A Revelation Mechanism for Shared Conditional Preferences in Multi-Attribute Negotiation

Buffett, Scott

A Revelation Mechanism for Shared Conditional Preferences in Multi-Attribute Negotiation

Scott Buffett
National Research Council Canada
46 Dineen Drive
Fredericton, New Brunswick, Canada
E3B 9W4
scott.buffett@nrc.gc.ca

ABSTRACT

Agents who negotiate over a space of multi-attribute agreements where conditional preferences may be present can encounter difficulties in converging toward Pareto-efficient outcomes. This is because of the fact that, while both agents may have strategic incentives for keeping their own preferences private, there may be a number of attributes for which, under certain conditions, the two agents have the same preference. If they could work together to discover such instances, and agree to eliminate a portion of the space of agreements that both dislike, it would greatly increase the probability and speed of reaching a mutually favourable deal. We present a negotiation mechanism for agents to eliminate portions of the agreement space that are mutually non-beneficial. The mechanism enables the semi-truthful revelation of conditional preferences in such a way that encourages agents to make progress towards finding non-Pareto efficient outcomes. We demonstrate the protocol for such negotiations and outline the set of strategies that (1) agents have incentive to follow and (2) will result in mutually favourable elimination. We also empirically measure the effectiveness of such agreement space reduction in terms of utility achieved.

Categories and Subject Descriptors

I.2.11 [**Distributed Artificial Intelligence**]: Intelligent agents

General Terms

Economics,Theory

Keywords

Negotiation mechanisms, Pareto efficiency, preferences, game theory

1. INTRODUCTION

While single-attribute negotiation introduces a host of difficult problems solve on its own, including deciding how much to offer, whether to accept an offer, or whether to walk away from negotiations altogether, to name a few, multi-attribute negotiation can introduce new problems that are even more complex. Negotiating over n attributes or issues is not equivalent to n single-attribute negotiations, for two reasons. First, a negotiator's decision making for one attribute might depend on the utility achieved from the negotiated result of other attributes. For example, a negotiator might be more likely to accept an unfavourable agreement for one attribute if it was able to reach a desirable agreement on another. Second, one's preference for various outcomes for one attribute might depend on the outcomes themselves (as opposed to the value of the outcomes) negotiated for other attributes. For example, for two attributes A and B, with possible outcomes $\{a_1, a_2\}$ and $\{b_1, b_2\}$, respectively, one might prefer b_1 over b_2 if a_1 is the negotiated outcome for A, and prefer b_2 otherwise, resulting in so-called "conditional" preferences. Thus individual attributes typically cannot be negotiated independently in such cases.

A consequence of the phenomena described above is that non-Pareto efficient outcomes, which are outcomes for which there exist other outcomes that are more desirable for some parties and at least as desirable for all, are introduced. As the fundamental objective of any rational party in a negotiation is utility maximization, whether it is one's own utility, utility of a coalition, entire group, etc., it is always mutually beneficial to avoid non-Pareto efficient agreements. In the typical negotiation scenario where one participant knows very little about the other's preferences, it can be very difficult to ensure that these outcomes are avoided. In the case where conditional preferences are present, the negotiating parties may sometimes even be drawn to such outcomes. Since there is little information available on the other's preferences, previous research has suggested analyzing a negotiator's previous offers in the current negotiation as a way of learning the types of things that it finds desirable. However, if conditional preferences are at work, using this information could be counterproductive. Consider the conditional preference mentioned above. One negotiator may propose an offer including attribute values a_1 and b_1. The other party, in an effort to reach an agreement within a reasonable time, concedes on a number of issues in order to present an offer that the opponent is likely to accept. One such issue might be the value b_1 for attribute B, thinking the opponent prefers this value, with no way of knowing that this preference was conditional on a_1 being present. If a_1 is not included, then the negotiator essentially makes a concession that is preferred less by its opponent, resulting in a non-Pareto efficient outcome being proposed.

