

# Distributed Fair Allocation of Indivisible Goods\*

Yann Chevaleyre  
LAMSADE  
Université Paris-Dauphine  
yann.chevaleyre@dauphine.fr

Ulle Endriss  
ILLC  
University of Amsterdam  
ulle.endriss@uva.nl

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

29 October 2009

## Abstract

Distributed mechanisms for allocating indivisible goods are mechanisms lacking central control, in which agents can locally agree on deals to exchange some of the goods in their possession. We study convergence properties for such distributed mechanisms when used as fair division procedures. Specifically, we identify sets of assumptions under which any sequence of deals meeting certain conditions can be shown to converge to a proportionally fair allocation and to an envy-free allocation, respectively. We also introduce an extension of the basic framework where agents are vertices of a graph limiting which agents can interact with each other and prove a similar convergence result for envy-freeness in this context. Finally, when not all assumptions guaranteeing envy-freeness are satisfied, we may want to minimise the degree of envy exhibited by an outcome. To this end, we introduce a generic framework for measuring the degree of envy in a society and establish the computational complexity of checking whether a given scenario allows for a deal that is beneficial to every agent involved and that will reduce envy.

## 1 Introduction

The problem of fairly dividing a number of goods between a number of agents has been studied in a variety of different settings. First, we may distinguish allocation problems for *divisible* and for *indivisible goods*. The literature on cake-cutting, for instance, is concerned with divisible goods [5, 27]. As for indivisible goods, we can distinguish *assignment problems* [1], where each agent can consume at most a single good, from more general settings where each agent may receive a set (or bundle) of goods [2, 31]. If agents can receive sets of goods, and if their preferences over the goods are not additively separable, then fair division becomes a *combinatorial optimisation problem*. We may also distinguish whether or not to allow for *monetary side payments* to be added to the bundles allocated to the agents, and if so we have to decide what assumptions to make regarding the agents' appreciation

---

\*Parts of the material in this paper has previously been presented at the 20th International Joint Conference on Artificial Intelligence (IJCAI-2007) in Hyderabad, the 22nd AAAI Conference on Artificial Intelligence (AAAI-2007) in Vancouver, the Dagstuhl Seminar on Computational Issues in Social Choice in October 2007, the Multiagent Resource Allocation Workshop held in Amsterdam in June 2008, as well as at departmental seminars at Delft University, Université Paris-Dauphine, Stanford University, and the University of Amsterdam. The authors would like to thank the participants of these meetings and the reviewers of IJCAI-2007 and AAAI-2007 for the useful feedback received.

of money (such as *quasi-linearity*, for instance). Finally, there are many different ways in which to interpret the term *fairness* itself [32]. For instance, we may be interested in equitable allocations, in proportional allocations, or in allocations where agents do not envy each other. In this paper, we adopt a model where sets of indivisible goods need to be allocated to a number of agents; agents express their preferences in terms of valuation functions over sets of goods; and side payments are possible and agents have quasi-linear preferences regarding money. In terms of fairness criteria, we shall concentrate on proportionality and envy-freeness (these terms will get defined formally later on).

While fair division is a problem originating in Economics and Political Science, with a number of important contribution by mathematicians, it has recently begun to also attract the attention of researchers in Artificial Intelligence, Multiagent Systems, and Theoretical Computer Science [7, 23, 14, 3, 6, 10]. The reason for this trend is twofold: on the one hand, concepts from fair division are immediately relevant to these disciplines (e.g., finding acceptable agreements in a multiagent systems) and the tools and techniques of these disciplines can shed new light on previously unexplored aspects of fair division (e.g., by applying ideas from complexity theory).

There is a large variety concerning the *procedures* that have been proposed to address fair division problems. Some work emphasises that such procedures should be simple enough to be executable by humans. Others have concentrated on devising computationally efficient algorithms. Our own approach is to use a *distributed* mechanism in which agents can locally agree on deals to exchange some of the goods in their possession. In this approach, we initially choose a random allocation of goods to agents and then let the agents negotiate freely to, hopefully, find a superior allocation. Each agent's preferences are defined using a (not necessarily additive) valuation function. Agents agree on a sequence of deals to exchange goods, which may be combined with monetary side payments. Once negotiation is underway, there is no central control regulating the process and agents are assumed to make deals according to their own preferences. Standard cake-cutting procedures that have been proposed for the fair partitioning of a single divisible good [5], for instance, are not distributed in this sense. While some of them, like e.g. the well-known Banach-Knaster "last diminisher procedure", do not require a referee to execute the procedure, they nevertheless proceed according to a fixed (global) protocol specifying which agent has to perform which type of action at any given time. The same is true for the "descending demand procedure" of Herreiner and Puppe [19], a procedure for finding equitable allocations of indivisible goods.

The kind of distributed approach followed here has recently been studied by several authors [29, 14, 12]. It is particularly attractive for large systems of autonomous software agents that need to negotiate an allocation of resources or agree on a distribution of tasks amongst themselves. Being able to distribute the computation of an optimal allocation of goods over the agents is useful in many applications where no central authority can be relied upon to eventually decide on the allocation. This may, for instance, be the case in view of computational limitations of a potential centre, or in view of its trustworthiness. In such a context the use of combinatorial auctions or similar mechanisms [9], which in principle could be employed to find an optimal allocation of indivisible goods, may be considered problematic. Then, the system designer will want to try to set up the system in such way that it nevertheless guarantees certain desirable properties, but without directly interfering into the process of negotiation.

