of bottom-up approaches to Pareto optimality.

2.1. Theoretical Foundations

We start by providing some basic definitions and notation that are needed to introduce the bottom-up Pareto optimality proof. Let $\mathcal{A} = \{a_1, \dots, a_n\}$ be a set of agents where k is the index of agent a_k and let $\mathcal{A}' = \{a_1, \dots, a_m\}$ be a superset of \mathcal{A} , i.e. $\mathcal{A} \subset \mathcal{A}'$ and $m > n$. \mathcal{O} is the set of all possible solutions in a given domain. By \succeq_i we represent agent's a_i *preference profile* over outcomes $o \in \mathcal{O}$. If $o \succeq_i o'$ then agent a_i likes o at least as well as o' , we write $o \succ_i o'$ to denote a strict preference for o and $o = o'$ to denote indifference. We assume that the agents' preference profiles are strict, transitive and complete.

An outcome o^* is Pareto optimal with respect to \mathcal{A} and \mathcal{O} , denoted by

$po(o^*, \mathcal{A}, \mathcal{O})$ if and only if

$$\nexists o \in \mathcal{O} \exists j \leq n \bigwedge_{i=1}^n o \succeq_i o^* \wedge o \succ_j o^*.$$

We denote the set of all Pareto optimal outcomes over \mathcal{A} by $\mathcal{O}_{\mathcal{A}}^* = \{o^* \in \mathcal{O} \mid po(o^*, \mathcal{A}, \mathcal{O})\}$.

Theorem 1. Given a set of outcomes \mathcal{O} . For all two sets of agents \mathcal{A} and \mathcal{A}' , if $\mathcal{A} \subset \mathcal{A}'$, then $\mathcal{O}_{\mathcal{A}}^* \subset \mathcal{O}_{\mathcal{A}'}^*$.

Proof. Let us assume by reductio ad absurdum that $\mathcal{A} \subset \mathcal{A}'$, but $\mathcal{O}_{\mathcal{A}}^* \not\subset \mathcal{O}_{\mathcal{A}'}^*$. This means there exists an $o^* \in \mathcal{O}_{\mathcal{A}}^*$ such that $o^* \notin \mathcal{O}_{\mathcal{A}'}^*$. Expanding the definition of Pareto optimal outcomes, we have

$$o^* \notin \{o \in \mathcal{O} \mid \nexists o' \in \mathcal{O} \exists k \leq m, \bigwedge_{i=1}^m o' \succeq_i o \wedge o' \succ_k o\}.$$

Recalling that n and m are the index of the last agent in \mathcal{A} and \mathcal{A}' respectively, this means that there must exist an $o \in \mathcal{O}$ and a $k \leq m$ such that $\bigwedge_{i=1}^m o \succeq_i o^* \wedge o \succ_k o^*$. We consider two scenarios: either $a_k \in \mathcal{A}$ or $a_k \notin \mathcal{A}$.

–If $a_k \in \mathcal{A}$ then o is an outcome that dominates o^* over \mathcal{A} , which is not possible as o^* is Pareto optimal over \mathcal{A} .

–Otherwise, $k > n$, so we have $\bigwedge_{i=1}^n o \succeq_i o^*$. In that case, as o^* is Pareto optimal over \mathcal{A} , the condition is only true if all of the agents in \mathcal{A} are indifferent between o and o^* . As preferences are strict, that cannot be true either.

Since both sides lead to a contradiction, we have proven the theorem. \square

Note that while the bottom-up Pareto optimality proof assumes that agents' preferences are strict, it may be the case that in real world domains an agent may be indifferent between multiple outcomes. Yet despite this, an outcome that is Pareto optimal in a subgroup with non-necessarily strict preferences is usually also Pareto optimal in a larger group containing the subgroup. For an outcome to be Pareto optimal in a subgroup but not Pareto optimal in the larger group, which we will refer from now as a false positive, all of the agents in the subgroup should be indifferent between such outcome and another Pareto optimal outcome. Then, in the larger group, at least one of the agents in the group is not indifferent between those outcomes. We believe that this situation is rare in practice. Its improbability stems from the fact that finding a situation where all of the agents in a subgroup are indifferent between two Pareto optimal outcomes becomes more and more unlikely as the size of the subgroup increases since the more agents the more probable it is that their opinions will be different.

Hypothesis 1. Overall, the likeliness of finding a false positive Pareto optimal outcome in a subgroup is small in practice due to the unlikeliness of finding a situation where agents are indifferent between two Pareto optimal outcomes.

We will confirm that this situation is generally rare, thus confirming **H1** in practice, in our experimental approach in Section 3.

So far, we have proved that an outcome that is Pareto optimal in a subgroup will also be Pareto optimal in any group containing that subgroup. As a

result, one can calculate Pareto optimal outcomes for a group by calculating outcomes that are Pareto optimal in the subgroup, restricting interactions to that subgroup, hence making it feasible to adopt a bottom-up approach to Pareto optimality. In addition to this, it is also interesting to study a few additional properties that stem from the definition of Pareto optimality and the previous theorem. These properties are helpful to understand how Pareto optimality in a group behaves, and the implications for bottom-up approaches. First, we bring about a well-known result from the literature: when two agents have completely opposite preferences then all of the outcomes are Pareto optimal.

Lemma 1. Given an outcome space \mathcal{O} and two agents a_i and a_j . If for all $o, o' \in \mathcal{O}$ it holds that $o \succ_i o' \iff o' \succ_j o$, then $\mathcal{O}_{\{a_i, a_j\}}^* = \mathcal{O}$.

Proof. The proof holds for any outcome space $\mathcal{O} = \{o_1, \dots, o_n\}$ with $o_k \succ_i o_l$ and $o_l \succ_j o_k$ for any $k < l$. Let us consider any outcome o_k , a_i can only improve its own utility by choosing an outcome o_r where $r < k$, as $o_k \succ_i o_l$ for any $k < l$. However, in any case, a_j will end up in a less beneficial position as $o_l \succ_j o_k$ for any $k < l$. Similarly, a_j can only improve its own utility by choosing an outcome o_r where $k < r$, which ends up in a less beneficial position for a_i . Therefore, any outcome is Pareto optimal in that situation. \square

Given Lemma 1 and Theorem 1, one can reach the following in a straightforward way:

Corollary 1. If there exist two agents $a_i, a_j \in \mathcal{A}$ such that

$$\forall o, o' \in \mathcal{O} \ o \succ_i o' \iff o' \succ_j o,$$

then $\mathcal{O}_{\mathcal{A}}^* = \mathcal{O}_{\{a_i, a_j\}}^* = \mathcal{O}$.

Proof. This follows directly: as any two agents with opposing preferences define a Pareto optimal frontier consisting of all possible outcomes, and the Pareto optimal space defined in the group is a superset of the Pareto optimal outcomes defined by that pair of agents, then all of the outcomes will be Pareto optimal in the group. \square

There are a few direct consequences to this corollary. First of all, there is no point in filtering Pareto optimal outcomes in scenarios where agents' have completely opposite preferences, as any outcome chosen by the group will be Pareto optimal by definition. A hypothesis that one can derive from this theoretical result is that in scenarios with a high degree of conflict, most of the outcomes are Pareto optimal for particular subgroups. The prospective lack of use for bottom-up approaches to Pareto optimality in high-conflict scenarios is not a shortcoming of the approach, but a direct consequence of the definition of Pareto optimality. There may still be use for bottom-up approaches for filtering outcomes according to stricter definitions of optimality like k-optimality (Corne and Knowles, 2007). As part of the study of the applicability of bottom-up approaches, we will test and analyze the relationship between the degree of conflict in a group and the number of outcomes that are Pareto optimal in the experiments described in Section 3.

2.2. A general scheme for bottom-up approaches to Pareto optimality

We showed that any outcome that is Pareto optimal in a subgroup with strict preferences, is also Pareto optimal in any extended group. Next, we discuss how this result may be applied in practical scenarios. Note that the scheme proposed in this subsection is generally applicable and aims to provide general guidelines for how the paradigm may be employed in practice. Figure 1 illustrates the general scheme for a group decision making setting.

There are three main steps to apply bottom-up approaches in group decision making:

1. *Subgroup formation*: The first step involved in any bottom-up approach is dividing the group of agents into subgroups. These subgroups will contribute to the group's Pareto optimal set by discovering the outcomes that are Pareto optimal in the subgroup. There are multiple ways of dividing agents into subgroups. Subgroups can either be overlapping or non-overlapping, and they can include all agents or just include a handful of agents. Even if the types of subgroups that are formed are limited, the strategies followed to devise those types of subgroups are almost unlimited in nature, and they may follow different criteria (e.g., trust among agents, similarity, agents' roles, etc.).
2. *Subgroup cooperation*: After subgroups have been identified and formed, it is time for members of subgroups to cooperate with the goal of finding out outcomes that are Pareto optimal. One of the assumptions underlying bottom-up approaches to Pareto optimality is the existence of a certain degree of trust among the agents that are part of a subgroup. The type of cooperative approach taken by the subgroup may differ according to the degree of trust in the subgroup. The simplest mechanism entails sharing the full preference profile among subgroup members and then aggregating these preference profiles (Ramanathan and Ganesh, 1994; Conitzer, Lang and Xia, 2011) to calculate Pareto optimal deals in the subgroup, but it also assumes complete trust among subgroup members. Depending on the degree of trust, there may be other mechanisms that can be applied such as mechanisms that only assume sharing partial information (Hu, Wang and Di Paolo, 2013) or negotiation mechanisms that assume sharing more limited information (Jonker, Robu and Treur, 2007). The key with this task is that it is possible to apply a more cooperative mechanism in the subgroup, than it is in the whole group.
3. *Group decision making*: Once the subgroups have applied a cooperation mechanism to calculate Pareto optimal outcomes in the subgroup, all the Pareto optimal outcomes identified in the subgroup are aggregated and presented to the whole group. At that point, it is the time to decide on an appropriate outcome for the whole group, which depends on the specific domain and the degree of cooperation vs. competition among agents. In any case, note that, in this process, the group has reduced from a potentially large space of outcomes to a more reduced space of outcomes that is known to be Pareto optimal. Hence, the subgroup cooperation mechanism acts as a filtering process that keeps relevant choices in the decision making process.

