On Maximal Classes of Utility Functions for Efficient one-to-one Negotiation

Yann Chevaleyre
LAMSADE
Université Paris-Dauphine
chevaley@lamsade.dauphine.fr

Ulle Endriss
Department of Computing
Imperial College London
ue@doc.ic.ac.uk

Nicolas Maudet
LAMSADE
Université Paris-Dauphine
maudet@lamsade.dauphine.fr

Abstract

We investigate the properties of an abstract negotiation framework where agents autonomously negotiate over allocations of discrete resources. In this framework, reaching an optimal allocation potentially requires very complex multilateral deals. Therefore, we are interested in identifying classes of utility functions such that any negotiation conducted by means of deals involving only a single resource at at time is bound to converge to an optimal allocation whenever all agents model their preferences using these functions. We show that the class of modular utility functions is not only sufficient but also maximal in this sense.

1 Introduction

The problem of *discrete* resource allocation has recently received much attention from the artificial intelligence community. A large amount of this work is focused on *combinatorial auctions* [Cramton *et al.*, 2005]. In this case, the allocation procedure is centralised, and the so-called *winner determination problem* consists in determining the allocation of resources maximising the auctioneer's revenue.

A different perspective is taken when one assumes that the allocation process is truly distributed, in the sense that agents autonomously negotiate over the bundles of resources they hold. This assumption is justified in many applications where no central authority can be relied on to decide upon the allocation of resources. In this case, the system designer will typically seek to set up the system in such way that it guarantees desirable properties, without directly interfering in the negotiation process itself. In this paper we will make use of such an abstract negotiation framework investigated by a number of authors [Sandholm, 1998; Endriss *et al.*, 2003; Dunne *et al.*, 2005].

To make things more precise, we assume a set of negotiating agents populating the system, and we model their preferences (over different bundles of resources) by means of utility functions. In order to pursue their own interests, agents agree on deals benefitting themselves but without planning ahead (*i.e.* they are both rational and myopic), thereby modifying the allocation of resources. From a global point of view, the quality of an allocation reflects the overall performance of the system, and the designer will naturally seek to ensure that negotiation converges towards an optimal allocation.

Section 2 introduces the negotiation framework used in this paper. As we shall recall in Section 3, it is known that very complex multilateral deals are potentially required to reach an optimal allocation. When deals are restricted (*e.g.* to a limited number of resources), it is only possible to guarantee an optimal outcome by also restricting the negotiation process to agents whose preferences have certain properties. In this paper, we study the conditions under which negotiation conducted by means of the simplest deals, involving one item at a time (or *1-deal negotiation* for short) still allows us to reach an optimal allocation. Section 4 generalises a result from the literature and shows that modelling preferences with *modular utility functions* is a sufficient condition. However, modularity is not a *necessary* condition. This is demonstrated in Section 5 by means of a counterexample. We also show that there can be *no* condition on utility functions that would be both necessary and sufficient for optimal allocations to be negotiable by means of rational 1-deals. The main contribution of this paper is to show that the class of modular utility functions is *maximal*, in the sense that no class strictly including the modular utility functions would still be sufficient for 1-deal negotiation. The proof detailed in Section 6 shows that, given any non-modular utility function, it is always possible to construct a modular utility function and select a scenario where the optimal allocation cannot be reached by 1-deals. Section 7 concludes.

2 Myopic negotiation over resources

In this section, we introduce the decentralised negotiation framework used throughout this paper and report a number of known technical results. In this framework, a finite set of agents negotiate over a finite set of discrete (*i.e.* non-divisible) resources. A resource *allocation* is a partitioning of the resources amongst the agents (that is, every resource has to be allocated to one and only one agent). As an example, the allocation A defined by $A(i) = \{r_1\}$ and $A(j) = \{r_2, r_3\}$ would allocate resource r_1 to agent i, while resources r_2 and r_3 would be owned by agent j.

We are going to model the preferences of agents by means of *utility functions* mapping bundles of resources to real numbers. Assuming that agents are only concerned with resources they personally own, we will use the abbreviation $u_i(A)$ for

$u_i(A(i))$, representing the utility value assigned by agent i to the bundle it holds for allocation A. The parameters of a negotiation problem are summarised in the following definition:

Definition 1 (Negotiation problems) *A negotiation problem is a tuple $\mathcal{P} = \langle \mathcal{R}, \mathcal{A}, \mathcal{U}, A_0 \rangle$, where*

- *$\mathcal{R}$ is a finite set of indivisible resources;*
- *$\mathcal{A} = \{1, \ldots, n\}$ is a finite set of agents ($n \geq 2$);*
- *$\mathcal{U} = \langle u_1, \ldots, u_n \rangle$ is a vector of utility functions, such that for all $i \in \mathcal{A}$, u_i is a mapping from $2^{\mathcal{R}}$ to $\mathbb{R}$;*
- *$A_0 : \mathcal{A} \to 2^{\mathcal{R}}$ is an (initial) allocation.*

Agents may agree on a *deal* to exchange some of the resources they possess. It transforms the current allocation of resources A into a new allocation A'; that is, we can define a deal as a pair $\delta = (A, A')$ of allocations (with $A \neq A'$).

