Reaching Envy-free States in Distributed Negotiation Settings

Yann Chevaleyre
LAMSADE
Univ. Paris-Dauphine
France

Ulle Endriss
ILLC
University of Amsterdam
The Netherlands

Sylvia Estivie
LAMSADE
Univ. Paris-Dauphine
France

Nicolas Maudet
LAMSADE
Univ. Paris-Dauphine
France

Abstract

Mechanisms for dividing a set of goods amongst a number of autonomous agents need to balance efficiency and fairness requirements. A common interpretation of fairness is envy-freeness, while efficiency is usually understood as yielding maximal overall utility. We show how to set up a distributed negotiation framework that will allow a group of agents to reach an allocation of goods that is both efficient and envy-free.

1 Introduction

There are often two opposing requirements that we apply when judging the quality of an allocation of goods to a number of autonomous agents: *efficiency* and *fairness* [Moulin, 1988]. While the quest for economically efficient outcomes is well aligned with the highly successful approach of modelling agents as rational players in the game-theoretical sense, fairness is much more elusive and has generally achieved far less attention in the AI community (but recent exceptions include, for instance, the work of Bouveret and Lang [2005], Endriss *et al*. [2006], and Lipton *et al*. [2004]).

In this paper, we interpret fairness as *envy-freeness*. An allocation of goods is envy-free if no agent would prefer to receive the bundle assigned to one of its peers rather than its own [Brams and Taylor, 1996]. Efficiency is interpreted as maximising utilitarian *social welfare*, *i.e.* as maximising the sum of the utilities of the individual agents [Moulin, 1988]. We are interested in identifying those efficient allocations of goods to agents that are also envy-free. More specifically, we are interested in reaching such allocations by means of a *distributed negotiation* process [Sandholm, 1998; Endriss *et al.*, 2006; Dunne *et al.*, 2005]. That is, rather than developing a "centralised" algorithm to compute such an allocation, we want agents to be able to reach the desired state in an interactive manner by means of a sequence of deals they negotiate locally, driven by their own interests and without regard for the global optimisation problem.

The paper is organised as follows. In Section 2 we review the distributed negotiation framework adopted. Section 3 formally defines the concept of envy-freeness and gives a proof for the existence of efficient envy-free states in the presence of money, for the case of supermodular valuations. Section 4 discusses how to reach such states by means of distributed negotiation. We prove a general convergence theorem for supermodular valuations, and show that very simple deals over one resource at a time suffice to guarantee convergence in the modular case. Next, to be able to study how envy evolves over the course of a negotiation process, we introduce different metrics for the degree of envy in Section 5. We have used these metrics in a number of experiments aimed at analysing how envy evolves in modular domains. The results of these experiments are reported in Section 6. Section 7 concludes.

2 Distributed Negotiation Settings

In this section, we briefly review the distributed negotiation framework we adopt in this paper [Endriss *et al.*, 2006].

2.1 Basic Definitions

Let $\mathcal{A} = \{1..n\}$ be a finite set of *agents* and let $\mathcal{R}$ be a finite set of indivisible *resources* (which we also refer to as *goods*). An *allocation* is a partitioning of the items in $\mathcal{R}$ amongst the agents in $\mathcal{A}$ (*i.e.* each good must be owned by exactly one agent). As an example, allocation A, defined via $A(i) = \{r_1\}$ and $A(j) = \{r_2, r_3\}$, would allocate r_1 to agent i, and r_2 and r_3 to agent j. The interests of individual agents $i \in \mathcal{A}$ are modelled using *valuation functions* $v_i : 2^{\mathcal{R}} \to \mathbb{R}$, mapping bundles of goods to real numbers. Throughout this paper, we shall make the assumption that all valuations v_i are *normalised* in the sense that $v_i(\{\ \}) = 0$. Technically, this is not a significant restriction, but it greatly eases presentation. We sometimes use $v_i(A)$ as a shorthand for $v_i(A(i))$, the value agent i assigns to the bundle received in allocation A.