We consider the setting of bilateral, alternating-offers agent negotiation over multi-attribute outcomes, and offer a mechanism for agents to work together to *reduce* the search space at the same time by eliminating non-Pareto efficient outcomes. This is done via a negotiation mechanism that enables each agent to learn under which conditions they may share the same preference for a given pair of attribute values. This mechanism takes the form of a secondary bilateral alternating-offers negotiation where agents declare conditional preferences and propose eliminating offers from the space of allowable agreements that are deemed less preferable by virtue of the conditional preferences. The protocol for these negotiations offers a strict set of rules that promote progress toward finding mutually agreeable eliminations, even when agents prefer to act dishonestly to mislead or keep preference information from the opponent. We discuss the set of equilibrium strategies and prove that a rational agent will participate in the prescribed manner and the objectives will be achieved when certain conditions are met.

This work advances the state of the art in two ways. First, it is the first work that enables agents to learn common conditional preferences during a negotiation session without revealing their entire preference models to each other or a third party mediator. Second, to the best of our knowledge this is first negotiation mechanism that strives to find agreements on portions of the offer space that can be *eliminated*, helping the agents find better deals, more quickly.

2. BACKGROUND

2.1 Multi-Attribute Negotiation

Much of the work done in multi-attribute negotiation (often alternatively referred to as multi-issue negotiation) revolves around determining agendas for a number of simpler single-issue negotiations [7, 8]. Dependencies among issues are determined and an ordering in which to independently tackle issues that minimizes the impact of these dependencies is declared. There is also a great deal of work in the areas of multi-issue negotiation when preferences are incomplete [9, 11] and when deadlines are imposed [10]. Boutilier *et al.* [1, 2] discuss the associated problems when preferences in the multi-attribute domain are conditional, and explore a representation referred to as a conditional preference network for structuring user preferences. Ito *et al.* [12] mitigate the problems introduced by conditional preferences by proposing a mechanism where each agent finds and declares regions of the outcome space that hold high utility. A moderator then chooses the final agreement deemed to have the highest social welfare. Robu *et al.* [13] take a learning-based approach to the problem, and attempt to model the opponent's preferences using utility graphs. These graphs are then used by a negotiator to find and propose offers that are likely to be Pareto efficient. As conditional preferences are of concern in this paper, we borrow the representation of Boutilier *et al.* and adapt it to our model, which we formalize later in the paper.

2.2 Predicting Opponent Preferences

When attempting to negotiate over the set of attributes as a whole, it can be difficult to determine good strategic offers to make that will not only represent a preferable outcome for the offerer, but that will also be seen as an improvement to the receiver and thus make progress toward the ultimate goal of reaching an agreement. Faratin *et al.* [6] propose a method for making tradeoffs in one's previous offer to generate new offers that have roughly equal value to the proposer, by choosing new attribute values, some of which are more preferred, some less, in such a way that the changes balance and approximately the same utility is achieved. To maximize the likelihood of acceptance on the part of the receiver, tradeoffs are performed in such a way as to generate an offer that is highly similar to previous offers that were received from the opponent. The theory is that if an offer is similar to outcomes that the opponent has itself proposed, then the opponent is more likely to find the offer acceptable. Coehoorn and Jennings [5] propose non-parametric methods for approximating the opponent's importance for each issue by considering data available on the opponent's previous negotiations. Buffett *et al.* [4] propose a similar strategy for negotiating the exchange of private information. In this scenario, where an entity representing a website requires some personal information from a visitor who is unwilling to freely give it up, it is nearly impossible to know what the opponent in the negotiation values most, since different websites likely have unique reasons for asking for certain information, and users likely have even more unique reasons for protecting certain information. In this work, offers are chosen by looking at one's top n offers that are deemed allowable by the protocol, and choosing the one deemed to be the most similar to the opponent's previous offers. A protocol is proposed to place a strict set of rules that guarantees convergence under this strategy to a mutually acceptable exchange, if one exists. Buffett *et al.* [3] propose a method for learning opponent preferences in a similar negotiation model by predicting which opponent offers represent a concession, as opposed to an attempt at a trade-off, by analyzing how it reacts to each offer received in the negotiation.

3. LEARNING CONDITIONAL PREFERENCES

While attempting to formulate offers that are similar to the opponent's previous offers has been shown to be effective in many situations, this strategy becomes less effective and even counterproductive when conditional preferences are present. This is because an attribute value may be introduced into the negotiation by an agent only because the conditions under which the value is preferred are met. The opposing agent will then typically not recognize the condition and, trying to achieve maximum similarity, offer the attribute value back possibly without the condition being met. If both agents attempt to maximize similarity, negotiations could get "stuck" on attribute values that are not highly valued by either party, either wasting time before better offers enter the conversation, or worse, resulting in sub-optimal agreements.

