How Equitable is Rational Negotiation?

Sylvia Estivie[1], Yann Chevaleyre[1], Ulle Endriss[2] and Nicolas Maudet[1]
[1]LAMSADE, Université Paris-Dauphine, France
[2]ILLC, University of Amsterdam, The Netherlands
[1]{estivie,chevaley,maudet}@lamsade.dauphine.fr, [2]ulle@illc..uva.nl

ABSTRACT

Notions of fairness have recently received increased attention in the context of resource allocation problems, pushed by diverse applications where not only pure utilitarian efficiency is sought. In this paper, we study a framework where allocations of goods result from distributed negotiation conducted by autonomous agents implementing very simple deals. Assuming that these agents are strictly self-interested, we investigate how equitable the outcomes of such negotiation processes are. We first discuss a number of methodological issues raised by this study, pertaining in particular to the design of suitable payment functions as a means of distributing the social surplus generated by a deal amongst the participating agents. By running different experiments, we finally identify conditions favouring equitable outcomes.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Multiagent systems*; J.4 [**Social and Behavioral Sciences**]: Economics

General Terms

Economics, Theory, Experimentation

Keywords

Multiagent resource allocation, Negotiation, Social welfare, Fair division

1. INTRODUCTION

In many applications of multiagent systems it is important to achieve a *fair*, or *equitable*, allocation of the available resources amongst the agents in the system. However, this fairness requirement will often compete with the *rational* interests of individual agents. If agents are allowed to freely negotiate their preferred allocations of resources, the emerging allocations will typically be good on average, but give a rather low payoff to an unfortunate minority. In this paper, we investigate to what degree it is possible to reconcile rationality and equitability requirements. We do this by running a number of simulations where agents negotiate in a rational manner and where we track the evolution of allocations in view of their equitability.

AAMAS'06 May 8–12 2006, Hakodate, Hokkaido, Japan.

Negotiation Problems

Take a set of (indivisible) goods $\mathcal{R}$ and a set of agents $\mathcal{A}$. A resource *allocation* is a partitioning of the resources of the set $\mathcal{R}$ amongst the n agents in $\mathcal{A}$ (that is, every resource has to be allocated to a unique agent). As an example, allocation A, defined through $A(i) = \{r_1\}$ 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. Each agent has individual preferences regarding the bundles of resources it may hold. We are going to model the preferences of agents by means of *utility functions* mapping bundles of resources to real numbers. Assuming that agents are only concerned with resources they personally own, we will use the abbreviation $u_i(A)$ for $u_i(A(i))$, representing the utility value assigned by i to the bundle it holds for allocation A.

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

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

From a designer's perspective, we will be interested in assessing the well-being of the whole society, or its *social welfare* [1], which is often defined as the sum of utilities of all the agents:

$$sw_u(A) \;=\; \sum_{i \in \mathcal{A}} u_i(A) \tag{1}$$

This is the *utilitarian* definition of social welfare (sw_u), but other criteria of assessment exist as well. In the next section, for instance, we are going to introduce the notion of *egalitarian social welfare*, which is a measure for the equitability of a given allocation. In this paper, we focus in general on equitability criteria. To decide how to (best) allocate the resources amongst the agents in a system is to solve a resource allocation problem.

Resource allocation problems are typically addressed from a centralized point of view (this is the case for instance in combinatorial auctions [4]). When no central authority is available, or when the computational burden would overwhelm a single agent, this approach is not well suited. A different approach, advocated by several authors [15, 7, 5], consists of distributing this process.

Distributed Approach

In the distributed perspective, instead of trying to centrally compute the optimal allocation, we shall progress in a stepwise manner towards an optimum. To do so, agents may agree on a *deal* to exchange some of the resources they possess. It transforms the current allocation of resources A into a new allocation A'; that is, we can define a deal as a pair $\delta = (A, A')$ of allocations (with $A \neq A'$). When such a deal improves the utilitarian social welfare (*i.e.* when $sw_u(A') > sw_u(A)$), we say that it is *socially beneficial*. Note that a single deal may involve the displacement of any number of resources between any number of agents. This level of generality is hardly realistic in practice. Sandholm [15] 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. In this paper, we shall only consider the simplest type of deals (1-deals), *i.e.* those involving only a single resource (and thereby only two agents). The set of agents involved in the deal δ will be denoted as $\mathcal{A}^\delta$.

The above is a condition on the *structure* of a deal. Other conditions relate to the *acceptability* of a deal to a given agent. We assume that agents are *rational* in the sense of aiming to maximise their individual welfare. Furthermore, agents are assumed to be *myopic*. This means that agents will not accept deals that would reduce their level of welfare, not even temporarily, because they are either not sufficiently able to plan ahead or not willing to take the associated risk (see also [15] for a justification of such an agent model). Also agents will not behave strategically by, for instance, postponing a rational deal in the hope of finding an even better opportunity. We will, however, permit agents to enhance deals with *monetary side payments*, in order to compensate for a possible loss a utility. This is modelled using a so-called *payment function* $p : \mathcal{A} \to \mathbb{R}$ satisfying $\sum_{i \in \mathcal{A}} p(i) = 0$. A positive value $p(i)$ indicates that agent i *pays* money, while a negative value means that that agent *receives money*. In summary, the following rationality criterion will define the acceptability of deals:

Definition 2 (Rationality). *A deal $\delta = (A, A')$ is individually rational iff there exists a payment function p such that $u_i(A') - u_i(A) > p(i)$ for all $i \in \mathcal{A}$, except possibly $p(i) = 0$ for agents i with $A(i) = A'(i)$.*

The following result, due to Endriss et al. [7], makes the connection between the local decisions of agents and the global behaviour of the system explicit:

Lemma 1. *A deal $\delta = (A, A')$ is individually rational iff $sw_u(A) < sw_u(A')$.*

A corollary of Lemma 1, originally due to Sandholm [15], is that *any* sequence of rational deals will eventually result in an allocation of resources with maximal social welfare. However, deals involving any number of agents and resources may be required to do so [15, 7]. In particular, 1-deals alone are certainly *not* sufficient for negotiation between agents that are not only rational but also myopic. Of course, for some particular negotiation problems, rational 1-deals *will* be sufficient. This is, for instance, the case when all agents are using *modular* utility functions [7].[1]

Paper Overview

The remainder of this paper is structured as follows. Section 2 presents the challenges posed by the study of egalitarian outcomes of distributed negotiation. Different types of payment functions are then introduced (Section 3), and we discuss how the choice of a particular function would affect the egalitarian nature of negotiation outcomes. Our experimental setup is introduced in Section 4. We then report results on different series of tests, both in modular (Section 5) and non-modular domains (Section 6). Section 7 concludes.

2. EGALITARIAN FRAMEWORK

While the utilitarian interpretation of the concept of social welfare is the definition usually adopted in the multiagent systems literature [16, 17], there are also several other formal tools for assessing the level of welfare of a society of agents that have been developed in the socio-economic sciences [1, 3, 12] and that have potential applications in the context of multiagent resource allocation as well [7].

Efficiency and Fairness

The well-known Pareto condition and the utilitarian social welfare ordering address the *efficiency* of an allocation of resources. In many situations, however, the efficiency of an outcome alone is not sufficiently accurate a criterion to assess the quality of an allocation of goods. Rather, the quest for efficiency needs to be balanced with certain *fairness* requirements.

This is, for instance, the case when a multiagent resource allocation system is used to automatically allocate airport slots to different airlines [9]: we want to minimise the overall delay of all flights (efficiency), without risking to completely ignore the requests of a particular flight that may be difficult to integrate into the current schedule (fairness). Another example are reverse combinatorial auctions for industrial procurement, where a single buyer solicits offers for, say, the parts required to assemble 500 cars [8]. Here it may be in the interest of the buyer to try to contract deals with more than one seller, even when that agreement does not entail the lowest possible price, in order to avoid being dependent on the good will of a single provider. This kind of safety constraint (from the viewpoint of the buyer) may also be interpreted as a fairness requirement (from the viewpoint of the competing sellers). The case of earth observation satellites discussed by Lemaître et al. [10] offers a particularly interesting and original example. In this application, several stakeholder (e.g. countries) are jointly funding a satellite that is circling the earth and can take photos, subject to various technical constraints. Different stakeholders request different (and often incompatible) photos, and a solution needs to be found that is both efficient (in the sense of

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

exploiting the satellite as well as possible) and fair (in the sense of honouring an appropriate share of the requests of each of the stakeholders). Finally, Porter et al. [13] justify their study of fairness by invoking the following application: In many countries, in crisis situations, the state can commandeer diverse resources (e.g. aircrafts) from companies to cope with the emergency. However, this should be done in such a way that it minimizes the total true cost incurred by airline carriers (and by the state), and distribute this cost fairly amongst all carriers (so as not to penalize a carrier).

What these examples show is the relevance of notions such as *fairness*, *equitability*, or *envy-freeness* in many (diverse) applications. This paper will focus on what may be considered the most "basic" measure of the equitability of given allocations of goods, known as the *egalitarian social welfare*. The egalitarian social welfare is given by the utility of the agent that is currently worst off [1, 12]. That is, maximising egalitarian social welfare amounts to maximising the minimum utility in the system. This *maxmin principle* is also called 1-efficiency in the work of Porter et al. [13].

Optimal Outcomes and Local Deals

We are going to be interested in mechanisms that allow agents to reach an allocation with high (and ideally, maximal) egalitarian social welfare by means of a sequence of locally agreed deals. These deals should be *local* also and particularly in the sense that the agents should be able to decide locally within the group of agents about to agree on a given deal whether or not that deal should be deemed acceptable. In other words, we are not looking for just any algorithm to optimise with respect to egalitarian social welfare, but for a distributed negotiation scheme.

The first kind of deal we are going to consider are the *individually rational* deals introduced earlier. It is easy to show that individually rational deals alone are not sufficient to allow agents to find an allocation with maximal egalitarian social welfare in all cases (while, on the contrary, this is possible in the case of utilitarian social welfare [15, 7]). This can be seen in the following example. Assume there are two agents and two possible allocations (*i.e.* there may be just a single resource):

$$\begin{array}{ll} u_1(A) = 5 & u_1(A') = 2 \\ u_2(A) = 1 & u_2(A') = 3 \end{array}$$

Then A' has maximal egalitarian social welfare, but in case A is the initial allocation there is no individually rational deal leading from A to A' (because the amount of utility lost by agent 1 would be greater than the utility gained by agent 2).

