Resource Allocation in Egalitarian Agent Societies

Ulrich Endriss
ue@doc.ic.ac.uk

Nicolas Maudet
maudet@doc.ic.ac.uk

Fariba Sadri
fs@doc.ic.ac.uk

Francesca Toni
ft@doc.ic.ac.uk

Department of Computing, Imperial College London
180 Queen's Gate, London SW7 2AZ
United Kingdom

Résumé :
Cet article présente une notion de société d'agents égalitaire et étudie comment il est possible au sein d'une telle société d'atteindre, par la négociation, une situation optimale d'allocation des ressources.

Mots-clés : Systèmes multi-agents, négociation, allocation de ressources, bien-être social égalitaire

Abstract:
We introduce the notion of an egalitarian agent society and study the problem of finding an optimal allocation of resources by means of negotiation amongst the agents inhabiting such a society.

Keywords: Multiagent systems, negotiation, resource allocation, egalitarian social welfare

1 Introduction

The basic concept of *resource allocation by negotiation* (as well as the related concept of task allocation by negotiation) have received much attention in the recent literature on multiagent systems ; see for instance [6, 9, 11]. A number of variants of this problem have been studied before ; here we consider the case of an artificial society of agents where, to begin with, each agent is in possession of a collection of discrete (i.e. non-divisible) resources. Agents may then negotiate with each other in order to agree on the redistribution of some of these resources to benefit either individual agents or society as a whole. In this paper, we focus on suitable mechanisms for the allocation of resources in agent societies that are governed by *egalitarian* principles.

To motivate this undertaking, we start by briefly reviewing previous work on the general framework of resource allocation by negotiation and then argue why an egalitarian variant of this framework may be the appropriate choice for certain applications of multiagent systems.

1.1 Resource allocation by negotiation

In a recent paper [3], we have analysed negotiation scenarios where self-interested agents exchange resources in order to increase their respective individual welfare. Rather than studying specific negotiation protocols or even strategies that agents may follow in order to further their interests, we were only concerned with the patterns of exchanges agents could possibly agree on and to what extent these patterns are sufficient or necessary to guarantee optimal outcomes of negotiations. One central assumption that we have made with respect to the strategies that agents follow, however, is that they are *individually rational* in the sense of never accepting a disadvantageous deal. To assess whether a given allocation of resources should be considered optimal, we have borrowed concepts from welfare economics, in particular the idea of maximising the sum of the utility values ascribed by all agents to the resources they hold in a given situation. A similar framework has been studied by Sandholm in [11] and elsewhere, mostly in the context of agents negotiating in order to reallocate tasks.

As shown in [3], the kinds of deals that individually rational agents are prepared to accept are *sufficient* to guarantee outcomes of negotiation that are Pareto optimal (i.e. society will never get stuck in a local minimum).[1] In an extended framework, where agents may use money to compensate other agents for (otherwise) disadvantageous deals, it is even possible to guarantee an outcome where the sum of utilities of the agents in the society is maximal.[2] Further results show that any deal (or pattern of resource exchanges) that is acceptable in these frameworks is also *necessary* in the sense that there are instances of the resource allocation problem where an optimal outcome is only possible if that particular deal is used at some point during negotiation.

Under a more general perspective, such results

[1] An allocation of resources is *Pareto optimal* iff there is no other allocation where some agents would be happier without any of the others being worse off.

[2] The result for the framework with money is essentially equivalent to Sandholm's main result on sufficient contract types in task-oriented domains [11].

may be interpreted as the emergence of a particular *global* behaviour (at the level of society) in reaction to *local* behaviour governed by some acceptability criterion of deals for individual agents.

The most widely studied mechanisms for the reallocation of resources in multiagent systems are *auctions* [14]. We should point out that our scenario of resource allocation by negotiation is *not* an auction. Auctions are mechanisms to help agents agree on a price at which an item (or a set of items) is to be sold [5]. In our work, on the other hand, we are not concerned with this aspect of negotiation, but only with the patterns of resource exchanges that agents actually carry out. On top of that, the egalitarian framework we are going to present in this paper does not involve a monetary component, i.e. there is no notion of a 'price' as such either. Typically, an auction involves a single auctioneer selling goods to a number of potential buyers. In contrast to this, our negotiation scenario is *symmetric* (there is no distinction between 'sellers' and 'buyers') and we specifically address the issue of multiple agents negotiating over multiple goods at the same time. The latter aspect has, to a certain degree, also been addressed in work on more complex auction mechanisms, in particular simultaneous auctions [10], combinatorial auction [12], and sequential auctions [2]. While it may be possible to use a combination of such auction mechanisms to negotiate the precise conditions accompanying a deal in our scenario (at least if we include a monetary component), in the present paper we are only concerned with the structure of these deals themselves, i.e. auctions and similar mechanisms are not of an immediate relevance.