The guidelines provided in the scheme can be applied to a wide variety of scenarios. For instance, the application to group decision making processes (Zhenh et al., 2015; Aydoğan, Hindriks and Jonker, 2014; Esparcia, Sanchez-




Fig. 1. A general workflow of steps to apply bottom-up approaches to Pareto optimality into a variety of scenarios

Anguix and Aydođan, 2013; Hara and Ito, 2013) is direct, with a pre-negotiation phase where subgroups are identified, and a negotiation phase where Pareto optimal outcomes in subgroups are identified, aggregated, and then negotiated as an entire group. Another prospective application of the scheme is negotiation teams (Sanchez-Anguix, Aydođan, Julian and Jonker, 2014; Sanchez-Anguix, Julian, Botti and Garcia-Fornes, 2013*b*; Sanchez-Anguix, Julian, Botti and Garcia-Fornes, 2013*a*; Sanchez-Anguix, Julian, Botti and Garcia-Fornes, 2012; Sanchez-Anguix, Dai, Semnani-Azad, Sycara and Botti, 2012), where a multi-individual party negotiates with one or several opponents. In this setting, bottom-up approaches to Pareto optimality may be applied before the negotiation to discover outcomes that are Pareto optimal among team members. These outcomes can later be used as a basis to negotiate with the opponent.

3. Experiments

In Section 2, we showed that it is possible to obtain a part of the Pareto optimal frontier of a group of agents by calculating the Pareto optimal frontier in a subgroup. As a consequence, it is possible to employ a bottom-up approach to obtain part of the Pareto optimal frontier in three steps: (1) divide the group into one or several subgroups, (2) calculate the Pareto optimal outcomes in each subgroup, and (3) use the obtained outcomes as part of the final Pareto optimal frontier of the entire group.

However, this does not provide any reasonable indication for how the approach would perform in practice, as there are many questions that should be answered before claiming that a bottom-up approach to Pareto optimality is feasible. For instance, what ratio of the final Pareto optimal frontier is obtained using this approach? What is the quality of the Pareto optimal frontier achievable in subgroups? In order to study this in more detail, we employ an experimental approach that aims to assess the performance of bottom-up approaches in practical scenarios.

This section is structured as follows. First, we introduce the wide range of decision making domains used in this experimental setting. Then, we move to analyze the applicability of bottom-up approaches to Pareto optimality by studying the relationship between the degree of conflict and the ratio of Pareto optimal outcomes in the outcome space. After that, and linked to studying the applicability of bottom-up approaches, we study the relationship between group size and the ratio of Pareto optimal outcomes in the outcome space. As a final step to study the applicability of bottom-up approaches to realistic scenarios, and due

to the fact that preferences are not guaranteed to be strict in real scenarios, we analyze the ratio of outcomes detected as Pareto optimal in subgroups that are not Pareto optimal in the whole group. Finally, we analyze the achievable performance of bottom-up approaches under two possible families of subgrouping strategies with regards to the final ratio of the Pareto optimal frontier calculated in the subgroups, and the relative joint utility performance.

3.1. Domains

In order to assess the experimental performance of bottom-up approaches, one needs to select a set of scenarios to test on. We classify the scenarios into synthetic and real decision-making domains. As a prior step to carry out our analyses, we decided to further characterize domains by their degree of conflict. For that, we employed the following metrics inspired by popular metrics used to compare rankings of items/outcomes:

- Average Spearman rank correlation over pairs of preference profiles in the domain.
- Average Kendall Tau rank correlation over pairs of preferences profiles in the domain.
- The average precision@10%, defined as the ratio of overlapping between the top 10% outcomes between pairs of preference profiles.

3.1.1. Synthetic domains

There are aspects that can affect the performance of bottom-up approaches to Pareto optimality. Based on the hypothesis raised in Section 2, one can reasonably think that the degree of conflict among agents shapes the Pareto optimal frontier and the number of Pareto optimal outcomes. In fact, as discussed, a group where at least two agents' preferences who are strict and are in complete opposition will result in a group where all of the outcomes are Pareto optimal. Therefore, it is suggested that the degree of preferential conflict may be an interesting factor to study, and a wide range of degrees of conflict should be covered. For that reason, we generated synthetic domains with varied degrees of conflict.

For generating scenarios with different degrees of conflict, we rely on the idea of clusters of preferences. A cluster of preferences is a group of preference profiles that are close to each other in a multidimensional space defined by the utility given by an agent to each outcome¹. Intuitively, we believe that the higher the number of clusters of preferences exist in a domain, the more different the preferences of the agents in a group should be overall. With that assumption in mind, we generated synthetic domains following the descriptive procedure below:

- A synthetic domain consists of $m = \{10, 100, 1000\}$ different outcomes, $k = \{1, 2, 3, 4, 5, \dots, 17, 18, 19, 20\}$ clusters of preferences and 100 preference profiles.

¹ This is similar to the classic machine learning notion of cluster, where the space has as many dimensions as outcomes in the domain, and a point represents the evaluation of an agent for each of the outcomes in the domain




Fig. 2. The figure shows the average Spearman rank correlation, Kendall Tau rank correlation, and Precision@10% for synthetic domains generated with a number of preferential clusters ranging from 2 to 20

- A cluster is defined by a multivariate m -dimensional isotropic Gaussian distribution that defines the utility vector provided by an agent to each outcome in the domain $\mathcal{U} = \langle u_1, \dots, u_m \rangle$, where u_i is the utility provided by the i -th outcome in the domain. Therefore, a cluster of preferences can be defined as those agents whose utility vectors have been sampled from $\mathcal{U} \sim \mathcal{N}(\mu = \langle \mu_1, \dots, \mu_m \rangle, \langle \sigma_1^2 \dots \sigma_m^2 \rangle \times I)$, with μ_i is uniformly sampled between between 0 and 10, and the standard deviation of each component σ_i is uniformly sampled between 0.1 and 2.0. A preference profile associated to a cluster consists of randomly sampling a point from the multivariate Gaussian distributions.
- For each possible combination of m and k , 20 possible domains are generated arbitrarily.

As a result, a total of $3 \times 20 \times 20 = 1,200$ synthetic domains were generated with different sizes and different number of preferential clusters. For all of these domains, we calculated the average Spearman rank correlation, the average Kendall Tau correlation, and the average precision@10% . In order to test the validity of the domain generation process, we aggregate the aforementioned metrics on the number of preferential clusters as depicted in Figure 2. It can be observed that when we increase the number of clusters it lowers the average Spearman rank correlation, the average Kendall Tau rank correlation, and the Precision@10%; this validates the generation process that aimed to generate domains with different degrees of conflict.

3.1.2. Real domains

The real domains employed in the experimental evaluation represent a variety of problems like deciding on a movie to be watched by a group, deciding on the best path to be taken by a robot explorer, or deciding on the specific details with regards to a party. The description of the 21 real domains selected for the experimental study can be found in Table 1.

It should be highlighted that the **real domains contain preference pro-**

Name	Domain size	# of Profiles
Sushi (Kamishima, 2003)	10	5,000
AGH (Skowron, Faliszewski and Slinko, 2013)	6	153
Book (Ziegler, McNee, Konstan and Lausen, 2005)	23	7
Movielens (Miller, Albert, Lam, Konstan and Riedl, 2003)	298	10
Holiday (Lin, Kraus, Baarslag, Tykhonov, Hindriks and Jonker, 2014)	1,024	9
Symposium (Lin et al., 2014)	2,305	9
Party (Lin et al., 2014)	3,072	24
Jester (Goldberg, Roeder, Gupta and Perkins, 2001)	100	7,200
Debian Leader (Mattei and Walsh, 2013)	9	430
Debian Logo (Mattei and Walsh, 2013)	8	125
ERS 1 (Mattei and Walsh, 2013)	10	252
ERS 2 (Mattei and Walsh, 2013)	19	80
Mariner (Mattei and Walsh, 2013)	32	10
Minneapolis Parks Election (Mattei and Walsh, 2013)	477	556
Minneapolis Tax Election (Mattei and Walsh, 2013)	379	723
Glasgow Election (Mattei and Walsh, 2013; O'Neill, 2013)	13	2,704
Skate Championship (Mattei and Walsh, 2013)	30	10
T-shirt (Mattei and Walsh, 2013)	11	30
Tram (Lin et al., 2014)	972	9
University (Lin et al., 2014)	2,250	9
Zone Planning (Lin et al., 2014)	448	9

Table 1. The real domains employed in the experimental section

files that are indifferent among different outcomes. This will allow us to test the applicability of the bottom-up proof in scenarios where strict preferences are not present.

For these domains, we also measure the average Spearman rank correlation, the average Kendall Tau rank correlation, and the Precision@10% in the domain. Then, we group the synthetic domains by the number of preferential clusters in them, and produce a centroid for each group to represent the group's average Kendall Tau rank correlation, average Spearman rank correlation, and average Precision@10%. Then, we make an educated guess on the prospective number of preferential clusters for real domains by assigning the same number of clusters than the closest² centroid calculated in the previous step. It should be highlighted that with this categorization we do not imply that the real number of clusters in the real domains is the one associated by this clustering process. The sole goal of this categorization is providing a rough idea on the degree of conflict in real domains, and being able to sort real domains by the degree of conflict that they present. The only claim that can be done on these real domains is that their conflict characteristics are similar to those outlined by the predicted category. This categorization will help us to analyze experimental results when analyzing results arising from real domains later in the experimental section. The results of this clustering process can be observed in the following list, which separates domains according to their expected levels of conflict (i.e., number of preferential clusters):

- Low conflict real domains: It includes those real domains whose number of assigned preferential clusters falls below the first quartile. More specifically, we have the Skate (1 cluster), Movielens (3), Minneapolis Park (3), Minneapolis Tax (3), and the Tee (3) domain.
- Mild conflict real domains: It includes those real domains whose number of assigned preferential clusters falls between the first and third quartile. More specifically, we have the Debian Leader (4 clusters), AGH (4), Mariner (4), Sushi (5), Debian Logo (6), ERS 1 (7), Holiday (8), University (9), Glasgow election (8), and the Jester (9) domain.
- High conflict real domains: It includes those real domains whose number of assigned preferential clusters above the third quartile. Therefore, we have the

² Euclidean distance




Fig. 3. The graph shows the relationship between the average Spearman Rank correlation in the domain and the ratio of outcomes that are Pareto optimal in a group of seven members

Zoning (10 clusters), Symposium (10), ERS 2 (12), Party (15), and Tram (20) domain.