Agents negotiate sequences of deals to improve their own welfare. A *deal* $\delta = (A, A')$ is a pair of allocations (with $A \neq A'$), specifying the situation before and afterwards. Observe that a single deal may involve the reassignment of any number of goods amongst any number of agents. A *1-deal* is a deal involving only a single resource (and hence only two agents). The set of agents involved in the deal δ is denoted as $\mathcal{A}^{\delta}$. Deals may be accompanied by monetary side payments to allow agents to compensate others for otherwise disadvantageous deals. This is modelled using so-called *payment functions* (PFs): $p : \mathcal{A} \to \mathbb{R}$, which are required to satisfy $\sum_{i \in \mathcal{A}} p(i) = 0$. A positive value $p(i)$ indicates that agent i *pays* money, while a negative value means that the agent *receives* money. We associate each allocation A that is reached in a sequence of deals with a function $\pi : \mathcal{A} \to \mathbb{R}$ mapping agents to the sum of payments they have made so far, *i.e.* we also have $\sum_{i \in \mathcal{A}} \pi(i) = 0$. A *state* of the system is a pair (A, π) of an allocation A and a *payment balance* π.

Each agent $i \in \mathcal{A}$ is also equipped with a *utility function* $u_i : 2^{\mathcal{R}} \times \mathbb{R} \to \mathbb{R}$ mapping pairs of bundles and previous payments to real numbers. These are fully determined by the valuation functions: $u_i(R, x) = v_i(R) - x$. That is, utilities are *quasi-linear:* they are linear in the monetary component, but the valuation over bundles of goods may be any set function. For example, $u_i(A(i), \pi(i))$ is the utility of agent i in state (A, π), while $u_i(A(j), \pi(j))$ is the utility that i would experience if it were to swap places with j (in terms of both the bundle owned, and the sum of payments made so far).

2.2 Individual Rationality and Efficiency

Agents are assumed to only negotiate *individually rational* (IR) deals, *i.e.* deals that benefit everyone involved:

Definition 1 (IR deals) *A deal $\delta = (A, A')$ is called individually rational iff there exists a payment function p such that $v_i(A') - v_i(A) > p(i)$ for all agents $i \in \mathcal{A}$, except possibly $p(i) = 0$ for agents i with $A(i) = A'(i)$.*

While negotiation is driven by the individual preferences of agents, we are interested in reaching states that are attractive from a "social" point of view. A common metric for efficiency is (utilitarian) *social welfare* [Moulin, 1988]:

Definition 2 (Social welfare) *The social welfare of an allocation A is defined as $sw(A) = \sum_{i \in \mathcal{A}} v_i(A(i))$.*

We also speak of the social welfare of a *state* (A, π). As the sum of all $\pi(i)$ is always 0, the two notions coincide, *i.e.* it does not matter whether we define social welfare in terms of valuations or in terms of utilities. A state/allocation with maximal social welfare is called *efficient*. A central result in distributed negotiation is due to Sandholm [1998]:

Theorem 1 (Convergence) *Any sequence of IR deals will eventually result in an efficient allocation of goods.*

This result guarantees that agents can agree on any sequence of deals meeting the IR condition without getting stuck in a local optimum; and there can be no infinite sequence of IR deals. On the downside, this result presupposes that agents are able to negotiate complex multilateral deals between any number of agents and involving any number of goods.

A valuation function v is *modular* iff $v(R_1 \cup R_2) = v(R_1) + v(R_2) - v(R_1 \cap R_2)$ for all bundles $R_1, R_2 \subseteq \mathcal{R}$. In modular domains, we can get a much stronger convergence result [Endriss *et al.*, 2006]:

Theorem 2 (Simple convergence) *If all valuation functions are modular, then any sequence of IR 1-deals will eventually result in an efficient allocation of goods.*

2.3 Payment Functions

Requiring deals to be IR puts restrictions on what deals are possible at all and it limits the range of possible payments, but it does not determine the precise side payments to be made. This is a matter to be negotiated by the participating agents. Estivie *et al.* [2006] introduce several concrete payment functions that agents may choose to adopt. The two simplest ones are the *locally uniform payment function* (LUPF) and the *globally uniform payment function* (GUPF).

