

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/304749404>

Can We Reach Pareto Optimal Outcomes Using Bottom-Up Approaches?

Conference Paper in Lecture Notes in Computer Science · August 2016

DOI: 10.1007/978-3-319-57285-7_2

CITATIONS

0

READS

44

4 authors:



Victor Sanchez-Anguix

Florida Universitaria

40 PUBLICATIONS 217 CITATIONS

[SEE PROFILE](#)



Tim Baarslag

Centrum Wiskunde & Informatica

69 PUBLICATIONS 669 CITATIONS

[SEE PROFILE](#)



Catholijn M. Jonker

Delft University of Technology

541 PUBLICATIONS 6,283 CITATIONS

[SEE PROFILE](#)



Reyhan Aydogan

Ozyegin University

47 PUBLICATIONS 214 CITATIONS

[SEE PROFILE](#)

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



Complex Negotiation [View project](#)



COREDEMA 2017 - 3rd International Workshop on Conflict Resolution in Decision Making In conjunction with the International Joint Conference on Artificial Intelligence (IJCAI-2017) [View project](#)

Can we reach Pareto optimal outcomes using bottom-up approaches?

Victor Sanchez-Anguix¹ and Reyhan Aydoğan^{2,4} and Tim Baarslag³ and Catholijn M. Jonker⁴

¹ Coventry University, Coventry, United Kingdom, ac0872@coventry.ac.uk

² Özyeğin University, Istanbul, Turkey, reyhan.aydogan@ozyegin.edu.tr

³ Centrum Wiskunde & Informatica, Amsterdam, Netherlands, T.Baarslag@cwi.nl

⁴ Technical University of Delft, Delft, Netherlands, C.M.Jonker@tudelft.nl

Abstract. Traditionally, researchers in decision making have focused on attempting to reach Pareto Optimality using horizontal approaches, where optimality is calculated taking into account every participant at the same time. Sometimes, this may prove to be a difficult task (e.g., conflict, mistrust, no information sharing, etc.). In this paper, we explore the possibility of achieving Pareto Optimal outcomes in a group by using a bottom-up approach: discovering Pareto optimal outcomes by interacting in subgroups. We analytically show that Pareto optimal outcomes in a subgroup are also Pareto optimal in a supergroup of those agents in the case of strict, transitive, and complete preferences. Then, we empirically analyze the prospective usability and practicality of bottom-up approaches in a variety of decision making domains.

Keywords: pareto optimality, agreement technologies, group decision making, multi-agent systems, artificial intelligence

1 Introduction

Group decision making has been studied within different disciplines with aim of reaching a mutually acceptable outcome. One of the desired properties of that outcome is Pareto optimality. However, reaching Pareto optimal agreements is not straightforward in practice. In open and dynamic environments, decision makers may not know each other's preferences completely. It may even be the case that it becomes more complicated to find Pareto optimal solutions when the number of participants increases, as the number of interactions required to achieve an optimal deal for the group may increase due to internal conflicts or lack of trust.

A number of works in the field focus on finding a global Pareto optimal solution by involving all agents at the same time [11, 16, 10, 29], which may lead to complicated 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 knowing or predicting the preferences of all agents. In essence, solving the Pareto optimal set problem in a smaller group may be less complicated than in larger groups (e.g., less privacy concerns, less interactions needed, more willingness to cooperate, etc.) and it may provide a relatively important ratio of the final Pareto Optimal outcomes. Such kind of property can be used in some complex group decision making scenarios. Imagine that a group of agents is negotiating in unison with an unknown opponent [22, 27, 25]. If the agents can find the Pareto optimal outcomes within the team, they may use these outcomes in their bidding strategy to reach a Pareto optimal agreement with their opponent.

In this paper we explore bottom-up strategies. For that, first we 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 are strict linear order. Second, we empirically simulate how bottom-up approaches may perform in realistic scenarios. More specifically, we show that we can obtain a reasonable ratio of the Pareto optimal outcomes within a group of agents by only finding the Pareto optimal outcomes within the subgroup of these agents.