1.2 Utilitarianism versus egalitarianism

The idea of aiming at maximising the sum of all utilities of the members of a society is a *utilitarian* concept. This interpretation of social welfare is often taken for granted in the multiagent systems literature. This is not the case in welfare economics and social choice theory though, where different notions of social welfare are being considered and compared with each other. Here, the concept of *egalitarian social welfare* takes a particularly prominent role [7, 13].

In an egalitarian system one would consider any differences in individual welfare unjust unless removing these differences would inevitably result in reducing the welfare of the agent who is currently worst off even further. (This is Rawls' so-called *difference principle* [8].) In other words, the first and foremost objective of such a society would be to maximise the welfare of its weakest member.

Clearly, the common utilitarian conception of multiagent systems is appropriate for many applications, particularly so if these applications have a commercial aspect of some sort. To demonstrate that there are other applications where the notion of an egalitarian agent society is more appropriate, let us consider the 'resource allocation problem' faced at regular intervals by the community of lecturers at a university department. We can imagine a multiagent system where each agent acts on behalf of one of the lecturers and negotiates over the allocation of resources (or tasks, which may be considered resources with negative utility values) before the beginning of term. Negative resources could be different courses to be taught, administrative tasks or supervision duties. Positive resources could be offices of different sizes or tutorial helpers. The values assigned to these resources may vary depending on the other resources held by the same agent. For example, the importance of being supported by a tutorial helper will usually depend on the size of the class you have to teach. Furthermore, different agents may assign different utility values to the same resources. In this context, the aim of the system designer (the head of department) would be to ensure that the least happy lecturer is as happy as possible (to preserve a good working atmosphere in the long-term), without having to regulate every single detail herself.

The question what social order is 'better' has concerned philosophers for a long time. A famous example is Rawls' *veil of ignorance* [8]. To determine what social principles are *just*, he suggests a thought experiment where rational agents have to choose the social principles governing a society before entering that society and without knowing their own position within that society. Behind this 'veil of ignorance', subjects must choose principles the consequences of which they are prepared to accept whatever their role in society may turn out to be. Rawls argues that under these circumstances the (egalitarian) principle of difference will be found to be just. Others have found similar arguments in defence of utilitarianism [4].

Moulin [7] analyses these arguments as follows. Someone who would prefer the egalitarian society is *risk-averse* ; they fear to end up as the weakest member of society and consequently opt for a social order based on egalitarian principles. Those favouring the utilitarian society, on the other hand, may be understood as *maximising their expected utility*.

Of course, either argument may be refuted (perhaps rather crudely) on the grounds of simply being too abstract a mental construction to yield any reliable ethical or social guidelines. In the context of a multiagent system, however, this kind of construction can become a very concrete issue. Before agreeing to be represented by a software agent in an artificial society, one would naturally want to know under what principles this society operates. If the agent's objective is to negotiate on behalf of its owner, then the owner has to agree to accept whatever the outcome of a specific negotiation may be. In many cases, there may be only very little information available regarding the agent's starting position ; that is, one would have to agree on accepting results—behind a veil of ignorance—on the basis of the principles governing negotiation alone. Clearly, for many sensitive domains, a risk-averse attitude would be appropriate, i.e. an egalitarian agent society would be a suitable option.

1.3 Paper overview

The remainder of this paper is structured as follows. In Section 2 we give a formal account of the basic framework of resource allocation by negotiation and in Section 3 we review the definitions of some well-known concepts from welfare economics, in particular the notion of egalitarian social welfare. A local acceptability criterion for agents matching this global measure of welfare will be introduced in Section 4. The main technical results of this paper are proved in Section 5. We show that our local criterion is sufficient to guarantee an optimal outcome of negotiation (in the egalitarian sense), but also that any admissible pattern of resource exchanges may be necessary to reach these optimal allocations. We conclude in Section 6 with a suggestion for future work in this area.

2 Negotiating over resources

Our basic scenario is that of an artificial society populated by a number of agents, each of which initially holds a set of resources. Agents may then engage into a negotiation process in order to agree on the reallocation of some of the resources. In this section we are going to formally define the various parameters of such a negotiation scenario.

2.1 Negotiation scenarios

Negotiation over resources takes place in a *system* $(\mathcal{A}, \mathcal{R})$, where $\mathcal{A}$ is a finite set of (at least two) *agents* and $\mathcal{R}$ is a finite set of *resources*. We will think of $\mathcal{A}$ as an initial segment of the natural numbers, that is, we identify agent names with numbers from 1 to n for some number $n \in \mathbb{N}$.

A particular *allocation* is a partitioning of the available resources $\mathcal{R}$ amongst the agents in $\mathcal{A}$.

Definition 1 (Allocations) *An allocation of resources is a function A from agents in $\mathcal{A}$ to subsets of $\mathcal{R}$ with $A(i) \cap A(j) = \{\,\}$ for $i \neq j$ and $\bigcup_{i \in \mathcal{A}} A(i) = \mathcal{R}$.*

Agents will usually ascribe different values to different sets of resources. The value an agent $i \in \mathcal{A}$ ascribes to a particular set of resources R will be modelled by means of a *utility function*, that is, a function from sets of resources (subsets of $\mathcal{R}$) to real numbers. This could really be *any* such function, that is, the utility ascribed to a set of resources is not just the sum of the values ascribed to its elements. This allows us to model the fact that utility may strongly depend on context, i.e. what other resources the agent holds at the same time.