Choosing a PF amounts to choosing how to distribute the *social surplus* $sw(A') - sw(A)$ generated by a deal $\delta = (A, A')$. It is known that the social surplus is positive iff δ is IR [Endriss *et al.*, 2006]. The LUPF divides this amount equally amongst the *participating* agents $\mathcal{A}^\delta$; the GUPF divides it equally amongst *all* agents $\mathcal{A}$:

LUPF: $p(i) = [v_i(A') - v_i(A)] - [sw(A') - sw(A)]/|\mathcal{A}^\delta|$ if $i \in \mathcal{A}^\delta$ and 0 otherwise

GUPF: $p(i) = [v_i(A') - v_i(A)] - [sw(A') - sw(A)]/n$

3 Envy-freeness

An important aspect of fairness is the absence of *envy* [Brams and Taylor, 1996]. In this section, we give two definitions of envy-freeness: the first applies to allocations of goods alone, while the second also takes previous payments into account. While the first may be in conflict with efficiency requirements, for the second definition there always exists a negotiation state that is both efficient and envy-free.

3.1 Envy-free Allocations of Goods

In the context of our negotiation framework, the concept of envy-freeness can be formally defined as follows:

Definition 3 (Envy-freeness) *An allocation A is called envy-free iff $v_i(A(i)) \geq v_i(A(j))$ for all agents $i, j \in \mathcal{A}$.*

Observe that, were we not to require that *all* goods be allocated, it would be easy to find an envy-free allocation: simply allocate the empty bundle to each agent. On the other hand, it is well-known that if we either require all goods to be allocated (so-called *complete* allocations), or if we restrict interest to allocations that are Pareto optimal, then an envy-free allocation may not always exist (just think of the case of two agents and a single good). Bouveret and Lang [2005] have investigated the computational complexity of checking whether a given negotiation problem admits an allocation that is envy-free. One of the simpler results states that checking whether there exists a complete and envy-free allocation is NP-hard, even if there are only two agents and these agents use the same dichotomous and monotonic valuation function.[1]

3.2 Envy-freeness in the Presence of Money

Definition 3 is the standard definition of envy-freeness, which applies to domains without transferable utility; that is, it does not (yet) take the monetary component of our negotiation framework into account. Next, we define envy-freeness with respect to full negotiation states:

Definition 4 (EF states) *A state (A, π) is called envy-free iff $u_i(A(i), \pi(i)) \geq u_i(A(j), \pi(j))$ for all agents $i, j \in \mathcal{A}$.*

A state that is both efficient and envy-free is called an *EEF state*. This corresponds to Pareto optimal and envy-free allocations without money.

[1] A valuation v is dichotomous iff $v(R) = 0$ or $v(R) = 1$ for any $R \subseteq \mathcal{R}$; v is monotonic iff $v(R_1) \leq v(R_2)$ whenever $R_1 \subseteq R_2$.

How does the move to a definition that explicitly accounts for payments affect the existence of envy-free states? In fact, in the presence of money, states that are both efficient and envy-free are known to always exist [Alkan *et al.*, 1991]. While the proof given by Alkan *et al*. is rather involved, in the case of *supermodular* valuations we can give a very simple existence argument. A valuation v is supermodular iff $v(R_1 \cup R_2) \geq v(R_1) + v(R_2) - v(R_1 \cap R_2)$ for all bundles $R_1, R_2 \subseteq \mathcal{R}$. We show a proof here, as this will be helpful in following some of the material later on (cf. Section 4.2).

Theorem 3 (Existence of EEF states) *If all valuations are supermodular, then an EEF state always exists.*