The remainder of this paper is organized as follows: first we present a proof of how Pareto optimal solutions in subgroups are also Pareto optimal in larger groups when agents have strict, transitive, and complete preferences. Section 3 discusses some of the implications of the proof, and how it can be applied to solve a wide variety of problems in multi-agent systems. In Section 4, we empirically validate the theory in practice and analyze empirically compare the ratio of Pareto optimal outcomes within subgroups to the Pareto optimal outcomes within the entire group in a wide variety of real domains. After discussing the related work, we finally conclude the paper with future lines of work.

2 Pareto optimality in subgroups

In this section we prove that any Pareto optimal outcome in a subgroup of agents is also Pareto optimal in any group of agents containing the subgroup. First, we provide some of the necessary definitions and introduce some notation.

Let $\mathcal{A} = \{a_1, \dots, a_n\}$ be a set of agents where k is the index of agent a_k and $\mathcal{A}' = \{a_1, \dots, a_m\}$ be a superset of \mathcal{A} , $\mathcal{A} \subset \mathcal{A}'$ where $m > n$. \mathcal{O} is the set of all possible solutions in a given domain, and $o \in \mathcal{O}$ represents a possible solution in the domain. We assume that \succeq_i represents agent's a_i preference relation over outcomes 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 relations are strict, transitive and complete.

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

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

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

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

This means 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.

At this point the reader may be wondering how the theorem above behaves in a scenario where agents' preferences are not strict. As we will discuss later, the likeliness of such a scenario is small, but the conclusion of the theorem above may in fact not hold in that case. Basically, an outcome that is Pareto optimal in a subgroup \mathcal{A} may not be Pareto optimal in the group \mathcal{A}' when all of the agents in \mathcal{A} are indifferent between such outcome and another Pareto optimal outcome. Then, one of the two outcomes may not be Pareto optimal with \mathcal{A}' when one of the agents in the group is not indifferent between those outcomes. Nevertheless, as we shall outline in Section 4.2, such situations are rare in practice, as all of the agents need to be indifferent between outcomes. This becomes increasingly unlikely as the group size grows and thus, for large enough groups, we can consider that the theorem is true for practically any scenario.

3 Prospective applications

In Section 2 we have demonstrated that an outcome that is Pareto optimal in a subgroup of agents will also remain Pareto optimal in a larger group⁵. It should be highlighted that we are not depicting achieving Pareto optimality as a simple task. However, there is value in computing Pareto optimality in smaller groups as long as we are able to use those solutions in more challenging scenarios:

⁵ For strict, transitive, and complete preferences

- **Negotiation teams:** In this scenario, a group of individuals negotiate as a party with opponent(s) to achieve a deal [22, 27, 26, 25, 23]. In that case, finding the outcomes that are Pareto optimal within the team may play in favor of the team as (i) if the team sticks to these outcomes while negotiating with opponents, it can ensure efficiency in the final outcome, (ii) the set of calculated deals may be reused in multiple negotiations with different opponents as they remain Pareto optimal, and (iii) finding Pareto optimal outcomes once may reduce the time spent in negotiation threads as the team exactly knows which outcomes are more beneficial for team members. On top of that, one can also assume that team members may be more willing to share information with teammates, which may make easier the search for Pareto optimal outcomes inside the team.
- **Multi-party negotiations:** Some participants in a multi-party negotiation [29, 6, 3, 8, 10, 11] may decide to collude and bias the agreement with their preferences. For that, the subgroup of agents may calculate Pareto optimal outcomes within the subgroup, and decide on the Pareto optimal outcomes that they plan to use in the upcoming multi-party negotiation. This way, there may be higher probabilities for the negotiation to finish with an outcome that satisfies the subgroups’ interests and that is efficient. Another possible application for this proof in multi-party settings is precisely the idea of looking for Pareto optimal agreements within subgroups of agents. For instance, agents with high degrees of trust may decide to share some information that facilitates the search of Pareto optimal outcomes within the subgroup. Then, once outcomes are found in subgroups, these may be shared among all of the agents, and the whole group may need to decide on the most appropriate Pareto optimal outcome.
- **Decision making in open environments:** Open multi-agent systems [2, 24, 12] have the particularity of being systems where agents enter and leave the system dynamically. In such environments, decision making tasks may suffer from the same characteristic and agents may enter and leave decision making tasks as needed, resulting in a real time problem. For those situations, agents in a decision making task may benefit from a continuous search for Pareto optimal outcomes. As new agents join the task, those Pareto optimal outcomes calculated should be kept as they will remain Pareto optimal in the new group. When agents leave, remaining group members can get rid of some outcomes that have become dominated in the new setting.

