Bargaining in-Bundle over Multiple Issues in Finite-Horizon Alternating-Offers Protocol

Francesco Di Giunta and Nicola Gatti [1]

Abstract. This paper provides an algorithm to compute the subgame perfect equilibrium strategies in perfect information finite-horizon alternating-offers bargaining in-bundle over multiple issues. We show that the agreement is achieved immediately and is Pareto efficient. We make a novel use of backward induction for multiple issues and we prove that, for linear multi-attribute utility functions, the problem of computing the equilibrium is tractable and the related complexity is polynomial with the number of issues and linear with the deadline of bargaining.

1 Introduction

Negotiation is the process whereby individuals try to solve disputes and reach mutually beneficial agreements communicating and compromising [12]. The interest of artificial intelligence research in negotiation lies in the possible exploitation of negotiation models and techniques to solve coordination and cooperation problems among rational agents [8]. Negotiation techniques are thus employed to address a number of classic problems, such as data allocation in information servers, resource allocation, and task distribution.

The theory of negotiation is also employed in artificial intelligence to address the challenge of automatizing some typically human negotiations, such as commercial negotiations. In this domain intelligent self-interested software agents negotiate with other intelligent agents on behalf of users for buying and selling services and goods. As underlined by Sandholm in [15], this automation, apart from saving labor time of human negotiators, can lead to more effective negotiations because software agents can enumerate and evaluate potential outcomes faster than humans and are more prone than humans to follow game-theoretic prescriptions.

Among the negotiation settings for commercial transactions, a very common one is *bargaining* [10, 11]: a buyer and a seller try to agree on the choice of the values of some parameters of the transaction they are carrying out together; if the parameter is only one (typically the price of the good to be sold) we have a *one-issue* bargaining; if there is more than one parameter (e.g., the price of more than one good or the price and the quality level of a service) we have a *multi-issue* bargaining.

The formalized study of negotiation (and therefore of bargaining) is commonly carried out with game-theoretic tools. In this approach one distinguishes the negotiation *protocol* and the negotiation *strategies*: a protocol is a set of rules that defines the possible ways the negotiation process can be led, specifying which actions are allowed and when [13]; a strategy is a set of actions, allowed by the protocol, that defines an agent's possible specific behavior in the negotiation. Given a protocol, the game-theoretic approach prescribes that rational agents would or should employ strategies which are somehow in *equilibrium*. The exact notion of equilibrium depends on the particular settings of the problem and is not always clear, but commonly *Nash equilibrium* and its refinements are used (see for example [3]).

The best known and perhaps the most elegant protocol for bilateral bargaining is the *alternating-offers* protocol, which comes in many variations. Basically, a player starts by offering a value for the parameter of the bargaining (say, a price) to her opponent. The opponent can accept the offer or exit the negotiation or make a counteroffer. If a counteroffer is made, the process is repeated until one of the players accepts or exits the negotiation. The study of bargaining over a single issue in an alternating fashion has been pioneered by Ståhl [16]. Ståhl analyzes bargaining games with finitely many possible agreements and a finite time horizon (i.e., a known negotiation deadline) assuming that players do not increase their demand during the game (*good-faith* assumption). He uses *backward induction* to identify optimal strategies for rational players: starting at the last stage of the game and then inductively working backwards to the beginning of the game. Rubinstein in [14] proposes a variation of Ståhl's alternating-offers protocol in which there are infinitely many possible agreements (the value of a parameter in $[0, 1]$), the time horizon is infinite, and the time preferences are stationary (i.e., the preference of getting x at time t over getting y at time $t+1$ does not depend on t). In his model Rubinstein identifies an unique *subgame perfect equilibrium* (see [4] for the definition of subgame perfect equilibrium) and the equilibrium is such that an agreement is immediately achieved.

Although much economics and computer science literature study the alternating-offers protocol (see [8]), several issues are still to be addressed before it can be usefully employed in automated negotiation. The two main open problems concern *incomplete information* and *multi-issue* bargaining in presence of rational agents. Easy and general solutions are available only when every pertinent information is common knowledge between the two players and the bargaining is only on one issue. Both assumptions are unrealistic or restrictive; e.g., it is very unlikely that one knows the other player's reservation price or her possible timeout; and it is very likely that one negotiates not only on the price of a good or service but also on its quantity or quality. The problem of incomplete knowledge in alternating-offers bargaining (see [1]) is a hard one and the relevant literature faces only very narrow problems or makes unrealistic assumptions.

Also alternating-offers bargaining on multiple issues is considered a hard problem to address (a literature review can be found in [9]). This is mainly due to difficulties in finding computationally tractable negotiation mechanisms that produce Pareto efficient agreements.

[1] Artificial Intelligence and Robotics Laboratory, Dipartimento di Elettronica e Informazione, Politecnico di Milano, Piazza Leonardo da Vinci 32, 20133, Milan, Italy, email: {digiunta, ngatti}@elet.polimi.it

Several procedures exist, the main ones being: *in-bundle* (i.e., all the issues are negotiated together) and *issue-by-issue* (i.e., the issues are negotiated one by one). In-bundle procedure is more efficient than issue-by-issue [6, 9] (i.e., issue-by-issue procedure does not assure the achievement of Pareto efficient agreement). However, as underlined in [12] by Raiffa, the in-bundle procedure requires complex computations to be carried out and it has not been deeply analyzed up to now. Some authors (e.g., Lai *et al.* in [9]) also say that finding a rigorous solution to the in-bundle procedure is an intractable problem. Conversely, issue-by-issue procedure has found a satisfactory discussion in literature and the main current open problem refers to the determination of the optimal *agenda* (i.e., the sequence of issues over which the bargaining is carried on); interesting works concerning issue-by-issue procedure can be found in [2, 9].

In this paper we describe the first step towards the development of an algorithm based on backward induction to determine the equilibrium strategies for fully rational agents in incomplete information, finite-horizon, multi-issue in-bundle alternating-offers bargaining. The incomplete information problem will be the subject of subsequent work. The step here developed concerns, instead, the multi-issue problem. The algorithm we propose makes a novel use of backward induction together with tools of mathematical optimization. And we prove that the problem of computing the equilibrium for multi-issue in-bundle is tractable when the (commonly employed) additive multi-attribute utility functions are used.

The paper is structured as follows. In the next section the bargaining protocol is formally described, some basic notations, concepts and techniques are described, and the one-issue solution is revised. In Sections 3 and 4 the equilibrium strategies for the multiple issue case are discussed. Section 5 concludes the paper. In Appendix A the proofs of the main propositions are given.

2 One-Issue Finite-Horizon Alternating-Offer Protocol

We consider the Rubinstein-Ståhl alternating-offers bargaining model [14, 16] enriched with reservation prices (i.e., the maximum price at which the buyer would buy the item and the minimum price at which the seller would sell it) and deadlines (i.e., time points after which the buyer or the seller have no more interest in bargaining). Two agents – a buyer agent b and a seller agent s – have strictly opposite interests on one attribute x of an item, which is a real number (typically the price), and bargain to reach an agreement. The agents alternatively act making or accepting offers and counteroffers or stopping negotiation without agreement. Each agent i has an utility function $U_i(x,t)$ that tells how much she gains from an agreement on the value x reached at time t. The utility of the seller increases linearly with x, while the utility of the buyer decreases linearly. Both utilities decrease exponentially as time passes by. The rules of the bargaining and the utilities of the agents are common knowledge (complete information hypothesis). In the next subsections we provide the exact model and revise its equilibrium analysis.