In this paper we study *convergence properties* for distributed allocation mechanisms when used

as fair division procedures. Specifically, we identify sets of assumptions under which any sequence of deals meeting certain conditions can be shown to converge to a proportionally fair allocation and to an envy-free allocation, respectively. We also introduce an extension of the basic framework where agents are vertices of a graph limiting which agents can interact with each other. The concept of envy-freeness is very naturally extended to such a setting: an agent will not envy another agent if they either believe that their own bundle is more valuable or if they cannot see that other agent, because there is no edge between them in the network. This allows us to prove a similar convergence result for envy-freeness in the context of *fair division on a graph*. Finally, when not all assumptions guaranteeing envy-freeness are satisfied, it is interesting to study the *degree of envy* exhibited by an outcome. To this end, we introduce a generic framework for defining metrics for assessing the degree of envy in a society and establish the *computational complexity* of checking whether a given scenario allows for a deal that is beneficial to every agent involved and that will reduce envy.

The remainder of this paper is organised as follows. Section 2 introduces the framework of distributed negotiation over indivisible goods we have adopted and recalls two results from the literature regarding convergence to an *efficient* allocation. Section 3 studies the case of proportionally fair outcomes. We show how to adapt the classical Knaster procedure to our setting and obtain a distributed procedure that will guarantee convergence to a proportional outcome. In Section 4 we carry out a similar programme, this time choosing envy-freeness as our goal. Envy-freeness is considerably more demanding a criterion than proportionality, so the convergence results that we obtain rely on somewhat stricter assumptions. The framework is extended to fair division on graphs in Section 5. We discuss possible extensions of the notions of efficiency, proportionality and envy-freeness to this richer setting, and arrive at a particularly satisfying definition as well as an attractive convergence result in the case of envy-freeness. We also argue that proportionality and efficiency are not naturally generalised to this setting. Section 6 proposes a typology of metrics for measuring the degree of envy in a society. We then establish a computational complexity result that shows that identifying a deal that is acceptable to each agent involved and that reduces envy is NP-complete (but not harder than finding a deal that is simply acceptable to the agents involved). This result applies for any of the measures of envy defined earlier. Section 7 concludes.

## 2 Distributed mechanisms for allocating indivisible goods

In this section we define the class of distributed mechanisms that we want to employ to negotiate fair allocations of goods and briefly recall a number of known results [14, 29].

### 2.1 Agents, goods, allocations, valuations

Let  $\mathcal{N} = \{1, \dots, n\}$  be a finite set of *agents* and let  $\mathcal{G}$  be a finite set of indivisible *goods*. An *allocation*  $A : \mathcal{N} \rightarrow 2^{\mathcal{G}}$  is a partitioning of the items in  $\mathcal{G}$  amongst the agents in  $\mathcal{N}$  (i.e., each good must be owned by exactly one agent). As an example, allocation  $A$ , defined via  $A(i) = \{g_1\}$  and  $A(j) = \{g_2, g_3\}$ , would allocate  $g_1$  to agent  $i$ , and  $g_2$  and  $g_3$  to agent  $j$ . We assume that the initial allocation has been chosen randomly, so agents cannot derive any rights from this allocation (as far as the definition of

fairness is concerned).

The interests of individual agents  $i \in \mathcal{N}$  are modelled using *valuation functions*  $v_i : 2^{\mathcal{G}} \rightarrow \mathbb{R}$ , mapping bundles of goods to the reals. To simplify presentation, throughout this paper, we shall make the assumption that all valuations  $v_i$  are *normalised*, in the sense that  $v_i(\emptyset) = 0$ , and *non-negative*, meaning that  $v_i(S) \geq 0$  for all  $S \subseteq \mathcal{G}$ . 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  $A$ .

Some of our results only apply to scenarios where all valuation functions meet additional properties. A valuation function is called *modular* (or *additive*) if  $v(S_1 \cup S_2) = v(S_1) + v(S_2) - v(S_1 \cap S_2)$  for all bundles  $S_1, S_2 \subseteq \mathcal{G}$ . Similarly,  $v$  is called *supermodular* if  $v(S_1 \cup S_2) \geq v(S_1) + v(S_2) - v(S_1 \cap S_2)$ .

## 2.2 Deals, payments, utility

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 good (and hence only two agents). The set of agents involved in the deal  $\delta$  is denoted as  $\mathcal{N}^\delta$ . (In Section 5 we shall introduce a refinement of the basic model, where agents can only make deals with agents to which they are connected in an underlying negotiation topology.)

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*:  $p : \mathcal{N} \rightarrow \mathbb{R}$ , which are required to satisfy  $\sum_i p(i) = 0$ . 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  $\pi : \mathcal{N} \rightarrow \mathbb{R}$  mapping agents to the sum of payments they have made so far, i.e., we also have  $\sum_i \pi(i) = 0$ . A *state* of the system is a pair  $(A, \pi)$  of an allocation  $A$  and a *payment balance*  $\pi$ . It is possible to impose an *initial payment* on each agent, at the time of awarding them the bundle they receive in the initial allocation. Payment function and initial payments together are referred to as the *payment scheme*.

Each agent  $i \in \mathcal{N}$  is also equipped with a *utility function*  $u_i : 2^{\mathcal{G}} \times \mathbb{R} \rightarrow \mathbb{R}$  mapping pairs of bundles and previous payments to the reals. These are fully determined by the valuation functions:  $u_i(S, x) = v_i(S) - x$ . 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  $(A, \pi)$ , 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).