We should stress that this is a *multilateral* negotiation framework. A single deal may involve the displacement of any number of resources between any number of agents. An actual implementation of this abstract framework may, however, not allow for the same level of generality. Sandholm [1998] has proposed a typology of different types of deals, such as *swap deals* involving an exchange of single resources between two agents or *cluster deals* involving the transfer of a set of items from one agent to another. The simplest type of deals are those involving only a single resource (and thereby only two agents).

Definition 2 (1-deals) *A 1-deal is a deal $\delta = (A, A')$ resulting in the reallocation of exactly one resource.*

The above is a condition on the *structure* of a deal. Other conditions relate to the *acceptability* of a deal to a given agent. We assume that agents are *rational* in the sense of aiming to maximise their individual welfare. Furthermore, agents are assumed to be *myopic*. This means that agents will not accept deals that would reduce their level of welfare, not even temporarily, because they are either not sufficiently able to plan ahead or not willing to take the associated risk (see also [Sandholm, 1998] for a justification of such an agent model in the context of multiagent resource allocation). We will, however, permit agents to enhance deals with *monetary side payments*, in order to compensate other agents for a possible loss in utility. This can be modelled using a *payment function* $p : \mathcal{A} \to \mathbb{R}$. Such a function has to satisfy the side constraint $\sum_{i \in \mathcal{A}} p(i) = 0$, *i.e.* the overall amount of money in the system remains constant. If $p(i) > 0$, then agent i *pays* the amount of $p(i)$, while $p(i) < 0$ means that it *receives* the amount of $-p(i)$. The following rationality criterion will define the acceptability of deals:

Definition 3 (Individual rationality) *A deal $\delta = (A, A')$ is rational iff there exists a payment function p such that $u_i(A') - u_i(A) > p(i)$ for all $i \in \mathcal{A}$, except possibly $p(i) = 0$ for agents i with $A(i) = A'(i)$.*

From a system designer's perspective, we are interested in assessing the well-being of the whole society, or *social welfare* [Arrow *et al.*, 2002], which is often defined as the sum of utilities of all the agents.

Definition 4 (Social welfare) *The social welfare $sw(A)$ of an allocation A is defined as follows:*

$$sw(A) = \sum_{i \in \mathcal{A}} u_i(A)$$

This is the *utilitarian* definition of social welfare. While this is the definition usually adopted in the multiagent systems literature [Wooldridge, 2002], we should stress that also several of the other notions of social welfare developed in the social sciences (e.g. egalitarian social welfare [Arrow *et al.*, 2002]) do have potential applications in the context of multiagent resource allocation.

We conclude this background section by recalling two important results [Sandholm, 1998; Endriss *et al.*, 2003]: the first one makes explicit the connection between the local decisions of agents and the global behaviour of the system, and the second one is the fundamental convergence theorem for this negotiation framework.

Lemma 1 (Individual rationality and social welfare) *A deal $\delta = (A, A')$ is rational iff $sw(A) < sw(A')$.*

Theorem 1 (Maximising social welfare) *Any sequence of rational deals will eventually result in an allocation of resources with maximal social welfare.*

The main significance of this latter result, beyond the equivalence of rational deals and social welfare-increasing deals stated in Lemma 1, is that *any* sequence of deals satisfying the rationality criterion will eventually converge to an optimal allocation. There is no need for agents to consider anything but their individual interests. Every single deal is bound to increase social welfare and there are no local minima.

3 Negotiating over one item at a time

While Theorem 1 shows that, in principle, it is always possible to negotiate an allocation of resources that is optimal from a social point of view, deals involving any number of agents and resources may be required to do so [Sandholm, 1998; Endriss *et al.*, 2003]. In particular, the most basic type of deal, which involves moving a single resource from one agent to another and which is the type of deal implemented in most systems realising a kind of *Contract Net* protocol [Smith, 1980], is certainly *not* sufficient for negotiation between agents that are not only rational but also myopic.