2.1 Bargaining Model

Two players exist, the buyer b and the seller s. They can act at integer times $t = 0, 1, 2, \ldots$. We denote by $\iota(t)$ the agent that acts at time t; function $\iota : \mathbb{N} \rightarrow \{b, s\}$ is called *player function*, and must be such that $\iota(t) \neq \iota(t+1)$.[2] We denote by σ_i^t the action performed by agent $i \in \{b, s\}$ at time t if $i = \iota(t)$. The possible values of $\sigma^0_{\iota(0)}$ are:

- $\mathit{offer}(\overline{x})$, where $\overline{x} \in \mathbb{R}$;
- exit.

If $t \neq 0$, the possible values of $\sigma^t_{\iota(t)}$ are:

- $\mathit{offer}(\overline{x})$, where $\overline{x} \in \mathbb{R}$;
- accept;
- exit.

If $\sigma^t_{\iota(t)} = \mathit{offer}(\overline{x})$, then the bargaining goes on to the next time point. If $\sigma^t_{\iota(t)} = \mathit{accept}$, then the bargaining stops and its *outcome* is $(\overline{x}, t)$, where $\overline{x}$ is the number such that $\sigma^{t-1}_{\iota(t-1)} = \mathit{offer}(\overline{x})$. If $\sigma^t_{\iota(t)} = \mathit{exit}$, then the bargaining stops and its outcome is $\mathit{NoAgreement}$.

The utility of player i, which is a function of the bargaining outcome, $U_i : (\mathbb{R} \times \mathbb{N}) \cup \{\mathit{NoAgreement}\} \rightarrow \mathbb{R}$, depends on three parameters:

- the *reservation price* $RP_i \in \mathbb{R}^+$;
- the *temporal discount factor* $\delta_i \in (0, 1]$;
- the *deadline* $T_i \in \mathbb{N}$, $T_i > 0$.

Exactly, if the outcome of the bargaining is an agreement (x, t), then the utility functions are, for the buyer b:

$$U_b(x,t) = \begin{cases} (RP_b - x)\delta_b^t & \text{if } t \leq T_b \\ -1 & \text{otherwise} \end{cases}$$

and for the seller s:

$$U_s(x,t) = \begin{cases} (x - RP_s)\delta_s^t & \text{if } t \leq T_s \\ -1 & \text{otherwise} \end{cases}$$

If the outcome is $\mathit{NoAgreement}$ the utilities are given by:

$$U_b(\mathit{NoAgreement}) = U_s(\mathit{NoAgreement}) = 0$$

Notice that the assignment of the value 0 to the utility given by $\mathit{NoAgreement}$ and of the value -1 to the utility given by any agreement beyond the deadline allows to effectively model the rational behaviour of the agents in presence of deadlines: once the deadline of an agent has expired, the agent prefers to make exit, being 0 the utility of $\mathit{NoAgreement}$, than reaching any agreement, being -1 the utility of any agreement beyond the deadline. Furthermore, we make the following four standard assumptions:

complete information: the protocol of the bargaining and the utility functions of the two agents (including the values of RP_i, δ_i and T_i) are common knowledge between the two agents;
feasibility: $RP_b \geq RP_s$;
rationality: it is common knowledge that each agent will act in order to maximize her utility;
benevolence: it is common knowledge that when an agent can choose between two outcomes which are indifferent for her but not for her opponent, she will choose the one that is better for her opponent.

The complete information assumption will be at least partially removed in subsequent work. The feasibility assumption is used to avoid a trivial situation. The rationality assumption is a standard one and will be *a posteriori* justified by the tractability of the problem of finding the equilibrium strategies (such that there is no need to take into account bounded rationality issues). Also the benevolence assumption is standard and, apart from being reasonable, is necessary to break ties and avoid multiple equilibria.

[2] The value of $\iota(0)$, i.e., the agent that starts bargaining, must be specified in an instantiation of the protocol.

2.2 Equilibrium Analysis

Finding the game-theoretic solution of the above model is an easy exercise, but we will explain it in some detail to introduce ideas and state notations useful for the subsequent less trivial situation of Sections 3 and 4.

An appropriate notion of solution for a complete information extensive form game like the one we are dealing with is *subgame perfect equilibrium* [4]: informally, a strategy of b and a strategy of s are a subgame perfect equilibrium if they are a Nash equilibrium in every possible subgame; i.e., also if the agents deviated for some time from the equilibrium, in following times it is still rational to follow the equilibrium.

In finite games it is possible to find subgame perfect equilibria by *backward induction*: one computes optimal actions for the last stage of the game, when it is known what the outcomes of the actions are; therefore one knows the outcomes for the actions at the stage before the last and can compute optimal ones; this process goes on recursively until all the stages of the game have been explored.

We remark that the protocol above described is not, rigorously speaking, a finite game; the deadlines are not in the protocol but in the agent's utility functions and the agents are allowed to offer and counteroffer also after the deadlines are expired. Nevertheless, it is essentially finite: a rational agent will give up bargaining after her deadline. It is therefore possible to use backward induction to solve it.

Informally, the agent that acts at the deadline of the bargaining $\overline{T}$ – let's say s – would accept any offer which has non-negative utility for her. Her opponent b knows that, and at time $\overline{T}-1$ she can safely offer RP_s (which would be accepted by s) or accept any possible previous offer x which is not worse for her than offering RP_s (i.e., $U_b(x,\overline{T}-1) \geq U_b(RP_s,\overline{T})$). Agent s knows that and at $\overline{T}-2$ would offer the maximum x such that $U_b(x,\overline{T}-1) \geq U_b(RP_s,\overline{T})$, which is the x such that $U_b(x,\overline{T}-1) = U_b(RP_s,\overline{T})$, or accept any possible previous offer which is not worse than that for her. This reasoning can be inductively carried on until the beginning of the game, finding an offer that the first player would do and her opponent would accept.

In other words, at each time point $\bar{t}$, starting from the time before the first deadline, it is possible to know which offer $x_{\bar{t}}$ would be made by agent $\iota(\bar{t})$ if she would make an offer; but she would accept and not make an offer if her opponent would have offered something not worse than $x_{\bar{t}}$ for $\iota(\bar{t})$ at time $\bar{t}-1$. The key feature of the calculation is, therefore, finding an offer $x_{\bar{t}-1}$ at time $\bar{t}-1$ which is the best for agent $\iota(\bar{t}-1)$ among those that would be accepted by $\iota(\bar{t})$ at time $\bar{t}$. This offer $x_{\bar{t}-1}$ is easily seen being the one such that $U_{\iota(\bar{t})}(x_{\bar{t}-1},\bar{t}) = U_{\iota(\bar{t})}(x_{\bar{t}},\bar{t}+1)$. On the space (x,t), the value $x_{\bar{t}-1}$ is therefore at the intersection of $t=\bar{t}$ with the level curve of $U_{\iota(\bar{t})}$ that goes through point $(x_{\bar{t}},\bar{t}+1)$. Because of the stationarity of the utility functions we employ, $x_{\bar{t}-1}$ is also at the intersection of $t=\bar{t}-1$ with the level curve of $U_{\iota(\bar{t})}$ that goes through point $(x_{\bar{t}},\bar{t})$.[3] We say that $x_{\bar{t}-1}$ is the one-step *backward propagation* of $x_{\bar{t}}$ along the level curve of $U_{\iota(t)}$.