Definition 2 (Utility functions) *The utility function u_i of each agent $i \in \mathcal{A}$ is a function from subsets of $\mathcal{R}$ to real numbers.*

Given an allocation A, $A(i)$ is the set of resources held by agent i in that situation. We are going to abbreviate $u_i(A) = u_i(A(i))$ for the utility value assigned to that set by agent i.

In summary, the kind of negotiation scenario we are interested in is characterised by four components : the set of agents $\mathcal{A}$, the set of resources $\mathcal{R}$, a collection $\{u_i : 2^{\mathcal{R}} \to \mathbb{R} \mid i \in \mathcal{A}\}$ of utility functions, and an initial resource allocation $A_0 : \mathcal{A} \to 2^{\mathcal{R}}$ for the system $(\mathcal{A}, \mathcal{R})$.

2.2 Deals

Agents can negotiate *deals* to exchange resources. An example would be : "I give you r_1 if you give me r_2". This would be a particularly simple deal, which only involves two agents and two resources. In general, any numbers of agents and resources could be involved in a single deal. For instance, an agent may only agree to exchange r_1 for r_2, if it can obtain another set of resources from a third agent during the same transaction. From an abstract point of view, a deal takes us from one allocation of resources to the next. That is, we may characterise a deal as a pair of allocations.

Definition 3 (Deals) *A deal is defined as a pair $\delta = (A, A')$ where A and A' are allocations of resources with $A \neq A'$.*

The intended interpretation of this definition is that the deal $\delta = (A, A')$ is only applicable for allocation A and will result in allocation A'. It thereby specifies for each resource in the system whether it is to remain where it has been before the deal or where it is to be moved to, respectively.

An agent may or may not find a particular deal acceptable. For instance, a selfish agent i may only agree to a deal $\delta = (A, A')$ iff $u_i(A) < u_i(A')$, that is, iff δ would strictly increase its individual welfare. This may be an appropriate policy for agents populating a society governed by utilitarian principles.[3] An acceptability criterion that is suitable for an egalitarian agent society will be introduced in Section 4.

The set of possible deals may also be restricted by the negotiation protocol in operation. Such a protocol may, for instance, only allow for deals that do not involve more than two agents at a time.[4]

3 Egalitarian welfare orderings

In this section we introduce two *social welfare orderings* over allocations of resources. Given the preference profiles of the individual agents in a society (which, in our scenario, are represented by means of their utility functions), a social welfare ordering formalises the notion of a society's 'preferences' [1]. We are going to make use of the egalitarian maximin- and the leximin-orderings, both of which are standard concepts in social choice theory and welfare economics (see, for instance, Moulin [7]).

3.1 Egalitarian social welfare

The first aim of an egalitarian society should be to maximise the welfare of its weakest member. In that sense, we can measure social welfare by measuring the welfare of the agent who is (currently) worst off. This idea leads to the definition of the following egalitarian social welfare function.

Definition 4 (Egalitarian social welfare) *The egalitarian social welfare $sw_e(A)$ of an allocation A is defined as follows :*

$$sw_e(A) \quad = \quad \mathit{min}\{u_i(A) \mid i \in \mathcal{A}\}$$

The function sw_e gives rise to a social preference ordering over different allocations of resources : allocation A' is strictly preferred over A iff $sw_e(A) < sw_e(A')$. This ordering is sometimes called the *maximin-ordering*.

An allocation A is said to have *maximal egalitarian social welfare* iff there is no other allocation A' such that $sw_e(A) < sw_e(A')$. The main technical objective of this paper is to investigate what are sufficient and necessary conditions for an agent society to be able to reach an allocation of resources that has maximal egalitarian social welfare by means of negotiation, that is, by agreeing on a sequence of deals to exchange resources.

3.2 The leximin-ordering

The maximin-ordering induced by sw_e only takes into account the welfare of the currently weakest agent, but is insensitive to utility fluctuation in the rest of society. To allow for a finer distinction of the social welfare of different allocations we introduce the so-called *leximin-ordering*.

Let $n = |\mathcal{A}|$ be the number of agents in the system $(\mathcal{A}, \mathcal{R})$. Then every allocation A determines

[3] In what sense such a policy (or slight variations of it) would be appropriate in a utilitarian negotiation framework has been discussed in [3].