In the previous bulleted list, it is possible to observe that most of the **real world domains tend to present characteristics that are similar to the synthetic domains with a low number of preferential clusters.**

3.2. Applicability: Analyzing the relationship between the degree of conflict and Pareto optimality in a group

From the definition of Pareto optimality, one can easily observe that when the preferential rankings of two agents are completely opposite, all of the outcomes are Pareto optimal. Apart from that, little is known with regards to the type of relationship between other degrees of conflict in a group and the number of outcomes that are Pareto optimal in that group. Understanding its relationship is important, as it may determine the kind of groups that may benefit from bottom-up approaches to Pareto optimality. Guided by this theoretical result, we formulate the following hypothesis:

Hypothesis 2. Overall, domains with a higher degree of conflict result in groups with a higher ratio of Pareto optimal outcomes.

More interestingly, we also want to study the type of relationship between the degree of conflict and the number of outcomes that are Pareto optimal in a group. For that, we carry out some experimental simulations with our synthetic domains. For each possible number of preferential clusters in the domain, we select 5 random domains with a domain size equal to 1,000 outcomes. Then, for each domain we calculate the average Spearman rank correlation in the domain, generate 100 random groups of 7 members, and calculate the ratio of outcomes that are Pareto optimal in each group.

Figure 3 shows the results of the simulation. It is easily observable that **H2** can be confirmed. Additionally, one can observe that there is a non-linear rela-




Fig. 4. The joint effect of group size and preferential conflict on the ratio of Pareto optimal outcomes

relationship between the degree of conflict in a domain and the ratio of outcomes that are Pareto optimal in a group. The higher the degree of conflict in the domain, the more rapidly the ratio of outcomes that are Pareto optimal increase. Despite the fact that the ratio of Pareto optimal outcomes does not even reach 50% in the graph, this non-linear relationship indicates that bottom-up approaches (or any other approach) to Pareto optimality may not be appropriate for domains where the degree of preferential conflict is very high, confirming our initial concern raised in Section 2. The problem is not caused by bottom-up approaches, but by the fact that Pareto optimality loses meaning and significance when almost all of the outcomes are Pareto optimal. Nevertheless, when practitioners employ bottom-up approaches to achieve Pareto optimality, they are advised to carry out a prior study to ensure that the degree of conflict in the domain does not tend to be high. Despite this result, and arising from the experiments in 3.1.2, it is suggested that most real and practical domains seem to have a low or moderate degree of conflict, high degrees of conflict being more unlikely.

3.3. Applicability: Analyzing the effect of the group size on Pareto optimality

There are still other aspects that we need to analyze to determine the applicability of bottom-up approaches to real scenarios. Another factor that may determine the applicability of bottom-up approaches to Pareto optimality is the size of the group. In (O'Neill, 1981), O'Neill estimated the number of Pareto optimal outcomes that one can expect in a domain with m outcomes and n agents:

$$E(K_{m,n}) = -\sum_{i=1}^m (-1)^i \binom{m}{i} \frac{1}{i^{n-1}} \quad (1)$$

To put it simply, the author proved that the number of Pareto optimal outcomes grows exponentially with the number of agents in the group, with the assumption that all preference profiles are equally likely. The author also showed that the size of the domain had an effect on the number of outcomes that are Pareto

optimal: larger outcome spaces tend to slow down the exponential growth of the Pareto optimal set, although the growth is still exponential. In order to draw that conclusion, the author assumed that all preference profiles are equally probable. We argue that, in practice, not all of the outcomes may be equally feasible (e.g., high prices in a team of buyers, popular choices in movies, popular choices in travel destinations, etc.). Therefore, our hypothesis is:

Hypothesis 3. The exponential relationship between group size and the ratio of outcomes that are Pareto optimal in a group will have a slower growth for scenarios where conflict is low, and its speed will increase as conflict increases in the domain.

This has a direct translation to the number of preferential clusters in the domain, as we have shown that there is an exponential relationship between the number of preferential clusters and the degree of conflict in a domain. Hence, the ratio of outcomes that are Pareto optimal should grow more slowly in domains with fewer preferential clusters (i.e., lower degree of conflict) than in domains with a higher number of preferential clusters (i.e., higher degree of conflict). The rationale behind this initial hypothesis is simple: given a group of agents, the more likely it is for all preferential clusters to be represented in a group when the number of clusters is low. Therefore, adding new agents to the group should not introduce significant conflict, and, therefore, keep the ratio of Pareto optimal outcomes almost invariant. A direct consequence of this hypothesis is that, if found true, the expected number of outcomes that are Pareto optimal in a domain may deviate from O’Neill’s formula. A question that stems from this consequence is the expected deviation that one can expect from the aforementioned formula when working with real domains with different degrees of conflict.

Firstly, we test whether or not the degree of preferential conflict has an effect on the exponential relationship between group size and the ratio of outcomes that are Pareto optimal. In order to isolate the effect of group size and the degree of conflict on the ratio of Pareto optimal outcomes, we focus on synthetic domains with a domain size equal to 1,000 outcomes. For each number of preferential clusters in the domain, we selected 5 random synthetic domains and generated random groups of sizes ranging from 3 to 9 members. More specifically, for each possible domain and group size we generated 100 random groups. Then, we calculated the ratio of Pareto optimal outcomes in each group. Figure 4 shows the results of this simulation. As it can be observed in Figure 4, there is an increasing relationship between the number of group members and the ratio of outcomes that are Pareto optimal. This is aligned with both (O’Neill, 1981) and our initial intuition. However, it can also be appreciated that, on average, the speed by which the ratio of Pareto optimal increases is different according to the number of preferential clusters in the domain. Higher numbers of preferential clusters (i.e., higher conflict) tend to increase the speed by which the ratio of Pareto optimal outcomes increases. The exponential relationship between both is more acute for larger numbers of preferential clusters, while it tends to flatten as the number of preferential clusters decreases. These results support **H3**.

After this analysis with synthetic data, we replicate a similar experiment with our real world domains. The goal of this experiment is twofold. First of all, we seek to analyze if bottom-up approaches to Pareto optimality are applicable to real world domains by providing some useful filtering when selecting Pareto optimal outcomes. Secondly, we also desire to analyze the differences between the exponential expression provided by (O’Neill, 1981) and the effect observed in real




Fig. 5. The effect of group size on the ratio of outcomes that are Pareto optimal for real domains with low degree of preferential conflict (top), mild degree of preferential conflict (middle), and high degree of preferential conflict (bottom)

domains with different degrees of conflict. For this experiment, we calculate the ratio of Pareto optimal outcomes for each real domain and group sizes ranging from 3 to 9 members. A selection of the results provided by this experiment can be observed in Figure 5. The figure shows the average ratio of outcomes that are Pareto optimal for different group sizes and domains, with the top row showing domains with a low degree of preferential conflict, the middle row showing domains with a mild degree of preferential conflict, and the bottom row showing domains with a high degree of preferential conflict. In these graphs, we represent the average ratio calculated in real scenarios (triangles) and the theoretical estimation provided by (O’Neill, 1981) (dots) for domains of the same size. In addition to this, for each data point we provide the total number of cases that provided the aggregate value³ that are considered for calculating the average.

Similarly to our previous experiment, it is possible to observe the increasing relationship between group size and the ratio of outcomes that are Pareto optimal in all of the graphs shown in Figures 5. It is also observable that the growth in the number of outcomes that are Pareto optimal is usually slower in real domains than in the theoretical estimation provided by (O’Neill, 1981). In our experiments including all of the real domains, only the symposium and the tram domain, two

³ The total number is $\min(1000, \binom{m}{n})$

of our domains with high conflict, show a similar or more exponential behavior to that of the theoretical case. The rest of the domains deviate from the theoretical behavior sooner or later, showing a less acute exponential relation. The reason behind this may be explained due to the fact that not all preference profiles are equally likely in real domains.

Another important appreciation from this experiment is that in no case the ratio of outcomes that are Pareto optimal was 100%. Even in the case of the largest groups (i.e., 9 members) and smaller domains, the overall ratio of Pareto optimal outcomes never saturated and reached all of the outcomes. This result is important, as it indicates that the application of bottom-up approaches to Pareto optimality would filter out outcomes. Of course, the exact number of outcomes that are discarded depends on the particular domain and its inherent characteristics (e.g., degree of conflict, domain size, etc.), and the group size. Therefore, it is difficult to provide a strict rule to determine the group size at which all of the outcomes will become Pareto optimal.