As the reader may have noticed, the range of applications where this approach could be applied is varied. We are not claiming that those are the sole applications for this approach, and there may be others in domains like social choice, group recommendations, and so forth.

4 Experimental study

Section 2 shows theoretically that Pareto optimal outcomes within a group of agents having complete, transitive and strict preferences are still Pareto optimal

when the group size increases with incoming agents. Even when preferences are non strict, we expect for the theorem to hold in most of the cases. In this section, we empirically analyze the prospective performance and applicability of bottom-up approaches. For that purpose, we selected a variety of domains:

- **Sushi domain:** 5000 preference profiles over 10 types of sushi [15].
- **AGH course selection:** 153 students’ preferences over 6 courses offered by AGU University of Science and Technology in 2004 [28].
- **Book crossing domain:** The original dataset contains preferences of 278,858 users that produced 1,149,780 ratings over 271,379 books [30]. In order to calculate Pareto optimality, we require preferences to be complete on at least a subset of items. We kept 7 users that had rated 23 books in common.
- **Movielens domain:** The original dataset contains 138,000 users that provided ratings over 27,000 movies [19]. As we require complete preferences, we picked 10 preference profiles that had rated a total of 298 movies in common.
- **Holiday domain:** A multi-party negotiation domain available in Genius [18]. In this scenario, participants need to decide on the details of a holiday trip. In total, 9 preference profiles over 1024 possible outcomes are available. These preferences have been elicited from TU Delft computer science students, but not with serious plans for a joint holiday in mind.
- **Symposium domain:** Another multi-party negotiation domain that is available in Genius [18] concerning the organization of a conference. There are 9 preference profiles over 2305 possible outcomes. These preferences have been elicited from faculty members in computer science of TU Delft experienced in organizing conferences, but not having a specific conference in mind.
- **Party domain:** Another multi-party negotiation domain, where agents decide on the details of a party [18]. We elicited preferences from students in a Master level AI course. Students were asked to input their real preferences via Genius based on their tastes for hosting parties. In total, we elicited 24 real preference profiles over 3072 outcomes.

From a global perspective, the sushi, agh, and book domain are *small* attending to the number of outcomes. These domains correspond to decision making domains where outcomes are non customizable objects (e.g., a movie, a book, a course, etc.). The data in the Movielens domain is less sparse and we were able to find 10 users that had rated 298 outcomes in common. This is again a domain where outcomes are non customizable, but the size of the domain is one order of magnitude larger than that of the small domains. The three remaining multi-party negotiation domains (i.e., holiday, symposium, and party domain) represent scenarios where the final outcome can be customized via the negotiable issues. As a result, the number of possible outcomes is larger. We consider these domains and the Movielens domain as the *large domains* in our study.

4.1 Validation and Performance Analysis

Our performance metric is the ratio of the Pareto optimal outcomes within a subgroup with a size of $\{2, \dots, n-1\}$ to the Pareto optimal outcomes within the

n-sized group. If the ratio remains low even for large subgroups, then this means that the performance of our theoretical finding may be of little value in practice, as only a small ratio of the final Pareto outcomes may be achievable. However, if the ratio is large, then it may indicate that bottom-up approaches may be valuable. Additionally, common sense indicates that, the larger the subgroup, the higher the ratio of final Pareto optimal outcomes that may be obtained in the subgroup. However, one question that arises is the actual speed by which the ratio of final Pareto outcomes increases, and whether or not subgroups may be able to calculate a respectable ratio of the final Pareto optimal outcomes.