This has first been shown by Sandholm [1998] and is best explained by means of an example.[1] Let $\mathcal{A} = \{1, 2, 3\}$ and $\mathcal{R} = \{r_1, r_2, r_3\}$. Suppose the utility functions of these agents are defined as follows (over singleton sets):

$$\begin{array}{lll} u_1(\{r_1\}) = 5 & u_1(\{r_2\}) = 1 & u_1(\{r_3\}) = 0 \\ u_2(\{r_1\}) = 0 & u_2(\{r_2\}) = 5 & u_2(\{r_3\}) = 1 \\ u_3(\{r_1\}) = 1 & u_3(\{r_2\}) = 0 & u_3(\{r_3\}) = 5 \end{array}$$

Furthermore, for any bundle R not listed above, suppose $u_i(R) = 0$ for all $i \in \mathcal{A}$. Let A_0 with $A_0(1) = \{r_2\}$, $A_0(2) = \{r_3\}$ and $A_0(3) = \{r_1\}$ be the initial allocation,

[1] A methodology for constructing such examples is easily generated from the proof of the result on the insufficiency of any kind of structurally limited class of deals given by Endriss *et al.* [2003].

i.e. $sw(A_0) = 3$. The optimal allocation would be A^* with $A^*(1) = \{r_1\}$, $A^*(2) = \{r_2\}$ and $A^*(3) = \{r_3\}$, which yields a social welfare of 15. All other allocations have lower social welfare than A^*. Hence, starting from A_0, the deal $\delta = (A_0, A^*)$ would be the only deal increasing social welfare. By Lemma 1, δ would also be the only rational deal. This deal, however, involves all three resources and affects all three agents. In particular, δ is not a 1-deal. Hence, if we choose to restrict ourselves to *rational* deals, then 1-deals are not sufficient to negotiate allocations of resources with maximal social welfare.

Of course, for some particular negotiation problems, rational 1-deals *will* be sufficient. The difficulty lies in recognising the problems where this is so. Closely related to this issue, Dunne *et al.* [2005] have shown that, given two allocations A and A' with $sw(A) < sw(A')$, the problem of checking whether it is possible to reach A' from A by means of a sequence of rational 1-deals is NP-hard in the number of resources in the system.

The structural complexity of deals required to be able to guarantee socially optimal outcomes partly stems from the generality of the framework. In particular, so far we have made no assumptions on the structure of utility functions used by the agents to model their preferences. By introducing restrictions on the class of admissible utility functions, it may be possible to ensure convergence to an allocation with maximal social welfare by means of simpler deals. In this paper, we are interested in characterising more precisely those classes of utility functions that permit 1-deal negotiation.

Definition 5 (1-deal negotiation) *A class $\mathcal{C}$ of utility functions is said to permit 1-deal negotiation iff any sequence of rational 1-deals will eventually result in an allocation of resources with maximal social welfare whenever all utility functions $\{u_1, \ldots, u_n\}$ are drawn from $\mathcal{C}$.*

Under this perspective, a relevant result is due to Endriss *et al.* [2003], who show that rational 1-deals are sufficient to guarantee outcomes with maximal social welfare in case all agents use *additive* utility functions.[2] We are going to prove a slight generalisation of this result in the next section.

4 Modular functions are sufficient

We are now going to define the class of *modular* utility functions. This is an important (see e.g. [Rosenschein and Zlotkin, 1994]), albeit simple, class of functions that can be used in negotiation domains where there are no synergies (neither complementaries nor substitutables) between different resources.

Definition 6 (Modular utility) *A utility function u is modular iff the following holds for all bundles $R_1, R_2 \subseteq \mathcal{R}$:*

$$u(R_1 \cup R_2) \quad = \quad u(R_1) + u(R_2) - u(R_1 \cap R_2) \qquad (1)$$

The class of modular functions includes the aforementioned additive functions. This may be seen as follows. Let R be any non-empty bundle of resources and let $r \in R$. Then equation (1) implies $u(R) = u(R \setminus \{r\}) + [u(\{r\}) - u(\{\,\})]$.

[2] A utility function is additive iff the utility assigned to a set of resources is always the sum of utilities assigned to its members.

If we apply this step recursively for every resource in R, then we end up with the following equation:

$$u(R) \quad = \quad u(\{\,\}) + \sum_{r \in R} [u(\{r\}) - u(\{\,\})] \qquad (2)$$

That is, in case $u(\{\,\}) = 0$, the utility assigned to a set will be the sum of utilities assigned to its members (*i.e.* u will be additive). Clearly, equation (2) also implies equation (1), *i.e.* the two characterisations of the class of modular utility functions are equivalent.