Proof. There clearly always exists an allocation that is efficient: *some* allocation must yield a maximal sum of individual valuations. Let A^* be such an efficient allocation. We show that a payment balance π^* can be arranged such that the state (A^*, π^*) is EEF. Define $\pi^*(i)$ for each agent i:

$$\pi^*(i) \quad = \quad v_i(A^*) - sw(A^*)/n$$

First, note that π^* is a valid payment balance: the $\pi^*(i)$ do indeed sum up to 0. Now let $i, j \in \mathcal{A}$ be any two agents in the system. We show that i does *not* envy j in state (A^*, π^*). As A^* is efficient, giving both $A^*(i)$ and $A^*(j)$ to i will not increase social welfare any further:

$$v_i(A^*(i)) + v_j(A^*(j)) \quad \geq \quad v_i(A^*(i) \cup A^*(j))$$

As v_i is assumed to be supermodular, this entails:

$$v_i(A^*(i)) + v_j(A^*(j)) \quad \geq \quad v_i(A^*(i)) + v_i(A^*(j))$$

Adding $sw(A^*)/n$ to both sides of this inequation, together with some simple rearrangements, yields $u_i(A^*(i), \pi^*(i)) \geq u_i(A^*(j), \pi^*(j))$, *i.e.* agent i does indeed not envy agent j. Hence, (A^*, π^*) is not only efficient, but also envy-free. □

This is an encouraging result: for the definitions adopted in this paper, efficiency and fairness *are* compatible. However, the mere *existence* of an EEF state alone is not sufficient in the context of negotiation amongst autonomous agents. Why should rational decision-makers accept the allocation and payments prescribed by the proof of Theorem 3? And even if they do, how can we compute them in practice? Just finding an efficient allocation is already known to be NP-hard [Rothkopf *et al.*, 1998]. Finally, as argued in the introduction, we are interested in a *distributed* procedure, where agents identify the optimal state in an interactive manner.

4 Reaching Efficient Envy-free States

In this section, we discuss to what extent EEF states can be reached by means of negotiation in a distributed manner.

4.1 Envy-freeness and Individual Rationality

In the previous section, we have seen that envy-freeness and efficiency *are* compatible in our framework. However, this does not necessarily mean that also envy-freeness and individual rationality will be compatible. And indeed, the following example (involving two agents and just a single resource) shows that this is not the case:

$$v_1(\{r\}) = 4 \quad v_2(\{r\}) = 7$$

Suppose agent 1 holds r in the initial allocation A_0. There is only a single possible deal, which amounts to passing r to agent 2, and which will result in the efficient allocation A^*. How should payments be arranged? To ensure that the deal is IR for both agents, agent 2 should pay agent 1 any amount in the open interval $(4, 7)$. On the other hand, to ensure that the final state is envy-free, agent 2 should pay any amount in the closed interval $[2, 3.5]$. The two intervals do not overlap. This means that, while we will be able to reach negotiation outcomes that are EEF, it is simply not possible in all cases to do so by means of a process that is fully IR.

For the procedure proposed in the sequel, we are going to circumvent this problem by introducing a one-off initial payment that may not be IR for all the agents involved:

$$\pi_0(i) \quad = \quad v_i(A_0) - sw(A_0)/n$$

That is, each agent has to first pay an amount equivalent to their valuation of the initial allocation A_0, and will then receive an equal share of the social welfare as a kick-back. We refer to this as making an *initial equitability payment*. Note that this payment does *not* achieve envy-freeness (and it does not affect efficiency at all). All it does is to "level the playing field". In the special case where none of the agents has an interest in the goods they hold initially ($v_i(A_0) = 0$ for all $i \in \mathcal{A}$), the initial payments reduce to 0.

4.2 Convergence in Supermodular Domains

We are now going to prove a central result on the reachability of EEF states by means of distributed negotiation. The result applies to supermodular domains and assumes that initial equitability payments have been made. It states that, if agents only implement deals that are individually rational (IR) and they use the globally uniform payment function (GUPF) each time to determine the exact payments, then negotiation will eventually terminate (*i.e.* no more such deals will be possible) and the final state reached will be both efficient and envy-free (EEF). Importantly, this will be the case whichever deals (that are meeting these conditions) the agents choose to implement, *i.e.* we can never get stuck in a local optimum. In short:

Theorem 4 (Convergence with GUPF) *If all valuations are supermodular and if initial equitability payments have been made, then any sequence of IR deals using the GUPF will eventually result in an EEF state.*