For testing the practical performance of our bottom-up approaches, we randomly generated groups of size n based on the preference profiles available for each domain. For each randomly generated group, we built all possible subgroups with varying sizes $\{2, \dots, n - 1\}$ and estimated the Pareto optimal set in each (sub)group. More specifically, for each domain we tested a maximum of 1000 groups⁶ of size $n = \{5, 7, 9\}$.




Fig. 1. Average ratio of the final Pareto optimal obtained in subgroups of different size

⁶ The total number is $\min(1000, \binom{m}{n})$, where m is the total number of available preference profiles and n is the size of the group

The results of this experiment can be observed in Figure 1. As expected, the results show that the larger the subgroup is, the larger the average ratio of the final Pareto Optimal set that we get. The increase is clearly continuous for all of the domains and group sizes. When we look at the results for groups of 7 and 9 members we observe a non-linear increase with the size of the subgroup. This non-linear increase is not as evident in the case of 5 members' groups, as in that case we only have 3 data points⁷.

One should highlight that for $n - 1$ agents in the subgroup, n being the total number of agents in the group, the average ratio of the Pareto optimal set obtained in the subgroup is always over 50% of the final set, being close to 80% in some cases (e.g, smaller domains, larger groups). This is a good result, especially for negotiation team scenarios [22, 27, 25], where the team could calculate the Pareto set inside the team and use those outcomes in the negotiation with an opponent. This is a clear case where a subgroup of size $n - 1$ can be formed (i.e., all of the team members) and, according to the experimental results, obtain a notably high ratio of final Pareto optimal outcomes. Consequently, they can propose Pareto optimal bids without knowing their opponent's preferences.

The result is also notable for smaller subgroups. For instance, in groups of size 5, we are able to obtain between an average of 68% of the final Pareto set for small domains and 32% for the larger domains with just about half of the group members (i.e., 3). In the case of groups of size 7, we get 68% of the final Pareto set for small domains and 28% for larger domains with just about half of the group members (i.e., 4). Similarly, for groups of size 9 we are able to obtain an average of 76% of the final Pareto set in small domains, and 30% in large domains with just about half of group members (i.e., 5).

The trends in the graphics and the results mentioned above may also suggest that larger domains may result in lower ratios of the final Pareto optimal set achievable by subgroups. Nevertheless, as we have been able to observe above, the results can still be considered as positive. Although the current results are promising, we would like to test a wider range of domains and domain sizes to strengthen the results of this study.

4.2 Applicability Analysis

There are still other aspects that we need to analyze to determine the applicability of bottom-up approaches in real situations. Even though considerable ratios of the final Pareto optimal set are obtainable within subgroups, this may be useless in practice if the total number of Pareto optimal outcomes is very close to all possible outcomes. In those cases, there would be no point in calculating Pareto optimal outcomes in subgroups, as almost any outcome would be Pareto optimal. Therefore, we are interested in checking that the set of final Pareto optimal outcomes does not dramatically approach the total number of outcomes. In [20], O'Neill studied how Pareto optimality was affected by the number of agents participating in a decision making process. To put it simply, the author

⁷ Even non-linear functions may look like linear when the number of points is reduced

proved that the number of Pareto optimal outcomes grows exponentially with the number of agents, with the assumption that all preference profiles are equally probable. Additionally, he proposed a formula to estimate the number of outcomes that are expected to be Pareto optimal based on the size of the domain m , and the number of agents in the group n : $E(K_{m,n}) = -\sum_{i=1}^m (-1)^i \binom{m}{i} \frac{1}{i^{n-1}}$.

He also stated that the size of the domain had an effect on the number of outcomes that were Pareto optimal: larger outcome spaces tend to slow down the exponential growth of the Pareto optimal set, although the growth is still exponential. Of course, for drawing such a conclusion, the author had to assume that all preference profiles were equally probable. We argue that, in practice, all preference profiles are not equally probable as in some domains 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.). Hence, we argue that the exponential growth may not be as fast as in the theoretical case, and bottom-up approaches may be applicable to more scenarios.