However, there are some trends and insights that one can employ as a practitioner. If we take into consideration the different degrees of preferential conflict in domains (i.e., number of clusters) we can also observe that domains with low, mild, and high levels of conflict behave differently. In domains with a low degree of conflict, it can be observed that the growth for the ratio of outcomes that are Pareto optimal tends to be slow compared to the theoretical case. On average, the difference in the ratio of outcomes that are Pareto optimal between experimental points in domains with a degree of conflict and the associated theoretical cases is 0.30. In practice, this means that larger group sizes would be needed to find situations where most of the outcomes are Pareto optimal. As we analyze mild conflict domains, we can observe that, overall, the growth in the number of outcomes that are Pareto optimal with the group size is closer to the theoretical case and, therefore, more exponential than domains with a low level of conflict. This difference can be numerically represented by the average difference in the ratio of outcomes that are Pareto optimal with respect to the theoretical case. In the case of mild conflict domains this difference is observed to be 0.17, lower than the low conflict case. However, the difference still suggests a relevant difference. This would suggest that, despite the point of saturation indicated by the formulation proposed by (O'Neill, 1981) for a domain of the same size, one can still form larger groups without the risk of all outcomes being Pareto optimal. Finally, when analyzing domains with a high degree of conflict, the growth in the ratio of outcomes that are Pareto optimal is the one that is closer to the exponential growth depicted by the theoretical case. The difference between the ratio of outcomes that are Pareto optimal in high conflict domains and the associated theoretical case is set at 0.07. A practical implication for domains with a high degree of conflict is that the group size at which one can expect for most of the outcomes to be Pareto optimal is close to the point indicated by O'Neill's equation.

The fact that, as we have shown, not all preference profiles are equally likely in practice makes bottom-up approaches more applicable to real life scenarios than the results depicted in theory (O'Neill, 1981). However, despite the fact that not all preference profiles are equally likely, and clusters of preferences do exist, the relationship between group size and the ratio of outcomes that are Pareto optimal still seems exponential. The exponential relationship between group size and the ratio of outcomes that are Pareto optimal, suggests that bottom-up approaches to Pareto optimality may not be useful for scenarios where very

large groups need to come to an agreement. Preferably, this approach should be taken for small and moderate group sizes.

3.4. Applicability: Determining the ratio of false positives in subgroups

As it was mentioned in Section 2, one can only guarantee with absolute certainty that a Pareto optimal outcome in a subgroup is also Pareto optimal in a larger subgroup when agents' preferences are strict. As we suggested before, in case of non-strict preferences, there are scenarios where an outcome that is Pareto optimal in a subgroup is not Pareto optimal in a larger group containing the subgroup. We refer to these outcomes that are Pareto optimal in the subgroup but not in the whole group as *false positives*. Even though the situation is possible, we argued that it may be unlikely in practice, as it requires all of the agents in a subgroup to be indifferent among the pair of outcomes. Taking that into consideration, we formulated the first hypothesis of this study (**H1**), which claims that the likeliness of finding a false positive Pareto optimal outcome in a subgroup is small in practice, due to the unlikeliness of finding a situation where agents are indifferent between two Pareto optimal outcomes

The study of the ratio of false positives raised when a bottom-up approach is applied is important, as a high ratio of false positives may deteriorate the efficiency of the final decision made by the group. Therefore, in this experiment we study the ratio of false positives found in subgroups of different sizes. We hypothesize the following:

Hypothesis 4. Overall, given a group of agents, the ratio of false Pareto optimal outcomes (i.e., false positives) discovered by bottom-up approaches is higher in smaller subgroups than it is in larger subgroups.

The rationale behind this hypothesis is that smaller groups of agents are more likely to agree on being indifferent between some outcomes than a larger group. As we are interested in studying this phenomenon in a realistic setting, we employ our real world domains.

For this experiment, we employ all of our real world domains, and we create 1,000 random groups of 5, 7, and 9 members for each domain. Then, for each group, we generate every possible subgroup of size ranging from 2 to the size of the group minus one. In each of these subgroups we calculate the Pareto optimal outcomes defined by the subgroup and then determine the percentage of false positives by comparing with the Pareto optimal outcomes defined by the whole group. Table 2 aggregates the results of this experiment by showing the average percentage and standard deviation (between brackets) of false positives generated in subgroups with different sizes and degrees of conflict.

As it can be observed in the table, the percentage of false positives generated in subgroups of all sizes is low, being approximately 1% at most. This suggests that, in practice, one can expect that if an outcome is Pareto optimal in a subgroup, it will also be Pareto optimal in any group containing that subgroup. **This result supports the applicability of bottom-up approaches to Pareto optimality, even in scenarios where preferences are not strict.** Our initial hypothesis **H4** is also confirmed by these experiments, as one can observe that smaller subgroups tend to generate higher percentages of false positives. Still, the percentages can be considered as low and it should not radically dis-

		Subgroup size						
Group size	2	3	4	5	6	7	8	
Low	5	0.31% (0.30)	0.20% (0.24)	0.10% (0.11)	-	-	-	-
	7	0.34% (0.42)	0.25% (0.31)	0.17% (0.21)	0.10% (0.13)	0.04% (0.06)	-	-
	9	0.32% (0.42)	0.25% (0.32)	0.18% (0.23)	0.12% (0.15)	0.08% (0.10)	0.04% (0.05)	0.01% (0.02)
Mild	5	0.25% (0.60)	0.09% (0.27)	0.02% (0.07)	-	-	-	-
	7	0.11% (0.28)	0.04% (0.13)	0.01% (0.04)	0.00% (0.00)	0.00% (0.00)	-	-
	9	0.00% (0.00)						
High	5	1.15% (2.44)	0.17% (0.38)	0.00% (0.00)	-	-	-	-
	7	0.31% (0.65)	0.02% (0.04)	0.00% (0.00)	0.00% (0.00)	0.00% (0.00)	-	-
	9	0.10% (0.20)	0.00% (0.00)	0.00% (0.00)	0.00% (0.00)	0.00% (0.00)	0.00% (0.00)	0.00% (0.00)

Table 2. Average percentage of false positives generated in subgroups of different sizes and degrees of conflict. The standard deviation of the average is found between brackets.

rupt the final decision made by the group. The results are also broken down according to the degree of conflict in the domain. However, we could not observe any pattern or notable difference arising from the individual analysis of the ratio of false positives in domains with different degrees of conflict.

3.5. Performance: Studying the ratio of the Pareto optimal outcomes achievable by bottom-up approaches

After studying the applicability of bottom-up approaches to Pareto optimality in scenarios with different characteristics, we aim to study the prospective performance of these approaches in real domains. First, we focus on the overall percentage of Pareto optimal outcomes that are achievable in subgroups⁴. If a higher percentage of the final Pareto optimal frontier is calculated, the whole group can be more flexible in the group negotiation. We argue that the prospective performance of bottom-up approaches depends on how the subgroups are formed. However, the subgroups can be formed by a very large number of strategies. As studying each of these strategies would be infeasible and it is out of the scope for this article, we focus on studying the performance achievable by bottom-up approaches when the subgrouping strategy ends up producing subgroup structures at extreme sides of a spectrum: one where a single subgroup is formed that does not include all agents, and another one where multiple subgroups are formed that contain all of the agents in the group. The performance achievable by specific subgrouping mechanisms should lie between both extremes:

- **Single maximum subgroup (SMS):** The general idea behind this family of scenarios is that only a single subgroup is formed, although it may be close to the whole group. Despite this, some agents are not included in the subgroup.

⁴ The number of Pareto optimal outcomes calculated in the subgroups compared to the total number of Pareto optimal outcomes in the whole group

- **Inclusive subgroups (IS):** In this case, multiple subgroups can be formed with varying size. The second family of scenarios ends up with each agent in a subgroup, leaving no agent left out of the bottom-up process.

By analyzing the prospective performance of bottom-up approaches in these two opposite scenarios, we aim to provide broad look at the performance achievable by the bottom-up paradigm. In addition to this, we also plan to provide recommendations on the outcomes that should be aimed by subgrouping strategies. Based on the description of SMS and IS, we can formulate the following hypotheses based on Theorem 1:

Hypothesis 5. In the *SMS* scenario, a higher ratio of the final Pareto optimal frontier are obtainable in larger subgroups.

Hypothesis 6. In the *IS* scenario, a higher ratio of the final Pareto optimal frontier are obtainable with fewer but larger subgroups.

Hypothesis 7. Higher ratios of the final Pareto optimal frontier are obtainable in *SMS* scenarios than in *IS* scenarios.

The reason for formulating both **H5** and **H6** is similar. An outcome is Pareto optimal due to the relationship between pairs of preference profiles and, thus, all the Pareto optimal outcomes in a group should be identified by the Pareto optimal outcomes arising from the interactions between all possible pairs of preference profiles in that group. Therefore, a larger subgroup should contain more pairs of preference profiles than smaller subgroups, and, then, it should produce a Pareto optimal frontier that is closer to the final Pareto optimal frontier in the whole group. In the case of **H7**, we believe that a larger subgroup should contain more pairs of agents than those contained by aggregating the pairs present in smaller subgroups.

In order to test these hypotheses, we run a series of experiments in our real domains. More specifically, for our *SMS* scenario, we create a maximum of 1,000 random groups for each combination of domain and group size, with the group sizes selected being 5, 7, and 9. For each random group we test all the possible subgroups of size ranging from 2 to the size of the group, and calculated the average ratio of outcomes of the final Pareto optimal frontier achievable in the subgroup. We follow a similar methodology for the *IS* scenario, whereby we create 1,000 random groups for each combination of domain, group size (5,7, and 9), and number of subgroups with equal size (from 2 subgroups to a setting where all agents are put in pairs). Then, we calculate the average ratio of the final Pareto optimal frontier that is achieved by aggregating the partial Pareto optimal frontiers found in the subgroups. The results of this experiment are gathered in Table 3. The table shows the average percentage of Pareto optimal outcomes calculated in case of ending up in a *SMS* scenario and a *IS* scenario for domains with a low, mild, and high degree of conflict respectively. The reader may also appreciate the standard deviation for each scenario within brackets.