## 2.3 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 (IR) if there exists a payment function  $p$  such that  $v_i(A') - v_i(A) > p(i)$  for all agents  $i \in \mathcal{N}$ , 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* [24]:

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

We also speak of the social welfare of a *state*  $(A, \pi)$ . 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 [29]:

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

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. In modular domains, we can get a much stronger convergence result [14]:

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

Theorems 1 and 2 are easy consequences of the fact that a move from one allocation to another is an IR deal if and only if that move increases social welfare (see [14] for a proof):

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

We shall frequently rely on this fact in the sequel, mostly without explicit reference to Lemma 1.

## 2.4 Specific 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. [15] 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 payment function amounts to choosing how to distribute the *social surplus*  $sw(A') - sw(A)$  generated by a deal  $\delta = (A, A')$ . By Lemma 1, the social surplus is positive if and only if  $\delta$  is IR. The LUPF divides this amount equally amongst the *participating agents*  $\mathcal{N}^\delta$ ; the GUPF divides it equally amongst *all agents*  $\mathcal{N}$ :

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

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

### 3 Proportional fairness and Knaster payments

Our first goal is to devise a distributed mechanism of the kind described in the previous section that can be used to negotiate outcomes that are *proportionally fair* (or simply *proportional*). An outcome, assigning each agent a bundle of goods and a payment balance, is proportional if the utility each agent assigns to their lot is at least  $1/n$ th of the value they place on the full set of goods (where  $n$  is the number of agents).

**Definition 3 (Proportionality)** *A negotiation state  $(A, \pi)$  is called proportional if  $u_i(A, \pi) \geq v_i(\mathcal{G})/n$  for all agents  $i \in \mathcal{N}$ .*

As a minimum requirement for this definition of fairness to be meaningful, we need to make the additional assumption that valuation functions are *monotonic*:  $S_1 \subseteq S_2$  implies  $v_i(S_1) \leq v_i(S_2)$ . Proportionality is the central fairness criterion applied in the cake-cutting literature, which addresses the problem of fairly dividing a single divisible good, when agents have *additive* preferences and monetary side payments are not possible [5, 27]. When agent valuations are *supermodular*, then it will often be impossible to obtain a proportional solution without using money (both for divisible and indivisible goods), because  $\frac{1}{n}$  of the value of the full set of goods would often be considered more than the value of  $\frac{1}{n}$  of the goods. For submodular valuations, on the other hand, proportionality would be a very undemanding fairness criterion, as it would often be easy to give an agent  $\frac{1}{n}$  of the value they ascribe to the full set by handing them a very small bundle. It is therefore questionable whether proportionality should be considered a relevant fairness criterion in case valuations may be submodular. This caution notwithstanding, due to the central importance of proportionality in the classical literature, we nevertheless choose to demonstrate here how to apply our methodology to this concept (but the reader should keep in mind that the resulting procedure is much more attractive in supermodular than in submodular domains).

To define a distributed procedure that can guarantee proportional outcomes we shall take inspiration from the classical fair division procedure proposed by Bronisław Knaster in the 1940s [30, 5, 26].<sup>1</sup> In this section, we recall the classical Knaster procedure, show how to adapt it to define a payment scheme, and prove a convergence theorem showing that this scheme guarantees outcomes that are both efficient and proportional.

#### 3.1 The Knaster procedure

Our exposition of the Knaster procedure follows Raith [26]. An important difference between the classical and our (more general) setting is that in the classical setting valuation functions are assumed to be modular (additive). This difference affects the first step of the procedure.

- (1) Compute an allocation  $A^*$  that maximises social welfare. (In the classical setting this step simplifies to assigning each good to the agent who values it the most.)
- (2) For each agent  $i$ , compute their *excess*  $ex_i(A^*) = v_i(A^*) - v_i(\mathcal{G})/n$  over their fair share. The *total excess* is defined as  $Ex(A^*) = \sum_i ex_i(A^*)$ .

---

<sup>1</sup>In an early paper, Knaster [21] introduced the concept of “excess” as used in Section 3.1 below, albeit in the context of cake-cutting (division of a single divisible good) without side payments. Steinhaus [30] later discussed a procedure for indivisible goods using this concept and attributed the idea for this procedure to Knaster.

- (3) Split the total excess evenly amongst all the agents; that is, require a payment of  $ex_i(A^*) - Ex(A^*)/n$  from each agent  $i$ .

Clearly, these payments sum up to zero, so are feasible. Note that an agent's individual excess  $ex_i$  could be negative. The total excess  $Ex(A^*)$ , however, will always be positive (or zero). This follows from the fact that  $Ex(A^*) = sw(A^*) - \frac{1}{n} \cdot \sum_i v_i(\mathcal{G})$  and that we must have  $sw(A^*) \geq v_i(\mathcal{G})$  for any  $i$  (because  $A^*$  is efficient by definition and hence at least as good as giving all goods to  $i$ ). After the procedure has been executed, agent  $i$  will enjoy utility  $v_i(A^*) - [ex_i(A^*) - Ex(A^*)/n] = v_i(\mathcal{G})/n + Ex(A^*)/n$ . Hence, the Knaster procedure does indeed produce a proportional outcome.

### 3.2 Definition of a payment scheme

As demonstrated by Theorem 1, our distributed negotiation procedure is capable of selecting an allocation  $A^*$  with maximal social welfare. We now want to devise a payment scheme that guarantees proportionality on top of efficiency. At the point of convergence we would like to obtain an outcome that is equivalent to the outcome we would have gotten by using the Knaster procedure.

