

Simple Negotiation Schemes for Agents with Simple Preferences: Sufficiency, Necessity and Maximality

Yann Chevaleyre · Ulle Endriss · Nicolas Maudet

Received: 31 July 2008 / Revised: 5 March 2009

Abstract We investigate the properties of an abstract negotiation framework where agents autonomously negotiate over allocations of indivisible resources. In this framework, reaching an allocation that is optimal may require very complex multilateral deals. Therefore, we are interested in identifying classes of valuation functions such that any negotiation conducted by means of deals involving only a single resource at a time is bound to converge to an optimal allocation whenever all agents model their preferences using these functions. In the case of negotiation with monetary side payments amongst self-interested but myopic agents, the class of modular valuation functions turns out to be such a class. That is, modularity is a sufficient condition for convergence in this framework. We also show that modularity is not a necessary condition. Indeed, there can be no condition on individual valuation functions that would be both necessary and sufficient in this sense. Evaluating conditions formulated with respect to the whole profile of valuation functions used by the agents in the system would be possible in theory, but turns out to be computationally intractable in practice. Our main result shows that the class of modular functions is maximal in the sense that no strictly larger class of valuation functions would still guarantee an optimal outcome of negotiation, even when we permit more general bilateral deals. We also establish similar results in the context of negotiation without side payments.

Y. Chevaleyre
LAMSADE, Université Paris-Dauphine, France
E-mail: yann.chevaleyre@lamsade.dauphine.fr

U. Endriss
ILLIC, University of Amsterdam, The Netherlands
E-mail: ulle.endriss@uva.nl

N. Maudet
LAMSADE, Université Paris-Dauphine, France
E-mail: maudet@lamsade.dauphine.fr

1 Introduction

The problem of allocating a set of indivisible resources to a set of self-interested agents has recently received much attention from the Artificial Intelligence and Multiagent Systems communities [7, 19, 27] and has important applications ranging from electronic commerce and industrial procurement [17], over the joint exploitation of earth observation satellites [21], to Grid computing [15]. Much work in this area has focussed on combinatorial auctions [10]. In this case, the allocation procedure is *centralised*, and the so-called winner determination problem is the problem of determining the allocation of resources maximising the sum of the prices associated with the bids satisfied by that allocation.

A different perspective is taken when one assumes that the allocation process is *distributed*, in the sense that agents autonomously negotiate over the bundles of resources they hold and that allocations emerge over time, as a consequence of a sequence of local deals agreed upon by the agents in the system. This assumption is justified in many applications where no central authority can be relied upon to decide on the allocation of resources. This may, for instance, be the case in view of computational limitations of any potential centre, or in view of its trustworthiness. In this case, the system designer will typically seek to set up the system in such way that it nevertheless guarantees certain desirable properties, but without directly interfering in the negotiation process itself. Similar considerations have led to research areas such as distributed constraint optimisation [14] or distributed mechanism design [25, 26]. In this paper we further analyse a framework for distributed negotiation over indivisible resources recently investigated by a number of authors [30, 13, 12, 29, 4].

We assume a set of negotiating agents populating the system, and we model their preferences (over different bundles of resources) by means of valuation 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 [30]), thereby modifying the allocation of resources. Negotiation is thus a *local* process; groups of individual agents come together to exchange resources according to their individual interests, without regard for the system as a whole. From a global point of view, on the other hand, a system designer may seek to ensure that negotiation converges towards an allocation that is “optimal” according to a suitable metric. This would typically be some form of aggregation of the individual agents’ preferences, *i.e.* we can employ well-known formal tools from welfare economics and social choice theory to characterise optimal allocations. Standard examples are the notions of Pareto optimality and social welfare [2, 22]. In this paper, we are mostly going to be interested in negotiating allocations of resources that maximise *utilitarian social welfare*, *i.e.* the sum of individual agent valuations.

The work described here is complementary to the large body of literature on mechanism design and game-theoretic models of negotiation in multiagent systems (see e.g. [19, 27, 25]). Rather than analysing the incentives of individual

agents in detail, here we simply assume that agents are rational in the sense of never accepting a disadvantageous deal, but we do not assume that they will necessarily attempt to negotiate the very best outcome for themselves in every single step of a negotiation process. In fact, we assume that the system designer does not know what the precise negotiation strategies used by the agents are. All that is known is that no agent will ever accept a deal that would decrease their utility. Our interests concentrate on the interplay of the local and the global level: assuming that agents only negotiate mutually beneficial deals, how does this affect social welfare for the system as a whole?

Section 2 introduces the negotiation framework used in this paper. We are going to distinguish two variants of the framework. In the first one, agents may enhance deals with monetary side payments (*framework with money*) to make them more attractive to agents who would otherwise have no incentive to agree to a particular deal. In the second variant of the negotiation framework, agents are assumed not to have the option of using side payments (*framework without money*). It is known that very complex multilateral deals are potentially required to reach an optimal allocation [30, 13]. 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 problem space to agents whose preferences have certain properties. These properties are modelled as restrictions to the class of valuation functions that agents may use to model their preferences over alternative bundles of resources. Section 3 defines two such restrictions. The first is the well-known class of *modular* (also known as *additive*) valuation functions. The second is a further restriction on modular functions, where the valuation an agent may place on any one resource is limited to a choice between three given values (representing positive preference, negative preference, and indifference). Each choice of values for positive and negative preference gives rise to a different class of restricted valuation functions (the value for indifference is 0 in all cases).

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. The focus on 1-deals is motivated by practical concerns: Negotiation protocols for 1-deals are certainly realisable in practice (witness the well-known *Contract-Net* protocol [33]), while implementing any class of deals that is significantly more complex is very challenging. Section 4 recalls (and slightly generalises) a result from the literature that shows that modelling preferences with modular valuation functions is a *sufficient* condition for reaching an optimal allocation by means of rational (*i.e.* mutually beneficial) 1-deals. This applies to the framework with money, and a similar result is proved for negotiation without money. 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 valuation functions that would be both necessary and sufficient for optimal allocations to be negotiable by means of rational 1-deals. These results pertain to conditions on single valuation functions, to be met by the functions of all the agents in the system. Section 5 concludes by showing

that describing such a necessary and sufficient condition over the *profile* of valuation functions of an entire agent society would be possible in theory, but that it would be computationally intractable to evaluate any such condition in practice. This result is obtained as a by-product of an NP-hardness proof for the problem of checking whether, for a given profile of valuation functions, any sequence of rational 1-deals is bound to result in an allocation with maximal social welfare, for a whole range of representation languages for encoding valuation functions.

The main contribution of this paper, as far as the framework with money is concerned, is to show that the class of modular valuation functions is *maximal*, in the sense that no class strictly including the modular valuations functions would still be sufficient for 1-deal negotiation. In fact, even if we allow for any kind of *bilateral deals* (each involving two agents and any number of resources), we cannot guarantee convergence anymore as soon as valuations are drawn from any superclass of the class of modular functions. The proof detailed in Section 6 shows that, given any non-modular valuation function, it is always possible to construct modular valuation functions for the other agents and select an initial allocation of resources such that the optimal allocation cannot be reached by means of rational bilateral deals alone. This is followed by a similar result for the framework without money, where each of the aforementioned restrictions on the class of modular valuation functions turns out to be both sufficient and maximal. We also show that when we change our requirements from obtaining outcomes with maximal social welfare to obtaining outcomes that are merely Pareto optimal, then we lose the maximality property for the framework without money. Finally, Section 7 concludes.

2 Myopic negotiation over indivisible resources

In this section, we introduce the decentralised negotiation framework used throughout this paper and report a number of known technical results [13, 30].

2.1 Negotiation problems and deals

In this framework, a finite set of agents negotiate over a finite set of indivisible 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 *valuation functions* mapping bundles of resources to real numbers. The parameters of a negotiation problem are summarised in the following definition:

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

- $\mathcal{N} = \{1, \dots, n\}$ is a finite set of agents ($n \geq 2$);
- $\mathcal{R} = \{r_1, \dots, r_m\}$ is a finite set of indivisible resources;
- $\mathcal{V} = \langle v_1, \dots, v_n \rangle$ is a profile of valuation functions, such that for all $i \in \mathcal{N}$, v_i is a mapping from $2^{\mathcal{R}}$ to \mathbb{R} .

Sometimes we will talk about negotiation problems together with an *initial allocation* $A_0 : \mathcal{N} \rightarrow 2^{\mathcal{R}}$.

Observe that the value of a valuation function only depends on the resources held by the agent in question, *i.e.* agents are not concerned with resources held by other agents or any other external factors. We are frequently going to use the abbreviation $v_i(A)$ for $v_i(A(i))$, representing the value assigned by agent i to the bundle it holds in allocation A .

Agents may agree on a *deal* to exchange some of the resources they possess. Such a deal 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 $A \neq A'$). When speaking of a *sequence* of deals $\langle \delta_1, \delta_2, \delta_3, \dots \rangle$, it is understood that the input allocation for δ_2 is the output allocation for δ_1 , the input allocation for δ_3 is the output allocation for δ_2 , and so forth.

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 [30] 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 $\delta = (A, A')$ reallocating exactly one resource: $\#\{r \in \mathcal{R} \mid \exists i, j \in \mathcal{N} \text{ such that } i \neq j \text{ and } r \in A(i) \cap A'(j)\} = 1$.

Another important class of deals are those that may involve any number of resources, but only two agents at a time.

Definition 3 (Bilateral deals) A bilateral deal is a deal $\delta = (A, A')$ between exactly two agents: $\#\{i \in \mathcal{N} \mid A(i) \neq A'(i)\} = 2$.

The above are conditions 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 at maximising their individual welfare (utility/valuation). 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's work [30] for a justification of such an agent model).

2.2 Negotiating with money

In the first variant of this framework, we will permit agents to enhance deals with *monetary side payments*, in order to compensate other agents for a

possible loss in valuation. This can be modelled using a *payment function* $p : \mathcal{N} \rightarrow \mathbb{R}$. Such a function has to satisfy the side constraint $\sum_{i \in \mathcal{N}} p(i) = 0$, *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 4 (Individual rationality) *A deal $\delta = (A, A')$ is individually rational (IR) if there exists a payment function p such that $v_i(A') - v_i(A) > p(i)$ for all $i \in \mathcal{N}$, 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* [2,22], which is often defined as the sum of valuations of all the agents.