Another approach would be to change the notion of individual rationality and to require agents to use a different criterion when deciding on the acceptability of a given deal. This idea has been followed in previous work where the class of so-called *equitable deals* has been proposed [7]. This criterion, which requires the minimum utility within the group of agents contracting a deal to increase, can guarantee convergence to an egalitarian optimum. However, as also pointed out in [7], this is not a precise criterion in the sense that it is still possible to contract further equitable deals after an allocation with maximal egalitarian social welfare has been found. In fact, it is *impossible* to design a deal acceptability criterion that would be precise in this sense. This is related to the fact that the egalitarian social welfare ordering is not *separable* [12], which means that it is not possible to decide whether or not egalitarian social welfare has increased simply by looking at the utility levels of the agents that have experienced a change in utility after a deal. Hence, it is also not possible for the agents involved in a deal to locally verify whether the deal they are about to implement would affect egalitarian social welfare in a positive manner.

Besides such technical considerations, of course, the individual rationality criterion is also much more realistic than the idealised concept of equitable deals. Therefore, in the present paper, we are going to follow a different approach. Rather than designing agents that are "good" for egalitarian social welfare, we are going to study just how "bad" the standard utility-maximising agents are from an egalitarian point of view. To put it differently, we are going to analyse how egalitarian social welfare develops in a society where agents negotiate a sequence of individually rational deals.

Egalitarian Social Welfare and Money

Before getting further in that direction, we need to address the first challenge this approach poses. Typically, a social welfare measure does *not* take a monetary component into account. There is an obvious reason for that when utilitarian social welfare is concerned: the payment function p has been defined in such a way that the sum of payments is always zero. Hence, side payments can never change the sum of utilities, *i.e.* they can never affect utilitarian social welfare. Things are different for egalitarian social welfare. Not taking into account monetary payments would amount to assessing buyers as being more satisfied than sellers, whatever the price they pay to get the resource(s) may be. In the most extreme case, an agent selling all its goods for very good (high) prices (and thereby certainly being "happy") would still provoke a very low egalitarian social welfare (because its utility function defined over bundles of resources alone would return a zero value).

To circumvent this problem, payments have to be included into agents' utility functions. Let the *balance* of agent i be defined as the sum of all the payments $p(i)$ of that agent paid out (or received, in the case of negative payments) during all previous deals. By Lemma 1, if agents only negotiate individually rational deals, no allocation can be visited more than once. So we can index an agent's balance with the name of the allocation in question: $bal^A(i)$ is the balance of agent i once allocation A has been reached. Then we define for each agent i a second utility function u'_i that combines the utility derived from the bundle of resources it currently holds and its monetary balance:

$$u'_i(A) \quad = \quad u_i(A) - bal^A(i) \tag{2}$$

In the literature, such utilities are often called *quasi-linear* utilities, in which case u would be referred to as a valuation and u' as the actual utility [4]. The egalitarian social welfare of an allocation with money is now simply defined as the egalitarian social welfare with respect to the utilities u'_i:

$$sw_e(A) \quad = \quad min\{u'_i(A) \mid i \in \mathcal{A}\} \tag{3}$$

3. PAYMENT FUNCTIONS

For any given deal that is individually rational, there will be a range of possible payment functions (in theory, there are infinitely many, in practice the number will depend on the granularity of the currency used). This raises the question what payments to implement in a simulation context, where agents do not actually negotiate prices with each other. Choosing a payment function p for an individually rational deal $\delta = (A, A')$ means deciding how to distribute the social surplus $sw_u(A') - sw_u(A)$ amongst the agents in the system. The only condition that the payment function has to meet is to ensure that side payments add up to zero.

Amongst these payment functions, some will ensure that any socially beneficial deal is indeed individually rational (and vice versa), that is that for each agent i involved in the deal, we indeed have $u_i(A') - u(A) > p(i)$: we shall call these *rational-compatible payment functions*.

Locally Uniform Payment Function

We begin with what we call the *locally uniform payment function*, which ensures an equal payoff amongst all the agents involved in the deal (and gives no payment to any of the non-involved agents):

$$p(i) \quad = \quad \begin{cases} u_i(A') - u_i(A) - \frac{sw_u(A') - sw_u(A)}{|\mathcal{A}^\delta|} & \text{if } i \in \mathcal{A}^\delta \\ 0 & \text{otherwise} \end{cases}$$

Note that these payments do indeed add up to 0 as required, and that each agent $i \in \mathcal{A}^\delta$ receives the same positive payoff $u_i(A') - u_i(A) - p(i) = \frac{sw_u(A') - sw_u(A)}{|\mathcal{A}^\delta|}$.