Recall that defining a payment scheme involves fixing an initial payment  $\pi_0(i)$  for each agent and choosing a payment function  $p$ . Our approach is simply to choose  $\pi_0$  and  $p$  such that at each stage  $(A, \pi)$  of the negotiation we will have  $\pi(i) = ex_i(A) - Ex(A)/n$ , thereby emulating the payments prescribed by the Knaster procedure (with the crucial difference that in the Knaster procedure payments are only required once, at the very end). Let  $A_0$  be the initial (randomly chosen) allocation. We must fix  $\pi_0$  as follows:

$$\pi_0(i) = ex_i(A_0) - Ex(A_0)/n \quad (1)$$

Here  $ex_i$  and  $Ex$  are defined as in Section 3.1. Of course, if  $A_0$  is sufficiently close to an efficient allocation, then these payments will already result in a proportional outcome. On the other hand, for a sufficiently inefficient  $A_0$  neither these nor any other payments could possibly ensure proportionality.

Next we choose a payment function  $p$ . We want to maintain  $\pi(i) = ex_i(A) - Ex(A)/n$  as an invariant for all negotiation states  $(A, \pi)$ . This forces the following payments for a deal  $\delta = (A, A')$ , with  $\pi$  being the payment balance for allocation  $A$  and  $\pi'$  being the payment balance for  $A'$ :

$$\begin{aligned} p(i) &= \pi'(i) - \pi(i) \\ &= [ex_i(A') - Ex(A')/n] - [ex_i(A) - Ex(A)/n] \\ &= [v_i(A') - v_i(A)] - [sw(A') - sw(A)]/n \end{aligned}$$

But this is exactly the GUPF, defined in Section 2.4. In particular, the payments prescribed by this choice of  $p$  are IR. In summary, the *Knaster payment scheme* is given by the GUPF together with the initial payments defined in Equation (1).

### 3.3 Convergence

We are now going to show that, indeed, negotiation using the Knaster payment scheme will always converge to a proportional outcome.

**Theorem 3 (Proportional outcomes)** *Under the Knaster payment scheme, any sequence of IR deals will eventually result in a state that is efficient and proportional.*

*Proof.* By Theorem 1, any sequence of IR deals must converge to an efficient state, independently from the specific payment scheme in use. As we are using the Knaster payment scheme, any state  $(A, \pi)$  reached during negotiation will satisfy  $\pi(i) = ex_i(A) - Ex(A)/n$  for each agent  $i$ . Hence, we are bound to converge to a state  $(A^*, \pi^*)$  with an efficient allocation  $A^*$  where  $\pi^*(i) = ex_i(A^*) - Ex(A^*)/n$ . This is exactly the kind of outcome we would have obtained had we used the Knaster procedure, which we have seen to always be proportional (see Section 3.1).  $\square$

If all agents have modular valuation functions (as is the case for the classical setting of the Knaster procedure), then we can guarantee convergence to a proportional state using a simpler negotiation protocol.

**Theorem 4 (Proportional outcomes by 1-deals)** *Under the Knaster payment scheme, if all valuation functions are modular, any sequence of IR 1-deals will eventually result in a state that is efficient and proportional.*

*Proof.* By Theorem 2, any sequence of IR 1-deals must converge to an efficient state whenever all valuation functions are modular. The rest of the argument is identical to the proof of Theorem 3.  $\square$

We conclude our treatment of proportionality by briefly discussing alternative payment schemes for guaranteeing proportional outcomes. As shown in Section 3.1, the Knaster procedure guarantees each agent  $i$  a marginal utility of  $Ex(A^*)/n$  above and beyond their fair share of  $v_i(\mathcal{G})/n$ . But any alternative division of the total excess  $Ex(A^*)$  will also satisfy proportionality. The equal division chosen by Knaster is a particularly natural choice, but it is not the only one. Raith [26], for instance, argues that a division of the total excess that is proportional to each individual's valuation of the full set of goods may be more appropriate and proposes an adjustment of the Knaster procedure to achieve this.

Let  $(\alpha_i)_{i \in \mathcal{N}}$  with  $\alpha_i \geq 0$  and  $\sum_i \alpha_i = 1$  be a vector specifying for each agent their entitlement to a proportion of the total excess. By revisiting the proofs given above, it is not difficult to see that Theorems 3 and 4 will continue to hold if the Knaster payment scheme is replaced by any payment scheme of the following form:

$$\begin{aligned}\pi_0(i) &= ex_i(A_0) - \alpha_i \cdot Ex(A_0) \\ p(i) &= [v_i(A') - v_i(A)] - \alpha_i \cdot [sw(A') - sw(A)]\end{aligned}$$

For the Knaster payment scheme we have  $\alpha_i = \frac{1}{n}$ , while for the payment scheme corresponding to the Raith-Knaster procedure [26] we have  $\alpha_i = \frac{v_i(\mathcal{G})}{\sum_j v_j(\mathcal{G})}$ . The two payment schemes coincide if  $v_i(\mathcal{G}) = 1$  for all  $i \in \mathcal{N}$ .

## 4 Envy-freeness and globally uniform payments

Besides proportionality, an important fairness criterion is *envy-freeness* [5]. An allocation of goods is envy-free if no agent would rather have the bundle held by one of the other agents. This is a more demanding requirement than proportionality. Indeed, if we require all goods to be allocated, then an envy-free allocation may not even exist at all, and the problem of checking whether or not an



envy-free allocation exists has been shown to be computationally intractable [3]. Fortunately, in the presence of money, when envy-freeness is defined in terms of both bundles of goods and payments received, the situation is more favourable and envy-free solutions do exist [2].

In this section we shall identify conditions under which we can guarantee convergence to an envy-free outcome.