Non-Pareto efficient outcomes could also result from mutual non-conditional preferences. For example, for an attribute A, a_1 could be preferred over a_2 for both parties, meaning that they should never agree on anything with a_2. However, since neither party has high value for a_2, it is less likely that a_2 would ever be introduced into negotiations. Thus we narrow our focus to the identification of common conditional preferences in the pursuit of eliminating non-Pareto efficient outcomes.

Previous work primarily solves the problem of eliminating such outcomes by forcing agents to participate in some sort of revelation process, where agents declare some or all preferences/utilities either to the opposing agent, or to a third party who keeps the preferences private and only declares which outcomes are necessarily non-optimal. The problem with this approach is there is insufficient incentive to be fully honest, since giving away one's preference information will give away any opportunity to exploit whatever information one may have had about the other's preferences. So it essentially levels the playing field, and would thus hurt whichever agent may have otherwise had an advantage. Even without any knowledge of the other's preferences, one might even be motivated to be dishonest merely to mislead the other party, somehow giving itself a particular advantage. Clearly, even though both parties may agree that the elimination of non-Pareto outcomes is mutually beneficially, often neither can be trusted to do their part to achieve it.

4. PROBLEM FORMALIZATION

4.1 Preference Model

Let X be a set of n attributes and let each attribute X^i consist of a set of possible values for that attribute. Then the set O of outcomes for the negotiation is a subset of $X^1 \times \ldots \times X^n$. Each agent j has a utility function $u_j : O \rightarrow \Re$ over the space of outcomes, which is consistent with its preference relation $\succ_j^i$ over the set of values for attribute i. Note that we write the preference operator simply as $\succ$ whenever the agent and attribute are obvious or unimportant. As in [1], we adopt the ceteris paribus (all else equal) principle for preferences on the attribute level, which means that $x_p^i \succ_j^i x_q^i$ implies that, for all pairs of outcomes (o_p, o_q) that have all attribute values in common except for x_p^i (for o_p) and x_q^i (for o_q) for attribute i, $u_j(o_p) > u_j(o_q)$. We also allow conditional preferences $c : x_p^i \succ x_q^i$ in our model, which indicate that the preference is true whenever the attribute values specified in c are present in the outcomes.

4.2 Negotiation Model

We consider two agent bilateral alternating-offers automated negotiation, where each proposal in the negotiation consists of a full agreement (i.e. a legal value for each attribute). After an initial offer is presented by an arbitrarily chosen participant, at each step the receiver of the previous offer has three choices, accept the previous offer, reject the previous offer and submit a counteroffer, or reject and quit. We allow any protocols that work within this model, with the added condition that any proposals that are deemed eliminated via outcome reduction agreements must then be disallowed during the primary negotiation.

4.3 Outcome Reduction Problem

We seek to reduce the space of allowable offers (and thus agreements) by providing a mechanism for agent revelation of mutual conditional preferences. A condition c in our model is formally represented as a set of attribute values, where each attribute in A is represented by at most one value. So a conditional preference $c : x_p^i \succ x_q^i$ means that for any outcome pair (o_p, o_q) that has all attribute values in common except for x_p^i (for o_p) and x_q^i (for o_q) for attribute i, and each outcome includes all atribute values specified in c, then $u_j(o_p) > u_j(o_q)$.

For each pair of values (where a preference exists), the set of conditions imposed on the preference is assumed to be exhaustive and non-conflicting. Let $C = C_p \cup C_q$ be the set of conditions that influence the preference over x_p^i and x_q^i, where C_p is the set of conditions where x_p^i is preferred and C_q is the set of conditions where x_q^i is preferred. Then, if C is non-empty, any pair of outcomes o_p and o_q that differ only in attribute i (where one has x_p^i and one has x_q^i) will satisfy a condition in exactly one of C_p and C_q. Note that C may be empty, indicating that the preference holds universally. The entire set C of conditions, with C_p and C_q indicated, is said to be a *partition* over the conditions for the preference over x_p^i and x_q^i.