It turns out that in domains where all utility functions are modular, it is always possible to reach a socially optimal allocation by means of a sequence of rational deals involving only a single resource each. This is a slight generalisation of a result proved by Endriss *et al.* [2003], and our proof closely follows theirs.

Theorem 2 (Negotiation in modular domains) *The class $\mathcal{M}$ of modular utility functions permits 1-deal negotiation.*

Proof. By Lemma 1, any rational deal results in a strict increase in social welfare. Together with the fact that the number of distinct allocations is finite, this ensures that there can be no infinite sequence of rational deals (termination). It therefore suffices to show that for any allocation that does not have maximal social welfare there still exists a rational 1-deal that would be applicable.

We are going to use the alternative characterisation of modular utility functions given by equation (2). For any allocation A, let f_A be the function mapping each resource r to the agent i that holds r in situation A. Then, for modular domains, the formula for social welfare (see Definition 4) can be rewritten as follows:

$$sw(A) \quad = \quad \sum_{i \in \mathcal{A}} u_i(\{\,\}) + \sum_{r \in \mathcal{R}} u'_{f_A(r)}(\{r\})$$

with $u'_i(R) = u_i(R) - u_i(\{\,\})$. Now assume we have reached an allocation of resources A that does not have maximal social welfare, *i.e.* there exists another allocation A' with $sw(A) < sw(A')$. Considering the above definition of social welfare and observing that $\sum_{i \in \mathcal{A}} u_i(\{\,\})$ is a constant that is independent of the current allocation, this implies that at least one resource r must satisfy the inequation $u'_{f_A(r)}(\{r\}) < u'_{f_{A'}(r)}(\{r\})$, *i.e.* the agent owning r in allocation A values that resource less than the agent owning it in allocation A'. But then the 1-deal consisting of passing r from agent $f_A(r)$ to agent $f_{A'(r)}$ would already increase social welfare and thereby be rational. □

Like Theorem 1, the above establishes an important convergence result towards a global optimum by means of decentralised negotiation between self-interested agents. In addition, provided all utility functions are modular, convergence can be guaranteed by means of a much simpler negotiation protocol, which only needs to cater for agreements on 1-deals (rather than multilateral deals over sets of resources).

5 Modular functions are not necessary

In the previous section we have introduced a class of utility functions (namely modular functions) such that it is possible to guarantee that sequences of rational 1-deals will converge to an allocation with maximal social welfare under the condition that *all* agents' utilities belong to this class. A natural question to ask would then be whether modularity is also a *necessary* condition in this sense.

It turns out that this is not the case. We demonstrate this by means of the following example. Suppose $\mathcal{R} = \{r_1, r_2\}$ and there are two agents with utility functions u_1 and u_2:

$$\begin{array}{lclllcl} u_1(\{\,\}) & = & 90 & & u_2(\{\,\}) & = & 90 \\ u_1(\{r_1\}) & = & 93 & & u_2(\{r_1\}) & = & 90 \\ u_1(\{r_2\}) & = & 95 & & u_2(\{r_2\}) & = & 90 \\ u_1(\{r_1, r_2\}) & = & 98 & & u_2(\{r_1, r_2\}) & = & 50 \end{array}$$

While u_1 is a modular function, u_2 is not. The optimal allocation is the allocation where agent 1 owns both items. Furthermore, as may easily be checked, any 1-deal that involves moving a single resource from agent 2 to agent 1 is rational. Hence, rational 1-deals are sufficient to move to the optimal allocation for this scenario, despite u_2 not being modular.

In fact, it is possible to show that there can be no class of utility functions that would be both sufficient and necessary in this sense. It suffices to produce two concrete utility functions u_1 and u_2 such that (i) both of them would guarantee convergence if all agents were using them, and (ii) there is a scenario where some agents are using u_1 and others u_2 and convergence is not guaranteed. This is so, because assuming that a necessary and sufficient class exists, (i) would imply that both u_1 and u_2 belong to that class, while (ii) would entail the contrary. We give two such functions for the case of two agents and two resources (the argument is easily augmented to the general case):

$$\begin{array}{lclllcl} u_1(\{\,\}) & = & 0 & & u_2(\{\,\}) & = & 0 \\ u_1(\{r_1\}) & = & 1 & & u_2(\{r_1\}) & = & 5 \\ u_1(\{r_2\}) & = & 2 & & u_2(\{r_2\}) & = & 5 \\ u_1(\{r_1, r_2\}) & = & 3 & & u_2(\{r_1, r_2\}) & = & 5 \end{array}$$