In order to examine this theoretical finding in practice, we calculated the ratio of the Pareto optimal outcomes to the total number of outcomes for each domain and group size. Figure 2 shows the average ratio of outcomes that are Pareto optimal for different groups sizes and domains. In these graphs, blue dots represent the average ratios calculated in real scenarios while green dots denote the theoretical estimation provided by [20] for domains of the same size. In addition to this, for each data point we provide the total number of cases⁸ that were considered for calculating the average. Numbers in red represent less than 30 samples and such averages should be ignored.

As it can be observed in Figure 2, the growth in the number of outcomes that are Pareto optimal is usually slower in real domains than in the theoretical estimation. Being more specific, we observe that only the symposium domain shows a similar behavior to that of the theoretical case. The rest of the domains deviate from the theoretical behavior sooner or later, showing a slower saturation. We can observe that this difference is specially acute in the Movielens, Book, Sushi, and Agh domain, which are the ones whose preferences have been rigorously elicited from real users (except for the party domain). This may reinforce our initial intuition, that, in real domains, the exponential growth on the number of Pareto optimal outcomes may not be as drastic as in the theoretical case. In other domains like the party and the holiday domain, the difference is less acute but still existent.

In fact, if one analyzes the proposed domains one by one, it is possible to realize that there are general preferential trends. This is clear in domains like Movielens or the Book domain, where we know that some movies and some books tend to be more popular than others. For instance, *The Shawshank Redemption* is one of the most popular movies of all times, and it has been able to obtain average ratings of 9.3 over 10 stars in sites like *IMDB*.⁹, where it has been voted

⁸ Again, the total number is $\min(1000, \binom{m}{n})$

⁹ <http://www.imdb.com>. Visited on 16th November 2015




Fig. 2. Average ratio of the final Pareto optimal set obtained with subgroups of different size

by more than 1 million users. Similarly, we can find books like *Harry Potter and the Deathly Hallows* that have received an average rating of 4.59 over 5 stars with more than 1 million ratings on sites like *GoodReads*¹⁰. Finding users that did not like these items has low odds, and as a consequence we can state that not all preference profiles are equally probable. Not only there are general trends in users preferences, but many times we find that there are clusters of users with similar preferences. For instance, in the book domain, we can expect that users that have liked *The Lord of the Rings* will also like other fantasy themed books like *Song of Ice and Fire*. This is the type of patterns exploited by recommender systems, and suggests that the number of likely preference profiles is even smaller.

With respect to the other small domains (e.g., AGH, Sushi), we analyzed the preferences of users. In fact, for analyzing the preferences of users on items we performed a Borda count with all of the preference profiles. We could observe that, in the Sushi domain, there are also some popular choices the *toro* (a total score of 39445) and some choices that are usually the least liked by users like the *kappa-maki* (a total score of 14928). In the case of the AGH domain, we could also observe that one of the courses (e.g., course 3) was the most preferred one with a score of 731, whereas the least preferred score had almost half the score.

¹⁰ <http://www.goodreads.com/>. Visited on 16th November 2015

This means that in these domains, preferences are not equally distributed and one should not expect such an exponential growth as in the theoretical case.

With respect to negotiation domains, we elicited real preferences from the Party domain, whereas we used the preference profiles provided by Genius in the Holiday and Symposium domain. Interestingly, we could observe that real users in the Party domain tend to consider the type of food, the type of drinks, and the music as the most important attributes. Even in some specific attributes, we could find that there were popular choices like for instance *Beer only* for drinks, and *Finger-food* and *Chips and Nuts* for food choices. With respect to the rest of negotiation domains, it has to be considered that they were not strictly and rigorously elicited like in the case of the party domain. Users were not contextualized in a specific scenario and their preferences were just elicited from their previous experiences in similar scenarios. In the case of the Holiday domain, we were able to observe some patterns like users considering the duration and the activities as the most important attributes. The users usually preferred longer durations to shorter durations, and we observed a slight positive inclination towards *Historical Places* and *Restaurants*. Even in the rest of less important attributes we were able to find some patterns like the fact that most users preferred *Miami* and *Amsterdam* as destinations. These patterns again show that not all preference profiles are equally likely, and that is reflected in the fact that the experimental growth depicted for Figure 2 is slower than the theoretical growth. On the other hand, the Symposium domain is closer to the theoretical expectation. This may be explainable due to the fact that the Symposium domain preferences were not elicited thinking on an specific symposium. In contrast to the Holiday domain, which did not follow a rigorous preference elicitation process either, in the Symposium domain it is harder to relate to the scenario, as it includes totally fictional speakers (e.g. *Mr. Talkolot*), whereas in the holiday domain one always can think about his/her own preferences on a trip. This may explain why the increase in the ratio of Pareto optimal outcomes is similar to the theoretical case where preference profiles are equally probable. It should be highlighted that in many negotiation domains, preferences are made different to test the performance of negotiation algorithms in conflicting scenarios.