Definition 5 (Social welfare) *The social welfare $sw(A)$ of an allocation A is defined as follows:*

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

This is the *utilitarian* definition of social welfare. While this is the definition usually adopted in the multiagent systems literature [31,36], we should stress that also several of the other notions of social welfare developed in the social sciences (e.g. egalitarian social welfare [2,22]) do have applications in the context of multiagent resource allocation [7].

We now recall two important known results. The first one makes the connection between the local decisions of agents and the global behaviour of the system explicit [13]:

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

That is, side payments can be arranged in such a way that a given deal is beneficial for all the agents involved if and only if that deal increases social welfare. The second result is the fundamental convergence theorem for this negotiation framework [30]:

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

That is, there can be no infinite sequence of IR deals, and any finite sequence that does not culminate in an allocation with maximal social welfare can still be extended with a further IR deal. The main significance of this result, above and 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 allocation that is socially optimal. 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 optima the system could get stuck in.

2.3 Negotiating without money

What happens if we do *not* allow agents to enhance deals with monetary side payments (or *explicit utility transfers*), in order to compensate other agents for a possible loss in value? In this context, we shall assume that agents are *cooperatively rational* in the sense of accepting deals that may not result in a *strict* increase in personal welfare, with the further condition that at least one agent will strictly benefit from the deal.

Definition 6 (Cooperative rationality) *A deal $\delta = (A, A')$ is called cooperatively (CR) rational if $v_i(A) \leq v_i(A')$ for all $i \in \mathcal{N}$ and there exists an agent $j \in \mathcal{N}$ such that $v_j(A) < v_j(A')$.*

Note that we have $sw(A) < sw(A')$ for any deal $\delta = (A, A')$ that is CR, but *not* vice versa. Clearly, in this general setting, it is not possible to guarantee that agents will eventually reach an allocation with maximal social welfare, in the sense previously defined. A result similar to Theorem 1 can be achieved, though, if we consider the Pareto optimality criterion instead (an allocation is Pareto optimal when no other allocation is strictly better for one agent without being worse for any of the others): *any* sequence of deals satisfying the cooperative rationality criterion will eventually converge to a Pareto optimal allocation [13].

3 Modular valuation functions and variants

In this section, we introduce several classes of restricted valuation functions. These types of functions will later be used to derive more specific results regarding negotiation when agents use such functions to model their preferences. We are first going to define the well-known class of *modular* functions. This is an important (see e.g. [27]), albeit simple, class of functions that can be used in negotiation domains where there are no synergies (either complementaries or substitutables) between different resources.

Definition 7 (Modular valuation) *A valuation function v is modular if the following condition is satisfied for all bundles $R_1, R_2 \subseteq \mathcal{R}$:*

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

The class of modular functions includes the class of *additive* functions.¹ This may be seen as follows. Let R be any non-empty bundle of resources and let $r \in R$. Then equation (1) implies $v(R) = v(R \setminus \{r\}) + [v(\{r\}) - v(\emptyset)]$. If we apply this step recursively for every resource in R , then we end up with the following equation:

$$v(R) = v(\emptyset) + \sum_{r \in R} [v(\{r\}) - v(\emptyset)] \quad (2)$$

¹ A valuation function is additive if the valuation assigned to a set of resources is always the sum of valuations assigned to its members.

That is, in case $v(\emptyset) = 0$, the value given to a set will be the sum of the values of its members (*i.e.* v will be additive). Clearly, equation (2) also implies equation (1), *i.e.* the two characterisations of the class of modular valuation functions are equivalent.

A further useful characterisation of the class of modular valuation functions is the following. A valuation function v is modular iff the following holds for every bundle $R \subseteq \mathcal{R}$ and every pair of resources $r_1, r_2 \in \mathcal{R} \setminus R$:

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

The equivalence to equation (1) is readily checked (one direction is obvious; the other can be established by a simple inductive argument).

We now introduce a restriction on the class of modular valuation functions, namely the classes of *modular functions with fixed α, β -values* (or $\mathcal{M}_{\alpha, \beta}$ for short). Intuitively, $\mathcal{M}_{\alpha, \beta}$ functions are suitable in domains where agents can only like, dislike, or possibly be indifferent towards any given resource in the system. The key point is that agents all agree on the intensities used to indicate positive and negative preferences for each single resource. $\mathcal{M}_{\alpha, \beta}$ functions hence define a set of classes, each class being parametrised by the range of possible values assigned to each resource.

Definition 8 (Modular functions with fixed α, β -values) *A profile of valuation functions $\langle v_1, \dots, v_n \rangle$ is called modular with fixed α, β -values if*

- there exists a list of coefficients $\alpha = (\alpha_{r_1}, \dots, \alpha_{r_{|\mathcal{R}|}}) \in (\mathbb{R}^+)^{|\mathcal{R}|}$,
- there exists a list of coefficients $\beta = (\beta_{r_1}, \dots, \beta_{r_{|\mathcal{R}|}}) \in (\mathbb{R}^-)^{|\mathcal{R}|}$, and
- for each valuation function v_i , there exist two sets $S_i^+, S_i^- \subseteq \mathcal{R}$ with $S_i^+ \cap S_i^- = \emptyset$, such that for all bundles $R \subseteq \mathcal{R}$, $v_i(R)$ can be computed as:

$$v_i(R) = v_i(\emptyset) + \sum_{r \in R \cap S_i^+} \alpha_r + \sum_{r \in R \cap S_i^-} \beta_r$$

Given two lists of coefficients $\alpha \in (\mathbb{R}^+)^{|\mathcal{R}|}$ and $\beta \in (\mathbb{R}^-)^{|\mathcal{R}|}$, the class $\mathcal{M}_{\alpha, \beta}$ is defined as the largest set of modular functions with fixed α, β -values.

Consider, for example, the class $\mathcal{M}_{(2,1),(-20,-7)}$ over resources $\{r_1, r_2\}$, and let v_1 and v_2 be two valuation functions drawn from this class such that $v_1(\emptyset) = v_2(\emptyset) = 0$, $S_1^+ = \{r_1\}$, $S_1^- = \{r_2\}$, $S_2^+ = \{r_2\}$, and $S_2^- = \emptyset$. Then, the values of v_1 and v_2 are as shown in the following table:

R	$v_1(R)$	$v_2(R)$
\emptyset	0	0
$\{r_1\}$	2	0
$\{r_2\}$	-7	1
$\{r_1, r_2\}$	-5	1

Observe that the union of all classes of modular functions with fixed α, β -values is the class of modular functions.

4 Sufficient classes of valuation functions

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 [13, 30]. 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 [33], is certainly *not* sufficient for negotiation between agents that are not only rational but also myopic.² This is best explained by means of an example. Let $\mathcal{N} = \{1, 2, 3\}$ and $\mathcal{R} = \{r_1, r_2, r_3\}$. Suppose the valuation functions of these agents are defined as follows (over singleton sets):

$$\begin{array}{lll} v_1(\{r_1\}) = 5 & v_1(\{r_2\}) = 4 & v_1(\{r_3\}) = 0 \\ v_2(\{r_1\}) = 0 & v_2(\{r_2\}) = 5 & v_2(\{r_3\}) = 4 \\ v_3(\{r_1\}) = 4 & v_3(\{r_2\}) = 0 & v_3(\{r_3\}) = 5 \end{array}$$

Furthermore, for any bundle R not listed above, suppose $v_i(R) = 0$ for all $i \in \mathcal{N}$. Let A_0 with $A_0(1) = \{r_2\}$, $A_0(2) = \{r_3\}$ and $A_0(3) = \{r_1\}$ be the initial allocation, *i.e.* $sw(A_0) = 12$. The optimal allocation is 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 both A_0 and 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 deal that is IR. 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 IR deals, then 1-deals are not sufficient to negotiate allocations of resources with maximal social welfare. Of course, for some particular negotiation problems, IR 1-deals *will* be sufficient. The difficulty lies in recognising the problems where this is so.

The structural complexity of deals required to be able to guarantee socially optimal outcomes partly stems from the generality of the framework. By introducing restrictions on the class of admissible valuation functions, it could conceivably be the case that it is possible to ensure convergence to an allocation with maximal social welfare by means of simpler deals. In other words, if all valuation functions belong to a certain restricted class of functions, then this may turn out to be a *sufficient* condition for convergence to a social optimum by means of structurally simple deals. In this section, we establish two such results for modular valuation functions and 1-deals.

4.1 Framework with money

It turns out that in domains where all valuation functions are modular, it is always possible to reach a socially optimal allocation by means of a sequence

² Dunne and colleagues [12, 11] have analysed the computational complexity of checking whether, for a given scenario, 1-deal negotiation is sufficient (a topic which we will also address in Section 5.3). Andersson and Sandholm [1] have run several experiments to analyse the quality of allocations that can be attained using 1-deals in practice.

of rational deals involving only a single resource each. This is a slight generalisation of a result originally proved for the case of additive valuations [13], and our proof closely follows that original proof.

Theorem 2 (Convergence in modular domains) *If all valuation functions are modular, then any sequence of IR 1-deals will eventually result in an allocation with maximal social welfare.*

Proof By Lemma 1, any IR 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 IR deals (termination). It therefore suffices to show that for any allocation that does not have maximal social welfare there still exists an IR 1-deal that would be applicable.

We are going to use the alternative characterisation of modular valuation 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 5) can be rewritten as follows:

$$sw(A) = \sum_{i \in \mathcal{N}} v_i(\emptyset) + \sum_{r \in \mathcal{R}} v'_{f_A(r)}(\{r\})$$

with $v'_i(R) = v_i(R) - v_i(\emptyset)$. 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 $\sum_{i \in \mathcal{N}} v_i(\emptyset)$ is a constant that is independent of the current allocation, this implies that at least one resource r must satisfy the inequality $v'_{f_A(r)}(\{r\}) < v'_{f_{A'}(r)}(\{r\})$, *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. \square

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 valuation 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).