Proof. We first show that the following invariant will be true for every state (A, π) and every agent i, provided that agents only negotiate IR deals using the GUPF:

$$\pi(i) \quad = \quad v_i(A) - sw(A)/n \qquad (1)$$

Our proof proceeds by induction over the number of deals negotiated. As we assume that initial equitability payments have been made, claim (1) will certainly be true for the initial state (A_0, π_0). Now let $\delta = (A, A')$ and assume (1) holds for A and the associated payment balance π. We obtain the payment balance π' associated with A' by adding the appropriate GUPF payments to π:

$$\begin{aligned} \pi'(i) \quad &= \quad \pi(i) + [v_i(A') - v_i(A)] - [sw(A') - sw(A)]/n \\ &= \quad v_i(A') - sw(A')/n \ \checkmark \end{aligned}$$

This proves our claim (1). Now, Theorem 1 shows that negotiation via IR deals (whichever PF is used) must eventually terminate and that the final allocation will be efficient. Let A^* be that allocation, and let π^* be the associated payment balance. Equation (1) also applies to (A^*, π^*). But as we have already seen in the proof of Theorem 3, the efficiency of A^* together with equation (1) applied to (A^*, π^*) implies that (A^*, π^*) must be an EEF state. $\square$

Theorem 4 may seem surprising. As pointed out elsewhere [Endriss *et al.*, 2006], it is not possible to define a "local" criterion for the acceptability of deals (which can be checked taking only the utilities of the agents involved into account) that would guarantee that a sequence of such deals always converges to an envy-free state. We circumvent this problem here by using the GUPF, which adds a (very limited) non-local element. Only the agents involved in a deal can ever be asked to give away money, and all payments can be computed taking only the utilities of those involved into account.

4.3 Convergence in Modular Domains

In modular domains, we can even guarantee convergence to an EEF state by means of 1-deals (over one item at a time):

Theorem 5 (Simple convergence with GUPF) *If all valuation functions are modular and if initial equitability payments have been made, then any sequence of IR 1-deals using the GUPF will eventually result in an EEF state.*

Proof. This works as for Theorem 4, except that we rely on Theorem 2 for convergence by means of 1-deals (in place of Theorem 1). Note that Theorem 3 still applies, because any modular valuation is also supermodular. $\square$

4.4 Related Work

There has been some work on procedures for finding EEF states in the social choice literature [Alkan *et al.*, 1991; Klijn, 2000; Haake *et al.*, 2002], albeit with little or no attention to computational issues. The work of Haake *et al.* is particularly interesting. They propose two variants of the same procedure, the first of which assumes that an efficient allocation is given *to begin with*. The actual procedure determines compensatory payments to envious agents such that an EEF state will eventually be reached. While their solution is very elegant and intuitively appealing, it fails to address the main issue as far as the *computational* aspect of the problem is concerned: by taking the efficient allocation as given, the problem is being limited to finding an appropriate payment balance. Certainly for supermodular domains, our proof of Theorem 3 shows that this is *not* a combinatorial problem.

The second procedure put forward by Haake *et al.* [2002] interleaves reallocations for increasing efficiency with payments for eliminating envy. However, here the authors also do not address a combinatorial problem, because they assume "exogenously given bundles". That is, negotiation relates only to who gets which bundle, but the composition of the bundles themselves cannot be altered. This is equivalent to allocating n objects to n agents—not an NP-hard problem.

5 Degrees of Envy

The convergence theorems of the previous section show that envy-freeness can be a guaranteed outcome of our negotiation process, under specific assumptions. But this is of course far from being the final word on this topic. The PF involved in proving Theorem 4 introduces a non-local element, in the sense that it redistributes the social surplus over the whole society. How would a local PF, like the LUPF, behave in comparison? Also, agents need to start negotiation with a non-IR payment. But how does this affect the results in practice? Lastly, the convergence theorems do not say how envy evolves over the course of negotiation. Because negotiation can become very long in practice, it is likely that the process will have to be be stopped before completion. In that event, it would be valuable to be able to guarantee some monotonicity properties—but with respect to what parameter?

To address such questions, not only the mere classification of a society as being envy-free or not is needed. It is important to investigate to what *degree* we may be able to approximate the ideal of envy-freeness. In the remainder of this section, we discuss several options for defining such a metric. We propose to analyse the degree of envy of a society as a three-level aggregation process, starting with envy between agents, over envy of a single agent, to eventually provide a definition for the degree of envy of a society.