As backward propagation of prices offered is extensively used in our work, we introduce a special notation. Given a price x, we denote by $x_{\leftarrow i}$ the backward propagation of x along the level curves of the utility of agent i, i.e., the price such that $U_i(x,t) = U_i(x_{\leftarrow i},t-1)$. If a price x is backward propagated n times along the level curves of agent i, we write $x_{\leftarrow n[i]}$. If a price is backward propagated along the level curves of more than one agent, we list them left to right in the subscript; for instance, $x_{\leftarrow i3[j]}$ is price x backward propagated along the level curves of agent i and subsequently three times along the curves of agent j.

We call $\overline{T} = \min\{T_b, T_s\}$ the *deadline of the bargaining*. We denote by $x^*(t)$ the price that backward induction prescribes would be offered at time $t < \overline{T}$ by agent $\iota(t)$ if she would make an offer. Recursively:

$$x^*(t-1) = \begin{cases} RP_{\iota(t)} & \text{if } t = \overline{T} \\ (x^*(t))_{\leftarrow \iota(t)} & \text{if } t < \overline{T} \end{cases}$$

The calculation of $x^*(\cdot)$ can be easily carried out recursively and its complexity is obviously linear with $\overline{T}$.[4]

We can now state the following result, whose rigorous proof is very easy but long and is therefore omitted (a sketch is provided in [3, 10]):

Proposition 2.1 *The bargaining game of Subsection 2.1 has one and only one subgame perfect equilibrium. The equilibrium strategies for* $t \leq \overline{T}$ *are given by:*

$$\sigma^t_{\iota(t)} = \begin{cases} accept & \text{if } \begin{cases} t > 0 \\ \iota(t) = b \\ \sigma^{t-1}_{\iota(t-1)} = \mathit{offer}(x) \text{ with } x \leq x^*(t-1) \end{cases} \\ accept & \text{if } \begin{cases} t > 0 \\ \iota(t) = s \\ \sigma^{t-1}_{\iota(t-1)} = \mathit{offer}(x) \text{ with } x \geq x^*(t-1) \end{cases} \\ \mathit{offer}(x^*(t)) & \text{otherwise} \end{cases}$$

The agreement is therefore achieved at time $t = 1$ *on the price* $x^*(0)$.

It can be also seen that, for $RP_s = 0$ and $RP_b = 1$, the solution converges to the Rubinstein solution [14] as $\overline{T}$ grows to infinity.

In Figure 1 we report on the (x,t) space a bargaining with $RP_b = 1$, $RP_s = 0$, $\delta_b = 0.7$, $\delta_s = 0.7$, $T_b = 9$, $T_s = 10$, and $\iota(0) = b$. The seller acts at the deadline $\overline{T} = 9$ of the game. At that time she is willing to accept any non-negative offer, so at $t = 8$ the buyer's possible offer would be $x^*(8) = 0$. At time $t = 7$ the seller's possible offer would be $x^*(7) = 0_{\leftarrow b}$ and so on. The values of $x^*(t)$ are highlighted with circles and denoted as backward propagations of the last value $x^*(8) = RP_s$. The agreement is achieved at time $t = 1$ on the price $x^*(0) = (RP_s)_{\leftarrow 4[bs]}$.

3 Bargaining over Multiple Issues In-Bundle

We take into account the scenarios in which the agents negotiate over several issues which can be either different attributes of one good or different goods. For instance, a buyer and a seller, trading on a single service, can negotiate over the price of the service, the quality of the service, the delivery time of the service, the guarantee expiration of the service, and so on; or the buyer and the seller can negotiate over the prices of different services.

In multi-issue negotiations the issues can be negotiated according to several procedures. The main procedures are: *in-bundle*: the agents negotiate all the issues together, namely, a bundle; *issue by issue*: the agents negotiate each issue separately.

The main problem to address in multi-issue bargaining is the development of mechanisms to produce Pareto efficient outcomes. In

[3] As it is better for visualization purposes, we will usually employ this second characterization of $x_{\bar{t}-1}$.

[4] Also closed form expressions for $x^*(t)$ can be given, but we omit them for the sake of brevity.

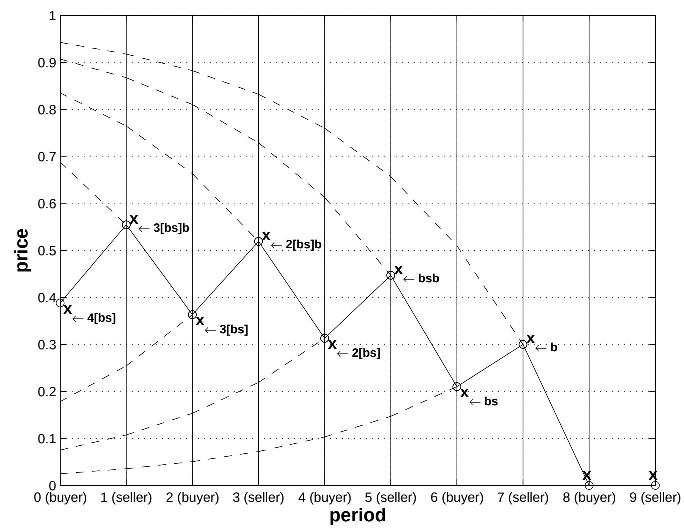

Figure 1. Backward induction construction for $RP_b = 1$, $RP_s = 0$, $\delta_b = 0.7$, $\delta_s = 0.7$, $T_b = 9$, $T_s = 10$, $\iota(0) = b$

single-issue bargaining the Pareto optimality of the outcome is trivial, since, by the opposite preferences of the agents, each agreement is Pareto efficient. In multi-issue bargaining, the achievement of Pareto efficient agreement is harder.

First of all we provide a formulation for the multi-issue utility functions of the agents in the case they bargain the issues in-bundle. We recall that in in-bundle bargaining the agents' strategies (i.e., offer, reject, and accept) are defined on the entire bundle of offers and not on just a portion of it (e.g., an agent cannot accept or reject just a part of the offers of the bundle, but the entire bundle itself). We consider multi-issue utility functions U given by the sum of single-issue utility functions. This kind of multi-issue utility functions, called *additive multi-attribute utility functions* is the most common one (see, e.g., [12]). The multi-issue utility function of agent i is:

$$U_i(x^1, \ldots, x^n, t) = \begin{cases} \sum_{j=1}^{n} U_i^j(x^j, t) & \text{if } \forall\, j, U_i^j(x^j, t) \geq 0 \\ -1 & \text{otherwise} \end{cases},$$

(notice that the -1 has been introduced to capture the case that the agreement has not been reached on every issue of the bundle) where:

$$U_i^j(x^j, t) = \begin{cases} u_i^j(x^j) \cdot (\delta_i^j)^t & \text{if } t \leq T_i^j \\ -1 & \text{otherwise} \end{cases},$$

(notice that the -1 has been introduced to capture the case that the agreement has been reached beyond the deadline) where:

- $\delta_i^j \in (0, 1]$;
- $T_i^j > 0$;
- u_i^j are continuous concave functions in $\mathbb{R}$;
- u_i^j are strictly monotonic;
- (*non-empty feasible agreement set*) $\forall\, j,\ \exists\, A^j \neq \emptyset, A^j = \{x^j \in \mathbb{R} : u_b^j(x^j) \geq 0 \text{ and } u_s^j(x^j) \geq 0\}$, we call $A = \times_{j=1}^{n} A^j$ feasible agreement set; notice that, by continuity and concavity of u_i^j, A^j are compact set;
- (*opposite preferences*) $\forall\, \overline{x}^j \in A^j : \exists\, \mu < 0$ such that $\dfrac{du_b^j}{dx^j}(\overline{x}^j) = \mu \cdot \dfrac{du_s^j}{dx^j}(\overline{x}^j)$.