4.2 Framework without money

We now prove a similar result for the variant of our framework that does not allow for monetary side payments.

Theorem 3 (Convergence in $\mathcal{M}_{\alpha,\beta}$ -domains) *Let $\mathcal{M}_{\alpha,\beta}$ be a class of modular functions with fixed α, β -values. If all valuation functions are drawn from $\mathcal{M}_{\alpha,\beta}$, then any sequence of CR 1-deals will eventually result in an allocation of resources with maximal social welfare.*

Proof By Definition 6, any CR 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 CR deals (termination). It therefore suffices to show that for any allocation that does not have maximal social welfare there still exists a CR 1-deal that would be applicable. If an allocation A does not have maximal social welfare then it must be the case that some agent i holds a resource r and that there is another agent j in the system such that $v_i(\{r\}) < v_j(\{r\})$. That is, either (1) some agent i holds a resource r with $v_i(\{r\}) - v_i(\emptyset) = 0$, and there is another agent j in the system with $v_j(\{r\}) - v_j(\emptyset) = \alpha_r$, or (2) some agent i holds a resource r with $v_i(\{r\}) - v_i(\emptyset) = \beta_r$, and there is another agent j in the system such that either $v_j(\{r\}) - v_j(\emptyset) = 0$, or $v_j(\{r\}) - v_j(\emptyset) = \alpha_r$. In each case, passing r from i to j would be a CR deal, so either negotiation has not yet terminated or we are in situation with maximal social welfare. \square

Theorem 3 generalises a previous result stating that any sequence of CR 1-deals will eventually result in an allocation with maximal social welfare in case all agents are using so-called 0-1 functions to model their preferences [13]. A valuation function v_i is a 0-1 function if it is additive and $v_i(\{r\}) = 0$ or $v_i(\{r\}) = 1$ for all $r \in \mathcal{R}$. That is, the class of 0-1 functions is like the class $\mathcal{M}_{(1,1,\dots,1),(0,0,\dots,0)}$, except that the former does not allow for non-zero values to be assigned to the empty set.

5 Necessity issues

As explained earlier, the convergence results of the previous section may be considered results on *sufficient* conditions on valuation functions for convergence by means of rational 1-deals. This naturally raises the question whether there may be a corresponding set of *necessary* conditions on valuation functions. In this section, we are going to give a negative answer to this question. We are first going to show that all valuation functions belonging to the class of modular functions (which we have shown to be a sufficient condition for convergence by means of 1-deals) is not a necessary condition. Then we are going to show that, in fact, there can be *no* condition on individual valuation functions that would be both necessary and sufficient in this sense. Of course, it *is* possible to formulate a necessary and sufficient condition for convergence by means of 1-deals for *profiles* of valuation functions (*i.e.* if we are allowed to give conditions on how the valuations of different agents should relate to each other). However, as we are going to show, evaluating such a condition would be computationally intractable.³

³ While the analysis in this section applies to the framework of negotiation with money, a similar case may be made for negotiation without money.

5.1 Modularity is not necessary

Modularity is not a necessary condition for convergence by means of IR 1-deals. We demonstrate this by means of the following example. Suppose $\mathcal{R} = \{r_1, r_2\}$ and there are two agents with valuation functions v_1 and v_2 defined as follows:

$$\begin{array}{llll} v_1(\emptyset) & = 0 & v_2(\emptyset) & = 0 \\ v_1(\{r_1\}) & = 10 & v_2(\{r_1\}) & = 10 \\ v_1(\{r_2\}) & = 10 & v_2(\{r_2\}) & = 10 \\ v_1(\{r_1, r_2\}) & = 0 & v_2(\{r_1, r_2\}) & = 0 \end{array}$$

These functions are not modular. The situation is as follows: each agent is willing to hold a single resource, and has no preference as to which resource it actually holds. The optimal allocations are those where each agent holds exactly one resource. Furthermore, as may easily be checked, IR 1-deals are sufficient to move to the optimal allocation for this scenario, despite the functions not being modular. Hence, modularity cannot be a necessary condition for convergence.

5.2 There is no sufficient and necessary class

In fact, it is possible to show that there can be no class of valuation functions that would be both sufficient and necessary. It suffices to produce two concrete functions v_1 and v_2 such that (i) each of them would guarantee convergence if all agents were to use it, and (ii) there is a scenario where some agents are using v_1 and others v_2 and convergence is not guaranteed. This is so, because assuming that a necessary and sufficient class exists, (i) would imply that both v_1 and v_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):

$$\begin{array}{llll} v_1(\emptyset) & = 0 & v_2(\emptyset) & = 0 \\ v_1(\{r_1\}) & = 1 & v_2(\{r_1\}) & = 5 \\ v_1(\{r_2\}) & = 2 & v_2(\{r_2\}) & = 5 \\ v_1(\{r_1, r_2\}) & = 3 & v_2(\{r_1, r_2\}) & = 5 \end{array}$$

The function v_1 is modular, *i.e.* all agents using that function is a sufficient condition for guaranteed convergence to an optimal allocation by means of IR 1-deals (Theorem 2). Clearly, convergence is also guaranteed if all agents are using v_2 . However, if the first agent uses v_1 and the second v_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 IR deal) would be to swap the two resources simultaneously. Hence, no condition on all agents' valuation functions can be both sufficient and necessary to guarantee convergence to an optimal allocation by means of IR 1-deals alone.

5.3 Evaluating conditions on profiles of valuations functions is intractable

Our argument for the non-existence of any necessary and sufficient condition for guaranteed convergence has directly exploited the fact that we were looking for a *single* condition to be met by the valuation functions of *all* agents. The problem could be circumvented by looking for suitable conditions on negotiation problems as a whole, where different valuation functions may meet different such conditions. That is, we could try to formulate a condition for *profiles* $\langle v_1, \dots, v_n \rangle$ of valuation functions. Clearly, such a condition *does* exist: trivially, $\langle v_1, \dots, v_n \rangle$ meets that condition iff convergence to an optimal allocation is guaranteed for any initial allocation when agents are using the valuation functions $\langle v_1, \dots, v_n \rangle$. Of course, phrased like this, this condition is of no practical use. What we would require is a “simple” condition on valuation profiles. The condition of modularity (for single valuations), for instance, clearly is such a simple condition.

To make this notion of simplicity precise, we first need to fix a *representation language* for valuation functions. Then we can ask whether or not it is the case that checking a particular condition for valuation profiles represented in this language is a tractable decision problem. As we are going to argue in this section, for a wide range of representation languages, there can be no condition on valuation profiles that would be simple in this sense.

Formally, a representation language over \mathcal{R} is defined in terms of a language $\mathcal{L} \subseteq \Sigma^*$, where Σ is some finite set of symbols, and a function $f : \mathcal{L} \times 2^{\mathcal{R}} \rightarrow \mathbb{Q}$ mapping each pair of a word (representing a valuation function) and a set of resources to a rational number. A representation language is said to be *polynomial-time computable* if the computation of $f(v, R)$ can be done in polynomial time with respect to the number of symbols in v and the size of R , for any $v \in \mathcal{L}$ and $R \subseteq \mathcal{R}$. We will continue to write $v(R)$ as a shorthand for $f(v, R)$. Most widely-used languages fit this definition. We will discuss several of them below. Note that there are also exceptions though, such as the so-called OR-language, mainly used in the combinatorial auction literature [24], which is not polynomial-time computable.

Checking whether a given profile of valuation functions guarantees convergence to an allocation with maximal social welfare by means of IR 1-deals is equivalent to a decision problem that has been referred to as 1-convergence (1-CONV) in the literature [7, 11]:

Definition 9 (1-convergence problem) *Given a negotiation problem $\langle \mathcal{N}, \mathcal{R}, \mathcal{V} \rangle$, with \mathcal{V} being represented in the representation language $\mathcal{R}\mathcal{E}\mathcal{P}$, the 1-convergence problem 1-CONV($\mathcal{R}\mathcal{E}\mathcal{P}$) is the problem of checking whether any sequence of IR 1-deals (originating from any initial allocation) will eventually result in an allocation with maximal social welfare.*

1-CONV has been shown to be coNP-complete in case valuations are represented using the so-called SLP form, which encodes valuation functions as *straight-line programs* [11]. This means that checking whether a profile of valuation functions given in SLP form satisfies the necessary and sufficient

condition for successful 1-deal negotiation is also coNP-complete, hence computationally intractable. In the sequel, we are going to extend this complexity result to a whole range of representation languages, which only have to meet very weak conditions for the hardness result to apply.

We are going to establish our hardness result via a reduction from the so-called *valuation improvement problem* (VIP) problem, which is closely related to the problem of finding the maximum of a valuation function. These problems are known to be NP-hard for a wide range of representation languages.

Definition 10 (Valuation improvement problem) *Given a valuation function v , expressed in a particular representation language $\mathcal{R}\mathcal{E}\mathcal{P}$, and a bundle $Q \subseteq \mathcal{R}$, the valuation improvement problem $\text{VIP}(\mathcal{R}\mathcal{E}\mathcal{P})$ is the problem of checking whether there exists a bundle $Q' \subseteq \mathcal{R}$ such that $v(Q) < v(Q')$.*

Our proof will apply to representation languages that are *r-composed*, where r is a particular resource. Roughly, a representation language $\mathcal{R}\mathcal{E}\mathcal{P}$ satisfies this property, if for any valuation function v that can be represented in $\mathcal{R}\mathcal{E}\mathcal{P}$ we can also represent certain functions that are like v , except that they either return 0 or the same value as v reduced by a particular constant, depending on whether or not r is included in the bundle being evaluated. (We stress that this is an artificial property of a language, specifically designed to suit the reduction to follow.)