The uniform payment function is not an arbitrary choice, but intended to be a reasonable approximation of actual payments that agents would agree upon in a real negotiation context (rather than a simulation). Equally dividing the social surplus amongst all the agents concerned also means maximising the product of individual payoffs. In the case of deals between just two agents, this corresponds to the result we would get if agents were to adopt the well-known Zeuthen strategy to negotiate the payments for a given deal [14]. This is true, at least, if we assume that agents calculate their "willingness to risk conflict" with respect to the payoff for the deal at hand (which is a reasonable assumption if agents have comparable levels of utility to begin with).

Globally Uniform Payment Function

The next type of payment that we introduce is one which ensures an equal payoff amongst *all* the agents within the society:

$$p(i) \quad = \quad u_i(A') - u_i(A) - \frac{sw_u(A') - sw_u(A)}{|\mathcal{A}|}$$

Observe that this function is also rational-compatible (for the same reason that the locally uniform one is), although agents involved in deals will typically receive a much more modest share of the social surplus. Note that this payment function is not inequality reducing, but not inequality increasing either. By distributing the social surplus equally amongst all agents of the society, everyone simply benefits with any deal and enjoys the same welfare increase.

It is easy to see that with such a payment function, the egalitarian social welfare depends both on the gain of utilitarian social welfare during the entire negotiation, and on the initial satisfaction of the agents. More precisely, consider a sequence $[A_0, A_1, \ldots, A_f]$ of allocations visited during a negotiation process. For all $t \in \{0..f{-}1\}$, we have (where $p_t(i)$ denotes the payment of agent i at time t):

$$u_i(A_{t+1}) - u_i(A_t) - p_{t+1}(i) = \frac{sw_u(A_{t+1}) - sw_u(A_t)}{\mid \mathcal{A} \mid}$$

Observe that:

$$u_i(A_f) - bal^{A_f}(i) = u_i(A_0) + \frac{sw_u(A_f) - sw_u(A_0)}{\mid \mathcal{A} \mid}$$

Thus, we finally have:

$$sw_e(A_f) = min_i \left\{ \frac{sw_u(A_f) - sw_u(A_0)}{\mid \mathcal{A} \mid} + u_i(A_0) \right\}$$

It is interesting to note that:

$$sw_e(A_f) - sw_e(A_0) = \frac{sw_u(A_f) - sw_u(A_0)}{\mid \mathcal{A} \mid}$$

This has two important consequences: first, the gain of egalitarian social welfare is proportional to the gain of utilitarian social welfare. Also, the optimal allocation(s) w.r.t. utilitarian social welfare is/are the *same* as the optimal allocation(s) w.r.t. egalitarian social welfare.

It should also be pointed out that in the case of two agents, locally uniform payment functions are of course equivalent to globally ones. As reported later on in this paper, we carried out experiments to find out from what number of agents on do locally uniform payment functions give rise to negotiations exhibiting significantly different outcomes (see Fig. 3).

Fully Locally Equitable Payment Function

In theory, it is possible to conceive an equitable payment function which would, at each step, compute a payment such that each agent (involved in the deal) would enjoy the same utility level after the deal has been achieved. This idea is very similar to that of *fair imposition* proposed in [13]. Such a payment would take the following form:

$$p(i) = \begin{cases} u_i(A') - bal^A(i) - \frac{\sum_{j \in \mathcal{A}^\delta}(u_j(A') - bal^A(j))}{|\mathcal{A}^\delta|} & \text{if } i \in \mathcal{A}^\delta \\ 0 & \text{otherwise} \end{cases}$$

Note that here, the balance $bal^A(i)$ refers to the balance of the previous allocation (A). The above must be read as follows: compute the restricted social welfare divided it by the number of agents involved (that would be the optimal repartition), then subtract the payment such that each agent reaches exactly that level of satisfaction. This is exemplified on the following example (assuming this is the first deal taking place, hence a balance of 0 for each agent):

$$\begin{array}{ll} u_1(A) = 3 & u_1(A') = 2 \\ u_2(A) = 1 & u_2(A') = 6 \end{array}$$

The final allocation exhibits a restricted (utilitarian) social welfare of 8, hence 4 for each agent would be the optimal situation. This is obtained by computing the payment function which pays +2 to agent 1 and -2 to agent 2. However, such a payment would not always be rational: in many situations, the payment would be such that at least a self-interested agent would have no interest in implementing the deal. This is for instance the case on the following example:

$$\begin{array}{ll} u_1(A) = 3 & u_1(A') = 2 \\ u_2(A) = 1 & u_2(A') = 3 \end{array}$$

In this case, the computed payment is +/- 0.5, and it is easy to see that agent 1 would not rationally accept such a deal. In general, a deal is only rational if every agent enjoys before the deal takes place a utility that is strictly lower than the mean of the restricted social welfare after the deal. This is then an example where a socially beneficial deal would not be individually rational for agents if they were to use this payment function, *i.e.* this proves that this payment function is not rational-compatible. The consequence is that it will change the structural properties of the framework: by restricting the number of feasible deals, certain sequences of deals (that could be rational with a different payment function) would not be explored. Clearly, this can be a good thing in some situations, as exemplified below:

$$
\begin{array}{lll}
u_1(\{r_1\}) = 2 & u_1(\{r_2\}) = 2 & u_1(\{r_1, r_2\}) = 104 \\
u_2(\{r_1\}) = 100 & u_2(\{r_2\}) = 0 & u_2(\{r_1, r_2\}) = 100 \\
u_3(\{r_1\}) = 0 & u_3(\{r_2\}) = 101 & u_3(\{r_1, r_2\}) = 101
\end{array}
$$

We first suppose that agents use *any* rational-compatible payment function, and that a_2 initially holds the bundle $\{r_1, r_2\}$ (initial social welfare is 100). Now the following deal sequence takes place: (i) a_2 gives r_2 to a_1 (leading to a social welfare of 102), (ii) a_2 gives r_1 to a_1 (leading to a social welfare of 104). No more rational deal is possible (giving any resource would result in a loss of 102).