Notice that the U_i^js are a generalization of the single issue utility functions introduced in the previous section. For a complete parallelism we denote by RP_i^j the (obviously unique) x such that $u_i^j(x) = 0$ and we call it *reservation value* of agent i over the issue j.

We give some notes on the above assumptions. Since u_i^j are continuous concave functions, then $\sum_{j=1}^{n} u_i^j$ is a continuous concave function (it can be trivially proved that the Hessian of $\sum_{j=1}^{n} u_i^j$ is a diagonal matrix with negative eigenvalues). Differently from the single issue case in which the utility of the buyer (seller) is decreasing (increasing) in x, in the multi-issue case the utility of a buyer (seller) can be either increasing or decreasing in x^j. For instance, if a buyer is negotiating on the price and the quality of a service, it desires to buy the service at the minimum price and at the maximum quality.

We call RP_{max}^j and RP_{min}^j the respectively highest and lowest values between the reservation values of the two agents. The feasible agreement set A can be simply described as: $\forall\, j, RP_{min}^j \leq x^j \leq RP_{max}^j$.

Note that the agents can have different deadlines for each single attribute. Actually, there are two common situations: (*1*) the agents negotiate on different attributes concerning an unique good, (*2*) the agents negotiate on attributes concerning different goods. In the first situation we expect that the deadlines depend only on the traded good itself, and not on its attributes; therefore, each attribute has the same deadline. In the second situation, instead, we expect that each single good has its own deadline, different from the deadline of the other issues. We say that an issue j is *negotiable* at a given time t if $t \leq T^j$. The main difference between the two situations discussed above is that in (*1*) the agents must agree on all the attributes (all the attributes are negotiable for all the bargaining length), while in (*2*) the agents can find partial agreements on the attributes that are negotiable at a certain time. Obviously, if the agreement is reached immediately, the agents agree on the values of all the issues; if the agreement is reached beyond some deadlines the agents agree on just a portion of the issues that were negotiable at the beginning of the bargaining. We call $\overline{T}^j$ the deadline of bargaining related to the issue j, i.e., $\overline{T}^j = \min\, \{T_b^j, T_s^j\}$. We call $\overline{T}$ the deadline of bargaining, i.e., $\overline{T} = \max\, \{\overline{T}^1, \ldots, \overline{T}^m\}$.

In the next section we study how an offer can be backward propagated and we give the subgame perfect equilibrium strategies in presence of multiple issues.

4 Backward Propagation with Multiple Issues

We modify our backward induction construction to address multi-issue negotiations in-bundle. As we will describe in what follows, the basic ideas behind the backward induction construction holds unaltered, the differences being: the construction will be built in a multi-dimensional space whose dimensions are exactly the issues, and the backward propagation of the offers with n issues will in general map a n-dimensional offer from the subspace at time t to a n-dimensional offer in the subspace at time $t-1$.

For simplicity, we initially consider the case in which the issues related to each single agent have the same deadline. Then we extend our solution to the more general case.

4.1 Issues with the Same Deadline

In the following, given an agent i we denote by $-i$ her opponent agent. Starting from an offer $\mathbf{x} = \langle x^1, \ldots, x^n \rangle$ at time $t+1$ we want to compute $\mathbf{x}_{\leftarrow -i}$ at time t given that $i = \iota(t)$. The backward propagation of an offer $\mathbf{x} = \langle x^1, \ldots, x^n \rangle$ to time t with $i = \iota(t)$ is tackled in two stages.

1. Determination of the set $X_{\leftarrow -i}$ of the offers $\mathbf{z}$ that give to agent $-i$ at time t the utility given by $\mathbf{x}$ at time $t+1$:

$$X_{\leftarrow -i} = \left\{ \mathbf{z} = \langle z^1, \ldots, z^n \rangle \in A : \sum_{j=1}^{n} \left[u^j_{-i}(z^j) \cdot (\delta^j_{-i})^t \right] = \right. \\ \left. = \sum_{j=1}^{n} \left[u^j_{-i}(x^j) \cdot (\delta^j_{-i})^{t+1} \right] \right\}, \quad (1)$$

   in general $X_{\leftarrow -i}$ is a geometric place of $n-1$ dimensions. Notice that all the offers belonging to $X_{\leftarrow -i}$ are indifferent for the agent $-i$ and $X_{\leftarrow -i}$ is a compact set since A is compact.
2. Determination of the set $\{\mathbf{x}_{\leftarrow -i}\}$ of the offers belonging to $X_{\leftarrow -i}$ that maximize the utility of agent i at time t:

$$\{\mathbf{x}_{\leftarrow -i}\} = \arg \max_{\mathbf{z} \in X_{\leftarrow -i}} \sum_{j=1}^{n} \left[u^j_i(z^j) \cdot (\delta^j_i)^t \right],$$

   this is due to the fact that among all the offers that $-i$ would accept at time $t+1$ (i.e., $X_{\leftarrow -i}$) the agent i makes at time t the offer that is the best for herself. In general $\{\mathbf{x}_{\leftarrow -i}\}$ can be a set of offers.

We provide a mathematical analysis of the multi-issue offer backward propagation. We prove that, given a multi-issue offer at time $t+1$, its backward propagation at time t is always made up of at least one offer and we prove that the offer prescribed by backward propagation is Pareto efficient in the subspace at time t.

Proposition 4.1 *For all $\mathbf{x} \in A$, for all $t \leq \overline{T}$, and for all i, $X_{\leftarrow -i}$ is always a non-empty set.*

We report the proof in Appendix A.

Proposition 4.2 *For all $\mathbf{x} \in A$, for all $t \leq \overline{T}$, and for all i, $\{\mathbf{x}_{\leftarrow -i}\}$ is always a non-empty set.*

Proof. Since U_i is continuous and concave in A and $X_{\leftarrow -i}$ is a compact set, there exists at least one global minimum by Weierstrass. □

Proposition 4.3 *For all $\mathbf{x} \in A$ and for all i, the agreement $(\mathbf{x}_{\leftarrow -i}, t)$ are Pareto efficient in subspace at time t.*

We report the proof in Appendix A.

Proposition 4.4 *The agreement $(\mathbf{x}^*(0), 1)$ prescribed by multi-issue backward induction is Pareto optimal.*

Proof. It trivially follows from Proposition 4.3 and from the discounting of the utility in time. □

As discussed in the proofs of the Proposition 4.3 the backward propagation of a multi-issue offer is the result a linear/convex programming problem. We recall that the linear programming complexity is polynomial [7]. This means that the complexity of the algorithm that produce the above multi-issue backward induction construction when the utility functions of the two agents are linear is linear with the deadline and polynomial with the number of issues. We report in Figure 2 an example of backward propagation with two issues. We report in Algorithm 1 the algorithm of backward propagation with multiple issues when the deadline over the issues is unique.