Definition 11 (*r*-composition) *Let $\mathcal{R}\mathcal{E}\mathcal{P}$ and $\mathcal{R}\mathcal{E}\mathcal{P}'$ be two representation languages for valuation functions and let $r \in \mathcal{R}$ be a resource. $\mathcal{R}\mathcal{E}\mathcal{P}$ is said to be an *r*-composition over $\mathcal{R}\mathcal{E}\mathcal{P}'$ if for all $v \in \mathcal{R}\mathcal{E}\mathcal{P}$ and for all $x \in \{v(R) \mid R \subseteq \mathcal{R}\}$, the functions \underline{v} and \bar{v} defined below belong to $\mathcal{R}\mathcal{E}\mathcal{P}'$ and the number of symbols needed to represent both functions is polynomial in $|v|$ (the number of symbols needed to encode v) and in the number of resources.⁴*

$$\underline{v}(R) = \begin{cases} v(R) - x & \text{if } r \in R \\ 0 & \text{otherwise} \end{cases} \quad \bar{v}(R) = \begin{cases} v(R) - x & \text{if } r \notin R \\ 0 & \text{otherwise} \end{cases}$$

For the sake of simplicity, we will say that a representation language $\mathcal{R}\mathcal{E}\mathcal{P}$ is *r-composed* (for some given resource r) if there exist some representation language $\mathcal{R}\mathcal{E}\mathcal{P}'$ over which $\mathcal{R}\mathcal{E}\mathcal{P}$ is an *r*-composition.

Our result will apply to any representation language $\mathcal{R}\mathcal{E}\mathcal{P}$ that is *r-composed* over some representation language $\mathcal{R}\mathcal{E}\mathcal{P}'$ for which the valuation

⁴ Note that encoding the function \underline{v} (or \bar{v}) as a sequence of bits requires encoding v , r , and x . Unfortunately, x may be a huge fractional number. Thus, if we choose to encode x directly as a sequence of bits, the total number of bits may not be polynomial in $|v|$ and $|\mathcal{R}|$. To overcome this problem, instead of encoding x , one may encode the bundle R whose value under v is x . This way, the size of \underline{v} remains polynomial in $|v|$ and $|\mathcal{R}|$. This will be of some importance in the proof to follow, as we will use a reduction to a decision problem in which the size of instances will also be polynomial in $|v|$ and $|\mathcal{R}|$. For representation languages $\mathcal{R}\mathcal{E}\mathcal{P}$ that encode the numbers used to make up the values of v directly (such as e.g. the k -additive form discussed in the sequel), this point is of course not an issue, as the complexity of representing those numbers is already being accounted for when we consider the complexity of representing v .

improvement problem $\text{VIP}(\mathcal{RE}\mathcal{P}')$ is either NP-hard or coNP-hard. As many interesting representation languages satisfy these properties, the results presented here are fairly general. Let us briefly mention some of the representation languages commonly used in multiagent resource allocation [7]. In the following, r is any resource chosen from \mathcal{R} .

- The *k-additive form* [18, 8] models a valuation function v as a set of coefficients α_T for $T \subseteq \mathcal{R}$ and $|T| \leq k$, and fixes $v(R) = \sum_{T \subseteq R} \alpha_T$. That is, α_T is the marginal value of obtaining all the items in T together, beyond the value associated with any of its strict subsets. VIP is NP-hard for the k -additive form for any $k \geq 2$ [34]. Also, for any $k \in \mathbb{N}$, the k -additive form is an r -composition over the $(k+1)$ -additive form. To see this, observe that given a representation $\{\alpha_T\}$ of v , we can represent \underline{v} by means of the coefficients $\{\alpha'_T\}$ defined as follows: $\alpha'_{\{r\}} := \alpha_\emptyset - x$, $\alpha'_{T \cup \{r\}} := \alpha_T$ for $T \neq \emptyset$, and $\alpha'_T := 0$ for all T not containing r (and similarly for \bar{v}).
- *Posiforms of degree k* [5], which are a generalisation of the k -additive form that also permit the definition of marginal values for *not* obtaining certain resources, are r -compositions over posiforms of degree $k+1$. The argument for r -composability is similar to the one given above, and NP-hardness of the VIP follows from the corresponding result for the k -additive form (or from observing that it is equivalent to a standard NP-hard combinatorial problem known as WEIGHTED MAX k -CONSTRAINT SAT [3]).
- The language of *weighted propositional formulas* associates resources with propositional variables and encodes valuation functions as sets of weighted formulas [34, 7]. The value of a bundle is then given by the sum of the weights of the formulas that are “true” over that bundle. This language is an r -composition over itself: to construct a representation of \underline{v} given a representation of v , we rewrite each weighted formula (φ, w) as $(\varphi \wedge r, w)$ and add the weighted formula $(r, -x)$. Clearly, VIP is NP-hard for this language, by virtue of an immediate reduction from SAT [3, 34]. Many sublanguages of this general framework will also satisfy r -composability and NP-hardness of the associated VIP [34]. These languages include the languages advocated by Boutilier and Hoos [6], amongst others, for the compact representation of bids in combinatorial auctions.
- The SLP (straight-line program) representation is an r -composition over itself. Intuitively, SLPs are computer programs without loops, taking a bundle of resources as input and returning the associated valuation as output. Here, VIP is also NP-hard [12].

We should also note that our result does not cover languages such as the so-called “bundle form” (which simply lists the valuation for each bundle in a large table) or the XOR-language used in combinatorial auctions [7, 24], which is like the bundle form, except that there is an implicit monotonicity assumption. These languages lack succinctness for most interesting classes of valuation functions, and as a consequence the associated VIP problem is linear in the size of the representation.

We are now ready to present our theorem on the intractability of 1-CONV. Our proof employs a polynomial-time Turing reduction (as opposed to the more widely used Karp reductions).⁵

Theorem 4 (Complexity of 1-convergence) *Let \mathcal{REP} be a polynomial-time computable representation language. If the following conditions hold, then $1\text{-CONV}(\mathcal{REP})$ is NP-hard under polynomial-time Turing reductions, even with only two agents:*

- there exist a resource $r \in \mathcal{R}$ and a representation language \mathcal{REP}' over which \mathcal{REP} is an r -composition;
- the VIP problem for this representation language \mathcal{REP}' is NP-hard; and
- the zero-valued valuation function $v_0 \equiv 0$ is representable in \mathcal{REP} .

Proof In the following, we will show NP-hardness using a Turing reduction from the VIP problem. More precisely, we will show that, given an oracle deciding the $1\text{-CONV}(\mathcal{REP})$ problem in a single time unit, we can build a deterministic algorithm solving the $\text{VIP}(\mathcal{REP}')$ problem in polynomial time. The proof is constructive, *i.e.* the algorithm will be given below.

Let $\langle v, Q \rangle$ be an instance of the VIP problem over the \mathcal{REP}' representation. Note that the number of bits required to encode such an instance is bounded from above by $O(|v| + |\mathcal{R}|)$. Let v_0 be the zero-valued valuation function, and let the set of agents be $\mathcal{N} = \{1, 2\}$. Now define two valuation functions (where r is the fixed resource mentioned in the theorem):

$$\underline{v}(R) = \begin{cases} v(R) - v(Q) & \text{if } r \in R \\ 0 & \text{otherwise} \end{cases} \quad \bar{v}(R) = \begin{cases} v(R) - v(Q) & \text{if } r \notin R \\ 0 & \text{otherwise} \end{cases}$$

Clearly, these two valuation function can be represented in \mathcal{REP} , because of the r -composition property. Based on these functions, we can now define two instances of 1-CONV using two agents and the same set of resources, as follows: Let $\underline{\mathcal{I}}$ be the 1-CONV instance for two agents with valuation functions \underline{v} and v_0 ; and let $\bar{\mathcal{I}}$ be the 1-CONV instance with \bar{v} and v_0 . We will show next that deciding whether these instances are 1-convergent or not will eventually give us the answer to the VIP problem $\langle v, Q \rangle$.

First of all, note that if the answer to the VIP problem over $\langle v, Q \rangle$ is NO (*i.e.* if Q is an optimal bundle), then both $\underline{\mathcal{I}}$ and $\bar{\mathcal{I}}$ are 1-convergent. The reason is that 0 will be the highest value the functions \underline{v} and \bar{v} can take in this case, and moving resource r from an agent to another in $\underline{\mathcal{I}}$ or $\bar{\mathcal{I}}$ will lead to an optimal allocation if the current allocation is not already optimal. By contraposition, we can deduce the following:

- If $\underline{\mathcal{I}}$ or $\bar{\mathcal{I}}$ are *not* 1-convergent, then the answer to the VIP problem is YES.

⁵ To be precise, we are using a polynomial-time truth-table reduction, which is a special case of a Turing reduction [20]. Recall that polynomial-time Turing (and truth-table) reductions do not allow us to distinguish between NP- and coNP-hardness [16,20], which explains why our NP-hardness result is not at odds with the known coNP-hardness result for the special case for the SLP form [11]

Algorithm: deciding VIP with the help of an oracle for 1-CONV

Input: valuation function v , bundle $Q \subseteq \mathcal{R}$

Output: YES or NO

- (1) if $\langle \underline{v}, v_0 \rangle$ is not in 1-CONV or $\langle \bar{v}, v_0 \rangle$ is not in 1-CONV, then return YES; else
- (2) if there exists an IR 1-deal improving allocation $\langle \emptyset, \mathcal{R} \rangle$ for $\langle \underline{v}, v_0 \rangle$, then return YES; else
- (3) if there exists an IR 1-deal improving allocation $\langle \mathcal{R}, \emptyset \rangle$ for $\langle \bar{v}, v_0 \rangle$, then return YES; else
- (4) return NO

Fig. 1 Algorithm summarising the reduction from 1-CONV to VIP

Consider now the case where both problems are 1-convergent. By definition of \underline{v} and \bar{v} , the answer to the VIP problem is YES *iff* there exists an allocation \underline{A} such that $\underline{v}(\underline{A}) > 0$ or there exists an allocation \bar{A} such that $\bar{v}(\bar{A}) > 0$. This in turn is equivalent to checking whether there exists an allocation \underline{A} such that $sw(\underline{A}) > 0$ for $\underline{\mathcal{I}}$ or there exists an allocation \bar{A} such that $sw(\bar{A}) > 0$ for $\bar{\mathcal{I}}$. Checking whether there exist allocations with social welfare exceeding 0 can be done the following way: First, identify allocations \underline{A}_0 with $sw(\underline{A}_0) = 0$ for $\underline{\mathcal{I}}$ and \bar{A}_0 with $sw(\bar{A}_0) = 0$ for $\bar{\mathcal{I}}$. Such allocations are easily constructed. Simply set $\underline{A}_0 = \langle \emptyset, \mathcal{R} \rangle$ and $\bar{A}_0 = \langle \mathcal{R}, \emptyset \rangle$, yielding $\underline{v}(\emptyset) + v_0(\mathcal{R}) = \bar{v}(\mathcal{R}) + v_0(\emptyset) = 0$. Then, because both problem instances are assumed to be 1-convergent, checking whether there exist allocations that are superior to \underline{A}_0 and \bar{A}_0 can be reduced to checking whether there are any IR 1-deals that would be applicable from these allocations. This can be checked in polynomial time. To summarise:

- If both $\underline{\mathcal{I}}$ and $\bar{\mathcal{I}}$ are 1-convergent, then consider allocations $\underline{A}_0 = \langle \emptyset, \mathcal{R} \rangle$ and $\bar{A}_0 = \langle \mathcal{R}, \emptyset \rangle$. If there exists an IR 1-deal from \underline{A}_0 for $\underline{\mathcal{I}}$ or if there exists an IR 1-deal from \bar{A}_0 for $\bar{\mathcal{I}}$, then the answer to the VIP problem is YES, otherwise it is NO.