EXAMPLE 1. *Let D, E, F and G be attributes, each with two values $\{d_1, d_2\}$, $\{e_1, e_2\}$, etc. If an agent was said to prefer f_1 over f_2 whenever either d_1 or e_1 are present, and f_2 over f_1 otherwise, then the conditions could be partitioned as (eliminating set brackets for individual conditions) $C_1 = \{d_1e_1, d_1e_2, d_2e_1\}$ and $C_2 = \{d_2e_2\}$.*

It is important to note that, as a consequence of C_p and C_q being exhaustive and non-conflicting, if $x_p^i \succ x_q^i$ holds under condition c, then the preference must hold under any superset of c. Let c_1 and c_2 be the conditions under which agent 1 and 2 respectively prefer $x_p^i \succ x_q^i$. Then by this axiom, if $c_1 \cup c_2$ does not violate the stipulation that each attribute be represented more than once in a condition, then agents 1 and 2 necessarily share the conditional preference $c_1 \cup c_2 : x_p^i \succ x_q^i$. This implies that any outcome containing all attribute values in $c_1 \cup c_2$ as well as x_q^i is considered inferior to the same outcome with x_q^i replaced by x_p^i by both agents, and are thus non-Pareto efficient. It is the goal then to encourage agents to seek and reveal these common conditional preferences, and to unanimously disallow the corresponding non-Pareto efficient outcomes from being proposed in the primary negotiation.

EXAMPLE 2. *Let agent R be the agent with conditional preferences outlined in Example 1. If agent I prefers f_1 over f_2 whenever g_1 is present, and f_2 over f_1 otherwise, then both agents must prefer f_1 when any conditions in $C_1' = \{d_1e_1g_1, d_1e_2g_1, d_2e_1g_1\}$ are satisfied, and f_2 when any conditions in $C_2 = \{d_2e_2g_2\}$ are satisfied.*

5. OUTCOME REDUCTION MECHANISM

5.1 Usage

Negotiating outcome reductions can take place at any time before or during the primary negotiation, however it is expected that it would typically be initiated during the negotiation, by the agent who is scheduled give the next offer. While negotiations can take place over the preference of any attribute values, one would expect that it would be most effective when the two agents share preference in a large number of situations. Predicting when this might be the case is, however, outside the scope of this paper, but is something we are currently working towards.

5.2 Protocol

Negotiations for outcome space reduction utilize a bilateral alternating offers protocol where, after an initial declaration, at each step each agent has the opportunity to propose a set of potential agreements to eliminate from the set of allowable offers, or to accept a previously proposed set. Negotiations continue until either an agreement is reached, or one agent does not respond within a previously agreed upon time limit, at which time the previous offer is considered mutually acceptable and enforced. Outcome reduction negotiations can be initiated by either agent at any time, either before the primary negotiations take place, or during a pause at some point within the primary negotiation, perhaps when one agent has a suspicion that a common conditional preference may cause sub-optimal outcomes, and thus should be resolved. In order to promote sufficient honesty and enforce convergence toward a mutually agreeable set of non-Pareto efficient agreements, communication is bound to the following protocol. It is perfectly legal for an agent to behave dishonestly, but it will be shown in the following section that there is sufficient incentive to act sufficiently honest to ensure favourable results.

Let there be two negotiating parties, (1) the agent who intitiates the outcome reduction negotiation, referred to as the initiator or agent I, and (2) the agent who responds, referred to as the responder or agent R. In the first step, the initiating agent declares a conditional preference of the form $a_i : b_j \succ b_k$, indicating that it prefers b_j over b_k for attribute B when a_i is present as the value for attribute A. This then consitutes a proposal that all agreements including both a_i and b_k be disallowed, on the grounds that, for any such agreement, replacing b_k with b_j might be preferred by both parties.

From this point, the following steps are repeated until an agreement is reached:

1. The responding agent can agree to the previous proposal offered by the initiating agent, thus ending the negotiation, or if it does not agree, can indicate a partition C_j and C_k over the set C of conditions under which b_j is preferred over b_k and b_k is preferred over b_j, respectively. As mentioned previously, it is assumed that this set of conditions will be exhaustive and nonconflicting. If the responding agent does not answer within the pre-specified time limit, the initiator's previous proposal will be considered mutually acceptable and enforced.