[4] However, as Theorem 2 in Section 5 will show, any such restriction to the protocol may prevent agents from being able to agree on an allocation of resources that is optimal (in a sense to be made precise in the next section).

a *utility vector* $\langle u_1(A), \ldots, u_n(A)\rangle$ of length n. If we rearrange the elements of that vector in increasing order we obtain the *ordered utility vector* for allocation A, which we are going to denote by $\vec{u}(A)$. We now declare a *lexicographic ordering* over vectors of real numbers (such as $\vec{u}(A)$) in the usual way : $\vec{x}$ lexicographically precedes $\vec{y}$ iff $\vec{x}$ is a (proper) prefix of $\vec{y}$ or $\vec{x}$ and $\vec{y}$ share a common (proper) prefix of length k (which may be 0) and we have $\vec{x}_{k+1} < \vec{y}_{k+1}$.

Definition 5 (Leximin-ordering) *The leximin-ordering $\prec$ over alternative allocations A and A' is defined as follows : $A \prec A'$ holds iff $\vec{u}(A)$ precedes $\vec{u}(A')$ in the lexicographic ordering over vectors of real numbers.*

We write $A \preceq A'$ iff either $A \prec A'$ or $\vec{u}(A) = \vec{u}(A')$ hold. An allocation A is called *leximin-maximal* iff there is no allocation A' such that $A \prec A'$ holds.

Let us note some simple consequences of Definitions 4 and 5. It is easily seen that $sw_e(A) < sw_e(A')$ implies $A \prec A'$, because the former requires already the element at the *first* position in the ordered utility vector of A to be smaller than that of the ordered utility vector of A'. Also note that $A \preceq A'$ implies $sw_e(A) \leq sw_e(A')$.[5] Finally, every leximin-maximal allocation has maximal egalitarian social welfare, but not vice versa.

4 Acceptable deals

In this section we are going to introduce a criterion that (egalitarian) agents may use to decide whether or not to accept a particular deal. We also show how this local acceptability criterion relates to the global notions of egalitarian social welfare and the leximin-ordering, respectively.

4.1 Equitability

Intuitively, agents operating according to egalitarian principles should help any of their fellow agents that are worse off than they are themselves (as long as they can afford to do so without themselves ending up even worse). This means, the purpose of any exchange of resource should be to improve the welfare of the weakest agent involved in the respective deal. This is precisely how we define our local acceptability criterion. We call a deal *equitable* iff it increases the minimum utility amongst the agents involved in it.

Definition 6 (Equitable deals) *Let $\delta = (A, A')$ be a deal and define the set of agents involved in δ as $\mathcal{A}^\delta = \{i \in \mathcal{A} \mid A(i) \neq A'(i)\}$. We call δ an equitable deal iff the following holds :*

$$\mathit{min}\{u_i(A) \mid i \in \mathcal{A}^\delta\} \; < \; \mathit{min}\{u_i(A') \mid i \in \mathcal{A}^\delta\}$$

Recall that, for $\delta = (A, A')$ to be a deal, we require $A \neq A'$, that is, $\mathcal{A}^\delta$ will not be the empty set.

4.2 Pigou-Dalton utility transfers

Our definition of equitable deals provides a criterion that allows agents to evaluate the acceptability of a particular deal at a local level. A related notion that can be found in the economic literature is the a so-called *Pigou-Dalton transfer* [7]. The Pigou-Dalton *principle* states that whenever a utility transfer between two agents takes place which reduces the difference in utility between the two, then that transfer should not be considered as reducing social welfare. Translating into our terminology, a Pigou-Dalton transfer can be characterised as a deal $\delta = (A, A')$ involving only two agents i and j that has the following properties :

(1) $u_k(A) = u_k(A')$ for all $k \in \mathcal{A}$ with $k \neq i$ and $k \neq j$,

(2) $u_i(A) + u_j(A) = u_i(A') + u_j(A')$, and

(3) $|u_i(A') - u_j(A')| < |u_i(A) - u_j(A)|$.

That is, while utility values of agents other than i and j as well as the sum of the utility values for all the agents in the society stay constant, the utility values of i and j move closer together. It is clear from this definition that any Pigou-Dalton transfer will also be an equitable deal, because it will always result in an improvement for the weaker one of the two agents concerned. The converse, however, does not hold (not even if we restrict ourselves to deals involving only two agents). In fact, equitable deals may even increase the inequality of the agents concerned, namely in cases where the happier agent gains more utility than the weaker does.

[5] Here is a proof : $A \preceq A'$ implies $A' \not\prec A$ (because $\preceq$ is a total order), which implies $sw_e(A') \not< sw_e(A)$ (by our earlier observation), which in turn implies $sw_e(A) \leq sw_e(A')$.

Agent 1	Agent 2	Agent 3
$A_0(1) = \{\}$	$A_0(2) = \{\}$	$A_0(3) = \{r_1\}$
$u_1(\{\}) = 0$ $u_1(\{r_1\}) = 5$	$u_2(\{\}) = 6$ $u_2(\{r_1\}) = 7$	$u_3(\{\}) = 8$ $u_3(\{r_1\}) = 9$

TAB. 1 – An example with a single resource

4.3 Local actions and their global effects

We are now going to prove two lemmas that provide the connection between the local acceptability criterion given by the notion of equitability and the two global notions of social welfare discussed in the previous section.