The fact that, as we have shown, not all preference profiles are equally likely makes bottom-up approaches more applicable to real life scenarios than the results depicted in theory[20]. However, it should be noted that, even though the growth is slower, the graphics still suggest an increase with the size of the group and eventually the proof may not be applicable for domains involving a large number of agents. These results raise an interesting trade-off that should be analyzed in the future: the relation between the performance of bottom-up approaches, which increases with the subgroup size, and its applicability, which decreases with the group size, as nearly all outcomes may be Pareto optimal.

There is another additional issue to be studied concerning the applicability of bottom-up approaches. As the reader may have guessed, the aforementioned domains do not guarantee strict preferences. Therefore, some Pareto optimal outcomes calculated in subgroups may not be Pareto optimal in the whole group

Group size	Subgroup size						
	2	3	4	5	6	7	8
5	7%	4%	1%	-	-	-	-
7	5%	2%	1%	0.7%	0.3%	-	-
9	4%	2%	1%	0.7%	0.4%	0.2%	0.07%

Table 1. Average % of false positives calculated in a subgroup

(we call these false positives). In order to study this, we measured the ratio of false positives in the previous experimental setting. The results are summarized in Table 1. As it can be observed, the percentage of false positives remains low for every possible scenario, and it tends to decrease with the size of the subgroup. This matches our initial intuition, and shows that the proof presented in this paper practically holds in every situation. Hence, this result supports the applicability of bottom-up approaches in practice.

5 Related work

As far as these authors are concerned, most of the studies have dedicated their efforts on reaching Pareto optimal solutions using a horizontal approach that involves interactions with all group members. In [16], the authors propose a general framework for bilateral negotiations where agents are able to reach near Pareto optimal outcomes by decomposing the negotiation process into iso-utility curves, from where outcomes are proposed based on the similarity to the last offer proposed by the opponent. Later, the authors extend their findings to a multilateral and multi-issue environment where convergence is guaranteed [29]. Ehtamo *et al.* [11] propose a centralized mechanism for achieving Pareto optimal outcomes based on real valued linear additive utility functions and information sharing. Amador *et al.* [1] propose a task allocation method for agents with temporal constraints that is capable of providing envy free and Pareto optimal solutions under specific conditions. Other works like [21] have extended the concept of Pareto optimality to argumentation frameworks. The authors study different agent attitudes, how they relate to the problem of efficiency in abstract argumentation dialogues, and define several situations and scenarios that lead to Pareto optimal arguments. Recently, Hara *et al.* [10] proposed a mediated mechanism based on genetic algorithms that is capable of achieving near Pareto optimal outcomes for multi-party negotiations where agents preferences present non-linear relationships and change over time. However, none of these works employ bottom-up approaches, which may prove more useful in some scenarios.

Another field related to our study is that of multi-objective optimization. Pareto optimality is a well-known efficiency measure in multi-objective optimization [13, 17, 14]. Similarly to our multiagent decision setting, researchers in centralized multi-objective optimization have noticed the exponential increase on the number of Pareto optimal outcomes with the number of objective functions [7, 5]. Due to this unfortunate property of Pareto optimality, some researchers have offered practical alternatives to the selection of Pareto optimal outcomes. Di Pierro *et al.* define the concept of k optimality for deciding over Pareto opti-

mal outcomes. Basically, 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. We want to highlight the practical usability of k -optimality on future decision making mechanisms for agents and how it complements our current findings. First, based on our proof, a subgroups of agents may calculate Pareto optimal outcomes on subgroups and communicate them to the rest of subgroups. Then, a mechanism may be devised to allow agents to select a k -optimal outcome over calculated Pareto optimal outcomes.