**Definition 4 (Envy-freeness)** *A negotiation state  $(A, \pi)$  is called envy-free if  $u_i(A(i), \pi(i)) \geq u_i(A(j), \pi(j))$  for all agents  $i, j \in \mathcal{N}$ .*

A state that is both efficient (in the sense of maximising social welfare) and envy-free will be referred to as an *EEF state*.

If all valuation functions are submodular (or modular), that is, when no agent ever likes the union of two disjoint bundles strictly more than the sum of the valuations of the two individual bundles, then any state that is envy-free is also proportional. The converse is not true: a proportional state is not necessarily envy-free. If valuation functions are supermodular, then envy-freeness does not guarantee proportionality, nor vice versa.

#### 4.1 Existence of envy-free states

In the presence of money, intuitively, chances for finding an envy-free solution are better than when there is no money. In particular, in the context of assignment problems, i.e., when each agent can only receive at most a single good, states that are both efficient and envy-free are known to always exist. This has been shown by Alkan et al. [1]. When agents can obtain (and express preferences over) bundles consisting of several goods, there is in general no such guarantee [31]. But the assumption of quasi-linearity of money is very important here: in this setting, the existence of EEF states can be proven by an argument similar to the one used by Alkan et al. [1], appealing to the Duality Theorem [17], as demonstrated by Beviá [2]. When all valuation functions are *supermodular*, then we can give a very simple existence argument for EEF states, which does not rely on the Duality Theorem. As we shall see later, this case is of special interest to us, which is why we present this argument here in some detail.

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  $\pi^*$  can be arranged such that the state  $(A^*, \pi^*)$  is EEF. Define  $\pi^*(i)$  for each agent  $i$ :

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

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

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

As  $v_i$  is assumed to be supermodular (and normalised), this entails:

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

Adding  $sw(A^*)/n$  to both sides of this inequality, 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,  $(A^*, \pi^*)$  is not only efficient, but also envy-free.

Of course, 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 above? And even if they do, how can we compute them in practice? Just finding an efficient allocation is already known to be NP-hard [28]. Finally, as argued in the introduction, we are interested in a *distributed* procedure, where agents identify the optimal state in an interactive manner.

## 4.2 Envy-freeness and individual rationality

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 good) shows that this is not the case:

$$v_1(\{g\}) = 4 \qquad v_2(\{g\}) = 7$$

Suppose agent 1 holds  $g$  in the initial allocation  $A_0$ . There is only a single possible deal, which amounts to passing  $g$  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.

We will be able to overcome this problem by choosing a suitable payment scheme. Namely, for the procedure proposed in the sequel, we use a payment scheme fixing the following initial payments:

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

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 choice of initial payment as an *initial equitability payment*. Note that this payment does *not* achieve envy-freeness (and it does not affect efficiency at all). 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  $i \in \mathcal{N}$ ), the initial payments reduce to 0.

## 4.3 Convergence in supermodular domains

We are now going to prove a 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.

**Theorem 5 (Envy-free outcomes)** *If all valuation functions 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  $(A, \pi)$  and every agent  $i$ , provided that agents only negotiate IR deals using the GUPF:

$$\pi(i) = v_i(A) - sw(A)/n \quad (2)$$

Our proof proceeds by induction over the number of deals negotiated. As we assume that initial equitability payments have been made, claim (2) will certainly be true for the initial state  $(A_0, \pi_0)$ . Now let  $\delta = (A, A')$  and assume (2) holds for  $A$  and the associated payment balance  $\pi$ . We obtain the payment balance  $\pi'$  associated with  $A'$  by adding the appropriate GUPF payments to  $\pi$ :

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

This proves our claim (2). Now, Theorem 1 shows that negotiation via IR deals (whichever payment function is used) must eventually terminate and that the final allocation will be efficient. Let  $A^*$  be that final allocation, and let  $\pi^*$  be the associated payment balance. Equation (2) also applies to the state  $(A^*, \pi^*)$ . But as we have already seen in Section 4.1, the efficiency of  $A^*$  together with Equation (2) applied to  $(A^*, \pi^*)$  implies that  $(A^*, \pi^*)$  must be an EEF state.  $\square$

Theorem 5 may seem surprising. As pointed out elsewhere [14], it is not possible to define a “local” criterion for the acceptability of deals (which can be checked taking only the valuation functions 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. It is limited, because only the agents involved in a deal can ever be asked to give away money, and all payments can be computed taking only the valuations of those involved into account.

A natural question that follows is whether Theorem 5 could be generalised, namely whether a more general class of valuations would still allow us to guarantee convergence to an envy-free state. The following result shows that this is not the case, at least not if we keep the payment scheme of Theorem 5 in place.

**Theorem 6 (Maximality of supermodularity)** *No class  $\mathcal{F}$  of valuation functions that strictly includes the class of supermodular functions can guarantee the following property: if all valuation functions are drawn from  $\mathcal{F}$  and initial equitability payments have been made, then any sequence of IR deals using the GUPF will eventually result in an EEF state.*

*Proof.* First, observe that under the payment scheme of Theorem 5, which is characterised by Inequality (2), the envy-freeness condition  $u_i(A(i), \pi(i)) \geq u_i(A(j), \pi(j))$  is equivalent to the following

inequality (which must hold for any two agents  $i$  and  $j$ ):

$$v_j(A(j)) \geq v_i(A(j)) \tag{3}$$

To prove the claim of the theorem, we give an example with two agents and two goods, where Inequality (3) fails to be satisfied for the unique efficient allocation. Hence, as negotiation by means of IR deals is bound to lead to that efficient allocation, and as the payment scheme is fixed, this will show that the final allocation cannot be envy-free. The first agent will have an arbitrary non-supermodular valuation function; for the second agent we will construct a supermodular valuation function. Thus, our counterexample is generic and will apply for any superclass  $\mathcal{F}$  of the supermodular functions (which must include at least one non-supermodular function and all supermodular functions). Showing failure of convergence for the simple case of two agents and two goods immediately implies failure of convergence for larger negotiation problems.