In the *SMS* scenario, the larger the subgroup is, the higher the percentage of the Pareto optimal frontier calculated in that subgroup. This trend just confirms **H5** and highlights the rationale provided above: Pareto optimality depends on the conflict between pairs of agents, so subgroups that include more pairs of agents (i.e., larger subgroups) should be able to achieve a higher percentage of the Pareto optimal frontier. On average, and regardless of the team size, one can calculate 65% of the Pareto optimal frontier in a subgroup of about half

		SMS						IS			
		Subgroup size						Number subgroups			
		2	3	4	5	6	7	8	2	3	4
Low	n=5	50.0 (14.3)	65.9 (13.1)	82.4 (8.9)	-	-	-	-	89.62 (10.2)	-	-
	n=7	40.0 (16.5)	51.9 (18.5)	63.8 (18.1)	75.8 (14.8)	87.9 (8.6)	-	-	88.9 (13.6)	86.5 (12.8)	-
	n=9	35.1 (17.7)	45.4 (20.3)	55.5 (21.0)	65.2 (19.6)	74.7 (16.4)	83.6 (11.8)	92.1 (6.3)	89.8 (14.1)	86.0 (16.7)	85.9 (13.6)
Mild	n=5	33.6 (16.5)	55.2 (16.9)	77.6 (10.9)	-	-	-	-	71.5 (21.2)	-	-
	n=7	24.9 (15.2)	40.5 (18.6)	56.3 (17.9)	71.5 (14.0)	86.1 (7.8)	-	-	75.7 (18.9)	60.9 (26.7)	-
	n=9	22.4 (14.0)	35.9 (18.3)	48.8 (19.4)	60.9 (18.1)	72.0 (15.1)	82.2 (10.9)	91.4 (5.8)	85.3 (16.6)	74.9 (25.0)	67.0 (28.7)
High	n=5	14.6 (11.3)	33.6 (13.1)	62.9 (9.4)	-	-	-	-	44.5 (21.5)	-	-
	n=7	8.2 (9.4)	18.4 (12.5)	34.3 (13.0)	54.3 (10.9)	76.7 (6.3)	-	-	52.9 (18.4)	30.8 (24.8)	-
	n=9	6.1 (7.9)	13.3 (11.1)	24.2 (12.5)	37.9 (12.3)	53.2 (10.6)	69.1 (7.7)	84.9 (4.0)	61.6 (15.1)	38.7 (22.4)	26.3 (25.7)

Table 3. Average percentage of the Pareto optimal frontier achieved and standard deviation (brackets) in low, mild, and high conflict real domains by the *single maximum subgroup* with different subgroup sizes, and the *maximum inclusive subgroups* for different number of subgroups

the size of the group (i.e., subgroup of size 3 for a group of size 5, subgroup of size 4 for a group of size 7, and a subgroup of size 5 for a group of size 9) when the domain has a low conflict, 58% when the domain as a mild conflict, and 35% in case of domains with high conflict. In all the three cases, a notable proportion of the Pareto optimal frontier can be calculated. Even in the case when the subgrouping ends up with a minimum subgroup (i.e. a pair of agents), the approach is capable of obtaining an average of 41%, and 27% for low and mild conflict domains. It is only in domains when just an average of 9% of the Pareto optimal frontier can be provided with a single subgroup of two agents.

In the *IS* scenario, it is observable that the fewer subgroups the higher ratio of the final Pareto optimal outcomes is obtained. If we follow an inclusive strategy that includes all of the agents, the fewer the number of subgroups, the largest those subgroups are. As mentioned above, the larger the subgroups the more pair of agents that are contained in that subgroup, and, therefore, the more Pareto optimal points we should discover in those subgroups. This result confirms *H5*. In addition to this, one can also observe that the average percentage of Pareto optimal outcomes calculated in subgroups tends to be notable for low and mild conflict domains. To support this claim, note that even in the case that smaller subgroups are formed (i.e., 2 subgroups in the case of teams of size 5, 3 subgroups in the case of teams of size 7, and 4 subgroups in the case of teams of size 9), an average of 87% of the Pareto optimal frontier can be calculated for domains with low conflict and 67% for domains with mild conflict. The performance decreases when the approach is employed in domains with high conflict. Yet, even in the worst case (i.e., largest number of subgroups), the approach can detect an average of 33% of the final Pareto optimal frontier.

Overall, a notable percentage of the final Pareto optimal frontier can be calculated in both scenarios. However, there are some observations that can be made. First of all, one can observe that for any team size, and almost any scenario, the best result achievable in a *SMS* scenario (i.e., one subgroup

formed by all group agents except for one) seems to outperform the best result in a *IS* scenario (i.e., 2 subgroups with half the group members each). The gathered average is always higher, also showing a lower standard deviation. The only exception is for domains with a low level of conflict. More specifically, the overall difference between the best result achievable in the *SMS* scenario and the best result achievable in the *IS* scenario is about 2% higher for the *IS* scenario in domains with low conflict. This advantage is lost when analyzing domains with mild and high conflict. The best performance in the *SMS* scenarios is 7.53% higher in the case of a domain with mild conflict, and 21.8% higher in domains with high conflict. These results partially support **H7. Subgrouping strategies may aim for outcomes closer to the *IS* scenario when domains present a low conflict**, as forming smaller subgroups should be easier than forming a single and large subgroup. However, **subgrouping strategies should aim for outcomes closer to the *SMS* scenario for domains with a mild or high level of conflict** as the ratio of the final Pareto optimal outcomes calculated by the single and large subgroup tends to be higher with a lower standard deviation than in the case of smaller subgroups.

3.6. Performance: Studying the joint utility achievable by bottom-up approaches

In the previous experiment we studied the percentage of the final Pareto optimal frontier that is achievable in subgroups. Even though the percentage of the Pareto optimal frontier obtained in subgroups is a relevant metric to characterize the performance of bottom-up approaches, this metric does not offer a full picture about the performance of bottom-up approaches. A question that should be answered is: How fair is the subset of Pareto optimal outcomes for the whole group? In order to answer that question, we study what the best joint utility⁵ achievable using a bottom-up approach is. As in the previous experiment, we aim to provide a broad look at the performance of bottom-up approaches by focusing in the two opposite scenarios proposed in the previous experiment: *SMS* and *IS*. This comparison should also help us to identify what type of outcomes should be aimed by subgrouping strategies. Attending to the observations raised in the previous experiment, we formulate the following hypothesis:

Hypothesis 8. The best joint utility achievable in *SMS* scenarios outperforms the best joint utility achievable in *IS* scenarios.

In this particular experiment, we repeated the same settings described in Section 3.5. However, this time we focus on the maximum joint utility observable in the outcomes calculated in the subgroup(s). The results of this experiment can be found in Table 4, which contains the performance of low, mild, and high conflict real domains. The results from the different scenarios and repetitions are aggregated by means of the 10% quantile. This means that in 90% of the situations it is expected that the joint utility will be at least equal or greater than the joint utility reported in the tables.

We start by describing the performance achievable in *SMS* scenarios. In general, the joint utility that is achievable in a small subgroup is low compared with

⁵ Product of utilities of the agents in the group

		<i>SMS</i>						<i>IS</i>			
		Subgroup size						Number subgroups			
		2	3	4	5	6	7	8	2	3	4
Low conflict	n=5	0.51	0.74	0.98	-	-	-	-	0.99	-	-
	n=7	0.39	0.61	0.79	0.97	1.00	-	-	1.00	0.99	-
	n=9	0.32	0.50	0.66	0.89	0.97	1.00	1.00	1.00	0.99	0.99
Mild conflict	n=5	0.41	0.9	1.00	-	-	-	-	0.98	-	-
	n=7	0.25	0.75	0.95	1.00	1.00	-	-	0.99	0.95	-
	n=9	0.17	0.61	0.89	0.98	1.00	1.00	1.00	1.00	0.99	0.96
High conflict	n=5	0.26	0.53	0.89	-	-	-	-	0.81	-	-
	n=7	0.20	0.38	0.61	0.9	1.00	-	-	0.92	0.79	-
	n=9	0.18	0.30	0.51	0.73	0.95	1.00	1.00	0.99	0.90	0.78

Table 4. 10% quantile for the normalized joint utility obtainable in subgroups formed in the *SMS* and the *IS* scenarios in real domains with low, mild and high degree of conflict.

the joint utility achievable from the whole Pareto optimal frontier. For instance, overall, the joint utility achievable in subgroups that are smaller than half of the group⁶ is only 0.50 for scenarios with a low conflict, 0.51 for scenarios with a mild conflict, and 0.30 for scenarios with high conflict. This means that 90% of the times we can only guarantee for the best joint utility in a subgroup of less than half the group size to be half as good as the best joint utility achievable in the whole group for domains with low and mild conflict, and approximately over one third as good as the best joint utility achievable in the whole group for domains with high conflict. However, the best joint utility found in a single subgroup rapidly grows from that point onwards. In fact, the average joint utility that one can expect 90% of the times for subgroups of about half the size⁷ is 0.81 for domains with low conflict, 0.94 for domains with mild conflict, and 0.62 for domains with high conflict. The best joint utility can be guaranteed in subgroups containing all agents but one in almost any possible scenario. Nevertheless, one should consider that the inherent trust in the domain must be high for all team members but one to be able to form a subgroup and share preferential information.

With respect to *IS* scenario, one can observe that even in the situations where the smallest subgroups are formed (i.e., pairs of agents) the best joint utility achieved in 90% of the cases is very close to the best joint utility found in the whole group. Overall, we can expect that 90% of the times the best joint utility achievable in the smallest subgroups formed by this approach to be 0.99 for domains with low conflict, 0.96 for domains with mild conflict, and 0.79 for domains with high conflict. These results are closer to the optimum than in the case of *SMS* scenarios. This results suggest that, **subgrouping strategies should aim for smaller but inclusive subgroups rather than a single and large subgroup**, which may leave some agents out of the subgroups, **when optimizing the best joint utility of the group**.