2. If a counter proposal is sent, the initiating agent has two options: (1) to enforce its previous proposal if one exists, where the initial declaration does not count as a proposal in this way, and thus this is not a valid option in the first iteration of this loop, or (2) to propose for elimination any subset of the conditions proposed in the previous step by the responder, by offering a special case C'_j of C_j and/or a special case C'_k of C'_k of the responder's conditions under which the initiator also prefers b_j over b_k, and/or b_j over b_k, respectively. Formally, a special case of a set of conditions C is a set C' such that for each $c' \in C$, there exists a $c \in C$ such that $c \subseteq c'$. So the conditions in C' are satisfied in a subset of the situations that satisfy C.

   For example, let the responder's partition be $C_j = \{a_1b_1, a_1b_2, a_2b_1\}$ for $b_j \succ b_k$ and $C_k = \{\{a_2b_2\}$ for $b_k \succ b_j$. If the initiator has the conditional preference that $d_1 : b_j \succ b_k$, then it could declare the special case $C'_j = \{a_1b_1d_1, a_1b_2, d_1\}$ as a special case of C_j, indicating some conditions under which both agents would prefer b_j.

   If each agent acts honestly and gives complete sets of conditions, then any agreements with b_k that meet conditions in C'_j and any agreements with b_j that meet conditions C'_k should be non-Pareto efficient. Thus the initiator is, in effect, proposing to eliminate these non-Pareto efficient outcomes.

   If the initiating agent does not answer within the pre-specified time limit, execution moves to the next iteration without a proposal, and the responder may offer a new proposal, re-offer the previous one, or close the negotiations without an agreement.

Thus in each iteration, the responder makes any proposal it wishes (regarding the attribute values specified in the declaration), with the knowledge that the initiator has the right to accept any part of that proposal in the future. However, because the initiator must wait until its next turn in the conversation to enforce the current proposal, meaning the responder has chance to make a new offer before enforcement is allowed, in the case that the initiator's choice is unfavourable, the responder also has the opportunity entice the initiator away from that choice by offering something that will be better (presumably for both parties) before choice is enforced. We will show in the next section that, when certain circumstances are achieved, the responder will have the incentive to constantly "sweeten the deal" in this way until the desired non-Pareto eliminations are acheived

5.3 Strategy

In this section we discuss equilibrium strategies by demonstrating optimal activity based the assumption that the opponent acts optimally. Initially, we discuss how each agent should behave at each step, and prove that each agent should behave in this manner below. The ultimate result is that, if an initiating agent can achieve its subgoals throughout the progress, the non-Pareto efficient outcomes that result from the initiator's initial declaration will be identified and eliminated, and that, while some Pareto efficient outcomes may also be eliminated, none of these eliminations will cause a non-Pareto efficient outcome to become Pareto efficient, thus leaving the Pareto frontier intact.

- *Initial declaration*: Agent I reveals a set C_d of conditions under which a preference over two attribute values x^i_p and x^i_q for attribute i holds, with the intention and priority of finding under which conditions I and R agree on the preference for x^i_p and x^i_q. We prove that (1) I should be truthful and (2) I should not reveal the full partition over the conditions under which x^i_p and x^i_q are preferred.

- *R's first move in the loop*: Agent R partitions the set of conditions. R could have incentive to exercise as much dishonesty as it wants. Even if R completely agrees with I, it is understood that it could indicate the opposite, hoping nothing (or relatively little) gets settled, thus capitalizing on the opportunity to learn about I's preferences without giving away its own. Thus R is free to act in any strategic manner it deems best in this step, and we make no claims regarding optimal actions.

- *I's first response to R*: I should declare a special case of the conditions in the partition expressed by R in the previous step, such that (1) all conditions where the two sides agree on the preferences are included, and (2) some conditions where I disagrees with R are included, which implies that I should lie and claim to prefer a value under a particular condition (in agreement with R) when it truly prefers the other under that condition. The purpose of these latter declarations is to identify dishonesties on the part of agent R, and to "call its bluff", so to speak. The intended effect of this strategy is demonstrated by R's activity in the next step.