Algorithm 1 MULTI_ISSUE_BACKWARD_PROPAGATION $(\mathbf{x}, t)$

```
1: determine the feasible agreement set A
2: determine X_{←−ι(t)}
3: determine {x_{←−ι(t)}}
4: return an element of {x_{←−ι(t)}}
```

4.2 Issues with Different Deadlines

We explore the situation of different deadlines related to each single issue to negotiate. The backward propagation construction is altered as follows:

- we consider the set of the issues negotiable at the deadline of the bargaining $\overline{T}$;
- at $\overline{T}$ the agent $\iota(\overline{T})$ would accept a bundle of offers such that it concerns exclusively the issues negotiable at $\overline{T}$ and the offer on each singular issue gives to her non-negative utility, i.e., the value of the offer for each negotiable issue is just the reservation price $RP^j_{\iota(\overline{T})}$;
- at $\overline{T}-1$ the agent $\iota(\overline{T}-1)$ would offer a bundle of offers composed of $RP^j_{\iota(\overline{T})}$ for each j such that the issue j is negotiable at $\overline{T}$;
- we backward propagate this bundle of offers according to the construction presented in the previous section until the deadline of bargaining of other issues is reached;
- as the deadline of bargaining of an issue (or several issues) is reached, the issue becomes negotiable. Suppose that in backward propagation we reach the deadline $\overline{T}^k$ of just the issue k: at $\overline{T}^k$ the set of negotiable issues is enriched by the issue k. The agent $\iota(\overline{T}^k)$ would accept at time $\overline{T}^k$ the bundle composed of *(a)* the bundle of offers that she would make at $\overline{T}^k$ backward propagated to the time $\overline{T}^k - 1$ (notice that a such bundle does not comprise any offer on the issue k) and *(b)* the offer $RP^k_{\iota(\overline{T}^k)}$ on the issue k. The situation in which more than one issue becomes negotiable at $\overline{T}^k$ is tackled similarly;
- from $\overline{T}^k - 1$ the backward propagation will be accomplished on the entire set of the issues negotiable at $\overline{T}^k$;
- the above procedure is repeated for each issue or set of issues that has deadlines between $\overline{T}$ and 0.

Summarily, the backward propagation construction has origin at $\overline{T}-1$ with the set of all negotiable issues at time $\overline{T}$, then during the construction every time $\overline{T}^k$ an issue k becomes negotiable the offer prescribed by backward induction for a such issue at time $\overline{T}^k - 1$ is the reservation price $RP^k_{\iota(\overline{T}^k)}$, the offer prescribed for the other issues is given by the backward propagation along the issues that are negotiable at $\overline{T}^k + 1$. Then the backward propagation continues from $\overline{T}^k - 1$ considering all the issues negotiable at time $\overline{T}^k$. Notice that the backward induction construction in the case the issues have different deadlines is a subconstruction of the construction accomplished in the case the issues have the same deadlines. As a result, the two properties of the backward induction construction accomplished with the same deadline for all the issues hold also in this case: (i) the agreement is Pareto efficient, (ii) the computational complexity of the construction is linear with the deadline of bargaining and polynomial with the number of issues. We report in Figure 3 an example of backward propagation in which the deadlines on the two issues are different. We report in Algorithm 2 the algorithm of backward

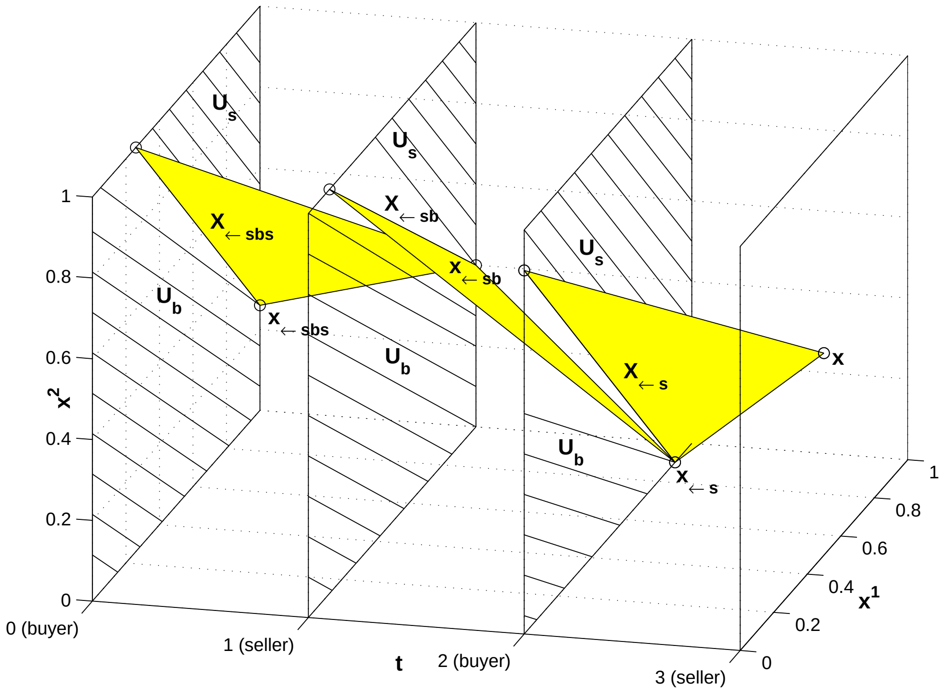

Figure 2. An example of backward propagation with two issues from $(\mathbf{x}, 3)$: $RP_b^1 = 0, RP_b^2 = 0, \delta_b^1 = 0.8, \delta_b^2 = 0.7,$ $RP_b^1 = 1, RP_b^2 = 1, \delta_b^1 = 0.9, \delta_b^2 = 0.9, \iota(0) = b$

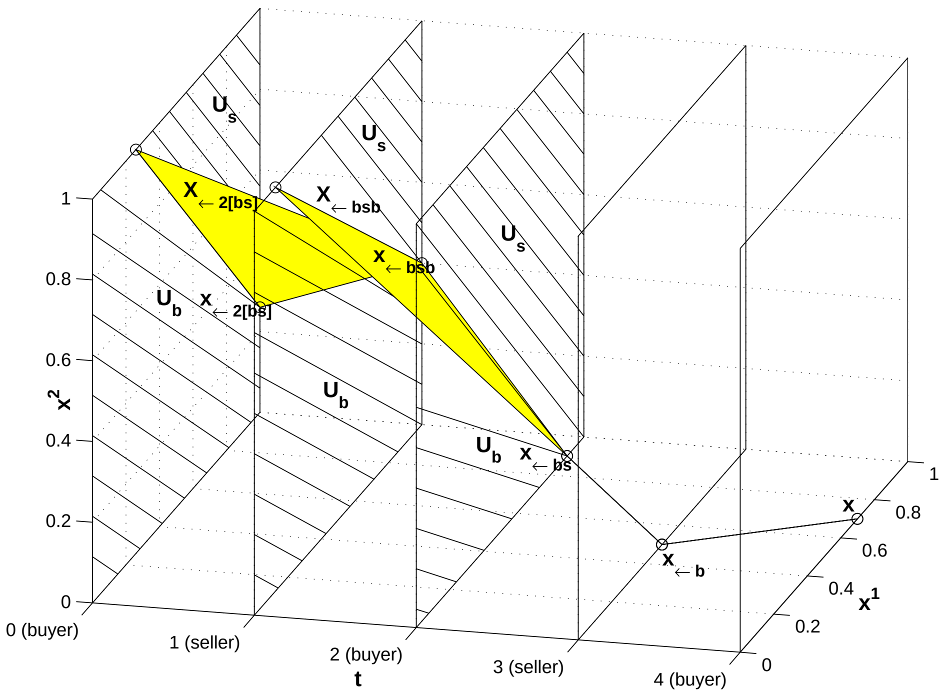

Figure 3. A example of backward propagation with two issues from $(\mathbf{x}, 4)$ in which the deadline of the issue x^2 is at $t = 3$: $RP_b^1 = 0, RP_b^2 = 0, \delta_b^1 = 0.8, \delta_b^2 = 0.7, RP_b^1 = 1, RP_b^2 = 1, \delta_b^1 = 0.9, \delta_b^2 = 0.9, \iota(0) = b$

propagation with multiple issues when the deadlines over the issues are different.