This covers all possible cases. Hence, given our oracle for 1-CONV, we can always decide VIP in polynomial time. Our reduction is summarised by the algorithm shown in Fig. 1. \square

We stress that the decision problem 1-CONV, for which we have given a complexity analysis here, is *not* intended to be viewed as a problem to be solved by any of the participating parties during negotiation. Instead, this is a (hypothetical) problem that would have to be solved by someone who wants to check whether an agent society with a given profile of valuation functions could be expected to solve any given resource allocation problem in an optimal manner by means of 1-deals alone. By establishing NP-hardness, we have shown that this would be an infeasible prospect.

To conclude this section, we recall that the NP-hardness result for 1-CONV entails that, for many common representation languages, checking whether profiles of valuation functions will guarantee convergence by means of IR 1-deals is intractable. Therefore, searching for a condition over profiles of valuation functions is unrealistic. As we have shown that no condition on *single* valuation functions is both sufficient and necessary, and as, on top of that, we have just argued that checking conditions on *profiles* of functions is intractable

in most cases, the best thing we can do is to investigate whether some restricted classes of valuation functions can be identified as being *maximal*. This is what we shall do in the next section.

6 Maximal classes of valuation functions

Recall that valuation functions being modular is a sufficient condition for convergence by means of IR 1-deals. In previous work [9] we have shown that the class of modular valuation functions are also *maximal* with respect to the class of 1-deals. This is to say that for no class of valuation functions strictly including the class of modular functions we would still be able to ensure the same kind of convergence result. The significance of this (surprising) result can only be fully appreciated when considered together with the “negative” result on necessary and sufficient conditions discussed in the previous section.

In this section, we are going to substantially strengthen the known maximality theorem by proving maximality with respect to the much larger class of *bilateral deals*, and we are also going to establish two related results for the framework without money.

6.1 Framework with money

The class of bilateral deals includes all 1-deals. Hence, by Theorem 2, if all valuation functions are modular, then negotiation by means of IR bilateral deals guarantees an outcome with maximal social welfare. We now show that this ceases to be the case as soon as just a single valuation is non-modular. That is, the class of modular valuation functions is maximal with respect to the class of bilateral deals if we wish to ensure convergence for the framework with money. This result applies to negotiation problems with three or more agents (naturally, for two agents the class of bilateral deals is sufficient).

Theorem 5 (Maximality wrt. bilateral deals) *Let \mathcal{M} be the class of modular valuation functions. Then for any class of valuation functions \mathcal{F} such that $\mathcal{M} \subset \mathcal{F}$, there are a negotiation problem with valuation functions drawn from \mathcal{F} and an initial allocation such that no sequence of IR bilateral deals will lead to an allocation with maximal social welfare.*

Proof Observe that for $|\mathcal{R}| \leq 1$, any valuation 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.

We first prove the claim for the case of three agents 1, 2, 3 and two resources r_1, r_2 and then argue how this immediately generalises to the general theorem. Let v_1 be any valuation function over $\{r_1, r_2\}$ that is not modular. We will show how to construct modular valuation functions v_2 and v_3 and an initial allocation A_0 such that moving to an optimal allocation is not possible by

means of bilateral deals alone. v_1 can be expressed in k -additive form using coefficients $a, b, c, d \in \mathbb{R}$:⁶

$$v_1 = a + b.r_1 + c.r_2 + d.r_1.r_2$$

Given that v_1 is non-modular we know that $d \neq 0$. We distinguish two cases:

- (1) First, suppose $d > 0$. We define v_2 and v_3 as follows:

$$\begin{aligned} v_2 &= (b + \frac{1}{3}d).r_1 + (c + \frac{1}{3}d).r_2 \\ v_3 &= (b + \frac{1}{3}d).r_1 + (c + \frac{1}{3}d).r_2 \end{aligned}$$

Furthermore, let A_0 be such that agent 2 holds r_1 and agent 3 holds r_2 . Hence, $sw(A_0) = a + b + c + \frac{2}{3}d$. Let A^* be the allocation where both resources are held by agent 1. We have $sw(A^*) = a + b + c + d > sw(A_0)$. The direct deal $\delta = (A_0, A^*)$ would not be bilateral. In fact, we can easily check that there are no IR bilateral deals from A_0 : First, as agents 2 and 3 have identical modular valuation functions, no deal between them would be IR. Second, the deal where agent 2 is selling r_1 to agent 1 is reducing social welfare by $\frac{1}{3}d$ and similarly for the deal where agent 3 is selling r_2 to agent 1. There are no other bilateral deals from A_0 . Hence, A_0 is not optimal, but negotiation by means of IR bilateral deals is stuck.

- (2) Now suppose $d < 0$. Choose a very large $\Omega \in \mathbb{R}$ (in fact, any $\Omega > |b + \frac{2}{3}d|$ will do) and define v_2 and v_3 as follows:

$$\begin{aligned} v_2 &= (b + \frac{2}{3}d).r_1 \\ v_3 &= (c + \frac{1}{2}d).r_2 - \Omega.r_1 \end{aligned}$$

Furthermore, let A_0 be such that agent 2 holds r_1 and agent 1 holds r_2 . Hence, $sw(A_0) = a + b + c + \frac{2}{3}d$. Let A^* be the allocation where agent 1 holds r_1 and agent 3 holds r_2 , *i.e.* $sw(A^*) = a + b + c + \frac{1}{2}d$. This sum is greater than $sw(A_0)$, because d is negative. That is, there exists a better allocation than A_0 and that allocation is not reachable by means of a single bilateral deal. It remains to be checked that no bilateral deal applicable in A_0 would be IR: First, the only possible deal between agents 1 and 2 would be to give r_1 to agent 1, which would decrease social welfare by $|\frac{1}{3}d|$. Second, the only possible deal between agents 1 and 3 would be to give r_2 to agent 3, which would decrease social welfare by $|\frac{1}{2}d|$. Third, the only possible deal between agents 2 and 3 would be to give r_1 to agent 3, which would also reduce social welfare (due to our choice of Ω).

⁶ The k -additive form has been introduced in Section 5.3. Here (and in the sequel) we use a simplified notation: an agent with valuation function v_1 assigns value a to the empty bundle, increases their valuation of the bundle by b on receiving r_1 , by c on receiving r_2 , and by an additional amount of d if both r_1 and r_2 are being received.

According to the characterisation of modular functions given by equation (3) in Section 3, if there are more than two resources in the system, then v_1 must still be non-modular with respect to two of the resources and a particular fixed allocation as far as the other resources are concerned. We can distribute these other resources amongst agents 1 and 2 according to this particular partial allocation and define v_2 such that agent 2 places prohibitively high positive or negative values on them so that these additional resources will never be involved in any IR deal. For example, if $v_1 = 2.r_1 + 3.r_2 - 1.r_1.r_2 + 5.r_1.r_2.r_3$, then v_1 is non-modular with respect to $\{r_1, r_2\}$ for the fixed partial allocation where agent 1 receives r_3 . Then, if we add the term $-100.r_3$ to the definition of v_2 , our proof given above (with $d = 5 - 1$) still applies. Clearly, this is a general approach that applies to any non-modular function v_1 . Hence, it is sufficient to consider the case of two resources, as done above. \square

Observe that the theorem only makes an existential statement, so there is no need to prove anything for more than three agents, although examples for larger numbers of agents can be constructed in the same manner as done in the proof. Clearly, for the special case of just two agents, bilateral deals *do* guarantee convergence—whatever the valuation functions may be.

Theorem 5 entails our earlier result [9] on the maximality of modular valuation functions with respect to 1-deals, because every 1-deal is also a bilateral deal. We state this result here as a corollary.

Corollary 1 (Maximality wrt. 1-deals) *Let \mathcal{M} be the class of modular valuation functions. Then for any class of valuation functions \mathcal{F} such that $\mathcal{M} \subset \mathcal{F}$, there are a negotiation problem with valuations drawn from \mathcal{F} and an initial allocation such that no sequence of IR 1-deals will lead to an allocation with maximal social welfare.*

The original direct proof of Corollary 1 [9] has been considerably more complex than our proof of Theorem 5. We reproduce that proof in Appendix A, because beyond of what follows from Theorem 5 that proof also shows that the class of modular valuation functions is maximal with respect to 1-deals *even for problems with only two agents*.