Of course, a social welfare of 201 would have been attainable, by simply offering r_2 to a_3 as a first deal. Equitable deals (payment 0.5) would indeed have discarded the deals implemented in the suboptimal sequence, leaving only this deal possible. Whether this is a good thing or not in general, *i.e* whether this payment function behaves as a good heuristic in this context will be experimentally evaluated later on in this paper (see Fig. 2 and 4).

Rational Locally Equitable Payment Function

What would be the "most equitable" payment function that would still be rational-compatible? A slight modification of the previous payment function would in fact suffice to make it rational-compatible: compute the payment function such that it makes every agent marginally better off, then allocate the remaining payments induced by the deal so as to reduce inequalities as much as possible (that is, give the remaining social surplus to the agent currently worse off).

For a deal $\delta = (A, A')$, formally, this payment can be defined as the function p maximizing

$$min_{i \in \mathcal{A}^\delta} \left\{ u_i(A') - bal^A(i) - p(i) \right\}$$

where $bal^A(i)$ is the balance of agent i before the deal, and under the two following constraints: p has to be a payment function ($\sum_{i \in \mathcal{A}} p(i) = 0$) and the deal must be rational ($\forall i \in \mathcal{A}^\delta, u_i(A') - u_i(A) > p(i)$). Obtaining an analytic formulation for p in the general case is difficult. However, in the case of bilateral negotiation, the solution to the above constraint equation for any socially beneficial deal $\delta = (A, A')$ involving agents i and j, can easily be computed. First note that in case the deal is rational with the fully locally equitable payment, this defines our payment function. Otherwise, observe that there is necessarily a single agent, say i, such that $u_i(A) \geq (u_i(A') + u_j(A'))/2$. This is so, because if both agents met this condition, then the deal would be fully equitable, while if neither of them met the condition, the deal would clearly not be be individually rational. Now we fix the payment for that agent i, in such a way that this agent is just marginally better off, in order to make the deal rational. That is, $p(i) = u_i(A') - u_i(A) - \epsilon$ for some suitably small constant $\epsilon > 0$. Note that $p(i)$ could be a positive or negative value; the payoff for i will be $+\epsilon$ either way. Then we define the other payment as $p(j) = -p(i)$.

We take again our previous example (that proved fully equitable payment functions not to be rational-compatible), to show how this works. Agent 1 has been better off before the deal than the average utility after the deal. Its payment will $p(1) = 2 - 3 - \epsilon$, *i.e.* it *receives* an amount of $1 + \epsilon$ from agent 2. After the deal, we have the following utilities (assuming there have been no earlier payments):

$$
\begin{array}{l}
u'_1(A') = 3 + \epsilon \\
u'_2(A') = 2 - \epsilon
\end{array}
$$

Whether this payment function outperforms other functions (as far as egalitarian social welfare is concerned) will be evaluated experimentally in this paper (see Fig. 1 and 5).

4. EXPERIMENTAL METHODOLOGY

Our experimental results are based on the following methodology: we first *create the system* (agents' utility functions and the initial allocation of goods), we then *compute the optimal utilitarian social welfare*, and we finally *run negotiations* (that is, let agents negotiate until no more deals are possible). Agents typically desire to hold several bundles (the details of how these utility functions are generated depends on the domain studied, and will be given later on). To avoid favouring a specific agent, the negotiation protocol is randomized: we pick an agent at random, and try to identify one rational deal involving this agent. Remember that as soon as one such rational deal meets the rationality constraints of the seller and the buyer, the deal is implemented and the allocation is updated. The results are averaged over 500 negotiation runs.

The reason why we compute the optimal utilitarian social welfare is that it serves to derive a (very generous) upper bound on optimal egalitarian social welfare: The best you can do is to optimize overall utility ($=sw_u$) and then use payments to distribute the overall wealth as evenly as possible (it is then simply $1/n$ times the maximal utilitarian social welfare). In the restricted negotiation framework that we investigate here, this optimal value will seldom be attained. What will be highly significant will be the value of the final utilitarian social welfare, which would represent an upper bound under the constraints of the negotiation framework (*e.g.* restrictions to 1-deals, rationality constraints). Note that affecting the class of acceptable deals by employing a non-rational-compatible payment function falls into that category. Again, distributing the overall wealth as evenly as possible would be the optimal egalitarian outcome (recall that a *theoretical* equitable payment function would allow to attain this value).