5.1 Envy between Agents

How much does agent i envy agent j in state (A, π)? Generally speaking, when assessing the degree of envy between two agents, we can either focus on agents that we indeed envy (positive envy, e^{pos}), or consider both agents that appear to be better off, of course, but also agents we believe to be in a situation worse than ours (total envy, e^{tot}).

$$
\begin{aligned}
e^{tot}(i,j) &= u_i(A(j), \pi(j)) - u_i(A(i), \pi(i)) \\
e^{pos}(i,j) &= max\{e^{tot}(i,j), 0\}
\end{aligned}
$$

While the latter option (e^{pos}) seems to fit best with our intuitions about envy, the former (e^{tot}) may be justifiable if we want agents to be able to compensate for envy.

5.2 Degree of Envy of a Single Agent

How envious is a single agent i in state (A, π)? This notion considers the agent in relation with many other agents. A first option would be to count the *number of agents* that i envies. This arguably only makes sense when e^{pos} is used to assess the envy between two agents.

$$
e^{0-1,pos}(i) = \begin{cases} 1 & \text{if } \exists j:\ e^{pos}(i,j) \neq 0 \\ 0 & \text{otherwise} \end{cases}
$$

More fine-grained, quantitative measures would also account for *how* envious an agent is. At least two obvious options for doing this come to mind. Firstly, we could measure how much our agent envies the agent it envies *the most* ($e^{max,op}$). Or we could compute the *sum* of envies it experiences with respect to all the agents in the system ($e^{sum,op}$).

$$
\begin{aligned}
e^{max,op}(i) &= \max_j e^{op}(i,j) \text{ where } op \in \{pos, tot\} \\
e^{sum,op}(i) &= \textstyle\sum_j e^{op}(i,j) \text{ where } op \in \{pos, tot\}
\end{aligned}
$$

Note that these two options can be of interest whatever the operator chosen to compute the degree of envy between agents.

5.3 Degree of Envy of a Society

Given the degrees of envy of each individual agent in a system for a given negotiation state, we can now define suitable aggregation operators to yield the degree of envy for the agent society as a whole. As for the aggregation of individual preferences to obtain a social preference ordering [Moulin, 1988], there are a multitude of different options available for doing this. Here we just mention two options that appear particularly appropriate in our context: the max- and sum-operators. Using the max operator must be interpreted as focusing on the most envious agent of the society (whatever operator was chosen to measure that), while the sum (or average) provides a more global picture of the situation. Lipton *et al.* [2004], for instance, use $e^{max,max,pos}$ to describe the degree of envy of a society. This is the maximum envy experienced by any agent, where each agent measures their envy as the difference to the opponent they envy the most.

5.4 Discussion

The different measures of envy introduced here are by no means exhaustive. Other options for aggregating degrees of envy could be used, e.g. based on the *leximin-ordering* [Moulin, 1988]. However, by just combining the options listed here, we obtain several candidate measures of envy. Some of them are of course much more natural than others, as our discussion already suggested. While this question would certainly necessitate a deeper discussion, we only provide here some basic justification that supports the choice of the measure we decided to use in this paper. Note first that a basic requirement of these measures would be that the degree of envy should be ≤ 0 iff that allocation is envy-free. This already rules out any measure based on e^{tot}. As for the rest, it is mostly a matter of what aspect you want to study as a priority. In what follows, we use two measures: the number of envious agents ($e^{sum,0-1,pos}$), because it gives a good snapshot of the distribution of envy within the society; and the sum of the sum of envies ($e^{sum,sum,pos}$), because it also reflects the overall magnitude of envy, which the previous one misses.

6 Analysis of Modular Domains

In this section, we report on a first experimental analysis of how close we can get to the ideal of an EEF state, if we do not make the initial equitability payment required by our convergence theorems; that is, if we use a framework that is fully IR. We restrict ourselves to the case of modular domains.