This finding is in contrast with our findings in Section 3.5, where we highlighted that, in the best case, higher percentages of the final Pareto optimal set are obtainable in *SMS* scenarios. Despite providing a Pareto optimal subset that is less flexible for negotiations (i.e., a lower percentage of the final Pareto optimal frontier), the frontier obtainable in *IS* scenarios contains the best possible joint

⁶ Subgroups of size 2 for teams of size 5, subgroups of size 2, and 3 for groups of size 7, and subgroups of size 2, 3, and 4 for subgroups of size 9

⁷ Subgroups of size 3, 4, and 5 for groups of size 5, 7 and 9 respectively.

utility outcome, or at least an outcome that is very close to it. This rejects our initial hypothesis **H8** and raises a need for system designers to trade-off between both type of scenarios when creating/deploying subgrouping strategies.

3.7. Discussion

In the previous experiments we have studied both the applicability and prospective performance of bottom-up approaches to Pareto optimality. As other approaches, bottom-up approaches to Pareto optimality have both their strengths and limitations. As a consequence of the experiments that we have carried out, we have identified the scenarios where these approaches may be the most useful. This information is useful for researchers aiming to develop new group decision mechanisms based on this philosophy. Next, we highlight the most important results of our experiments:

- Firstly, we confirmed that the ratio of outcomes that are Pareto optimal increases in a non-linear way with the conflict in the scenario. As the whole objective of bottom-up approaches is pre-selecting outcomes that are Pareto optimal and filtering out the rest, it may not be useful if most of the outcomes are already Pareto optimal. This growth is consequence of the definition of Pareto optimality itself. Therefore, bottom-up approaches, and any other approach to reach Pareto optimality, should not be employed in domains where there is a high degree of conflict. The results suggest that they are **more useful when applied to scenarios with a low and mild level of conflict**.
- A partial consequence of the definition of Pareto optimality, and its relationship with the degree of conflict, is the effect of the group size. As a rule of thumb, the more group members, the more outcomes that are Pareto optimal. Again, this limits the applicability and meaning of Pareto optimality to particular scenarios. When the degree of conflict in the domain is low, it is meaningful to apply bottom-up approaches to Pareto optimality even for large groups. Even though in the experiments we never found scenarios and groups for which all of the outcomes were Pareto optimal, **we only suggest the application of this approach to small and medium groups in case of high and mild degrees of conflict**.
- Despite the fact that one can only guarantee that an outcome that is Pareto optimal in a subgroup is also Pareto optimal in the whole group in case of strict preferences, **in practice an outcome that is Pareto optimal in a subgroup will also be Pareto optimal in the whole group regardless the strictness of preferences**. This supports the applicability of bottom-up approaches to Pareto optimality.
- We have identified that, given the appropriate circumstances, **bottom-up approaches can provide with notable percentages of the final Pareto optimal frontier and good quality Pareto optimal outcomes for the group**.
- When **aiming to maximize the percentage of the Pareto optimal outcomes** calculated in subgroups, the **subgrouping strategies should aim to produce a larger and single subgroup**.
- When **optimizing the quality of the Pareto optimal outcomes for the group**, the **subgrouping strategies should aim to include as many agents as possible in subgroups**.

The aforementioned recommendations are guidelines that should provide researchers with indication on how and when to apply bottom-up approaches to Pareto optimality, and how to create subgrouping mechanisms that maximize the goals of the system. Of course, they are not strict rules that will determine the exact performance of bottom-up approaches in every single domain. Every domain has its own characteristics (e.g., conflict degree, group size, trust, etc.) and these characteristics should be properly analyzed before deciding on the application or rejection of bottom-up approaches to Pareto optimality.

4. Related Work

Since its introduction by the Italian mathematician Vilfred Pareto, Pareto optimality has been an efficiency and stability concept that has had an impact on many disciplines and areas of knowledge. Not only it has been studied in mathematics, but Pareto optimality has been considered a cornerstone concept in some computer science areas like artificial intelligence, especially in those fields concerned with making decisions by means of automated software (e.g., multi-agent systems, automated negotiation, etc.).

In automated negotiation, Pareto efficiency is a central quality measure of the negotiated outcome, and in particular to quantify the success in estimating the opponent's preferences (Baarslag, Hendrikx, Hindriks and Jonker, 2016). There have been several successful approaches proposing mechanisms that guarantee Pareto optimal or near Pareto optimal outcomes in negotiation processes.

For instance, Ehtamo *et al.* (Heiskanen et al., 2001) propose a centralized and mediated mechanism for achieving Pareto optimal outcomes in groups of agents. The negotiation model assumes that the domains are solely composed of multiple real valued negotiation issues, the utility function of agents is linear and additive, and that the agents have agreed on a set of feasible agreements resulting of the aggregation of the agents' constraints. The mediator proceeds by requiring agents to inform about the gradient to be followed to increase one's utility. The mediator then chooses a compromise direction and proposes a tentative agreement in that direction. The agents in the process then inform the mediator on an agent in the same compromise direction that improves the utility of the tentative agreement. With this information, the mediator chooses a new tentative agreement and the process is repeated until no further improvements can be done. The authors proved that by the end of the process, the achieved outcome is Pareto optimal. Differently to this work, bottom-up approaches are more general as they do not assume any particular domain values, or any particular type of utility function.

Luo *et al.* (Luo, Jennings, Shadbolt, Leung and Lee, 2003) propose a semi-cooperative negotiation model for buyers and sellers in electronic commerce. The agents represent their preferences over multiple issues by means of fuzzy constraints, and negotiate with each other in a bilateral process that gradually converges towards a Pareto optimal outcome in case that it exists. Similarly to this work setting, we also assume that the environment is semi-cooperative as some agents in the group are willing to engage in a more cooperative process with some group members. However, we again do not assume any particular type of preferences or utility functions. In addition to this, we consider group decision settings instead of bilateral processes.

In (Lai and Sycara, 2009), the authors propose a general framework for multi-issue bilateral negotiation that reaches Pareto optimal or near Pareto optimal

outcomes. The framework assumes convex utility functions, and a time constrained bilateral negotiation. The authors propose two proposing mechanisms that can be applied to different situations. First, the authors propose an offer proposal mechanism for scenarios where agents' know their own utility function but they do not have information about their opponents' utility function. In that situation, agents build their own iso-utility curve and chooses the offer from the iso-utility curve that is the most similar to the best offer proposed by the opponent in the previous round. Then, if more offers need to be sent in that round, offers in the neighborhood of the chosen offer are randomly selected. Experimentally, it is shown that, when both agents follow this strategy, final outcomes tend to be close to the Pareto optimal frontier. In the second mechanism, it is assumed that agents do not have an elicited utility function, but they can compare a handful of offers. The mechanism assumes the existence of a non-biased mediator that works by dividing the negotiation space into base lines. Agents choose a base offer from the base line and then a process is started to find a point in the base line that improves the utility of both agents with respect to the base offer. If the selected offer is rejected, the mediator updates the base line and the same process is repeated until an agreement is found. The experiments show that near Pareto optimal outcomes are achieved. Bottom-up approaches to Pareto optimality guarantee Pareto optimality, while also not making implicit assumptions about the agents' utility functions.

Lou *et al.* (Lou and Wang, 2016) propose a mediated multi-party negotiation mechanism that finds an approximation of the whole Pareto optimal frontier in a decentralized way. The negotiation model assumes a negotiation domain composed by real-valued issues and utility functions that are strictly convex. Under this assumption, the authors proposed a negotiation mechanism based on an iterative process that employs the weighted sum method and subgradient optimization. By controlling the number of iterations of the iterative process, it is possible to gradually converge towards the real Pareto optimal frontier. The main difference striving from this work and our paradigm is that we do not assume any particular type of utility function or domain. In addition to this, under a bottom-up approach to Pareto optimality it is assumed that a Pareto optimal solution is reached.

The aforementioned approaches normally make several assumptions with regards to the type of decision making domain, the way preferences are represented, and in many cases also with regards to the information that is shared in the process. One may be tempted to think that either bottom-up approaches to Pareto optimality are not necessary, or that the previous works will be automatically substituted by bottom-up approaches to Pareto optimality. The reality is that both are complementary and employable in different scenarios, as the underlying assumptions are different. Many important optimality metrics that employ the notion of Pareto efficiency can be used in conjunction with our work, including the distance to the Pareto frontier (Baarslag, Hendriks, Hindriks and Jonker, 2012; Lau, Li, Song and Kwok, 2008; Niemann and Lang, 2009; Rahman, Bahgat and Farag, 2011; Rau, Tsai, Chen and Shiang, 2006), correctly estimated Pareto outcomes (Baarslag, Hendriks, Hindriks and Jonker, 2013), and the distance to a fair solution (located on the Pareto frontier), such as the Nash solution (Hindriks, Jonker and Tykhonov, 2009; Mok and Sundarraj, 2005) or Kalai-Smorodinsky (Baarslag et al., 2012; Hindriks et al., 2009). Bottom-up approaches to Pareto optimality assume that there is some degree of trust and willingness to cooperate among some of the members of the group. It is this will-

ingness to cooperate and trust the one that allows for the formation of subgroups and the application of a more cooperative mechanism in subgroups. We expect for this cooperative mechanism to be less complex and costly than, for instance, the ones described in this literature review. Then, after the use of a cooperative mechanism, Pareto optimal outcomes found in subgroups can be aggregated to make a decision in the group, again with a mechanism that may be less complex than the ones described in the literature review. This inherent sense of trust and cooperation among some group members is not necessarily present in the approaches described above. Thus, bottom-up approaches to Pareto optimality are complementary to existing work.