To sum up, the experimental results reported in this paper will typically mention the following outcome values:

(1) the value of the *optimal utilitarian welfare* (sw_u^{opt}/n), upper bound in unrestricted negotiation frameworks;

(2) the value of the *final utilitarian welfare* (sw_u/n), upper bound under the constraints imposed during this negotiation;

(3) the value of the *final egalitarian welfare* (sw_e), after negotiation took place.

The difference between (1) and (2) is an indicator for what we shall call the *protocol bias*. The difference between (2) and (3), more interesting to us, permits to appreciate the *payment function bias*. Comparison with the initial value of the egalitarian welfare proved to present very little interest in our experiments: it has been omitted for the sake of readability.

Figures presented in the next two sections exhibit curves defined by the following parameters: (1) the type of social welfare measured (`SWe` or `SWu/n`); (2) whether the social welfare shown is the `optimal` value or the `final` value; (3) the payment function used; and (4) the type of distribution used.

5. MODULAR DOMAINS

We start with very simple domains of negotiation, where no synergies can occur between different items of the bundles agents hold. In this case agents' utility functions are said to be *modular*. We generate these functions by simply picking 50 resources for each agent, and assigning random coefficients between 0 and 100. Note that in this case the optimal utilitarian social welfare can be easily computed: It suffices to assign each item to the agent who values it the most. (This allows in particular to run experiments with a large number of items). A recent (somewhat surprising) complexity result, due to Bouveret and Lang [2], shows that the difficulty of the corresponding problem largely differs for egalitarian frameworks *without money:* Optimising the egalitarian social welfare is an NP-hard problem, even when agents use modular functions (this is even true when they all use the same utility function). This is *not* the case in our setting though, because of the use of money (which allows us to separate the process of maximising overall utility and then distributing that utility evenly amongst the agents). A final important point to note is that the complexity of the overall negotiation process remains polynomial in this case, because the length (number of deals) of rational negotiations in modular domains is known to be linear [6].

Locally Uniform vs. Rational Locally Equitable

The first set of experiments has been run with 5 agents using modular utility functions, deals enhanced with rational-compatible payment functions. The number of resources varies between 50 and 150. Recall that in this case, we know that the optimal utilitarian social welfare will be reached (*i.e.* there is no "protocol bias"): the experiment will really measure the "payment function bias". To put it another way, if there were such a thing as a rational-compatible fully equitable payment function, both curves would be equivalent.

Fig. 1 shows the final egalitarian social welfare attained when locally uniform or rational equitable payments are used. It is striking that, first, both payment functions perform reasonably well (the ratio (sw_e/sw_u^{opt}) is constantly higher than 70% and improves when the number of resources increases). The rational equitable payment is, as expected, significantly above the locally uniform one (84% vs. 70% for 60 resources, 93% vs. 80% for 120 resources, 96% vs. 83% for 150 resources). However, given that this latter payment is, as argued before, the "most equitable" conceivable

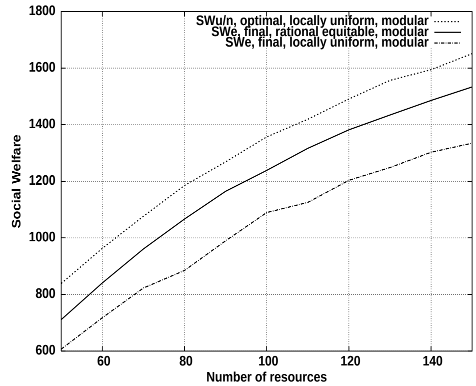

Figure 1: Modular Domains (I)

rational-compatible payment function, it is still surprising to see how close the "strategically justified" locally uniform payment stands.

Fully Locally Equitable

The second series of experiments (Fig. 2) carried out in modular domains tests the assumption that fully equitable deals can be a good heuristic in our negotiation setting, despite reducing the number of feasible deals. The results are unambiguous: in modular domains, this is far from being the case. The protocol bias (which exists in this case), is illustrated by the ratio between the optimal value and the final sw_u/n reached when allowing only fully equitable deals. For 50 resources already, the figure is around 75%, and drops to 50% for 150 resources. This really measures the loss due to the restriction put on the framework. Worse than that, the final sw_e exhibits a figure (sw_e/sw_u^{opt}) of 45% for 50 resources, shrinking down to 16% for 150 goods. Actually, the final egalitarian social welfare hardly increases in this setting. This really proves that this restriction is simply too hard to be met, and almost freezes the negotiation process.

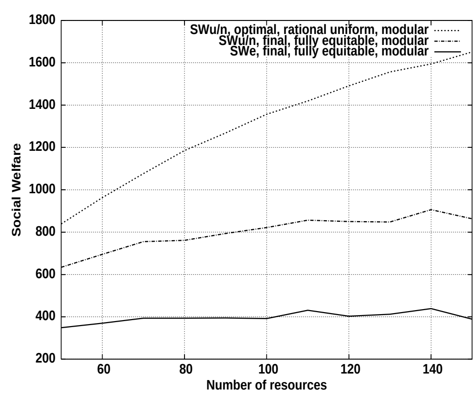

Figure 2: Modular Domains (II)

6. NON-MODULAR DOMAINS

In many situations, the assumption that resources have no synergies between them will be highly unrealistic. It is then very important to investigate the outcome of negotiation in such (non-modular) domains. The first experiments that we envisaged were produced using a uniform distribution for the generation of the structure of utility functions. The results (not reported here), regarding the final egalitarian welfare, turned out to be very good. This simply reflected the fact that uniform distributions tended to distribute evenly the participants' needs: agents were seldom in conflict.