6.1 Experimental Setup

The experiments discussed in this section have been carried out using a simple simulation platform. After creating a number of virtual agents and resources, we randomly distribute the resources amongst the agents and randomly generate valuation functions for them. Negotiation is simulated by going for each (randomly chosen) agent through all the goods currently held by that agent and trying to find a potential partner with whom an IR 1-deal over that good can be agreed upon. This process is repeated until no more IR 1-deals are possible. As we restrict ourselves to modular valuations, the final allocation is known to always be efficient (cf. Theorem 2).

For all the experiments reported, the number of agents is $n = 20$. The number $|\mathcal{R}|$ of resources in the system varies between 50 and 500, depending on the experiment. For each agent, we randomly choose 50 distinguished resources. To all the other resources, the agent in question assigns valuation 0. For each of the distinguished resources, we randomly draw a weight from the uniform distribution over the set $\{1..100\}$ to determine the (modular) valuation function of that agent.

6.2 Envy over Time

The first set of experiments is aimed at clarifying how envy evolves over the course of negotiation. Before turning to the experiments, we observe the following fact: If an agent with a modular valuation function makes an IR bilateral deal, this will reduce its envy towards its trading partner. This means that the envy *between* the two agents cutting a deal will always reduce. However, non-concerned agents may become more envious than before, because one of those two agents may have improved its lot according to their own valuation.

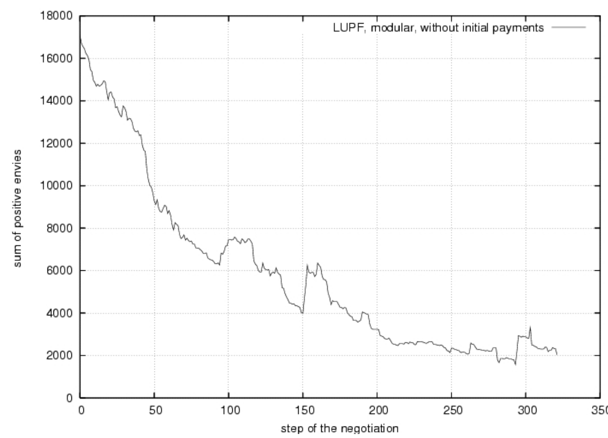

Figure 1: Evolution of envy over time (150 goods)

This makes it difficult to predict how overall envy will be affected by a deal. A typical example is given in Figure 1. The figure shows the sum of positive envies after each deal, for one particular negotiation process. The payment function used is the LUPF (without any initial payments). As we can see, after about 320 deals negotiation terminates and we have reached an allocation that is almost envy-free. Figure 1 shows that the degree of envy does not diminish strictly monotonically. On the other hand, the graph suggests (and more extensive experiments, not shown here, clearly demonstrate), that envy does diminish monotonically *on average*. Further experiments show that the GUPF (also without initial payments) performs slightly better than the LUPF (faster convergence and lower overall envy in the end).

6.3 Impact of Number of Resources

While for the example shown in Figure 1 we do not reach an EEF state, it seems nevertheless remarkable that we get rather close. It turns out that *how* close we can get depends very strongly on the number of goods in the system. This issue is analysed in more detail in our second set of experiments. An example is given in Figure 2. It shows the number of envious agents in the final state of negotiation, as a function of the number of goods in the system (average of 100 runs each). For this experiment, the standard deviation ranges from (almost) 0 to 4.2, depending on the number of resources.