Finally, economic and theoretical studies are also a source of related work. As introduced in the text, [20] 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 [20]. 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 [4, 9], but their focus of study has not been on the exploration of bottom-up approaches for reaching Pareto optimality.

6 Conclusion

In this paper, we have explored the applicability and performance of bottom-up approaches for reaching Pareto optimal outcomes in groups. Our analysis shows that Pareto optimal outcomes in a group remain optimal when increasing the number of agents in the group in many practical scenarios. This has implications for bottom-up approaches, as Pareto optimal outcomes may be calculated in subgroups first, and then be used in scenarios involving the whole group.

We performed experimental analysis on preferences elicited from users in real-life scenarios and validated that this principle can be applied to a wide range of domains. Our results on performance and applicability indicate that we are able to calculate a significant ratio of the final Pareto optimal frontier within subgroups. Conversely, we analyzed the applicability of our approach by considering how the ratio of Pareto outcomes increases with the size of the group. Our findings highlight that this increase is not as abrupt as expected in theoretical studies, as not all preference profiles are equally likely in many real-life domains. Still, the increase of the ratio of final Pareto optimal outcomes points to a necessary trade off in practice, which we plan to analyze in the future.

Additionally, as a future work, we plan to design novel negotiation approaches for intra-team negotiations that benefit from our findings. In particular, we plan to design a negotiation strategy for negotiation teams, which first calculate the Pareto optimal solutions within the team using our approach, and then target that set of Pareto optimal proposals when negotiating with the opponent.

References

1. S. Amador, S. Okamoto, and R. Zivan. Dynamic multi-agent task allocation with spatial and temporal constraints. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*, pages 1495–1496. International Foundation for Autonomous Agents and Multiagent Systems, 2014.
2. E. Argente, V. Botti, C. Carrascosa, A. Giret, V. Julian, and M. Rebollo. An abstract architecture for virtual organizations: The thomas approach. *Knowledge and Information Systems*, 29(2):379–403, 2011.
3. R. Aydođan, K. V. Hindriks, and C. M. Jonker. Multilateral mediated negotiation protocols with feedback. In *Novel Insights in Agent-based Complex Automated Negotiation*, pages 43–59. Springer, 2014.
4. A. Bogomolnaia and H. Moulin. Size versus fairness in the assignment problem. *Games and Economic Behavior*, 90:119–127, 2015.
5. D. W. Corne and J. D. Knowles. 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, pages 773–780, New York, NY, USA, 2007. ACM.
6. D. de Jonge and C. Sierra. \ hbox {NB} ^ {3}: a multilateral negotiation algorithm for large, non-linear agreement spaces with limited time. *Autonomous Agents and Multi-Agent Systems*, 29(5):896–942, 2015.
7. F. di Pierro. *Many-objective evolutionary algorithms and applications to water resources engineering*. PhD thesis, University of Exeter, 2006.
8. S. Esparcia, V. Sanchez-Anguix, and R. Aydođan. A negotiation approach for energy-aware room allocation systems. In *Highlights on Practical Applications of Agents and Multi-Agent Systems*, pages 280–291. Springer, 2013.
9. J. Garca-Segarra and M. Gines-Vilar. The impossibility of paretian monotonic solutions: A strengthening of roths result. *Operations Research Letters*, 43(5):476–478, 2015.
10. K. Hara and T. Ito. A mediation mechanism for automated negotiating agents whose utility changes over time. In *Twenty-Seventh AAAI Conference on Artificial Intelligence*, 2013.
11. P. Heiskanen, H. Ehtamo, and R. P. Hamalainen. Constraint proposal method for computing pareto solutions in multi-party negotiations. *European Journal of Operational Research*, 133(1):44–61, 2001.
12. C. Hewitt. Open information systems semantics for distributed artificial intelligence. *Artificial Intelligence*, 47(1-3):79–106, 1991.
13. J. Horn, N. Nafpliotis, and D. E. Goldberg. A niched pareto genetic algorithm for multiobjective optimization. In *Evolutionary Computation, 1994. IEEE World Congress on Computational Intelligence., Proceedings of the First IEEE Conference on*, pages 82–87. Ieee, 1994.
14. X.-B. Hu, M. Wang, and E. Di Paolo. Calculating complete and exact pareto front for multiobjective optimization: a new deterministic approach for discrete problems. *IEEE Transactions on Cybernetics*, 43(3):1088–1101, 2013.
15. T. Kamishima. Nantonac collaborative filtering: recommendation based on order responses. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 583–588. ACM, 2003.
16. G. Lai, C. Li, and K. Sycara. Efficient multi-attribute negotiation with incomplete information. *Group Decision and Negotiation*, 15(5):511–528, 2006.