Realistic Distributions

In order to get realistic agent preferences we decided to make use of the CATS bid generator [11], primarily designed for the use with combinatorial auctions. The output of the CATS software is a set of bids, taking the form of XOR demands $\langle b_1, p_1 \rangle \ldots, \langle b_n, p_n \rangle$, where b_i stands for the bundle of resources, and p_i stands for the value the agent assigns to this bundle. One problem that we faced when using CATS is the fact the output file contains only anonymous bids. This is perfectly suited for the case of studies of algorithms for the Winner Determination Problem in combinatorial auctions [4], because then you are not really concerned with associating resources with agents who claimed them, but only with computing the optimal solution. What we did to deal with this was to generate a (large) number of bids in a single shot (interacting in the same domain, it is likely that agents share a common dependency structure between goods). However, by allowing dominated bids to be generated, we permit agents to express different degree of preferences over the same bundles of goods. And finally, by picking (at random) only a subset of these generated bids, we allow agents to demand different bundles and admit that they may only partially share this common structure. The experiments reported in this section have been conducted using the CATS *arbitrary* distribution.

Locally Uniform vs. Globally Uniform

Fig. 3 shows the egalitarian social welfare reached when an increasing number of agents negotiate over 15 resources using locally uniform payments. Each agents claims 10 bundles. Recall that, as mentioned in the previous section, when there are no more than two agents, the locally uniform payment and the globally uniform payment coincide.

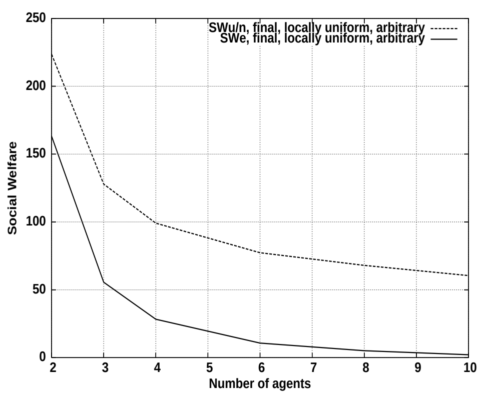

Figure 3: Non-Modular Domains (I)

First notice that, as expected, the final sw_u/n decreases exactly as a $1/n$ curve. Intuitively, we could have expected the final sw_e to decrease at the same speed. This turns out not to be the case: from 2 to 3 agents, sw_e decreases by 65%, whereas a $1/n$ decrease rate would have been of 33%. This observation can be explained in the light of the locally uniform payment function singularity: when there are no more than two agents, the locally and globally uniform payment are identical, in which case the inequalities between both agents will not increase during the negotiation. In other words, as explained in the previous section, the gain of sw_e between the beginning and the end of the negotiation will be equal to the gain of sw_u divided by n. This explains why the final sw_e for two agents is close to the final sw_u.

On the contrary, for more than 2 agents, each bilateral deal will only favor the two agents involved in it, and will thus potentially increase inequalities between agents contracting many deals, and others. This allows us to explain the rapid decrease of sw_e at the beginning of the curve.

Fully Locally Equitable vs. Locally Uniform

Next we compare fully locally equitable payment functions and locally uniform payment functions (Fig. 4). The results show that negotiation using fully equitable payments is notably less equitable than negotiation with uniform payments. This is so because the criterion is, again, too restrictive and prevents agents from contracting deals that could be rational (and hence improve social welfare). The fact that this function helps to select highly beneficial deals does not compensate for the loss incurred by these missed opportunities. This completes the answer to the question we left pending in Section 5: this payment function does not behave as a good heuristic, in this type of domains neither. The reason why the difference is far less important with uniform payment functions in these domains lies mainly in the fact that, typically, much fewer deals are contracted, and that values assigned to bundles are typically higher than those assigned to single resources (augmenting the probability to contract a fully equitable deal).

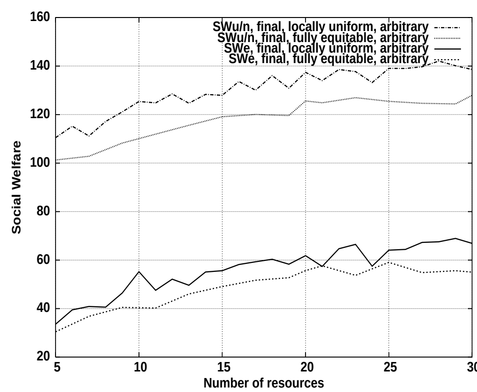

Figure 4: Non-Modular Domains (II)

Rational Locally Equitable vs. Locally Uniform

In this last series of experiments, we show that rational equitable payments slightly outperform uniform payments as far as the satisfaction of the poorest agent is concerned. (Remember that feasible deals being the same for these two types of payments, and payment functions adding up to zero at the end of a given negotiation, the fact that final values of averaged utilitarian welfare are similar should come as no surprise.) In Fig. 5, we observe that the difference with uniform payments is marginal for small number of resources, but becomes slightly larger when this number increases. This is explained as follows: when there are few resources to be distributed, the number of deals that will take place during a negotiation will typically be very small. Each contracted deal increases the probability of redistributing some wealth to the poorest agent, the more deals contracted the happier that agent should get (on average).