The first lemma shows how global changes are reflected locally. If a deal happens to increase (global) egalitarian social welfare, that is, if it results in a rise with respect to the maximin-ordering, then that deal will in fact be an equitable deal.

Lemma 1 (Maximin-rise implies equitability) *Let A and A' be resource allocations with $sw_e(A) < sw_e(A')$. Then $\delta = (A, A')$ is an equitable deal.*

Proof. Let A and A' be allocations with $sw_e(A) < sw_e(A')$ and let $\mathcal{A}^\delta$ be the set of agents involved in deal $\delta = (A, A')$ as defined in Definition 6. Any agent with minimal utility for allocation A must be involved in δ, because social welfare, and thereby their individual utility, is higher for allocation A'. That is, we have $\mathit{min}\{u_i(A) \mid i \in \mathcal{A}^\delta\} = sw_e(A)$. Furthermore, because of $\mathcal{A}^\delta \subseteq \mathcal{A}$, we certainly have $sw_e(A') \leq \mathit{min}\{u_i(A') \mid i \in \mathcal{A}^\delta\}$. Together with our original assumption of $sw_e(A) < sw_e(A')$ we now get $\mathit{min}\{u_i(A) \mid i \in \mathcal{A}^\delta\} < \mathit{min}\{u_i(A') \mid i \in \mathcal{A}^\delta\}$, i.e. δ will indeed be an equitable deal. □

Observe that the converse does not hold ; not every equitable deal will necessarily increase egalitarian social welfare. This is for instance not the case if only agents who are currently better off are involved in a deal.

We will illustrate this point by means of an example. To make this example simpler, we will draw an intuitive picture where the only resource r_1 is a moderate amount of money. Then we will assume that agents have different backgrounds : agent 3 is very rich and is already happy without the money, agent 2 is pretty rich and happy, and agent 1 is poor and unhappy without the money. Having the money would make each of these agents happier. However, to make the picture even more striking (although this is not necessary to exemplify our case), we can assume that obtaining the sum of money would involve a higher utility gain for the poorest agent. The example is pictured in Table 1. The social welfare for this allocation of resources is 0. It is easy to see that passing r_1 from agent 3 to agent 1 (which is of course an equitable deal) would increase the social welfare to 5. But it is also an equitable deal for agent 3 to pass the resource r_1 to agent 2. This move, however, has no influence on the social welfare of our egalitarian agent society, as the poorest agent's utility remains unchanged.

In fact, there can be no class of deals (that could be defined without reference to the *full* set of agents in a society) that will always result in an increase in egalitarian social welfare. This is a consequence of the fact that the maximin-ordering induced by sw_e is not separable [7].[6]

To be able to detect changes in welfare resulting from an equitable deal we require the finer differentiation between alternative allocations of resources given by the leximin-ordering. In fact, as we shall see next, any equitable deal can be shown to result in a strict improvement with respect to the leximin-ordering.

Lemma 2 (Equitability implies leximin-rise) *Let $\delta = (A, A')$ be an equitable deal. Then $A \prec A'$ holds.*

Proof. Let $\delta = (A, A')$ be an equitable deal. We define the set $\mathcal{A}^\delta$ of agents involved in δ as in Definition 6 and set $\alpha = \mathit{min}\{u_i(A) \mid i \in \mathcal{A}^\delta\}$. The value α may be considered as partitioning the ordered utility vector $\vec{u}(A)$ into three subvectors. To begin with, $\vec{u}(A)$ has got a (possibly empty) prefix $\vec{u}(A)^{<\alpha}$ where all elements

[6] A social welfare ordering is called *separable* iff the effect of a local welfare redistribution with respect to that ordering (rise or fall) is independent of non-concerned agents.

are strictly lower than α. In the middle, it has got a subvector $\vec{u}(A)^{=\alpha}$ (with at least one element) where all elements are equal to α. Finally, $\vec{u}(A)$ has got a suffix $\vec{u}(A)^{>\alpha}$ (which again may be empty) where all elements are strictly greater than α.

By definition of α, the deal δ cannot affect agents whose utility values belong to $\vec{u}(A)^{<\alpha}$. Furthermore, by definition of equitability, we have $\alpha < min\{u_i(A') \mid i \in \mathcal{A}^\delta\}$, which means that all of the agents that *are* involved will end up with a utility value which is strictly greater than α, and at least one of these agents will come from $\vec{u}(A)^{=\alpha}$. We now collect the information we have on $\vec{u}(A')$, the ordered utility vector of A'. Firstly, it will have a prefix $\vec{u}(A')^{<\alpha}$ identical to $\vec{u}(A)^{<\alpha}$. This will be followed by a (possibly empty) subvector $\vec{u}(A')^{=\alpha}$ where all elements are equal to α and which must be strictly shorter than $\vec{u}(A)^{=\alpha}$. All of the remaining elements of $\vec{u}(A')$ will be strictly greater than α. It follows that $\vec{u}(A)$ lexicographically precedes $\vec{u}(A')$, i.e. $A \prec A'$ holds as claimed. $\Box$

Again, the converse does not hold, i.e. not every deal resulting in a leximin-rise is necessarily equitable. Counterexamples are deals where the utility value of the weakest agent involved stays constant, despite there being an improvement with respect to the leximin-ordering at the level of society.