Our counterexample is constructed as follows. Take  $v_1$  to be *any* non-supermodular valuation function defined over two goods  $g_1$  and  $g_2$ ; that is

$$v_1(\{g_1, g_2\}) = v_1(\{g_1\}) + v_1(\{g_2\}) - d,$$

where  $d > 0$ . Now we construct a second valuation function  $v_2$ , which we assume to be modular (so it is certainly supermodular), as follows:

$$\begin{aligned} v_2(\{g_1\}) &= v_1(\{g_1\}) - d/2 \\ v_2(\{g_2\}) &= v_1(\{g_2\}) - d/3 \end{aligned}$$

But now in a scenario with two agents with these two valuation functions, it can easily be checked that the (only) efficient allocation is to give good  $g_1$  to agent 1, and good  $g_2$  to agent 2, yielding a social welfare of  $v_1(\{g_1\}) + v_2(\{g_2\}) - d/3$ . However it is clear that this allocation does not satisfy Inequality (3). Indeed, agent 1 values  $g_2$  more than agent 2 does; that is, it is not the case that agent 2 values its own share at least as much as anyone else would.  $\square$

Theorem 6 is reminiscent of so-called *maximality theorems* in multiagent resource allocation [8]. Observe that the proof of Theorem 6 uses Equation (3), which is only mandatory if the payment scheme being used is the GUPF with initial equitability payments. In particular, it could be the case that for different payment functions, different valuation functions could work. However, as the GUPF is attractive due to its simplicity, Theorem 6 offers an interesting characterisation of a domain of valuation functions guaranteeing convergence to states without envy.

Vice versa, as we shall see next, if we fix the supermodularity condition, then the payment scheme of Theorem 5 turns out to be the only scheme (meeting certain mild conditions) that would allow us to obtain a convergence result. To make this claim precise, let us call a payment function  $p$  *reasonable* if there exists a function  $f : \mathcal{N} \times \mathbb{R}^n \rightarrow \mathbb{R}$  such that for any deal  $\delta = (A, A')$  the payment  $p(i)$  of any agent  $i$  is equal to  $f(i, \langle \Delta_1, \dots, \Delta_n \rangle)$ , where  $\Delta_j = v_j(A') - v_j(A)$ . That is, a payment function is reasonable if payments only depend on the agents' changes in valuation resulting from a deal. In particular, payments are not dependent on the position of the deal in the overall sequence. We call a payment scheme reasonable if the associated payment function is reasonable.

**Theorem 7 (Payment schemes)** *No reasonable payment scheme  $\Pi$  other than the GUPF with initial equitability payments can guarantee the following property: if all valuation functions are supermodular, then any sequence of IR deals using  $\Pi$  will eventually result in an EEF state.*

*Proof.* Observe that there are many scenarios where, for the final state to be EEF, the final payment balance must be defined by  $\pi^*(i) = v_i(A^*) - sw(A^*)/n$ , as in the proof of Theorem 5. Any situation where all agents have the same valuation function may serve as an example. Of course, any number of payment functions could achieve these final payment balances, as long as we can be sure that the rule applied during the very last deal is such that we get the correct values for  $\pi^*$ . We need to show that, amongst the *reasonable* payment schemes, the scheme of Theorem 5 is the only one with this property.

As reasonable schemes do not allow us to make payments dependent on where in the sequence of deals we currently are,  $\pi(i) = v_i(A) - sw(A)/n$  must hold after *every* deal. But this *forces* initial payments to be exactly as in Theorem 5 (initial equitability payments), because the initial allocation may already be efficient and hence final; and the only possible payment function is the GUPF, as it is precisely the function we obtain when we compute the difference of the payment balances for two consecutive negotiation states.  $\square$

In summary, Theorems 6 and 7 together show that our convergence result, Theorem 5, is in some sense as strong as possible: we can neither relax the range of valuation functions to which it applies nor the payment scheme for which it will go through.

#### 4.4 Convergence in modular domains

In modular domains, we can strengthen Theorem 5 and even guarantee convergence to an EEF state by means of 1-deals (over one item at a time):

**Theorem 8 (Envy-free outcomes by 1-deals)** *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 5, except that we rely on Theorem 2 for convergence by means of 1-deals (in place of Theorem 1). Note that the argument of Section 4.1 still applies, because any modular valuation function is also supermodular.  $\square$

#### 4.5 Related work

There has been some work on procedures for finding EEF states in the social choice literature [1, 20, 18], albeit with little or no attention to computational issues. The work of Haake et al. [18] is particularly interesting. These authors 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 elegant and intuitively appealing, it does not 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, as our discussion in Section 4.1 demonstrates that this is *not* a hard combinatorial problem: there is a simple procedure for choosing the payments.

The second procedure put forward by Haake et al. [18] interleaves reallocations for increasing efficiency with payments for eliminating envy. However, here the authors also do not address a hard 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 the assignment problem of allocating  $n$  objects to  $n$  agents, which, unlike the problem addressed by Theorem 5, is not an NP-hard problem [22].