The function u_1 is modular, *i.e.* all agents using that function is a sufficient condition for guaranteed convergence to an optimal allocation by means of rational 1-deals (Theorem 1). Clearly, convergence is also guaranteed if all agents are using u_2. However, if the first agent uses u_1 and the second u_2, then the allocation A with $A(1) = \{r_1\}$ and $A(2) = \{r_2\}$ is not socially optimal and the only deal increasing social welfare (and thereby, the only rational deal) would be to swap the two resources simultaneously. Hence, no condition on all agents' utility functions can be both sufficient and necessary to guarantee convergence to an optimal allocation by means of rational 1-deals alone.

Our argument for the inexistence of any such necessary and sufficient condition has directly exploited the fact that we were looking for a *single* condition to be met by the utility functions of *all* agents. The problem could be circumvented by looking for suitable conditions on negotiation problems as a whole, where different utility functions may meet different such conditions. Clearly, such a condition *does* exist. However, the aforementioned result of Dunne *et al.* [2005] on the NP-hardness of checking whether there exists a path of rational 1-deals between two given allocations immediately suggests that verifying whether a given negotiation problem meets any such condition would be intractable.

6 The modular class is maximal

We are now going to prove the main result of this paper, namely the surprising fact that the class of modular utility functions is not only sufficient for 1-deal negotiation but also *maximal* in the sense that no class of utility functions strictly including the modular functions would still be sufficient for 1-deal negotiation. The significance of this result can only be fully appreciated when considered together with the "negative" result on necessary and sufficient conditions discussed in the previous section.

Before stating the main result, we prove the following auxiliary lemma:

Lemma 2 (Alternative characterisation of modularity) *A utility function u is modular iff the following holds for all $R \subseteq \mathcal{R}$ and all $r_1, r_2 \in \mathcal{R}$ with $r_1, r_2 \notin R$ and $r_1 \neq r_2$:*

$$u(R \cup \{r_1, r_2\}) = u(R \cup \{r_1\}) + u(R \cup \{r_2\}) - u(R) \quad (3)$$

Proof. To show this, let us recall elementary facts about submodular functions. A function $v : \mathcal{R} \to \mathbb{R}$ is submodular iff $\forall R_1, R_2 \subseteq \mathcal{R}, v(R_1) + v(R_2) \geq v(R_1 \cup R_2) + v(R_1 \cap R_2)$. It is also known that v is submodular iff $v(R \cup \{r_1\}) + v(R \cup \{r_2\}) \geq v(R \cup \{r_1, r_2\}) - v(R)$ for any $R \subseteq \mathcal{R}, r_1, r_2 \in \mathcal{R} \backslash R$, with $r_1 \neq r_2$ [Nemhauser and Wolsey, 1988, p.662]. Because a function u is modular iff both u and $-u$ are submodular, the lemma holds. □

We are now in a position to prove our theorem on the maximality of the class of modular utility functions with respect to rational negotiation over one resource at a time:

Theorem 3 (Maximality) *Let $\mathcal{M}$ be the class of modular utility functions. Then for any class of utility functions $\mathcal{F}$ such that $\mathcal{M} \subset \mathcal{F}$, $\mathcal{F}$ does not permit 1-deal negotiation.*

Proof. First observe that for $|\mathcal{R}| \leq 1$, *any* utility function is modular, *i.e.* the theorem holds vacuously in these cases. Therefore, without loss of generality, from now on we assume that there are at least two distinct resources in the system.

The proof is constructive. We will show that for any non-modular utility function u_1 on m resources, it is possible to construct a modular utility function u_2 (with $u_i \equiv 0$ for all other agents i) and an initial allocation such that no optimal allocation can be reached by means of 1-deals. This implies that $\mathcal{M} \cup \{u_1\}$ does not permit 1-deals.

Because u_1 is non-modular, Lemma 2 can be applied in the following way: there exist a bundle X and distinct resources $r_1, r_2 \notin X$ such that ϵ, defined as follows, is not equal to 0:

$$\epsilon = u_1(X \cup \{r_1\}) + u_1(X \cup \{r_2\}) - u_1(X) - u_1(X \cup \{r_1, r_2\}) \quad (4)$$

From now on, $A_{12|}$, $A_{|12}$, $A_{1|2}$ and $A_{2|1}$ will refer to allocations in which r_1 and r_2 belong to one of the first two agents, and in which resources in X are owned by 1, and resources in $Y = \mathcal{R} \backslash (X \cup \{r_1, r_2\})$ by 2, as shown in the following table.