- *R's response to I's special case proposals*: If I is successful in identifying a case where R had been dishonest, then it must be the case where there is a conditional preference $c' : x_p^i \succ x_q^i$ upon which the two agents agreed, and thus all outcomes with both c' and x_q^i are proposed to be eliminated, but R was dishonest and actually prefers $c' : x_q^i \succ x_p^i$. Therefore, eliminating all outcomes with both c' and x_q^i in favor of keeping those with both c' and x_p^i is detrimental to R. R must then attempt to prevent I from enforcing this proposal by adjusting its partition so that (1) dishonesties that were identified are corrected, and that (2) I will be able to choose a new special case of the conditions that will be more favourable to itself than the current one, thus being less apt to enforce the current one. Since R knows which subset C'_p of the conditions C_p in the initial declaration are mutual (since R knows that it was optimal for I to have acted honestly in this declaration), if there is any condition that satisfies C'_p (and thus where both agents prefer x_p^i) that R has misreported in the previous as a condition under which he prefers x_q^i, R should adjust its partition so that at least one such case is rectified. I will be able to confirm agreement with this condition by adding it to its special case in the next step, and thus will choose not to enforce the proposal in the previous step.

  If I does not hit on any of R's dishonesties (perhaps if R had been entirely truthful, or if I was unlucky), then R will entirely agree with I's proposal and may choose to accept it. In this case, some non-Pareto efficient outcomes will be eliminated (i.e. those that I truly agreed with), and unfortunately some Pareto efficient outcomes would be eliminated (i.e. those that that I lied about and actually preferred). In this case, the two agents disagree, and one of them would have needed to concede anyway. This way, agent I essentially declares under which conditions (likely those that are very favourable to it) under which it will concede. So it is not disasterous, and is possibly even beneficial for I.

- *Remainder of the negotiation*: I should repeatedly try to identify dishonesties, as they will continue to improve the offers it gets from R. Once offers stop improving for I then R's offer will include all common conditions from C_d in the initial declaration. (as we will show), and will not include eliminating anything that both agents prefer (also to be shown).

First we show that the mechanism is not counterproductive, and thus that rational agents hoping to achieve better results have incentive to participate (individual rationality).

Assumption 1: For two attribute values x_p^i and x_q^i and condition c, whenever an agent prefers $c : x_p^i \succ x_q^i$, it will always act in favor of eliminating an outcome with c and x_q^i in favour of keeping one with c and x_p^i, all else equal.

THEOREM 1. *The elimination of a Pareto efficient outcome will never result in a non-Pareto efficient outcome becoming Pareto efficient in the new set of allowable outcomes.*

Proof: Suppose not. Then there is a conditional preference $c : x_p^i \succ x_q^i$ that the two agents agreed upon, even though they both truly preferred $c : x_q^i \succ x_p^i$. Thus Pareto efficient outcomes (i.e. those with c and x_q^i) were eliminated, making non-Pareto efficient outcomes (i.e. those with c and x_p^i) Pareto efficient. However, whenever I proposes eliminating such an outcome to R in favour of one that R prefers, then by Assumption 1, R will act to attempt to avoid this by rectifying his partition indicating that $x_q^i \succ x_p^i$ holds under condition c (perhaps also trying to entice I to offer something new). Since I also prefers $x_q^i \succ x_p^i$ under condition c, I will have the incentive adjust its proposal to not include the elimination of outcomes with c and x_q^i. □

Next we show that agent R will act in such a way as to ultimately reveal the conditional preferences it shares with I's initial declaration.

THEOREM 2. *Let C_d be the conditions declared in the initial step under which agent I prefers x_p^i over x_q^i, and let C'_d by the special case of C_d indicating the conditions under which both agents prefer x_p^i over x_q^i. If agent I can identify dishonesties in R's partition at each step, the mechanism is guaranteed to reveal C'_d, thus enabling the elimination of non-Pareto efficient offers that include x_q^i when conditions from C'_d are satisfied.*