## 5 Fair division on graphs

As argued in the introduction, there are good reasons for studying distributed mechanisms for fair division, such as a lack of confidence in —or the mere absence of— a central authority for regulating interaction. One even more compelling reason is that agents may be spatially distributed, with restricted interaction opportunities between them. In this context, the assumption of a fully connected graph (or global network) connecting agents quickly disappears. The pervasiveness of applications exhibiting underlying graph-like structures, like for instance small-worlds (be it in Grid computing, or in social networks) is thus a solid motivation to study distributed mechanisms for allocating goods. But it also necessitates making appropriate assumptions about agents being only able to act and perceive their environment *locally*.

To account for this, we shall assume that not every agent is able to “see” all of the other agents and formulate an appropriate generalisation of our model.<sup>2</sup> A *negotiation topology* is an undirected graph  $G = (\mathcal{N}, E)$ , the vertices of which are the agents in  $\mathcal{N}$ . Two agents  $i$  and  $j$  stand in the relation  $E$  if and only if they can see each other. This means in particular that  $i$  and  $j$  may engage in negotiation and exchange goods. We put the following structural restriction on deals: any deal is possible, as long as it only involves agents belonging to a common *clique* of the graph  $G$ . (A clique is a set of vertices  $C \subseteq \mathcal{N}$  such that  $(i, j) \in E$  for all distinct  $i, j \in C$ .) We call deals meeting this condition *clique-deals*. Our visibility relation  $E$  is symmetric (the graph is undirected); this is important to be able to define negotiation along graphs in a meaningful manner. Note that we shall use the same topological constraints in order to define what information is available to agents in the graph. For instance, agents may not be aware of the goods being held by agents outside of their scope of visibility. (In theory, different graphs could be used for these two aspects, but we shall not consider this here.)

As we shall see in greater detail below, it will not always be straightforward to adapt notions of fairness to the context of graphs. In particular, we will see that no definition of proportionality seems totally satisfying, while envy gives rise to a very natural interpretation.

---

<sup>2</sup>De Weerd et al. [11] explore a similar model in the context of task allocation problems.

## 5.1 Defining fairness on graphs

Defining proportional fairness on graphs is less straightforward than one might expect. This is so because we need to make more precise what information is available to agents located on the graph when they have to assess whether their share is proportionally fair. Proportionality is typically evaluated with respect to the full set of goods  $\mathcal{G}$  available in the system. In the distributed settings involving graphs discussed in this section, agents may not have full information regarding the current location of the goods, so it is important to define exactly what information is available to them. It is reasonable to assume that each agent must have at least access to its local neighbourhood (agents and goods that they can “see”). Then, depending on what we assume to be available to agents on top of this, we may favour different intuitive definitions of proportionality on a graph.

The first option is to assume that the full set of goods is known to the agents (even if their exact location is not known), and to define proportionality accordingly. A less demanding (and arguably more intuitive) definition is to have proportionality defined with respect to a subset of the full set of goods of the system, namely the goods an agent has seen so far. The rationale behind this approach is that agents start the process with very little knowledge regarding the goods being negotiated, and learn during interactions. Unfortunately, it is easy to see that, for both definitions, our framework could not possibly guarantee proportional outcomes to occur in all cases (even though the second option is more likely to reach such an outcome, of course). The reason is simple: even though “full” efficiency is not required as such to guarantee proportionality, allocations can be too inefficient to allow proportionality. But it is easy to design negotiation topologies so that outcomes are bound to be highly inefficient, due to agents with low valuations for particular goods acting as bottlenecks. A third alternative option is to define proportionality with respect to what you can currently see in the network. This solution is highly questionable though, because it means that agents must forget about the goods they have seen in previous states.

Unlike proportionality, the notion of envy happens to give rise to a very natural extension in the context of graphs. In this case indeed, the restriction to the information available to agents in their neighbourhood is very appealing: envy can only be experienced by agents with respect to the agents they can see on the network. For example, in the extreme case of a society of completely disconnected agents, no agent would ever be envious of the situation of any other agent. To account for a negotiation topology, we then propose the following modification of the definition of envy-freeness:

**Definition 5 (GEF states)** *A state  $(A, \pi)$  is called graph-envy-free (GEF) with respect to the graph  $G = (\mathcal{N}, E)$  if  $u_i(A(i), \pi(i)) \geq u_i(A(j), \pi(j))$  for all agents  $(i, j) \in E$ .*

As we have seen in Section 4, convergence to envy-freeness relies on the final allocation being efficient (or at least efficient “enough”). But as argued above, it is always possible to construct scenarios that would keep outcomes very far from full efficiency. In what follows, we see however that some weaker notion of efficiency can be satisfied, which will suffice to permit envy-freeness on graphs to hold.

## 5.2 Clique-wise efficiency

We now develop a notion of efficiency that takes the negotiation topology into account and relate this notion to IR negotiation when restricted to clique-deals.

**Definition 6 (Clique-variants)** *Let  $A$  be an allocation. Another allocation  $A'$  is called a clique-variant of  $A$  if and only if there exists a clique  $C \subseteq \mathcal{N}$  such that  $\bigcup_{i \in C} A(i) = \bigcup_{i \in C} A'(i)$  and  $A(i) = A'(i)$  for all  $i \notin C$ .*

Observe that  $A$  and  $A'$  are clique-variants of each other if and only if  $\delta = (A, A')$  is a clique-deal.