Why is Corollary 1 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 valuation functions for which such a guarantee can be given. However, our discussion in Section 5 has shown that there can be no class of valuation 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 valuation functions that do guarantee optimal outcomes by means of 1-deal negotiation. As shown by Theorem 2, the class of modular functions is such a class and it

is certainly a natural, albeit simple and limiting, 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. Corollary 1 settles this question by giving a negative answer: For any agent with a non-modular valuation function there exist modular valuation 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. This is the case even for small systems with just two agents.⁷

Theorem 5 substantially sharpens the previously known maximality result (here stated as Corollary 1). It says that even if we are able to implement a negotiation framework in which agents can successfully negotiate arbitrary bilateral deals (requiring a rich negotiation protocol, going well beyond of what is possible in many existing systems today), we cannot give guarantees for socially optimal outcomes, unless all agents have preferences that are expressible in terms of modular valuation functions alone. Given the very limited expressive power of modular preferences and the seemingly powerful concept of general bilateral negotiation, this result is both sobering and surprising.

Can Theorem 5 be generalised even further? As we shall argue next, it can but not in interesting ways. For instance, the class of modular valuation functions is also maximal with respect to the class of bilateral deals enriched with one specific trilateral deal, such as e.g. the deal that moves r_1 from agent 1 to agent 2 and at the same time r_2 from agent 1 to agent 3 (to see this, revisit the proof of Theorem 5 and observe that having this additional deal available would not have made a difference for the argument used). However, this specific extension is of course of no practical interest. Extending Theorem 5 to larger classes of deals that are natural and interesting does not seem possible. For instance, it is *not* the case that the class of modular valuations functions would be maximal with respect to the class of all bilateral and trilateral deals together. Put differently, at least for some superclasses of \mathcal{M} we can always guarantee convergence by means of IR deals that involve at most three agents each. This follows from known results: Observe that two resources suffice to construct a non-modular function. In a domain with two resources any deal can involve at most four agents. But in fact, a deal involving four agents must be the composition of two independent deals between two agents each, and it is known that such a deal can always be decomposed without violating IR [13].

⁷ As an aside we remark here that there may well be further such classes (that are both sufficient and maximal), but we have not been able to identify any such class that would also be *interesting*. An example for a class of valuation functions that is sufficient but *not* interesting is the class of *pseudo-constant* functions. A valuation function v_i is pseudo-constant iff there exists a $c \geq 0$ such that $v_i(R) = c$ for all $R \neq \emptyset$ and $v_i(\emptyset) = 0$. Clearly, if all agents use such a function, then IR 1-deals will be possible until an allocation with maximal social welfare has been reached—but this is of course a very limited model of agent preferences.

6.2 Framework without money

We now turn our attention to the framework without money and study maximality questions for the family of classes $\mathcal{M}_{\alpha,\beta}$ of modular valuation functions with fixed α, β -values introduced in Section 3. We shall prove two theorems. The first establishes maximality with respect to CR 1-deals. The second shows that when we lower our ambitions from negotiating allocations with maximal social welfare to negotiating Pareto optimal allocations, then we lose the maximality property.

As we have seen in the proof of Theorem 3, for any class $\mathcal{M}_{\alpha,\beta}$, all valuation functions being drawn from that class is a sufficient condition for convergence to an allocation with maximal social welfare by means of CR 1-deals. It turns out that any such class is also maximal.

Theorem 6 (Maximality wrt. 1-deals) *Let $\mathcal{M}_{\alpha,\beta}$ be a class of modular valuation functions with fixed α, β -values. Then for any class of valuation functions \mathcal{F} such that $\mathcal{M}_{\alpha,\beta} \subset \mathcal{F}$, there are a negotiation problem with valuation functions drawn from \mathcal{F} and an initial allocation such that no sequence of CR 1-deals will lead to an allocation with maximal social welfare.*

Proof Suppose agent 1 has a valuation function $v_1 \notin \mathcal{M}_{\alpha,\beta}$. We shall construct a negotiation problem with two further agents with valuation functions in $\mathcal{M}_{\alpha,\beta}$ and an initial allocation such that that allocation is not optimal, but no CR 1-deal can improve upon it. We distinguish two cases:

(1st case) Suppose there exist a bundle R and a resource r such that $v_1(R \cup \{r\}) - v_1(R) \notin \{0, \alpha_r, \beta_r\}$. Let $\lambda = v_1(R \cup \{r\}) - v_1(R)$. We define $v_3 = \sum_{x \in R \cup \{r\}} \beta_x \cdot x + \sum_{x \in \mathcal{R} \setminus (R \cup \{r\})} \alpha_x \cdot x$ and give all of $\mathcal{R} \setminus (R \cup \{r\})$ to agent 3. That is, agent 3 gets its single most preferred bundle and will not be willing to participate in any deals. We also allocate all of R to agent 1. It remains to define v_2 and to determine the location of r in the initial allocation A_0 . Let $v'_2 = \sum_{x \in R} \beta_x \cdot x$. An agent equipped with the valuation function v'_2 would dislike all the items in R and be indifferent towards all others. We will define v_2 of agent 2 in terms of v'_2 , by additionally specifying that agent's attitude towards r . The location of r and the value agent 2 places on r depend on λ :

- (1) If $\lambda > \alpha_r$, then we give r to agent 2 and define $v_2 = v'_2 + \alpha_r \cdot r$.
- (2) If $\alpha_r > \lambda > 0$, then we give r to agent 1 and define $v_2 = v'_2 + \alpha_r \cdot r$.
- (3) If $0 > \lambda > \beta_r$, then we give r to agent 2 and define $v_2 = v'_2 + \beta_r \cdot r$.
- (4) If $\beta_r > \lambda$, then we give r to agent 1 and define $v_2 = v'_2 + \beta_r \cdot r$.

In case (1), for example, social welfare would increase by $\lambda - \alpha_r$ if we were to give r to agent 1, but this deal is not CR. Indeed, there is no CR deal starting from A_0 . In case (2), social welfare would increase by $\alpha_r - \lambda$ if agent 1 were to give r to agent 2, but while this is the only rational 1-deal for agent 2, it is not rational for agent 1. The situation for the other two cases is similar. This concludes the proof for the first case.

(2nd case) Now suppose that $v_1(R \cup \{r\}) - v_1(R) \in \{0, \alpha_r, \beta_r\}$ for all bundles R and resources r . Given that $v_1 \notin \mathcal{M}_{\alpha, \beta}$, this means that v_1 cannot be modular. Hence, by equation (3), there exist a bundle R and resources $r_1, r_2 \notin R$ such that $v_1(R \cup \{r_1, r_2\}) \neq v_1(R \cup \{r_1\}) + v_1(R \cup \{r_2\}) - v_1(R)$.

We will again use agent 3 to allow us to focus on an interesting subset of resources: give all of $\mathcal{R} \setminus (R \cup \{r_1, r_2\})$ to agent 3 and define v_3 such that agent 3 perceives this as the best possible situation and will not participate in any deals. So below we only need to specify allocations and v_2 with respect to $R \cup \{r_1, r_2\}$. We will refer to the following four (partial) allocations:

	A	A_1	A_2	A_{12}
Agent 1	R	$R \cup \{r_1\}$	$R \cup \{r_2\}$	$R \cup \{r_1, r_2\}$
Agent 2	$\{r_1, r_2\}$	$\{r_2\}$	$\{r_1\}$	\emptyset

The marginal valuation of receiving $\{r_1, r_2\}$ will be either more or less than the sum of the marginal valuations of the two individual items:

- (1) Suppose $v_1(R \cup \{r_1, r_2\}) > v_1(R \cup \{r_1\}) + v_1(R \cup \{r_2\}) - v_1(R)$. Define $v_2 = [v_1(R \cup \{r_1\}) - v_1(R)] \cdot r_1 + [v_1(R \cup \{r_2\}) - v_1(R)] \cdot r_2 + \sum_{x \in R} \beta_x \cdot x$. Clearly, $v_2 \in \mathcal{M}_{\alpha, \beta}$. We have $v_1(A) + v_2(A) = v_1(R) + v_1(R \cup \{r_1\}) + v_1(R \cup \{r_2\}) - 2 \cdot v_1(R) = v_1(R \cup \{r_1\}) + v_1(R \cup \{r_2\}) - v_1(R)$. Also, $v_1(A_1) + v_2(A_1) = v_1(R \cup \{r_1\}) + v_1(R \cup \{r_2\}) - v_1(R) = v_1(A_2) + v_2(A_2)$ and $v_1(A_{12}) + v_2(A_{12}) = v_1(R \cup \{r_1, r_2\})$. Now, if A is made initial allocation, then there is no CR 1-deal available, despite A_{12} being socially preferable.
- (2) Suppose $v_1(R \cup \{r_1, r_2\}) < v_1(R \cup \{r_1\}) + v_1(R \cup \{r_2\}) - v_1(R)$. Define $v_2 = [v_1(R \cup \{r_1\}) - v_1(R)] \cdot r_1 + [v_1(R \cup \{r_1, r_2\}) - v_1(R \cup \{r_1\})] \cdot r_2 + \sum_{x \in R} \beta_x \cdot x$. Clearly, $v_2 \in \mathcal{M}_{\alpha, \beta}$. We have $v_1(A) + v_2(A) = v_1(A_1) + v_2(A_1) = v_1(A_{12}) + v_2(A_{12}) = v_1(R \cup \{r_1, r_2\})$ and $v_1(A_2) + v_2(A_2) = v_1(R \cup \{r_1\}) + v_1(R \cup \{r_2\}) - v_1(R)$. Hence, if the initial allocation is A_1 , then moving to A_2 would increase social welfare, but there is no possible CR 1-deal.

This covers all cases, and we are done. \square

Next we analyse the situation where we aim for a Pareto optimal outcome rather than an outcome maximising social welfare. As any allocation with maximal social welfare is also Pareto optimal, Theorem 3 shows that CR 1-deals will always lead to a Pareto optimal allocation if all valuation functions are drawn from the same $\mathcal{M}_{\alpha, \beta}$. Our next theorem shows that the $\mathcal{M}_{\alpha, \beta}$ classes are *not* maximal, however. We give a constructive proof by showing how $\mathcal{M}_{\alpha, \beta}$ can be extended without jeopardising convergence.⁸

Theorem 7 (Lack of maximality) *Let $\mathcal{M}_{\alpha, \beta}$ be a class of modular valuation functions with fixed α, β -values. Then there exists a class of valuation*

⁸ We label this theorem as a negative result (on the lack of a maximality property) rather than as a positive result (on the sufficiency of a superclass of $\mathcal{M}_{\alpha, \beta}$), because we believe that this highlights what is most interesting about it: The superclass used in the proof is hardly of practical interest, but the lack of maximality is insightful, given the strong affirmative maximality results proved earlier.

functions \mathcal{F} with $\mathcal{M}_{\alpha,\beta} \subset \mathcal{F}$, such that any sequence of CR 1-deals will eventually result in a Pareto optimal allocation, whenever all valuation functions are drawn from \mathcal{F} .