A well-known result in welfare economics states that every Pigou-Dalton utility transfer results in a leximin-rise [7]. Given that we have observed earlier that every deal that amounts to a Pigou-Dalton transfer will also be an equitable deal, this result can now be seen to also be a simple corollary to Lemma 2.

5 Optimal outcomes

In this section we are going to prove the two main technical results of this paper : (i) equitable deals are *sufficient* to guarantee outcomes of negotiation with maximal social welfare, and (ii) all deals are also *necessary*, provided only equitable deals are allowed.

5.1 Termination

We first show that, as long as agents only agree on deals that are equitable, negotiation will always terminate, i.e. after a finite number of equitable deals no further equitable deals will be possible.

Lemma 3 (Termination) *There can be no infinite sequence of equitable deals.*

Proof. Given that both the set of agents $\mathcal{A}$ as well as the set of resources $\mathcal{R}$ in a negotiation system are required to be finite, there can only be a finite number of distinct allocations. Furthermore, by Lemma 2, any equitable deal will result in a strict rise with respect to the leximin-ordering $\prec$. Hence, negotiation will have to terminate after a finite number of deals. $\Box$

5.2 Guaranteed optimal outcomes

The proof of the following theorem shows that equitable deals are sufficient for agents to reach an allocation of resources with maximal egalitarian social welfare. In fact, the result is even stronger than this : *any* sequence of equitable deals will eventually result in an optimal allocation. That is, agents may engage 'blindly' into negotiation. Whatever their course of action, provided they restrict themselves to equitable deals, once they reach an allocation where no further equitable deals are possible, that allocation is bound to have maximal welfare.

Theorem 1 (Maximal social welfare) *Any sequence of deals that are equitable will eventually result in an allocation of resources with maximal egalitarian social welfare.*

Proof. By Lemma 3, negotiation will eventually terminate if all deals are required to be equitable. So suppose negotiation has terminated and no more equitable deals are possible. Let A be the corresponding terminal allocation of resources. The claim is that A will be an allocation with maximal egalitarian social welfare. For the sake of contradiction, assume it is not, i.e. assume there exists another allocation A' for the same system such that $sw_e(A) < sw_e(A')$. But then, by Lemma 1, the deal $\delta = (A, A')$ will be an equitable deal. Hence, there is still a possible deal, namely δ, which contradicts our earlier assumption of A being a terminal allocation. This shows that A will be an allocation with maximal egalitarian social welfare, which proves our claim. $\Box$

After having reached the allocation with maximal egalitarian social welfare, it may be the case

Agent 1	Agent 2	Agent 3
$A_0(1) = \{\}$	$A_0(2) = \{\}$	$A_0(3) = \{r_1, r_2\}$
$u_1(\{\}) = 0$	$u_2(\{\}) = 6$	$u_3(\{\}) = 8$
$u_1(\{r_1\}) = 5$	$u_2(\{r_1\}) = 7$	$u_3(\{r_1\}) = 9$
$u_1(\{r_2\}) = 0$	$u_2(\{r_2\}) = 6.5$	$u_3(\{r_2\}) = 8.5$
$u_1(\{r_1, r_2\}) = 5$	$u_2(\{r_1, r_2\}) = 7.5$	$u_3(\{r_1, r_2\}) = 9.5$

TAB. 2 – An example with two resources

that still some equitable deals are possible, although they would not increase social welfare any further (but they would still cause a leximin-rise).

This, again, can be shown by means of a simple example. Let us reuse the setting of our earlier example (see Table 1) and slightly modify it by adding a second resource which is also initially allocated to agent 3 (say, a book that gives advice on how to avoid paying taxes). Both agents 2 and 3 would be happier with the book, but agent 1 does not care about this resource since it does not have to pay taxes anyway. This can be represented by appropriate utility functions, as shown in Table 2.

The social welfare for the initial allocation of resources A_0 is 0 and the corresponding ordered utility vector is $\vec{u}(A_0) = \langle 0, 6, 9.5\rangle$. Passing r_1 from agent 3 to agent 1 would lead to a new allocation with the ordered utility vector $\langle 5, 6, 8.5\rangle$ and increase the social welfare to 5, which is indeed the maximal social welfare that can be achieved by this particular society. However, there is still another equitable deal that can be processed from this latter allocation : agent 3 could offer the book to agent 2. Of course, this deal does not affect agent 1. The resulting allocation would then have the ordered utility vector $\langle 5, 6.5, 8\rangle$, which corresponds to the leximin-maximal allocation.