Proof: Since Agent I reveals honestly in the first step, C'_d is known to agent R. If I's last offer includes $x_p^i \succ x_q^i$ for all conditions in C'_d, then I can enforce this at the next step and the goal is achieved. Otherwise there exists a condition c' that satisfies C'_d that is not in I's last offer. So the conditional preference $c' : x_p^i \succ x_q^i$ holds for both agents, but I does not know it. Suppose I identifies a dishonest preference $c'' : x_p^i \succ x_q^i$ as proposed by R in its last step, and proposes removing all outcomes with x_q^i such that c'' is satisfied. Then by Assumption 1, R will have incentive to convince I to not eliminate these outcomes. It can do so only by proposing a new partition where (1) this false conditional preference no longer holds and (2) agent I will have incentive to propose a special case instead of enforcing the previous. The only way to ensure (2) is to adjust the partition so that $c' : x_p^i \succ x_q^i$ now holds. □.

Finally we show that I will act truthfully in the initial declaration.

THEOREM 3. *If agent I's key objective is to find under which conditions I and R agree on the preference for x_p^i and x_q^i, then I will truthfully reveal conditional preferences in the initial declaration. Moreover, I will not reveal the full partition of conditions for preferences over these two attribute values.*

Proof: Due to Theorem 2, agent I must be honest in order for agent R to know how to improve its partition towards

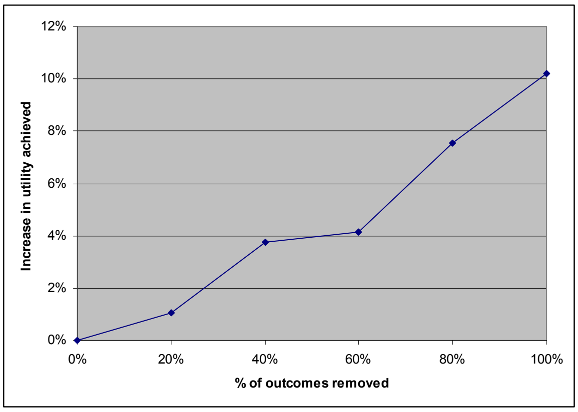

Figure 1: Increase in utility achieved by percentage of non-Pareto efficient outcomes eliminated

achieving the priority of finding where the two agents agree on the preference for x_p^i and x_q^i. If agent I revealed its full partition, it would hamper its leverage in calling agent R's bluff and encouraging it to offer a more desirable partition. □

Thus if I's top priority is finding under which conditions I and R agree on the preference for x_p^i and x_q^i, and can achieve its objective by acting in the prescribed manner provided R acts in the prescribed manner, and if Assumption 1 holds then R must act in the prescribed manner in response to agent I's actions, then the prescribed strategies constitute a Nash equilibrium, on the assumption that I can sufficiently identify dishonesties in R's responses, if and when R acts dishonestly.

6. EMPIRICAL MEASURES OF THE EFFECT OF OUTCOME REDUCTION

We complete the discussion with a demonstration of the effect that eliminating non-Pareto efficient outcomes can have on the overall utility achieved in the final result of the negotiation. We tested a simple scenario with 10 attributes, where each attribute had 2 values. Preferences and conditions were chosen randomly and independently for each agent. Agents used a simple similarity-maximization strategy, similar to those discussed above. Figure 1 shows the resulting average utility increase realized over a number of runs, for various degrees of non-Pareto efficient outcome elimination. One can see, in this particular scenario, that although the line is not is not smooth, utility achieved appears to increase linearly in the percentage of outcomes eliminated, up to a maximum of approximately 10% increase when all such outcomes are eliminated. So it does not take a large number of eliminations to have an effect; each step toward decreasing the space of non-Pareto efficient outcomes is an equal step toward improving the utility achieved by all agents.

7. CONCLUSIONS AND FUTURE WORK

Agents who negotiate over a space of multi-attribute agreements where conditional preferences may be present can encounter difficulties in converging toward Pareto-efficient outcomes. This is because of the fact that, while both agents may have strategical incentives for keeping their own preferences private, there may be a number of attributes for which, under certain conditions, the two agents have the same preference. Agents that work together to discover such instances, and agree to eliminate a portion of the space agreements that both dislike, can greatly increase the probability and speed of reaching a mutually favourable deal.