Algorithm 2 MULTI_ISSUE_BACKWARD_PROPAGATION ($\mathbf{x}$, t)

1: determine the feasible agreement set $A(t)$ on the issues negotiable at time t
2: for all j such that $t = \overline{T}^j - 1$ set $x^j = RP^j_{\iota(t+1)}$
3: for all j such that $t < \overline{T}^j - 1$ determine $X_{\longleftarrow -\iota(t)}$ for only the issues negotiable at time $t+2$
4: determine $\{\mathbf{x}_{\longleftarrow -\iota(t)}\}$ for only the issues negotiable at time $t+2$
5: return an offer composed of an element of $\{\mathbf{x}_{\longleftarrow -\iota(t)}\}$ and the values assigned to the x^js in the step 2 of the algorithm

4.3 Equilibrium Strategies with Multiple Issues

We denote by $\mathbf{x}^*(t)$ the offer prescribed by the backward induction constructions given above. In general, $\mathbf{x}^*(t) \in \mathbb{R}^m$, however, as discussed above, in the case the issues have different deadlines the offer $\mathbf{x}^*(t)$ is defined on just a subset of $\mathbb{R}^m$, i.e., the set of the negotiable issues. We can now state the following result as extension of the result given for the one-issue bargaining:

Proposition 4.5 *The bargaining game of Subsection 2.1 in which the utility functions are defined on multiple attributes has one and only one subgame perfect equilibrium. The equilibrium strategies for $t \leq \overline{T}$ are given by:*

$$\sigma^t_{\iota(t)} = \begin{cases} accept \text{ if } \begin{cases} t > 0 \\ \sigma^{t-1}_{\iota(t-1)} = offer(\mathbf{x}) \text{ with } U_{\iota(t)}(\mathbf{x}, t-1) \geq \\ \qquad U_{\iota(t)}(\mathbf{x}^*(t-1), t-1) \end{cases} \\ offer(\mathbf{x}^*(t)) \text{ otherwise} \end{cases}$$

where $\mathbf{x}$ is an offer defined on exclusively the issues negotiable at time t. The agreement is therefore achieved at time $t = 1$ on the price $\mathbf{x}^*(0)$.

Notice that a bundle comprising any non-negotiable issue is rejected independently from the values of the offers on the other issues.

5 Conclusions

Automated negotiations have been suggested as a way to facilitate increasingly efficient electronic trading. The computational speed of autonomous agent can significantly enhance negotiation, especially in presence of a combinatorial number of possible deals. In this paper we analyze the alternating-offers negotiation protocol. One of the main problems related to this protocol concerns its computational tractability. Economics and computer science literature lacks of studies on what is tractable and what is not in alternating-offers. In this paper we have provided an analysis of perfect information alternating-offers bargaining with finite-horizon on multiple issues in-bundle and we have shown that the determination of the equilibrium strategies – considered intractable in literature – is computationally tractable both when all the issues related to a single agent have the same deadline and when each single issue has its own deadline. In detail, when the utility functions are linear, the complexity is linear with the bargaining time and polynomial with the number of issues. Our proposal can be also employed when the utility functions are generically concave; in this case the computational complexity is still linear with the bargaining time, but it is not polynomial with the number of issues being the complexity of convex programming techniques.

Finally, the backward induction construction presented in this paper is a prominent technique to study efficient multi-issue negotiation. In particular, we will employ it in future to determine the the equilibrium strategies of the agents in the presence of incomplete information.

References

[1] P. C. Cramton, L. M. Ausubel, and R. J. Deneckere, *Handbook of Game Theory*, volume 3, chapter Bargaining with Incomplete Information, 1897–1945, Elsevier Science, 2002.
[2] S. S. Fatima, M. Wooldridge, and N. R. Jennings, 'An agenda-based framework for multi-issue negotiation', *Artificial Intelligence*, **152**, 1–45, (2004).
[3] D. Fudenberg and J. Tirole, *Game Theory*, The MIT Press, Cambridge, MA, USA, 1991.
[4] J. C. Harsanyi and R. Selten, 'A generalized nash solution for two-person bargaining games with incomplete information', *Management Science*, **18**, 80–106, (1972).
[5] J.-B. Hiriart-Urruty and C. Lemaréchal, *Convex Analysis and Minimization Algorithms I*, Springer-Verlag, Berlin, Germany, 1996.
[6] Y. In and R. Serrano, 'Agenda restrictions in multi-issue bargaining', *Journal of Economic Behavior and Organization*, **53**, 385–399, (2004).
[7] N. Karmarkar, 'A new polynomial-time algorithm for linear programming', *Combinatorica*, **4**(4), 373–395, (1984).
[8] S. Kraus, *Strategic Negotiation in Multiagent Environments*, The MIT Press, Cambridge, MA, USA, 2001.
[9] G. Lai, C. Li, K. Sycara, and J. Giampapa, 'Literature review on multi-attribute negotiations', Technical Report CMU–RI–TR–04–66, Carnegie Mellon University, (December 2004).
[10] S. Napel, *Bilateral Bargaining: Theory and Applications*, Springer-Verlag, Berlin, Germany, 2002.
[11] M. J. Osborne and A. Rubinstin, *Bargaining and Markets*, Academic Press, San Diego, CA, USA, 1990.
[12] H. Raiffa, *The Art and Science of Negotiation*, Harvard University Press, Cambridge, USA, 1982.
[13] J. S. Rosenschein and G. Zlotkin, *Rules of Encounter. Designing Conventions for Automated Negotiations among Computers*, The MIT Press, Cambridge, MA, USA, 1994.
[14] A. Rubinstein, 'Perfect equilibrium in a bargaining model', *Econometrica*, **50**(1), 97–109, (1982).
[15] T. Sandholm, 'Agents in electronic commerce: Component technologies for automated negotiation and coalition formation', *Autonomous Agents and Multi-Agent Systems*, **3**(1), 73–96, (2000).
[16] I. Stahl, *Bargaining Theory*, Stockholm School of Economics, Stockholm, Sweden, 1972.