	Agent 1	Agent 2
$A_{12\mid}$	$\{r_1, r_2\} \cup X$	Y
$A_{\mid 12}$	X	$\{r_1, r_2\} \cup Y$
$A_{1\mid 2}$	$\{r_1\} \cup X$	$\{r_2\} \cup Y$
$A_{2\mid 1}$	$\{r_2\} \cup X$	$\{r_1\} \cup Y$

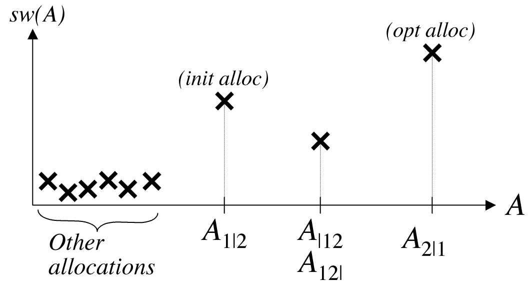

Figure 1: Values of sw for the four allocations (case $\epsilon > 0$).

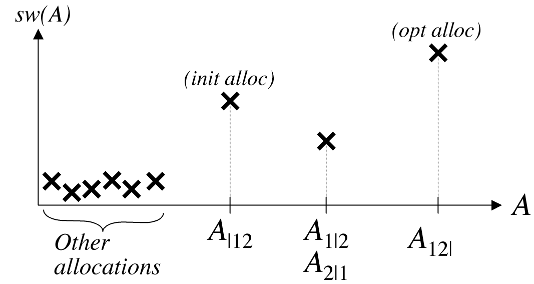

Figure 2: Values of sw for the four allocations (case $\epsilon < 0$).

Let us build a modular utility function u_2 defined as follows: $\forall R \in \mathcal{R}$,

$$u_2(R) = \sum_{r \in \{r_1, r_2\} \cap R} \alpha_r + \sum_{r \in R \cap Y} \omega - \sum_{r \in R \cap X} \omega \quad (5)$$

with $\omega = 14 \times \max|u_1| + 1$. Let $\Omega = u_2(Y) = |Y| \times \omega$. As the rest of the proof shall reveal, the value of ω has been chosen such that the social welfare of each of these four allocations is greater than that of any other allocation. Of course, this will imply that the optimal allocation has to be among these four. The values of α_{r_1} and α_{r_2} will be chosen later. The social welfare of each of these four allocations can then be written as follows:

$$\begin{aligned}
sw(A_{|12}) &= \Omega + \alpha_{r_1} + \alpha_{r_2} + u_1(X) \\
sw(A_{12|}) &= \Omega + u_1(X \cup \{r_1, r_2\}) \\
sw(A_{1|2}) &= \Omega + \alpha_{r_2} + u_1(X \cup \{r_1\}) \\
sw(A_{2|1}) &= \Omega + \alpha_{r_1} + u_1(X \cup \{r_2\})
\end{aligned}$$

It remains to be shown that depending on the value of ϵ, we can always choose an initial allocation among these four and values of α_{r_1} and α_{r_2} such that (1) this initial allocation does not have optimal social welfare, (2) there is only one rational deal from this allocation, (3) this deal leads to the optimal allocation but however (4) this rational deal would involve more than one resource. We will have to consider two cases for equation (4): the case of $\epsilon > 0$ and the case of $\epsilon < 0$.

(1st case) Suppose $\epsilon > 0$. Let us choose $\alpha_{r_1} = u_1(X \cup \{r_1\}) - u_1(X) - \frac{\epsilon}{4}$ and $\alpha_{r_2} = u_1(X \cup \{r_1, r_2\}) - u_1(X \cup \{r_1\}) + \frac{\epsilon}{4}$.

Let us first show that the four allocations have a greater social welfare than any other. With the help of equation (4), observe that both $|\alpha_{r_1}|$ and $|\alpha_{r_2}|$ are less than $3 \times \max|u_1|$. Thus, all four allocations have a social welfare of at least $\Omega - |\alpha_{r_1}| - |\alpha_{r_2}| - \max|u_1| \geq \Omega - 7 \times \max|u_1| > \Omega - \frac{\omega}{2}$. All other allocations have a social welfare lower than $\Omega - \omega + |\alpha_{r_1}| + |\alpha_{r_2}| + \max|u_1| \leq \Omega - \omega + 7 \times \max|u_1| < \Omega - \frac{\omega}{2}$. Thus, the social welfare of each of the four allocations is greater than that of any other allocation.