Proof (sketch) Fix some $r \in \mathcal{R}$. Define $\mathcal{F} := \mathcal{M}_{\alpha,\beta} \cup \{v\}$ with $v(R) = a$ whenever $r \in R$ and $v(R) = 0$ otherwise, where a is chosen to be any positive number different from α_r . (So we have $v \notin \mathcal{M}_{\alpha,\beta}$, as required.) Suppose all valuation functions are drawn from \mathcal{F} . Then it is not hard to show that any sequence of CR 1-deals will converge to a Pareto optimal allocation. The proof is similar to that of Theorem 3. Note that an agent with valuation function v is only interested in obtaining r , will never give away r (whatever is being offered in return), and will give away any other resources without expecting anything in return. \square

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 indivisible resources. We have shown that the use of *modular* valuation functions to model agent preferences is a *sufficient* condition, if side payments are allowed, to guarantee final allocations with maximal social welfare in case agents only negotiate *1-deals* (involving one resource each). When no side payments are permitted, any class of $\mathcal{M}_{\alpha,\beta}$ functions would be sufficient. We have then seen, however, that this is not a *necessary* condition for optimal outcomes and, indeed, there can be no condition on (individual) valuation functions that would be both necessary and sufficient in this sense. Furthermore, we have shown that, while a necessary and sufficient condition on profiles of valuation functions obviously does exist, checking such a condition is intractable for many representation languages used for encoding agent preferences. We have therefore concentrated on establishing *maximality* results. In the framework with money, there is no superclass of the class of modular valuation functions that would still ensure convergence by means of either 1-deals or the much larger class of bilateral deals. For the framework without money we have also proved maximality with respect to 1-deals. On the other hand, if we only require Pareto optimal outcomes, this kind of maximality property ceases to hold.

To summarise the main results more formally, let us introduce the $\text{CONV}(\cdot, \cdot)$ problem, a generalisation of the 1-convergence problem presented in Section 5.3. Given a negotiation problem where all valuation functions of the agents belong to some class \mathcal{F} , the convergence problem $\text{CONV}(\mathcal{F}, \mathcal{T})$ is the problem of checking whether any sequence of deals of type \mathcal{T} (originating from any initial allocation) will result in an allocation with maximal social welfare. When the valuation functions are encoded using a particular representation language $\mathcal{R}\mathcal{E}\mathcal{P}$, this problem is referred to as $\text{CONV}(\mathcal{R}\mathcal{E}\mathcal{P}, \mathcal{T})$. Regarding types of deals, we have considered 1-deals, bilateral deals, and arbitrary deals as far

as their structure is concerned, and we have distinguished IR deals (individually rational deals with side payments) from CR deals (cooperatively rational deals without side payments). Let \mathcal{M} be the class of modular functions, and let $\mathcal{M}_{\alpha,\beta}$ be the class of modular functions with fixed α, β -values for some given vectors α and β . Using this notation, the main results reported in this paper may be summarised as follows. We first give the results pertaining to the framework with money:

- **Theorem 1:** For any class \mathcal{F} , $\text{CONV}(\mathcal{F}, \text{IR deals})$ always answers YES. (This result is due to Sandholm [30].)
- **Theorem 2:** $\text{CONV}(\mathcal{M}, \text{IR 1-deals})$ always answers YES.
- **Section 5.1:** There exist classes $\mathcal{F} \not\subseteq \mathcal{M}$ such that $\text{CONV}(\mathcal{F}, \text{IR 1-deals})$ always answers YES.
- **Section 5.2:** There exists no class \mathcal{F} such that for all classes \mathcal{F}' we have that $\mathcal{F}' \subseteq \mathcal{F}$ if and only if $\text{CONV}(\mathcal{F}', \text{IR 1-deals})$ always answers YES.
- **Theorem 4:** $\text{CONV}(\mathcal{R}\mathcal{E}\mathcal{P}, \text{IR 1-deals})$ is NP-hard (under Turing reduction) for many representation languages $\mathcal{R}\mathcal{E}\mathcal{P}$.
- **Theorem 5:** There exists no class $\mathcal{F} \supset \mathcal{M}$ such that $\text{CONV}(\mathcal{F}, \text{IR bilateral deals})$ always answers YES.

In addition, the proof of Corollary 1 shows that the maximality result of Theorem 5 continues to hold even for the case of just two agents when we replace “bilateral deals” by the weaker “1-deals”. For the setting without money, we have obtained these results:

- **Theorem 3:** $\text{CONV}(\mathcal{M}_{\alpha,\beta}, \text{CR 1-deals})$ always answers YES, for any vectors α and β .
- **Theorem 6:** There exists no class $\mathcal{F} \supset \mathcal{M}_{\alpha,\beta}$ (for any vectors α and β) such that $\text{CONV}(\mathcal{F}, \text{CR 1-deals})$ always answers YES.

Finally, Theorem 7 demonstrates that maximality is not as pervasive a notion as our other results may suggest: for the framework without money, it fails to apply when we aim for Pareto optimality rather than maximal social welfare.

We consider these not only surprising results, but also useful characterisations 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 foreseen 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 the introduction, we have argued for *distributed* approaches to resource allocation. Of course it would be over-simplistic to merely distinguish a centralised approach, *à la* combinatorial auction, versus a purely distributed approach. The truth is that there exists an entire spectrum of approaches between these two extremes. Very often, for instance, even if an auctioneer is available, it is desirable anyway to reduce part of the burden of the computational task usually assigned to it. One way to do so is to delegate it to the bidders instead.

This distribution of the *computation* of the winner determination problem has been studied by some authors, including Vidal and colleagues [23,35]. What remains centralised in these approaches, however, is the fact that agents must be aware of the bids placed by other agents during the process to compute their new bid. That is, while the computation of the solution is indeed distributed, the process itself remains centralised. In our approach, on the other hand, we assume that agents agree *locally* on beneficial deals, without requiring any sort of global perception of the situation. Also, while the actual allocation is modified once and for all in the case of combinatorial auctions, our setting regards negotiation as a continuously evolving process with a sequence of local steps incrementally modifying the allocation of resources within the system. In that sense, it is closer to combinatorial exchanges [32]. But even these are centralised mechanisms where a central authority makes all the decisions and has to perform all the computation.

The present paper continues recent work on convergence properties for distributed negotiation schemes [30,13]. Some related work, notably by Dunne and colleagues [12,11], has concentrated on the computational complexity of decision problems arising in the context of such negotiation frameworks. The complexity of the most fundamental problem in resource allocation, namely social welfare optimisation, has been known for some time. The underlying decision problem, *i.e.* checking whether a given negotiation problem admits a solution where the sum of valuations exceeds a given value, is NP-complete [28].⁹ This can be shown, for instance, via a reduction from the well-known SET PACKING problem [16]. Dunne *et al.* [12] have introduced the 1-PATH problem and analysed its complexity: given two allocations A and A' with $sw(A) < sw(A')$, check whether it is possible to reach A' from A by means of a sequence of rational 1-deals. Dunne *et al.* [12] have shown that 1-PATH is NP-hard. Recent results have further strengthened this complexity bound and established that the problem is in fact PSPACE-complete [11]. The closely related problem 1-CONV [7,11], discussed earlier in the context of checking conditions for guaranteed convergence with respect to full valuation profiles, asks whether any sequence of individually rational 1-deals would culminate in an allocation with maximal social welfare. As we have seen, 1-CONV is intractable. This has been a known result for the case where valuation functions are represented as straight-line programs [11], which we have extended here to a wider class of preference representation languages.

Acknowledgements. We would like to thank the reviewers of this paper and its earlier incarnations for their very helpful feedback.

⁹ As explained in Section 5.3, complexity results apply to a particular decision problem with respect to a particular representation language for the input (here, the valuation functions of the agents). However, in practice, it is the case that most of the results mentioned here have been established for several of the most common representations, so we omit giving such details here.