To be able to detect situations where a social welfare maximum has already been reached but some equitable deals are still possible, and to be able to stop negotiation (assuming we are only interested in maximising sw_e as quickly as possible), however, we would require a *global* criterion.[7]

We could define a class of *strongly equitable* deals that are like equitable deals but on top of that require the (currently) weakest agent to be involved in the deal. This would be a sharper criterion, but it would also be against the spirit of distributivity and locality, because every single agent would be involved in every single deal (in the sense of everyone having to announce their utility in order to be able to determine who is the weakest).

[7] This is again a consequence of the fact that the maximin-ordering is not separable. No measure that only takes into account the welfare of agents involved in a particular deal could be strong enough to always tell us whether or not the deal in question will result in an increase in social welfare (see also our discussion after Lemma 1 in Section 4).

5.3 Necessity of complex deal types

As our second theorem will show, if we restrict the set of admissible deals to those that are equitable, then every single deal δ may be necessary to guarantee an optimal result (that is, no sequence of equitable deals excluding δ could possibly result in an allocation with maximal egalitarian social welfare). This emphasises the high complexity of our negotiation scenarios.

Theorem 2 (Necessity) *Let the sets of agents and resources be fixed. Then for every deal δ there are utility functions and an initial allocation of resources such that δ is necessary to reach an allocation with maximal egalitarian social welfare, provided only equitable deals are admitted.*

Proof. Given a set of agents $\mathcal{A}$ and a set of resources $\mathcal{R}$, let $\delta = (A, A')$ with $A \neq A'$ be any deal for this system. We need to show that there are a collection of utility functions for the agents in $\mathcal{A}$ as well as an initial allocation of resources such that δ is necessary for the agent society to be able to move to an allocation with maximal egalitarian social welfare.

As we have $A \neq A'$, there will be a (at least one) agent $j \in \mathcal{A}$ with $A(j) \neq A'(j)$. We use this particular j to fix suitable utility functions u_i for agents $i \in \mathcal{A}$ and sets of resources $R \subseteq \mathcal{R}$ as follows :

$$u_i(R) = \begin{cases} 2 & \text{if } R = A'(i) \\ 2 & \text{if } R = A(i) \text{ and } i \neq j \\ 1 & \text{if } R = A(i) \text{ and } i = j \\ 0 & \text{otherwise} \end{cases}$$

That is, for allocation A' every agent assigns a utility value of 2 to the resources it holds. The same is true for allocation A, with the sole exception of agent j, who only assigns a value of 1. For any other allocation, agents assign the value of 0 to their set of resources, unless that set is the same as for either allocation A or A'. Hence, for every allocation other than A or A' at least one agent will assign a utility value of 0 to its allocated set of resources. We get $sw_e(A') = 2$, $sw_e(A) = 1$, and $sw_e(B) = 0$ for every other allocation B, i.e. A' is the only allocation with maximal egalitarian social welfare.

The ordered utility vector of A' is of the form $\langle 2, \ldots, 2 \rangle$, that of A is of the form $\langle 1, 2, \ldots, 2 \rangle$, and that of any other allocation has got the form $\langle 0, \ldots \rangle$, i.e. we have $A \prec A'$ and $B \prec A$ for all allocations B with $B \neq A$ and $B \neq A'$. Therefore, if we make A the initial allocation of resources, then δ will be the only deal that would result in a rise with respect to the leximin-ordering. Thus, by Lemma 2, δ would also be the only equitable deal. Hence, if the set of admissible deals is restricted to equitable deals then δ is indeed necessary to reach an allocation with maximal egalitarian social welfare. $\Box$

An important consequence of this result is that there can be no simple class of deals (such as the class of deals only involving two agents at a time) that would be sufficient to guarantee an optimal outcome of negotiation.

5.4 Specific utility functions

In our previous work on resource allocation by negotiation in the utilitarian setting [3], we have shown that the optimal outcome of a negotiation process may be guaranteed even when we admit only very specific types of deals, provided that we put suitable restrictions on the class of utility functions that agents may use to represent their valuation of different sets of resources. In the egalitarian setting, to date, we have not been able to establish similar results.

Even the (arguably) strongest restrictions used in the utilitarian case do not allow us to eliminate any type of deal in the egalitarian framework. Let us consider the example of *0-1 additive utility functions*, where agents can only assign the values 1 or 0 to single resources (simply distinguishing whether or not they *need* a particular resource) and where the utility value for a set of resources is always the sum of the values assigned to the single resources in that set. As shown in [3], this restriction guarantees an optimal outcome of negotiation for the utilitarian framework, even when the only deals allowed are those where a single resource is being transferred from one agent to another (that is, no deal may involve more than two agents or more than one resource at a time).