Pareto optimality is not only important in multi-agent systems, but also in other research areas such as multi-objective optimization, where the solution to problems where different functions must be optimized at the same time. Normally, the maximization/minimization of one of the functions incurs in some loss for the other functions. Therefore, the importance to detect those solutions that are Pareto optimal. The importance of Pareto optimality in multi-objective optimization has given rise to a variety of centralized optimization methods to achieve Pareto optimality. For instance, specific genetic algorithms have been designed to seek approximations of the real Pareto optimal frontier (Deb, Pratap, Agarwal and Meyarivan, 2002; Seada and Deb, 2015). The idea behind these genetic algorithms is the preservation of those solutions that are non-dominated, and then the application of genetic operators taking these non-dominated solutions as parents. In addition to this, other specific algorithms have been designed to obtain Pareto optimal outcomes to optimization problems. For instance, Hu *et al.* (Hu et al., 2013) propose an iterative process to discover the Pareto optimal frontier in discrete domains. The algorithm relies on the existence of mechanisms that quickly determine the k best possible solutions to a single function. Generally, the algorithm compares the k -th best solution for a single objective j with the k -th best solution for the rest of objectives to determine if an objective needs to continue in the iterative process. k is progressively incremented and objectives are removed accordingly from the iterative process, unless no more objectives are left. At that point, the Pareto optimal frontier can be quickly determined.

The curse of the definition of Pareto optimality has also been documented in the multi-objective literature. For instance, Winkler (Winkler, 1985) documented that the number of non-dominated solutions increases as random objectives are added to an optimization problem. Due to this unfortunate property of the definition of Pareto optimality, some researchers have proposed practical alternatives to optimization with many problems with many objectives. For instance, some body of research has made successful attempts at simplifying a search space composed by many objectives into a space with fewer objectives that retain part of the information contained in the objectives removed (Deb and Saxena, n.d.; Murata and Taki, 2009). One could attempt to apply similar dimensionality reduction methods in our setting to reduce the number of agents participating in the decision-making process. However, there will always be some point at which most outcomes will become Pareto optimal or many preferences will not be represented in the decision making process. The solution to the curse of Pareto optimality is quite probably further guiding the decision making process into a particular subset of Pareto optimal outcomes. In this sense, the multi-objective optimization literature has already provided with some prospective efficiency metrics that could substitute Pareto optimality in cases where most outcomes are Pareto optimal (di Pierro, 2006; Corne and Knowles, 2007).

For example, Di Pierro *et al.* (di Pierro, 2006) define the concept of k optimality for deciding over Pareto optimal outcomes. A non-dominated outcome is defined as k -optimal when that outcome is non-dominated over every possible combination of k objectives. Thus, it results in a stronger concept of optimality that may help to choose a solution over a set of Pareto optimal outcomes. Despite the recent proposal in the multi-objective literature, as far as we know, the application of these stricter definitions of optimality has been largely overlooked in the decision-making literature. Due to the curse of Pareto optimality, it may be necessary to design group decision making algorithms that are guided by these metrics instead of Pareto optimality. The study of how bottom-up approaches can be applied to obtain other type of efficient outcomes is left as an interesting future line of study.

Finally, economic and theoretical studies are also a source of related work. For instance in (O'Neill, 1981), it is analyzed how the number of Pareto optimal outcomes exponentially increases with the number of agents by assuming that all preference profiles are equally probable. In our present study, we have, among other contributions, shown how real domains in practice behave with regards to Pareto optimality. More specifically, we have shown that, despite the increase in the number of Pareto optimal outcomes with the number of agents, the growth speed is not as quick as portrayed by (O'Neill, 1981). This is, as far as we know, our closest work in the study of the underlying properties of Pareto optimality. Of course, there have been other successful studies on Pareto optimality for specific domains and problems like characterizing fairness, or studying the relationship between monotonic solutions and Pareto optimality (Bogomolnaia and Moulin, 2015; García-Segarra and Ginés-Vilar, 2015), but their focus of study has not been on the exploration of bottom-up approaches for reaching Pareto optimality.

5. Conclusions & Future work

In this paper we have introduced a new paradigm to reach Pareto optimal outcomes in group decision making: bottom-up approaches to Pareto optimality. The paradigm is based on dividing groups into subgroups, and calculating portions of the final Pareto optimal frontier in the subgroups by means of a more cooperative mechanism than the one that would be applied in the whole group. It assumes the existence of certain trust and willingness to cooperate among some of the agents in the group. For the applicability of this paradigm, we have shown that an outcome that is Pareto optimal in a subgroup of agents, is also Pareto optimal in a group containing the aforementioned subgroup. This property holds always as long as agents' preferences are strict, but we have also shown how the property can hold in practice in case of non-strict preferences.

As far as we know, this is the first study on bottom-up approaches for finding Pareto optimal outcomes in a group of decision makers. This is a clear step away from the classic approach followed so far in the decision making literature, which consisted of the application of complex and specific mechanisms that guaranteed Pareto optimality by involving all of the agent at the same time. Being a new paradigm, in this article we have (i) proved the applicability of this paradigm from a theoretical perspective; (ii) discussed how bottom-up approaches can be applicable in practice (iii) studied what are the conditions that make the application of bottom-up approaches to Pareto optimality more convenient; (iv) studied the performance achievable by the paradigm under different conditions

and schemes; and (vi) identified goals that should be pursued by subgrouping strategies when optimizing different criteria. The experimental setting has proved both the applicability and achievable performance of bottom-up approaches in real world domains.

The bottom-up paradigm to Pareto optimality still remains largely unexplored and there is potential for further interesting research. We hope this work can further inspire research in the field of group decision making. For instance, given a decision making scenario that is prone to the application of the paradigm, one of the questions that needs to be analyzed is what are the optimal or near optimal processes that should be followed to divide agents into subgroups. As we have observed in the experiments, the performance of bottom-up approaches largely depends on this division mechanism. Moreover, some Pareto optimal outcomes may be ignored by a bottom-up approach. In those particular scenarios, it would be beneficial to research whether or not one can easily detect new Pareto optimal outcomes given an existing Pareto optimal subset. Finally, another interesting question is whether or not bottom-up approaches can be used to identify outcomes given stricter definitions of optimality, as Pareto optimality soon becomes meaningless as conflict and group size increase.

Acknowledgements

This work is part of the Veni research programme with project number 639.021.751, which is financed by the Netherlands Organisation for Scientific Research (NWO) and is supported by ITEA M2MGrids Project, grant number ITEA141011. It is also partly funded by the Spanish Ministry of Economy under the project reference TIN2015-65515-C4-1-R. We would like to thank the anonymous reviewers for their useful comments.

References

- Aydođan, R., Hindriks, K. V. and Jonker, C. M. (2014), Multilateral mediated negotiation protocols with feedback, *in* ‘Novel Insights in Agent-based Complex Automated Negotiation’, Springer, pp. 43–59.
- Baarslag, T., Hendriks, M. J., Hindriks, K. V. and Jonker, C. M. (2012), Measuring the performance of online opponent models in automated bilateral negotiation, *in* M. Thielscher and D. Zhang, eds, ‘AI 2012: Advances in Artificial Intelligence’, Vol. 7691 of *Lecture Notes in Computer Science*, Springer Berlin Heidelberg, pp. 1–14.
- Baarslag, T., Hendriks, M. J., Hindriks, K. V. and Jonker, C. M. (2013), Predicting the performance of opponent models in automated negotiation, *in* ‘International Joint Conferences on Web Intelligence (WI) and Intelligent Agent Technologies (IAT), 2013 IEEE/WIC/ACM’, Vol. 2, pp. 59–66.
- Baarslag, T., Hendriks, M. J., Hindriks, K. V. and Jonker, C. M. (2016), ‘Learning about the opponent in automated bilateral negotiation: a comprehensive survey of opponent modeling techniques’, *Autonomous Agents and Multi-Agent Systems* **30**(5), 849–898.
- Bogomolnaia, A. and Moulin, H. (2015), ‘Size versus fairness in the assignment problem’, *Games and Economic Behavior* **90**, 119–127.
- Conitzer, V., Lang, J. and Xia, L. (2011), Hypercubewise preference aggregation in multi-issue domains, *in* ‘IJCAI Proceedings-International Joint Conference on Artificial Intelligence’, Vol. 22, p. 158.
- Corne, D. W. and Knowles, J. D. (2007), Techniques for highly multiobjective optimisation: Some nondominated points are better than others, *in* ‘Proceedings of the 9th Annual Conference on Genetic and Evolutionary Computation’, GECCO ’07, ACM, New York, NY, USA, pp. 773–780.