References

1. M. Andersson and T. W. Sandholm. Time-quality tradeoffs in reallocative negotiation with combinatorial contract types. In *Proc. 16th National Conference on Artificial Intelligence (AAAI-1999)*. MIT Press, 1999.
2. K. J. Arrow, A. K. Sen, and K. Suzumura, editors. *Handbook of Social Choice and Welfare*, volume 1. North-Holland, 2002.
3. G. Ausiello, P. Crescenzi, G. Gambosi, V. Kann, A. Marchetti-Spaccamela, and M. Protasi. *Complexity and Approximation: Combinatorial Optimization Problems and their Approximability Properties*. Springer-Verlag, 1999.
4. Y. Bachrach and J. S. Rosenschein. Distributed multiagent resource allocation in diminishing marginal return domains. In *Proc. 7th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2008)*, pages 1103–1110. IFAA-MAS, 2008.
5. E. Boros and P. L. Hammer. Pseudo-boolean optimization. *Discrete Applied Mathematics*, 123(1–3):155–225, 2002.
6. C. Boutilier and H. H. Hoos. Bidding languages for combinatorial auctions. In *Proc. International Joint Conference on Artificial Intelligence (IJCAI-2001)*, pages 1211–1217. Morgan Kaufmann, 2001.
7. Y. Chevaleyre, P. E. Dunne, U. Endriss, J. Lang, M. Lemaître, N. Maudet, J. Padget, S. Phelps, J. A. Rodríguez-Aguilar, and P. Sousa. Issues in multiagent resource allocation. *Informatica*, 30:3–31, 2006.
8. Y. Chevaleyre, U. Endriss, S. Estivie, and N. Maudet. Multiagent resource allocation in k -additive domains: Preference representation and complexity. *Annals of Operations Research*, 163(1):49–62, 2008.
9. Y. Chevaleyre, U. Endriss, and N. Maudet. On maximal classes of utility functions for efficient one-to-one negotiation. In *Proc. 19th International Joint Conference on Artificial Intelligence (IJCAI-2005)*, pages 941–946. Morgan Kaufmann Publishers, 2005.
10. P. Cramton, Y. Shoham, and R. Steinberg, editors. *Combinatorial Auctions*. MIT Press, 2006.
11. P. E. Dunne and Y. Chevaleyre. The complexity of deciding reachability properties of distributed negotiation schemes. *Theoretical Computer Science*, 396(1–3):113–144, 2008.
12. P. E. Dunne, M. Wooldridge, and M. Laurence. The complexity of contract negotiation. *Artificial Intelligence*, 164(1–2):23–46, 2005.
13. U. Endriss, N. Maudet, F. Sadri, and F. Toni. Negotiating socially optimal allocations of resources. *Journal of Artificial Intelligence Research*, 25:315–348, 2006.
14. B. Faltings and M. Yokoo. Introduction: Special issue on distributed constraint satisfaction. *Artificial Intelligence*, 161(1–2):1–5, 2005.
15. A. Galstyan, K. Czajkowski, and K. Lerman. Resource allocation in the Grid with learning agents. *Journal of Grid Computing*, 3(1–2):91–100, 2005.
16. M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-completeness*. W. H. Freeman and Co., 1979.
17. A. Giovannucci, J. A. Rodríguez-Aguilar, A. Reyes, F. X. Noria, and J. Cerquides. Towards automated procurement via agent-aware negotiation support. In *Proc. 3rd International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2004)*, pages 244–253. ACM Press, 2004.
18. M. Grabisch. k -order additive discrete fuzzy measures and their representation. *Fuzzy Sets and Systems*, 92:167–189, 1997.
19. S. Kraus. *Strategic Negotiation in Multiagent Environments*. MIT Press, 2001.
20. R. E. Ladner, N. A. Lynch, and A. L. Selman. A comparison of polynomial time reducibilities. *Theoretical Computer Science*, 1(2):103–123, 1975.
21. M. Lemaître, G. Verfaillie, and N. Bataille. Exploiting a common property resource under a fairness constraint: A case study. In *Proc. 16th International Joint Conference on Artificial Intelligence (IJCAI-1999)*. Morgan Kaufmann Publishers, 1999.
22. H. Moulin. *Axioms of Cooperative Decision Making*. Cambridge University Press, 1988.
23. M. V. Narumanchi and J. M. Vidal. Algorithms for distributed winner determination in combinatorial auctions. In *Agent-Mediated Electronic Commerce. Designing Trading Agents and Mechanisms*, pages 43–56. Springer-Verlag, 2006.

24. N. Nisan. Bidding languages for combinatorial auctions. In P. Cramton et al., editors, *Combinatorial Auctions*. MIT Press, 2006.
25. D. Parkes and J. Shneidman. Distributed implementations of Vickrey-Clarke-Groves mechanisms. In *Proc. 3rd International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2004)*, pages 261–268. ACM Press, 2004.
26. A. Petcu, B. Faltings, and D. Parkes. MDPOP: Faithful distributed implementation of efficient social choice problems. In *Proc. 5th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2006)*, pages 1397–1404. ACM Press, 2006.
27. J. S. Rosenschein and G. Zlotkin. *Rules of Encounter*. MIT Press, 1994.
28. M. H. Rothkopf, A. Pekeč, and R. M. Harstad. Computationally manageable combinatorial auctions. *Management Science*, 44(8):1131–1147, 1998.
29. S. Saha and S. Sen. An efficient protocol for negotiation over multiple indivisible resources. In *Proc. 20th International Joint Conference on Artificial Intelligence (IJCAI-2007)*, pages 1494–1499, 2007.
30. T. W. Sandholm. Contract types for satisficing task allocation: I Theoretical results. In *Proc. AAAI Spring Symposium: Satisficing Models*, 1998.
31. T. W. Sandholm. Distributed rational decision making. In G. Weiß, editor, *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*, pages 201–258. MIT Press, 1999.
32. T. W. Sandholm, S. Suri, A. Gilpin, and D. Levine. Winner determination in combinatorial auction generalizations. In *Proc. 1st International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2002)*. ACM Press, 2002.
33. 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.
34. J. Uckelman, Y. Chevaleyre, U. Endriss, and J. Lang. Representing utility functions via weighted goals. *Mathematical Logic Quarterly*, 2009. In press.
35. J. M. Vidal. Multiagent coordination using a distributed combinatorial auction. In *Proc. AAAI Workshop on Auction Mechanisms for Robot Coordination*, 2006.
36. M. Wooldridge. *An Introduction to MultiAgent Systems*. John Wiley and Sons, 2002.

Appendix A: Direct proof of Corollary 1

We give a direct proof of Corollary 1 that demonstrates that this results applies even for systems with just two agents.

Proof The proof is constructive. We will show that for any non-modular valuation function v_1 on m resources, it is possible to construct a modular valuation function v_2 (with $v_i \equiv 0$ for all other agents i , if any) and an initial allocation such that no optimal allocation can be reached by means of IR 1-deals. This implies that $\mathcal{M} \cup \{v_1\}$ does not guarantee convergence by means of IR 1-deals. Because v_1 is non-modular, we can apply equation (3) to show that there must exist a bundle X and distinct resources $r_1, r_2 \notin X$ such that ϵ , defined as follows, is not equal to 0:

$$\epsilon = v_1(X \cup \{r_1\}) + v_1(X \cup \{r_2\}) - v_1(X) - v_1(X \cup \{r_1, r_2\}) \quad (4)$$

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, resources in X are owned by 1, and resources in $Y = \mathcal{R} \setminus (X \cup \{r_1, r_2\})$ by 2, as shown in the following table.

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

Let us build a modular valuation function v_2 defined as follows: $\forall R \subseteq \mathcal{R}$,

$$v_2(R) = \sum_{r \in \{r_1, r_2\} \cap R} \alpha_r + \sum_{r \in R \cap Y} \omega - \sum_{r \in R \cap X} \omega \quad (5)$$

with $\omega = 14 \times \max |v_1| + 1$. Let $\Omega = v_2(Y) = |Y| \times \omega$.¹⁰ 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:

$$\begin{aligned} sw(A_{1|2}) &= \Omega + \alpha_{r_1} + \alpha_{r_2} + v_1(X) \\ sw(A_{12|}) &= \Omega + v_1(X \cup \{r_1, r_2\}) \\ sw(A_{1|2}) &= \Omega + \alpha_{r_2} + v_1(X \cup \{r_1\}) \\ sw(A_{2|1}) &= \Omega + \alpha_{r_1} + v_1(X \cup \{r_2\}) \end{aligned}$$

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 (a) this initial allocation does not have optimal social welfare, (b) there is only one rational deal from this allocation, (c) this deal leads to the optimal allocation, but (d) 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 $\alpha_{r_1} = v_1(X \cup \{r_1\}) - v_1(X) - \frac{\epsilon}{4}$ and $\alpha_{r_2} = v_1(X \cup \{r_1, r_2\}) - v_1(X \cup \{r_1\}) + \frac{\epsilon}{4}$.

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 $3 \times \max |v_1|$. Thus, all four allocations have a social welfare of at least $\Omega - |\alpha_{r_1}| - |\alpha_{r_2}| - \max |v_1| \geq \Omega - 7 \times \max |v_1| > \Omega - \frac{\omega}{2}$. All other allocations have a social welfare lower than $\Omega - \omega + |\alpha_{r_1}| + |\alpha_{r_2}| + \max |v_1| \leq \Omega - \omega + 7 \times \max |v_1| < \Omega - \frac{\omega}{2}$. 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. More precisely, let us show that $sw(A_{1|2}) < sw(A_{12|})$, 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:

$$\begin{aligned} sw(A_{1|2}) &= \Omega + v_1(X \cup \{r_1, r_2\}) \\ sw(A_{12|}) &= \Omega + v_1(X \cup \{r_1, r_2\}) \\ sw(A_{1|2}) &= \Omega + v_1(X \cup \{r_1, r_2\}) + \frac{\epsilon}{4} \\ sw(A_{2|1}) &= \Omega + v_1(X \cup \{r_1\}) + v_1(X \cup \{r_2\}) - v_1(X) - \frac{\epsilon}{4} \\ &= \Omega + v_1(X \cup \{r_1, r_2\}) + \frac{3}{4}\epsilon \end{aligned}$$

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 $\delta(A_{1|2}, A_{1|2})$. These deals decrease social welfare, and thus are not IR 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 $\alpha_{r_1} = v_1(X \cup \{r_1\}) - v_1(X) - \frac{\epsilon}{4}$ and $\alpha_{r_2} = v_1(X \cup \{r_2\}) - v_1(X) - \frac{\epsilon}{4}$. Note that, again, both $|\alpha_{r_1}|$ and $|\alpha_{r_2}|$ are less than $3 \times \max |v_1|$. Thus, by the same argument as in the first case, the four allocations all have greater social welfare than any other allocation.

¹⁰ Here, $\max |v_1|$ is the maximal value that v_1 may take, for any bundle of resources.

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 $sw(A_{1|2}) < sw(A_{12|})$.

$$\begin{aligned} sw(A_{1|2}) &= \Omega + v_1(X \cup \{r_1\}) + v_1(X \cup \{r_2\}) - v_1(X) - \frac{\epsilon}{2} \\ sw(A_{12|}) &= \Omega + v_1(X \cup \{r_1, r_2\}) \\ &= \Omega + v_1(X \cup \{r_1\}) + v_1(X \cup \{r_2\}) - v_1(X) - \epsilon \\ sw(A_{1|2}) &= \Omega + v_1(X \cup \{r_1\}) + v_1(X \cup \{r_2\}) - v_1(X) - \frac{\epsilon}{4} \\ sw(A_{2|1}) &= \Omega + v_1(X \cup \{r_1\}) + v_1(X \cup \{r_2\}) - v_1(X) - \frac{\epsilon}{4} \end{aligned}$$

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