A Proofs

A.1 Proof of the Proposition 4.1

We prove that $X_{\longleftarrow -i}$ consists of at least one element. We consider $\mathbf{x} \in A, \mathbf{x} = \langle x^1, \ldots, x^n \rangle$. We call $X^*_{\longleftarrow -i} = \left\{ \mathbf{z} \in \mathbb{R}^n : \sum_{j=1}^n \left[u^j_{-i}(z^j) \cdot (\delta^j_{-i})^t \right] = \sum_{j=1}^n \left[u^j_{-i}(x^j) \cdot (\delta^j_{-i})^{t+1} \right] \right\}$, then $X_{\longleftarrow -i} = X^*_{\longleftarrow -i} \bigcap A$. $X^*_{\longleftarrow -i}$ is a non-empty set since $\sum_{j=1}^n \left[u^j_{-i}(x^j) \cdot (\delta^j_{-i})^{t+1} \right]$ is finite and U_{-i} is continuous, concave, and strictly monotonic. In addition, since $\sum_{j=1}^n \left[u^j_{-i}(x^j) \cdot (\delta^j_{-i})^{t+1} \right] \geq 0$ and δ^j_{-i}s are non-negative, there exists at least $\mathbf{z} \in X^*_{\longleftarrow -i}$ such that for all j we have $u^j_{-i}(z^k) \geq 0$ (i.e., $\mathbf{z} \in X^*_{\longleftarrow -i} \cap A_{-i}$). We need to prove that a such $\mathbf{z}$ satisfies, for all j, $u^j_i(z^k) \geq 0$ (i.e., $\mathbf{z} \in X^*_{\longleftarrow -i} \cap A_{-i} \cap A_i$). Starting from $\mathbf{x}$ it is possible to build a path that connects $\mathbf{x}$ to $\mathbf{z}$ just along the directions such that $\frac{du^j_{-i}}{dx^j} < 0$ and $u^j_{-i} \geq 0$. Given $\mathbf{x} = \langle x^1, \ldots, x^n \rangle$, we consider a movement on just the component x^1 such that $\mathbf{y}_1 = \langle y^1, x^2, \dots, x^n \rangle$. By backward propagation:

$$\sum_{j=2}^{n} \left[u_{-i}^j(x^j) \cdot (\delta_{-i}^j)^t \right] + u_{-i}^1(y^1) \cdot (\delta_{-i}^1)^t =$$

$$= \sum_{j=2}^{n} \left[u_{-i}^j(x^j) \cdot (\delta_{-i}^j)^{t+1} \right] + u_{-i}^1(x^1) \cdot (\delta_{-i}^1)^t,$$

we derive:

$$u_{-i}^1(y^1) - u_{-i}^1(x^1) = \sum_{j=2}^{n} \left[u_{-i}^j(x^j) \cdot (\delta_{-i}^j - 1) \right],$$

since $\sum_{j=2}^{n} \left[u_{-i}^j(x^j) \cdot (\delta_{-i}^j - 1) \right] < 0$, we have that $u_{-i}^1(y^1) < u_{-i}^1(x^1)$. We can move along x^1 direction until (1) is satisfied or $u^1(y^1) = 0$. If (1) is not satisfied, we can move along x^2, and so on. By continuity, concavity, and strictly monotonicity of U_i and by the positivity of $\sum_{j=1}^{n} \left[u_{-i}^j(x^j) \cdot (\delta_{-i}^j)^{t+1} \right]$, there exists at least a $\mathbf{z}$ such that $\forall\ j, u_{-i}^j(z^j) \leq u_{-i}^j(x^j)$ and (1) is satisfied. Then, by opposite preferences hypothesis, along each singular components x^j on which the path is built the corresponding utility u_i^j increases. Since, by hypothesis on $\mathbf{x}$, we have $\forall j, u_i^j(x^j) \geq 0$, then $\forall j, u_i^j(z^j) \geq u_i^j(x^j) \geq 0$. Follows that $\forall j, u_i^j(z^j) \geq 0, u_{-i}^j(z^j) \geq 0$. □

A.2 Proof of the Proposition 4.3

We recall the definition of Pareto efficiency, $(\mathbf{x}_{\leftarrow -i}, t)$ is Pareto efficient in t if:

$$\nexists \mathbf{x} : \begin{cases} U_i(\mathbf{x}, t) \geq U_i(\mathbf{x}_{\leftarrow -i}, t) \\ U_{-i}(\mathbf{x}, t) \geq U_{-i}(\mathbf{x}_{\leftarrow -i}, t) \end{cases}$$

We call $S_i(U,t)$ and $S_{-i}(U,t)$ the sub-level sets of U_i and U_{-i} at t defined as follows: $S_i(U,t) = \{\mathbf{x} \in A : U_i(\mathbf{x}) \geq U\}$ and $S_{-i}(U,t) = \{\mathbf{x} \in A : U_{-i}(\mathbf{x}) \geq U\}$. We identify the frontier of S_i and S_{-i} with $\overline{S}_i$ and $\overline{S}_{-i}$, respectively. The frontier $\overline{S}$ is a level set (i.e., $\overline{S}_{-i}(U,t) = \{\mathbf{x} \in A : U_{-i}(\mathbf{x}) = U\}$). We note that, given a $\mathbf{x} \in A$, for all $\mathbf{z} \in S_i(U_i(\mathbf{x}), t) \cap S_{-i}(U_{-i}(\mathbf{x}), t)$ we have that $(\mathbf{z}, t)$ Pareto dominates or is Pareto indifferent to $(\mathbf{x}, t)$. Moreover, given a $\mathbf{x} \in A$, for all $\mathbf{z} \in \overline{S}_i(U_i(\mathbf{x}), t) \cap \overline{S}_{-i}(U_{-i}(\mathbf{x}), t)$ we have that $(\mathbf{z}, t)$ is Pareto indifferent to $(\mathbf{x}, t)$.

By hypothesis $\mathbf{x}_{\leftarrow -i} = \max_{\mathbf{x} \in X_{\leftarrow -i}} U_{-i}(\mathbf{x}, t)$, what we want to prove is that $\mathbf{x}_{\leftarrow -i}$ is Pareto efficient in the subspace at time t. We consider a generic optimization problem of the form:

$$\begin{cases} \min\ f(\mathbf{z}) \\ \text{s.t. } g(\mathbf{z}) \leq 0 \\ \text{s.t. } h(\mathbf{z}) = 0 \end{cases} \tag{2}$$

By Karush-Kuhn-Tucker theorem [5], $\overline{\mathbf{z}}$ is a local minimum of the problem (2) if and only if f and g are convex and h linear and there exist $\lambda^0 \in \mathbb{R}$ and $\lambda \in \mathbb{R}^m$, $\mu \in \mathbb{R}^p$ such that $\lambda^0 > 0$, $\lambda \geq 0$, $\mu \neq 0$ and:

$$\begin{cases} \lambda^0 \nabla f(\overline{\mathbf{z}}) + \lambda \nabla g(\overline{\mathbf{z}}) + \mu \nabla h(\overline{\mathbf{z}}) = 0 \\ g(\overline{\mathbf{z}})\lambda = 0 \end{cases} \tag{3}$$

If h is not linear, Karush-Khun-Tucker expresses just a necessary condition [5]. In our case (2) can be written as:

$$\begin{cases} \min\ -U_i(\mathbf{z}, t) \\ \text{s.t. } g(\mathbf{z}) \leq 0 \\ \text{s.t. } U_{-i}(\mathbf{z}, t) - U_{-i}(\mathbf{x}, t+1) = 0 \end{cases} \tag{4}$$

where $\mathbf{x}$ and t are fixed, and $g(\mathbf{z}) < 0$ are of the form $z^j - RP_{max}^j \leq 0$ and $RP_{min}^j - z^j \leq 0$.

We analyze two cases: (*a*) the minimum $\mathbf{x}_{\leftarrow -i}$ is inner to A (i.e., $g(\mathbf{x}_{\leftarrow -i}) \neq 0$), (*b*) the minimum $\mathbf{x}_{\leftarrow -i}$ is on the frontier of A (i.e., $g(\mathbf{x}_{\leftarrow -i}) = 0$).