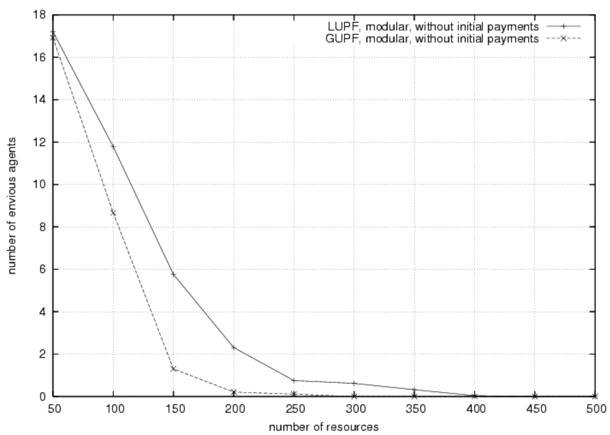

Figure 2: Number of envious agents in the final state

Note that for valuations generated using a uniform distribution, as is the case here, the *initial* degree of envy will always be rather high, independently of the number of goods (on average, 19 out of 20 agents are envious to begin with—data not shown here). This can be explained by the fact that the initial allocation is independent of agent preferences, so for any agent the chance of liking a particular bundle more than its own is already 50%. On the other hand, as our experiments show, the potential for rational negotiation to reduce or even fully eliminate envy improves dramatically as the number of goods increases. This can be explained, at least in part, by the fact that for a high number of goods the probability of an agent receiving several of their top goods in the efficient allocation is very high. For instance, if each agent has 50 distinguished goods, it will on average value 5 of them between 90 and 100. It is then very unlikely that any of the other 19 agents will also value these 5 goods that highly, when there are, say, 500 items in the system. So the fact that we are using a uniform distribution does go some way towards explaining these, nevertheless surprising, results. Finally, as Figure 2 shows, the GUPF does better than the LUPF.

7 Conclusion

This paper has addressed the question of the fairness (understood here as envy-freeness) of allocations obtained when one lets rational decision-makers negotiate over indivisible goods. We show in particular that EEF outcomes *can* be guaranteed in supermodular domains, provided agents initiate the process with some equitability payments that are not necessarily IR (but envy-freeness cannot be guaranteed with a fully rational process). To appreciate how this initial payment affects negotiation, we report on experimental results in modular domains where no such payments are made. The experiments demonstrate that envy decreases on average (although not monotonically) over a negotiation, and even tends towards zero at the end of the negotiation when the number of goods is large.

There are many potential extensions of this work. Perhaps the most obvious concerns the experimental study of non-modular domains. Theorem 4 still holds in supermodular domains, and it would be interesting to see how negotiation proceeds in these domains, when the efficient allocation may not be reached in practice. Another tempting avenue of research is the design of *local* PFs specially customised to help approach envy-freeness as much as possible. While it is known that no such perfect PF can exist, more systematic experiments could shed a light on what makes a PF well-behaved. At last, this work makes the assumption that each agent has full knowledge of the bundles held by the others; this can be unrealistic in many applications. There are different ways to alleviate this, for instance by assuming that agents have a partial (pre-defined or dynamic) perception of the society.

References

[Alkan *et al.*, 1991] A. Alkan, G. Demange, and D. Gale. Fair allocation of indivisible goods and criteria of justice. *Econometrica*, 59(4):1023–1039, 1991.

[Bouveret and Lang, 2005] S. Bouveret and J. Lang. Efficiency and envy-freeness in fair division of indivisible goods: Logical representation and complexity. In *Proc. IJCAI-2005*. Morgan Kaufmann, 2005.

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

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

[Endriss *et al.*, 2006] U. Endriss, N. Maudet, F. Sadri, and F. Toni. Negotiating socially optimal allocations of resources. *Journal of Artif. Intell. Res.*, 25:315–348, 2006.

[Estivie *et al.*, 2006] S. Estivie, Y. Chevaleyre, U. Endriss, and N. Maudet. How equitable is rational negotiation? In *Proc. AAMAS-2006*. ACM Press, 2006.

[Haake *et al.*, 2002] C.-J. Haake, M.G. Raith, and F.E. Su. Bidding for envy-freeness: A procedural approach to n-player fair-division problems. *Social Choice and Welfare*, 19:723–749, 2002.

[Klijn, 2000] F. Klijn. An algorithm for envy-free allocations in an economy with indivisible objects and money. *Social Choice and Welfare*, 17:201–215, 2000.

[Lipton *et al.*, 2004] R.J. Lipton, E. Markakis, E. Mossel, and A. Saberi. On approximately fair allocations of indivisible goods. In *Proc. EC-2004*. ACM Press, 2004.

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

[Rothkopf *et al.*, 1998] M.H. Rothkopf, A. Pekeč, and R.M. Harstad. Computationally manageable combinational auctions. *Management Science*, 44(8):1131–1147, 1998.

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