Now let us show that $A_{2|1}$ is the optimal allocation, as illustrated in Figure 1. More precisely, let us show that $sw(A_{|12}) < sw(A_{1|2})$, that $sw(A_{12|}) < sw(A_{1|2})$ and that $sw(A_{1|2}) < sw(A_{2|1})$. By substituting the values of α_{r_1} and α_{r_2} and using equation (4), the social welfare of each allocation can be written as follows:

$$\begin{aligned}
sw(A_{|12}) &= \Omega + u_1(X \cup \{r_1, r_2\}) \\
sw(A_{12|}) &= \Omega + u_1(X \cup \{r_1, r_2\}) \\
sw(A_{1|2}) &= \Omega + u_1(X \cup \{r_1, r_2\}) + \frac{\epsilon}{4} \\
sw(A_{2|1}) &= \Omega + u_1(X \cup \{r_1\}) + u_1(X \cup \{r_2\}) \\
&\quad -u_1(X) - \frac{\epsilon}{4} \\
&= \Omega + u_1(X \cup \{r_1, r_2\}) + \frac{3}{4}\epsilon
\end{aligned}$$

Here, $A_{2|1}$ is clearly the optimal allocation. If we choose $A_{1|2}$ as the initial allocation, then the only 1-deals involving resources r_1 or r_2 are $\delta(A_{1|2}, A_{12|})$ and $\delta(A_{1|2}, A_{|12})$. These deals decrease social welfare, and thus are not individually rational by Lemma 1. Thus, it is not possible to reach the optimal allocation $A_{2|1}$ starting from $A_{1|2}$ using only 1-deals.

(2nd case) Suppose $\epsilon < 0$. Let us choose $\alpha_1 = u_1(X \cup \{r_1\}) - u_1(X) - \frac{\epsilon}{4}$ and $\alpha_2 = u_1(X \cup \{r_2\}) - u_1(X) - \frac{\epsilon}{4}$.

Note that again, both $|\alpha_{r_1}|$ and $|\alpha_{r_2}|$ are less than $3 \times \max|u_1|$. Thus, by the same argument as in the first case, the four allocations all have greater social welfare than any other allocation.

The optimal allocation is now $A_{12|}$. To see this, let us show that $sw(A_{1|2}) < sw(A_{|12})$, that $sw(A_{2|1}) < sw(A_{|12})$, and that $sw(A_{|12}) < sw(A_{12|})$ as illustrated in Figure 2.

$$\begin{aligned}
sw(A_{|12}) &= \Omega + u_1(X \cup \{r_1\}) + u_1(X \cup \{r_2\}) \\
&\quad -u_1(X) - \frac{\epsilon}{2} \\
sw(A_{12|}) &= \Omega + u_1(X \cup \{r_1, r_2\}) \\
&= \Omega + u_1(X \cup \{r_1\}) + u_1(X \cup \{r_2\}) \\
&\quad -u_1(X) - \epsilon \\
sw(A_{1|2}) &= \Omega + u_1(X \cup \{r_1\}) + u_1(X \cup \{r_2\}) \\
&\quad -u_1(X) - \frac{\epsilon}{4} \\
sw(A_{2|1}) &= \Omega + u_1(X \cup \{r_1\}) + u_1(X \cup \{r_2\}) \\
&\quad -u_1(X) - \frac{\epsilon}{4}
\end{aligned}$$

Here, $A_{12|}$ is clearly the optimal allocation. If we choose $A_{|12}$ as the initial allocation, then the only 1-deals involving

r_1 or r_2 are $\delta(A_{|12}, A_{1|2})$ and $\delta(A_{|12}, A_{2|1})$. These deals decrease social welfare, and thus are not individually rational by Lemma 1. Thus, it is not possible to reach the optimal allocation $A_{12|}$ starting from $A_{|12}$ using only 1-deals. □

Why is this result significant? As argued earlier, while the full abstract negotiation framework introduced at the beginning of this paper would be difficult to implement, designing a system that only allows for pairs of agents to agree on deals over one resource at a time is entirely feasible. As we would like to be able to guarantee socially optimal outcomes in as many cases as possible, also for such a restricted negotiation system, we would like to be able to identify the largest possible class of utility functions for which such a guarantee can be given. However, our discussion in Section 5 has shown that there can be no class of utility functions that *exactly* characterises the class of negotiation problems for which negotiating socially optimal allocations by means of rational 1-deals is always possible. Still, there *are* classes of utility functions that permit 1-deal negotiation. As shown by Theorem 2, the class of modular functions is such a class and it is a very natural class to consider. An obvious question to ask is therefore whether this class can be enlarged in any way without losing the desired convergence property.