17. H. Li and Q. Zhang. Multiobjective optimization problems with complicated pareto sets, moea/d and nsga-ii. *Evolutionary Computation, IEEE Transactions on*, 13(2):284–302, 2009.
18. R. Lin, S. Kraus, T. Baarslag, D. Tykhonov, K. Hindriks, and C. M. Jonker. Genius: An integrated environment for supporting the design of generic automated negotiators. *Computational Intelligence*, 30(1):48–70, 2014.
19. B. N. Miller, I. Albert, S. K. Lam, J. A. Konstan, and J. Riedl. MovieLens unplugged: experiences with an occasionally connected recommender system. In *Proceedings of the 8th international conference on Intelligent user interfaces*, pages 263–266. ACM, 2003.
20. B. O’Neill. The number of outcomes in the pareto-optimal set of discrete bargaining games. *Mathematics of Operations Research*, pages 571–578, 1981.
21. I. Rahwan and K. Larson. Pareto optimality in abstract argumentation. In *AAAI*, pages 150–155, 2008.
22. V. Sanchez-Anguix, R. Aydoğlan, V. Julian, and C. Jonker. Unanimously acceptable agreements for negotiation teams in unpredictable domains. *Electronic Commerce Research and Applications*, 13(4):243–265, 2014.
23. V. Sanchez-Anguix, T. Dai, Z. Semnani-Azad, K. Sycara, and V. Botti. Modeling power distance and individualism/collectivism in negotiation team dynamics. In *45 Hawaii International Conference on System Sciences (HICSS-45)*, pages 628–637, 2012.
24. V. Sanchez-Anguix, A. Espinosa, L. Hernandez, and A. Garcia-Fornes. Mamsy: A management tool for multi-agent systems. In *7th International conference on practical applications of agents and multi-agent systems (PAAMS 2009)*, pages 130–139. Springer, 2009.
25. V. Sanchez-Anguix, V. Julian, V. Botti, and A. Garcia-Fornes. 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, 2012.
26. V. Sanchez-Anguix, V. Julian, V. Botti, and A. Garcia-Fornes. Studying the Impact of Negotiation Environments on Negotiation Teams’ Performance. *Information Sciences*, 219:17–40, 2013.
27. V. Sanchez-Anguix, V. Julian, V. Botti, and A. Garcia-Fornes. Tasks for agent-based negotiation teams: Analysis, review, and challenges. *Engineering Applications of Artificial Intelligence*, 26(10):2480–2494, 2013.
28. P. Skowron, P. Faliszewski, and A. Slinko. Achieving fully proportional representation is easy in practice. In *Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems*, pages 399–406. International Foundation for Autonomous Agents and Multiagent Systems, 2013.
29. C. N. D. T. Zhenh, Ronghuo and K. Sycara. Automated multilateral negotiation on multiple issues with private information. *INFORMS Journal on Computing*, page In Press, 2015.
30. C.-N. Ziegler, S. M. McNee, J. A. Konstan, and G. Lausen. Improving recommendation lists through topic diversification. In *Proceedings of the 14th international conference on World Wide Web*, pages 22–32. ACM, 2005.