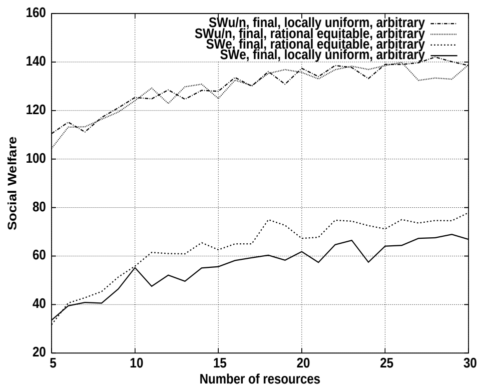

Figure 5: Non-Modular Domains (III)

7. CONCLUSION

This paper has investigated the egalitarian properties of some distributed resource allocation processes. An important aspect of the framework, that has radical consequences on egalitarian outcomes, is the choice of a specific payment function (*i.e.* the way to redistribute the social surplus induced by any given deal). We have discussed a variety of payment functions, and shown in particular that a perfectly equitable payment function is not rational-compatible. Our study has been complemented with several experiments, that have allowed us to identify conditions favourable to egalitarian outcomes. In particular, we should stress that:

- In modular domains, we observed that locally uniform payments give rise to reasonably equitable outcomes, which becomes comparatively more equitable as the number of resources grows. This remains significantly below (but closer than expected) to what could be achieved at best within the class of rational-compatible payments.
- It is typically not beneficial to enforce agents to obey to implement a fully equitable payment function, because this would conflict with rationality principles and prevent agents from conducting socially beneficial deals. This is striking in modular domains, less obvious in non-modular domains where there are typically fewer deals taking place during the negotiation.
- Locally uniform payment functions significantly differ from globally uniform payment functions for more than 2 agents. More generally, as the number of agents grows, we cannot expect any local payment to compensate for the inequalities.

A possible extension of this work will be to investigate how the conclusions drawn here will vary for other “standard” CATS distributions. A longer-term issue will be the study of other equitability measures: the *maximin* egalitarian principle is but one (admittedly rough, because it only takes into account the poorest agent) equitability measure. To refine our study, we plan to investigate other measures, such as Gini-like indexes for instance [12], which give a more global measure of a society’s wealth distribution.

8. REFERENCES

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

[2] S. Bouveret, H. Fargier, J. Lang, and M. Lemaître. Allocation of indivisible goods: A general model and some complexity results. In *Proc. AAMAS-2005*. ACM Press, 2005.

[3] S. J. Brams and A. D. Taylor. *Fair Division: From Cake-cutting to Dispute Resolution*. Cambridge University Press, 1996.

[4] P. Cramton, Y. Shoham, and R. Steinberg, editors. *Combinatorial Auctions*. MIT Press, 2006.

[5] P. E. Dunne, M. Wooldridge, and M. Laurence. The complexity of contract negotiation. *Artificial Intelligence*, 164(1–2):23–46, 2005.

[6] U. Endriss and N. Maudet. On the communication complexity of multilateral trading: Extended report. *Journal of Autonomous Agents and Multiagent Systems*, 11(1):91–107, 2005.

[7] U. Endriss, N. Maudet, F. Sadri, and F. Toni. Negotiating socially optimal allocations of resources. *Journal of Artificial Intelligence Research*, 2006. To appear.

[8] A. Giovannucci, J. A. Rodríguez-Aguilar, A. Reyes, F. X. Noria, and J. Cerquides. *i*Bundler: An agent-based decision support service for combinatorial negotiations. In *Proc. AAAI-2004*. AAAI Press, 2004.

[9] G. Jonker, J.-J. Meyer, and F. Dignum. Efficiency and fairness in air traffic control. In *Proc. 17th Belgian-Dutch Conf. on Artificial Intelligence*, 2005.

[10] M. Lemaître, G. Verfaillie, F. Jouhaud, J.-M. Lachiver, and N. Bataille. Selecting and scheduling observations of agile satellites. *Aerospace Sciences and Technology*, 6:367–381, 2002.

[11] K. Leyton-Brown, M. Pearson, and Y. Shoham. Towards a universal test suite for combinatorial auction algorithms. In *Proc. EC-2000*, 2000.

[12] H. Moulin. *Axioms of Cooperative Decision Making.* Cambridge University Press, 1988.

[13] R. Porter, Y. Shoham, and M. Tennenholtz. Fair imposition. *Journal of Economic Theory*, 118(2):209–228, 2004.

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

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

[16] T. W. Sandholm. Distributed rational decision making. In *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*. MIT Press, 1999.

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