We consider the setting of bilateral, alternating-offers agent negotiation over multi-attribute outcomes, and offer a mechanism for agents to work together *reduce* the search space at the same time by eliminating non-Pareto efficient outcomes. This is done via a negotiation mechanism that enables each agent to learn under which conditions they may share the same preference for a given pair of attribute values. This mechanism takes the form of a secondary bilateral alternating-offers negotiation where agents declare conditional preferences and propose eliminating offers from the space of allowable agreements that are deemed less preferable by virtue of the conditional preferences. The protocol for these negotiations offers a strict set of rules that promote progress toward finding mutually agreeable eliminations, even when agents prefer to act dishonestly to mislead or keep preference information form the opponent. We discussed the set of equilibrium strategies for the two agents and showed that (1) the mechanism is individually rational, meaning that the two agents benefit by participating, (2) if the strategies are followed, the initiating agent will discover under which conditions the two agents agree on the preference over two attribute values, when certain objectives are met in the process, and (3) that these prescribed strategies constitute a Nash equilibrium under certain reasonable conditions.

The key goal for future work is to devise methods for agents to recognize when a conditional preference may be at work causing inefficiencies in the negotiation, and thus this secondary negotiation can be used accordingly to identify the non-Pareto efficient outcomes that may be hampering the progress.

8. REFERENCES

[1] C. Boutilier, R. I. Brafman, C. Domshlak, H. H. Hoos, and D. Poole. CP-nets: A tool for representing and reasoning with conditional ceteris paribus preference statements. *Journal of Artificial Intelligence Research*, 21:135–191, 2004.

[2] C. Boutilier, R. I. Brafman, C. Geib, and D. Poole. A constraint-based approach to preference elicitation and decision making. *AAAI Spring Symposium on Qualitative Preferences in Deliberation and Practical Reasoning*, pages 19–28, 1997.

[3] S. Buffett, L. Comeau, B. Spencer, and M. W. Fleming. Detecting opponent concessions in multi-issue automated negotiation. In *Proc. of ICEC2006*, pages 11–18, Fredericton, Canada, 2006.

[4] S. Buffett, K. Jia, S. Liu, B. Spencer, and F. Wang. Negotiating exchanges of P3P-labeled information for compensation. *Computational Intelligence*, 20(4):663–677, 2004.

[5] R. M. Coehoorn and N. R. Jennings. Learning an opponent's preferences to make effective multi-issue negotiation tradeoffs. In *Proc. of the 6th International*

Conference on Electronic Commerce (ICEC2004), pages 113–120, Delft, The Netherlands, 2004.

[6] P. Faratin, C. Sierra, and N. Jennings. Using similarity criteria to make issue trade-offs in automated negotiations. *Artificial Intelligence*, 142:205–237, 2002.

[7] S. Fatima, M. Wooldridge, and N. Jennings. Optimal agendas for multi-issue negotiation. In *Proceedings of the second international joint conference on Autonomous agents and multiagent systems*, pages 129–136, Melbourne, Australia, 2003.

[8] S. Fatima, M. Wooldridge, and N. Jennings. An agenda-based framework for multi-issue negotiation. *Artificial Intelligence*, 152(1):1–45, 2004.

[9] S. Fatima, M. Wooldridge, and N. Jennings. Automated multi-attribute negotiation with efficient use of incomplete preference information. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2004)*, pages 1054–1061, New York, NY, 2004.

[10] S. Fatima, M. Wooldridge, and N. Jennings. Multi-issue negotiation with deadlines. *Journal of Artificial Intelligence Research*, 27(1):381–417, 2006.

[11] S. Fatima, M. Wooldridge, and N. R. Jennings. Optimal negotiation of multiple issues in incomplete information settings. In *Proc. 3rd Int. Conf. on Autonomous Agents and Multi-Agent Systems*, pages 1080–1087, New York, NY, 2004.

[12] T. Ito, H. Hattori, and M. Klein. Multi-issue negotiation protocol for agents: exploring nonlinear utility spaces. In *Proceedings of the 20th international joint conference on Artifical intelligence (IJCAI07)*, pages 1347–1352, Hyderabad, India, 2007.

[13] V. Robu, D. J. A. Somefun, and J. A. L. Poutré. Modeling complex multi-issue negotiations using utility graphs. In *Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems (AAMAS'05)*, pages 280 – 287, Utrecht, The Netherlands, 2005.