- Deb, K., Pratap, A., Agarwal, S. and Meyarivan, T. (2002), ‘A fast and elitist multiobjective genetic algorithm: Nsga-ii’, *IEEE transactions on evolutionary computation* **6**(2), 182–197.
- Deb, K. and Saxena, D. K. (n.d.), ‘On finding pareto-optimal solutions through dimensionality reduction for certain large-dimensional multi-objective optimization problems’.
- di Pierro, F. (2006), Many-objective evolutionary algorithms and applications to water resources engineering, PhD thesis, University of Exeter.
- Esparcia, S., Sanchez-Anguix, V. and Aydođan, R. (2013), A negotiation approach for energy-aware room allocation systems, in ‘Highlights on Practical Applications of Agents and Multi-Agent Systems’, Springer, pp. 280–291.
- García-Segarra, J. and Ginés-Vilar, M. (2015), ‘The impossibility of paretian monotonic solutions: A strengthening of roth’s result’, *Operations Research Letters* **43**(5), 476–478.
- Goldberg, K., Roeder, T., Gupta, D. and Perkins, C. (2001), ‘Eigentaste: A constant time collaborative filtering algorithm’, *information retrieval* **4**(2), 133–151.
- Hara, K. and Ito, T. (2013), A mediation mechanism for automated negotiating agents whose utility changes over time, in ‘Twenty-Seventh AAAI Conference on Artificial Intelligence’.
- Heiskanen, P., Ehtamo, H. and Hämäläinen, R. P. (2001), ‘Constraint proposal method for computing pareto solutions in multi-party negotiations’, *European Journal of Operational Research* **133**(1), 44–61.
- Hindriks, K. V., Jonker, C. M. and Tykhonov, D. (2009), The benefits of opponent models in negotiation, in ‘Proceedings of the 2009 IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology’, Vol. 2, IEEE Computer Society, pp. 439–444.
- Hu, X.-B., Wang, M. and Di Paolo, E. (2013), ‘Calculating complete and exact pareto front for multiobjective optimization: a new deterministic approach for discrete problems’, *IEEE Transactions on Cybernetics* **43**(3), 1088–1101.
- Jonker, C. M., Robu, V. and Treur, J. (2007), ‘An agent architecture for multi-attribute negotiation using incomplete preference information’, *Autonomous Agents and Multi-Agent Systems* **15**(2), 221–252.
- Kamishima, T. (2003), Nantonac collaborative filtering: recommendation based on order responses, in ‘Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining’, ACM, pp. 583–588.
- Lai, G., Li, C. and Sycara, K. (2006), ‘Efficient multi-attribute negotiation with incomplete information’, *Group Decision and Negotiation* **15**(5), 511–528.
- Lai, G. and Sycara, K. (2009), ‘A generic framework for automated multi-attribute negotiation’, *Group Decision and Negotiation* **18**(2), 169–187.
- Lau, R. Y., Li, Y., Song, D. and Kwok, R. C.-W. (2008), ‘Knowledge discovery for adaptive negotiation agents in e-marketplaces’, *Decision Support Systems* **45**(2), 310–323.
- Lin, R., Kraus, S., Baarslag, T., Tykhonov, D., Hindriks, K. and Jonker, C. M. (2014), ‘Genius: An integrated environment for supporting the design of generic automated negotiators’, *Computational Intelligence* **30**(1), 48–70.
- Lou, Y. and Wang, S. (2016), ‘Approximate representation of the pareto frontier in multi-party negotiations: Decentralized methods and privacy preservation’, *European Journal of Operational Research* **254**(3), 968–976.
- Luo, X., Jennings, N. R., Shadbolt, N., Leung, H.-f. and Lee, J. H.-m. (2003), ‘A fuzzy constraint based model for bilateral, multi-issue negotiations in semi-competitive environments’, *Artificial Intelligence* **148**(1-2), 53–102.
- Mattei, N. and Walsh, T. (2013), Preflib: A library of preference data [HTTP://PREFLIB.ORG](http://preflib.org), in ‘Proceedings of the 3rd International Conference on Algorithmic Decision Theory (ADT 2013)’, Lecture Notes in Artificial Intelligence, Springer.
- Miller, B. N., Albert, I., Lam, S. K., Konstan, J. A. and Riedl, J. (2003), MovieLens unplugged: experiences with an occasionally connected recommender system, in ‘Proceedings of the 8th international conference on Intelligent user interfaces’, ACM, pp. 263–266.
- Mok, W. W. H. and Sundarraj, R. P. (2005), ‘Learning algorithms for single-instance electronic negotiations using the time-dependent behavioral tactic’, *ACM Transactions on Internet Technology* **5**(1), 195–230.
- Murata, T. and Taki, A. (2009), Many-objective optimization for knapsack problems using correlation-based weighted sum approach, in ‘Evolutionary Multi-Criterion Optimization’, Springer, pp. 468–480.
- Niemann, C. and Lang, F. (2009), Assess your opponent: A bayesian process for preference observation in multi-attribute negotiations, in T. Ito, M. Zhang, V. Robu, S. Fatima and

- T. Matsuo, eds, 'Advances in Agent-Based Complex Automated Negotiations', Vol. 233 of *Studies in Computational Intelligence*, Springer Berlin Heidelberg, pp. 119–137.
- O'Neill, B. (1981), 'The number of outcomes in the pareto-optimal set of discrete bargaining games', *Mathematics of Operations Research* pp. 571–578.
- O'Neill, J. (2013), 'Open stv'.
URL: www.OpenSTV.org
- Rahman, S. A., Bahgat, R. and Farag, G. M. (2011), 'Order statistics bayesian-mining agent modelling for automated negotiation', *Informatica: An International Journal of Computing and Informatics* **35**(1), 123–137.
- Ramanathan, R. and Ganesh, L. (1994), 'Group preference aggregation methods employed in ahp: An evaluation and an intrinsic process for deriving members' weightages', *European Journal of Operational Research* **79**(2), 249–265.
- Rau, H., Tsai, M.-H., Chen, C.-W. and Shiang, W.-J. (2006), 'Learning-based automated negotiation between shipper and forwarder', *Computers & industrial engineering* **51**(3), 464–481.
- Sanchez-Anguix, V., Aydođan, R., Baarslag, T. and Jonker, C. M. (2017), Can we reach pareto optimal outcomes using bottom-up approaches?, in 'Conflict Resolution in Decision Making: Second International Workshop, COREDEMA 2016, The Hague, The Netherlands, August 29-30, 2016, Revised Selected Papers', Springer International Publishing, pp. 19–35.
- Sanchez-Anguix, V., Aydođan, R., Julian, V. and Jonker, C. (2014), 'Unanimously acceptable agreements for negotiation teams in unpredictable domains', *Electronic Commerce Research and Applications* **13**(4), 243–265.
- Sanchez-Anguix, V., Dai, T., Semnani-Azad, Z., Sycara, K. and Botti, V. (2012), Modeling power distance and individualism/collectivism in negotiation team dynamics, in '45 Hawaii International Conference on System Sciences (HICSS-45)', pp. 628–637.
- Sanchez-Anguix, V., Julian, V., Botti, V. and Garcia-Fornes, A. (2012), '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., Julian, V., Botti, V. and Garcia-Fornes, A. (2013a), 'Studying the Impact of Negotiation Environments on Negotiation Teams' Performance', *Information Sciences* **219**, 17–40.
- Sanchez-Anguix, V., Julian, V., Botti, V. and García-Fornes, A. (2013b), 'Tasks for agent-based negotiation teams: Analysis, review, and challenges', *Engineering Applications of Artificial Intelligence* **26**(10), 2480–2494.
- Seada, H. and Deb, K. (2015), U-nsga-iii: A unified evolutionary optimization procedure for single, multiple, and many objectives: Proof-of-principle results, in 'International Conference on Evolutionary Multi-Criterion Optimization', Springer, pp. 34–49.
- Skowron, P., Faliszewski, P. and Slinko, A. (2013), Achieving fully proportional representation is easy in practice, in 'Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems', International Foundation for Autonomous Agents and Multiagent Systems, pp. 399–406.
- Winkler, P. (1985), 'Random orders', *Order* **1**(4), 317–331.
- Zhenh, R., Chakraborty, Nilanjan, Dai, T. and Sycara, K. (2015), 'Automated multilateral negotiation on multiple issues with private information', *INFORMS Journal on Computing* p. In Press.
- Ziegler, C.-N., McNee, S. M., Konstan, J. A. and Lausen, G. (2005), Improving recommendation lists through topic diversification, in 'Proceedings of the 14th international conference on World Wide Web', ACM, pp. 22–32.

Author Biographies



Victor Sanchez-Anguix received a BSc. in Computer Science at Universitat Politècnica de València, Valencia, Spain. Later, he obtained a MSc. in Artificial Intelligence and Pattern Recognition by Universitat Politècnica de València. In 2013 obtained a PhD. in Computer Science at Universitat Politècnica de València, which was awarded with a Magna Cum Laude award and an award for being one of the the top 10% dissertations presented during 2013 at Universitat Politècnica de València. Victor Sanchez-Anguix has published several journal and conference articles on artificial intelligence. He is currently a lecturer at Florida Universitaria and Universidad Isabel I, Spain, as well as a guest member at Coventry University, UK. He carries out research on multi-agent systems, applied data mining, artificial intelligence, and automated negotiation.



Reyhan Aydoğan is an assistant professor at Özyegin University, Istanbul and guest researcher in Interactive Intelligence Group at Delft University of Technology, the Netherlands. She received her PhD in Computer Science from Boğaziçi University in 2011. After then she has worked at Delft University of Technology. As a visiting scholar, she has been at Massachusetts Institute of Technology, Norwegian University of Science and Technology and Nagoya Institute of Technology. Her research focuses on the modeling, development and analysis of agent technologies that integrate different aspects of intelligence such as reasoning, decision making and learning. She is applying artificial intelligence techniques such as machine learning and semantic reasoning in designing and developing agent-based decision support systems, particularly negotiation support systems and automated negotiation tools.



Tim Baarslag is a Scientific Staff Member at CWI (Centrum Wiskunde & Informatica) and an Assistant Professor at Utrecht University. He is a Visiting Associate Professor at Nagoya University of Technology and Visiting Fellow at the University of Southampton. Tim Baarslag graduated from Utrecht University with a MSc in Mathematics and a BSc in Computer Science. He obtained his PhD from Delft University of Technology in 2014 on the topic of intelligent decision support systems for automated negotiation. Tim Baarslag studies intelligent and autonomous systems that can collaborate through the process of joint decision making. His research interests include agent-based negotiation, coordination and cooperation, preference elicitation, decision making under uncertainty.



Catholijn Jonker is full professor of Interactive Intelligence (0.8 fte) at the Faculty of Electrical Engineering, Mathematics and Computer Science of the Delft University of Technology and full professor of Explainable Artificial Intelligence (0.2 fte) at the Leiden Institute for Advanced Computer Science of Leiden University. She received her PhD in computer science from Utrecht University in 1994. She chaired De Jonge Akademie (Young Academy) of the KNAW (The Royal Netherlands Society of Arts and Sciences) in 2005 and 2006, and she was a member from 2005 to 2010. She is a member of the Koninklijke Hollandsche Maatschappij der Wetenschappen and of the Academia Europaea. She was the president of the National Network Female Professors (LNVH) in The Netherlands from 2013 till 2016. Catholijn is EurAI Fellow since 2015, and EurAI board member since 2016.

Correspondence and offprint requests to: Victor Sanchez-Anguix, Florida Universitaria, Carrer del Rei en Jaume I, 2, 46470, Catarroja, Spain. Email: vsanchez@florida-uni.es