Our proof of Theorem 3 settles this question by giving a negative answer: For any agent with a non-modular utility function there exist modular utility functions (for the other agents) and an initial allocation such that rational 1-deals alone do not suffice to negotiate an allocation of resources with maximal social welfare. There may well be further such classes (that are both sufficient and maximal), but the class of modular functions is one that is particularly natural and useful in the context of modelling agent preferences.

7 Conclusion

This paper makes a contribution to the theoretical analysis of a negotiation framework where rational but myopic agents agree on a sequence of deals regarding the reallocation of a number of discrete resources. We have shown that the use of *modular* utility functions to model agent preferences is a *sufficient* condition to guarantee final allocations with maximal social welfare in case agents only negotiate *1-deals* (involving one resource each). We have then seen that this is, however, not a *necessary* condition for optimal outcomes and, indeed, there can be no condition on (individual) utility functions that would be both necessary and sufficient in this sense. We have therefore concentrated on showing that the class of modular functions is *maximal*, *i.e.* no strictly larger class of functions would still permit an optimal allocation to be found by means of rational 1-deals in all cases.

We consider this not only a surprising result, but also a useful characterisation of negotiation domains that can be handled reliably using simple negotiation protocols, catering only for *Contract Net*-like deals over single items between pairs of agents, rather than the full range of multilateral deals forseen in the abstract framework. Such theoretical results affect both the design of agents and of negotiation mechanisms. For instance, if a given mechanism can only handle 1-deals, then it may be inappropriate to design myopic agents with very rich preference structures to use such a mechanism.

In a companion paper [Chevaleyre *et al.*, 2005], we prove a generalisation of Theorem 1 which shows that rational deals involving at most k resources each are sufficient for convergence to an optimal allocation in case all utility functions are *additively separable* with respect to a common partition of $\mathcal{R}$ (*i.e.* synergies across different parts of the partition are not possible and overall utility is defined as the sum of utilities for the different sets in the partition [Fishburn, 1970]), and each set in this partition has at most k elements. The arguments against the existence of sufficient conditions for negotiation over k items at a time that are also necessary generalise in the expected manner. An important issue that remains to be investigated in the future therefore is to see whether it is possible to derive a similar maximality property as the one proved in this paper for this richer class of utility functions.

Another topic for future work would be to investigate what classes of utility functions are sufficient and maximal for negotiating socially optimal allocations by means of 1-deals without side payments.

References

[Arrow *et al.*, 2002] K. J. Arrow, A. K. Sen, and K. Suzumura, editors. *Handbook of Social Choice and Welfare*, volume 1. North-Holland, 2002.

[Chevaleyre *et al.*, 2005] Y. Chevaleyre, U. Endriss, J. Lang, and N. Maudet. Negotiating over small bundles of resources. In *Proc. AAMAS-2005*. ACM Press, 2005.

[Cramton *et al.*, 2005] P. Cramton, Y. Shoham, and R. Steinberg, editors. *Combinatorial Auctions*. MIT Press, 2005. To appear.

[Dunne *et al.*, 2005] P. E. Dunne, M. Wooldridge, and M. Laurence. The complexity of contract negotiation. *Artificial Intelligence*, 2005. To appear.

[Endriss *et al.*, 2003] U. Endriss, N. Maudet, F. Sadri, and F. Toni. On optimal outcomes of negotiations over resources. In *Proc. AAMAS-2003*. ACM Press, 2003.

[Fishburn, 1970] P. C. Fishburn. *Utility Theory for Decision Making*. John Wiley and Sons, 1970.

[Nemhauser and Wolsey, 1988] G. L. Nemhauser and L.A. Wolsey. *Integer and Combinatorial Optimization*. John Wiley and Sons, 1988.

[Rosenschein and Zlotkin, 1994] J. S. Rosenschein and G. Zlotkin. *Rules of Encounter*. MIT Press, 1994.

[Sandholm, 1998] T. W. Sandholm. Contract types for satisficing task allocation: I Theoretical results. In *Proc. AAAI Spring Symposium: Satisficing Models*, 1998.

[Smith, 1980] R. G. Smith. The contract net protocol: High-level communication and control in a distributed problem solver. *IEEE Transactions on Computers*, C-29(12):1104–1113, 1980.

[Wooldridge, 2002] M. Wooldridge. *An Introduction to MultiAgent Systems*. John Wiley and Sons, 2002.