This result does not hold anymore for egalitarian agent societies. Counterexamples can easily be constructed. Take, for instance, a scenario of three agents furnishing their flats. *Ann* needs a *picture* and has a *desk* which she does not need. *Bob* needs a *desk* and a *chair*, but only has the *chair*. *Carlos* needs a *picture*, a *chair*, and a *cushion*, and he only owns the *picture* and the *cushion* at the beginning of the negotiation process. The ordered utility vector for this allocation is $\langle 0, 1, 2 \rangle$. However, in the situation where *Ann* has the *picture*, *Bob* the *desk* instead of the *chair*, and *Carlos* the *chair* and the *cushion* is better ; the corresponding ordered utility vector would be $\langle 1, 1, 2 \rangle$. Unfortunately, only a very complex equitable deal (involving all three agents, namely *Carlos* giving the *picture* to *Ann*, *Ann* giving the *desk* to *Bob*, and *Bob* giving the *chair* to *Carlos*) would allow this agent society to reach the preferred allocation of resources.

6 Conclusion

In this paper we have argued that egalitarian social welfare can be an interesting concept in the context of negotiation in multiagent systems. Specifically, we have shown the the notion of equitability serves as a suitable acceptability criterion for agents operating in an egalitarian environment and proved sufficiency and necessity results along the lines of those established for a similar utilitarian framework by the present authors in [3] and by Sandholm in [11].

6.1 Welfare engineering

The approach followed in this paper may be regarded as a kind of '*welfare engineering*'. We have chosen a global welfare measure (such as utilitarian or egalitarian social welfare) appropriate for a given application domain and constructed local acceptability criteria accordingly, which allow an agent society to reach allocations considered optimal by means of negotiation and in a distributed fashion.

We conclude by introducing another interesting welfare measure that we hope to investigate in more detail in our future work.

6.2 Elitist societies

In this paper we have focussed on egalitarian social welfare orderings. This kind of social welfare function is actually a particular case of a class of functions sometimes called *k-rank* dictators [7], where a particular agent of the society (the one corresponding to the kth element in the ordered utility vector) is chosen to be the representative of the society. Amongst this class of functions, another particularly interesting case is where the welfare of society is evaluated on the basis of the happiest agent (as opposed to the unhappiest agent, as in the case of egalitarian welfare). In such *elitist* societies, agents would cooperate in order to support their champion (the happiest agent).

While such an approach to social welfare may seem somewhat unethical as far as human society is concerned,[8] we believe that it could indeed be very appropriate for certain societies of artificial agents. A typical scenario could be where a system designer launches different agents with the same goal, with the aim that *at least one* agent achieves that goal—no matter what happens to the others.

Acknowledgements. We would like to thank the two anonymous referees for their comments on the paper, and particularly Jérôme Lang for his many valuable suggestions and pointers to the literature on welfare economics. This research has been funded by the European Union as part of the SOCS project (Societies Of ComputeeS), IST-2001-32530.

References

[1] Kenneth J. Arrow. *Social Choice and Individual Values*. John Wiley and Sons, 2nd edition, 1963.

[2] Craig Boutilier, Moisés Goldszmidt, and Bikash Sabata. Sequential Auctions for the Allocation of Resources with Complementaries. In *Proceedings of the 16th International Joint Conference on Artificial Intelligence (IJCAI-1999)*, pages 527–523. Morgan Kaufmann, 1999.

[3] Ulrich Endriss, Nicolas Maudet, Fariba Sadri, and Francesca Toni. On Optimal Outcomes of Negotiations over Resources. In *Proceedings of the 2nd International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2003)*, 2003. To appear.

[4] John C. Harsanyi. Can the Maximin Principle Serve as a Basis for Morality ? A Critique of John Rawls' Theory. *American Political Science Review*, 69 :594–609, 1975.

[5] Gregory E. Kersten, Sunil J. Noronha, and Jeffrey Teich. Are All E-Commerce Negotiations Auctions ? In *Proceedings of the 4th International Conference on the Design of Cooperative Systems*, 2000.

[6] Sarit Kraus. *Strategic Negotiation in Multiagent Environments*. MIT Press, 2001.

[7] Hervé Moulin. *Axioms of Cooperative Decision Making*. Cambridge University Press, 1988.

[8] John Rawls. *A Theory of Justice*. Oxford University Press, 1971.

[9] Jeffrey S. Rosenschein and Gilad Zlotkin. *Rules of Encounter*. MIT Press, 1994.

[10] Michael H. Rothkopf. Bidding in Simultaneous Auctions with a Constraint on Exposure. *Operations Research*, 25(4) :620–629, 1977.

[11] Tuomas W. Sandholm. Contract Types for Satisficing Task Allocation : I Theoretical Results. In *AAAI Spring Symposium : Satisficing Models*, 1998.

[12] Tuomas W. Sandholm. Algorithm for Optimal Winner Determination in Combinatorial Auctions. *Artificial Intelligence*, 135 :1–54, 2002.

[13] Amartya K. Sen. *Collective Choice and Social Welfare*. Holden Day, 1970.

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

[8] But note that the elitist welfare ordering is not uncommon in the world of sports : During the *Tour de France*, for instance, cyclists are expected to support the member of their team who has currently the best chances to win the *maillot jaune*.