Case (*a*). We initially consider linear U_i and U_{-i}. Since $\mathbf{x}_{\leftarrow -i}$ is a minimum of (4), then, by Karush-Khun-Tucker, there exist $\lambda^0 \in \mathbb{R}$ and $\mu \in \mathbb{R}^p$ such that such that $\lambda^0 > 0$, $\mu \neq 0$ and:

$$\nabla U_i(\mathbf{x}_{\leftarrow -i}, t) = \frac{\mu}{\lambda^0} \nabla U_{-i}(\mathbf{x}_{\leftarrow -i}, t). \tag{5}$$

By opposite preferences and (5), $\nabla U_i(\mathbf{x}_{\leftarrow -i}, t)$ and $\nabla U_{-i}(\mathbf{x}_{\leftarrow -i}, t)$ are opposite. By linearity of U_i and U_{-i} we have that $\overline{S}_i$ and $\overline{S}_{-i}$ are hyper-planes and, in particular, $\overline{S}_i \equiv \overline{S}_{-i}$. By the fact that $\nabla U_i(\mathbf{x}_{\leftarrow -i}, t)$ and $\nabla U_{-i}(\mathbf{x}_{\leftarrow -i}, t)$ are opposite we have that the intersection between S_i and S_{-i} is just $\overline{S}_i \equiv \overline{S}_{-i}$ [5]. This means that in the subspace at time t: (*i*) all the points belonging to $\overline{S}_i \equiv \overline{S}_{-i}$ are Pareto-indifferent to $(\mathbf{x}_{\leftarrow -i}, t)$, (*ii*) any point outer $\overline{S}_i \equiv \overline{S}_{-i}$ are Pareto-dominated by $(\mathbf{x}_{\leftarrow -i}, t)$.

We consider concave U_i and U_{-i}. The gradients of U_i and U_{-i} in $\mathbf{x}_{\leftarrow -i}$ are opposite by Karush-Khun-Tucker and by hypothesis of opposite preferences. We consider the two gradients $\nabla U_i(\mathbf{x}_{\leftarrow -i}, t)$ and $\nabla U_{-i}(\mathbf{x}_{\leftarrow -i}, t)$ and the hyper-plane orthogonal to the gradients that passes in $\mathbf{x}_{\leftarrow -i}$. A such hyper-plane corresponds to the level sets of linear utility functions with the gradients in the same directions of $\nabla U_i(\mathbf{x}_{\leftarrow -i}, t)$ and $\nabla U_{-i}(\mathbf{x}_{\leftarrow -i}, t)$. Since U_i and U_{-i} are concave the sub-level sets S_i and S_{-i} are subsets of the the sub-level sets of linear functions with the gradients directed in the same way, then the intersection between S_i and S_{-i} is just $\overline{S}_i \equiv \overline{S}_{-i}$. This means that in the subspace at time t: (*i*) all the points belonging to $\overline{S}_i \cap \overline{S}_{-i}$ are Pareto-indifferent to $(\mathbf{x}_{\leftarrow -i}, t)$, (*ii*) any point outer $\overline{S}_i \cap \overline{S}_{-i}$ are Pareto-dominated by $(\mathbf{x}_{\leftarrow -i}, t)$. Note that if U_i and U_{-i} are strictly concave, then their level sets are strictly convex and $\overline{S}_i \cap \overline{S}_{-i} \equiv \mathbf{x}_{\leftarrow -i}$.

Case (*b*). We initially consider linear U_i and U_{-i} are linear. Since $\mathbf{x}_{\leftarrow -i}$ is a minimum of (4), then, by Karush-Khun-Tucker, then there exists $\lambda_i^0 \in \mathbb{R}$ and $\lambda_i \in \mathbb{R}^m$, $\mu_i \in \mathbb{R}^p$ such that $\left(\lambda_i^0, \lambda_i\right) > 0$ and $\left(\lambda_i^0, \lambda_i, \mu_i\right) \neq 0$ and:

$$\begin{cases} -\lambda_i^0 \nabla U_i(\mathbf{x}_{\leftarrow -i}, t) + \lambda_i \nabla g(\mathbf{x}_{\leftarrow -i}) + \mu_i \nabla U_{-i}(\mathbf{x}_{\leftarrow -i}, t) = 0 \\ g(\mathbf{x}_{\leftarrow -i}) \lambda_i = 0 \end{cases}$$

We consider the case in which just a single constraint $g^l(\mathbf{x}_{\leftarrow -i}) = 0$ is active, it is the case in which the minimum $\mathbf{x}_{\leftarrow -i}$ is placed on just one border. In the cases in which the minimum $\mathbf{x}_{\leftarrow -i}$ is placed on more than one border, (3) singularly holds for each border. We prove that also in this case $(S_i \cap S_{-i}) \equiv (\overline{S}_i \cap \overline{S}_{-i})$. From:

$$\lambda_i^0 \nabla U_i(\mathbf{x}_{\leftarrow -i}, t) = \lambda_i \nabla g(\mathbf{x}_{\leftarrow -i}) + \mu_i \nabla U_i(\mathbf{x}_{\leftarrow -i}, t)$$

since $\lambda_i \nabla g(\mathbf{x}_{\leftarrow -i})$ has at least one component equal to zero and U_i and U_{-i} have opposite preferences on such component, follows that $\mu_i < 0$. We consider now the optimization problem:

$$\begin{cases} \min\ -U_{-i}(\mathbf{z}, t) \\ \text{s.t. } g(\mathbf{z}) < 0 \\ \text{s.t. } U_i(\mathbf{z}, t) - U_i(\mathbf{x}_{\leftarrow -i}, t+1) = 0 \end{cases} \tag{6}$$

where $\mathbf{x}$ and t, and g is the same set of functions of the problem (4). We want to prove that $\mathbf{x}_{\leftarrow -i}$ is a minimum for (6). We have to find $(\lambda^0_{-i}, \lambda_{-i}) > 0$ and $(\lambda^0_{-i}, \lambda_{-i}, \mu_{-i}) \neq 0$ such that:

$$\begin{cases} -\lambda^0_{-i}\nabla U_{-i}(\overline{\mathbf{z}}, t) + \lambda_{-i}\nabla g(\mathbf{x}_{\leftarrow -i}) + \mu_{-i}\nabla U_i(\mathbf{x}_{\leftarrow -i}, t) = 0 \\ g(\mathbf{x}_{\leftarrow -i})\lambda_{-i} = 0 \end{cases}$$

We set $\lambda^0_{-i} = -\mu_i$, $\lambda_{-i} = \lambda_i$, and $\mu_{-i} = \lambda^0_i$. Notice that $\lambda^0_{-i} > 0$ and $\lambda_{-i} > 0$. Thus, by Karush-Khun-Tucker, $\mathbf{x}_{\leftarrow -i}$ is a minimum for U_{-i} along $U_i(\mathbf{z}, t) = U_i(\mathbf{x}_{\leftarrow -i}, t)$. Thus, the sub-level set S_i and S_{-i} have not common intersection in A but $\overline{S}_i \cap \overline{S}_{-i}$. This means that $\mathbf{x}_{\leftarrow -i}$ is not Pareto-dominated in A in the subspace at time t. Similarly to the case (*a*), if U_i and U_{-i} are concave, their sub-level sets are sub-sets of the sub-level sets of linear functions with the same gradients. Then $\mathbf{x}_{\leftarrow -i}$ is not Pareto-dominated in A in the subspace at time t. □