**Definition 7 (Clique-wise efficiency)** *An allocation  $A$  is called clique-wise efficient if  $sw(A) \geq sw(A')$  for every clique-variant  $A'$  of  $A$ .*

It should be noted that, in its own right, this notion of efficiency would only be of very limited interest. While our definition of envy with respect to a graph is very natural and reaching GEF states seems indeed desirable, it is questionable whether the standard notion of efficiency can be relativised with respect to a negotiation topology in a meaningful manner. Our interest in clique-wise efficiency stems from the fact that it will be helpful in characterising conditions under which convergence to a GEF state can be guaranteed (as will become clear in the sequel). But first we prove a convergence result for clique-wise efficiency:

**Lemma 2 (Clique-wise efficient outcomes)** *Any sequence of IR clique-deals will eventually result in a clique-wise efficient allocation of goods.*

*Proof.* There can be no infinite sequence of deals (clique-deals or otherwise), because any deal strictly increases social welfare and the set of possible allocations is finite. Now let  $A$  be the terminal allocation. For the sake of contradiction, suppose that  $A$  is *not* clique-wise efficient. Then there exists an allocation  $A'$  that is a clique-variant of  $A$  such that  $sw(A) < sw(A')$ . But then, by Lemma 1,  $\delta = (A, A')$  must be IR (besides being a clique-deal). This contradicts our assumption of  $A$  being a terminal allocation.  $\square$

Lemma 2 generalises Theorem 1. The latter corresponds to the case of a fully connected graph. While Lemma 2 guarantees clique-efficient outcomes, it does not say anything about *which* clique-efficient allocation will be reached. For most graphs there will be a range of clique-efficient allocation of varying quality in terms of global efficiency (in this sense the concept is similar to that of Pareto efficiency). There is no guarantee that we will end up with the “best” clique-efficient allocation. Nevertheless, as we shall see next, clique-wise efficiency is a sufficiently strong notion to serve as a basis for negotiating envy-free states.

### 5.3 Convergence to GEF states

We now prove a convergence theorem for GEF states, which extends Theorem 5 to the framework with a negotiation topology: we show that under the same conditions (on the valuation functions and for a particular choice of payment scheme), any sequence of IR deals that respect the negotiation topology will result in a GEF state.

**Theorem 9 (Convergence on graphs)** *If all valuations are supermodular and if initial equitability payments have been made, then any sequence of IR clique-deals using the GUPF will eventually result in a GEF state.*



*Proof.* Recall from the proof of Theorem 5 that the use of the GUPF together with initial equitability payments ensures that we get a payment balance satisfying  $\pi(i) = v_i(A) - sw(A)/n$  for every state  $(A, \pi)$  reached during negotiation and every agent  $i \in \mathcal{N}$ . By Lemma 2, negotiation will eventually terminate and the final allocation  $A^*$  will be clique-wise efficient. The associated payment balance will be  $\pi^*(i) = v_i(A^*) - sw(A^*)/n$ . We need to show that the state  $(A^*, \pi^*)$  must be GEF whenever all valuations  $v_i$  are supermodular. Let  $i$  and  $j$  be any two agents. If  $i$  cannot see  $j$  then we are done. Otherwise,  $i$  and  $j$  are part of a clique, and due to the clique-wise efficiency of  $A^*$ , giving both  $A^*(i)$  and  $A^*(j)$  to  $i$  will not increase the sum of valuations for this clique any further. Hence, by the argument familiar from Section 4.1,  $i$  does not envy  $j$  in state  $(A^*, \pi^*)$ .  $\square$

Theorem 9 means that agents can negotiate in a distributed manner, guided only by their own rational interests and limited to their “neighbourhoods” as given by the cliques of the negotiation topology, and—as long as all the side conditions are satisfied—a state that is envy-free according to all agents (whose vision is limited by the negotiation topology) will eventually emerge. In particular, agents can go ahead and negotiate any beneficial deals, without fear of getting stuck in a local optimum.

A critical point in Theorem 9 is the use of the GUPF, as this payment function does not respect the negotiation topology. However, it is easy to show that there can be *no clique-wise payment function* (a payment function giving non-zero payments only to agents belonging to a particular maximal clique within which the deal is taking place) that would allow us to achieve a convergence result for GEF states. To see this, consider the following example. Suppose there are three agents on a line (i.e.,  $E = \{(1, 2), (2, 1), (2, 3), (3, 2)\}$ ), the payment balance is currently 0 for all agents, and the valuation function of agent 3 is  $v_3(S) = 0$  for any  $S \subseteq \mathcal{G}$ . Then any deal between agents 1 and 2, where the former makes a non-zero payment to the latter, will render agent 3 envious of agent 2.

A variation of this example shows that even when there exists a clique-wise payment function leading to a GEF state, the exact amount of the payments to be made may depend on agents outside the clique where the deal is taking place. For the same negotiation topology as above, let again  $v_3(S) = 0$  for any  $S \subseteq \mathcal{G}$ , but now suppose that agent 3 has benefited from a previous deal in monetary terms, i.e.,  $\pi(3) = x$  for some  $x < 0$ . Then an envy-eliminating deal between agents 1 and 2 should be such that it brings the payment balance of agent 2 to at least  $x$  as well (which may or may not be possible, depending on the scenario at hand). This shows that the best possible clique-wise payment function may not be identifiable locally.

A positive point to be made about the GUPF is that the payments to non-involved agents (in particular those outside the clique where the deal is taking place) solely depend on the social surplus generated by the deal and the overall number of agents in the system. So agents do only need to be “aware” of agents they cannot “see” in so far as they need to know their overall number. This arguably corresponds well to human society: our sphere of influence may be very much restricted to a small section of society (negotiation topology), but we are still aware of some basic facts concerning society as a whole (such as the number of its members).

Finally, in analogy to Theorems 4 and 8, if all valuation functions are modular we can strengthen Theorem 9 and prove convergence by means of IR 1-deals between connected agents.

## 6 Degrees of envy and computational complexity

The convergence theorems of the previous two sections show that envy-freeness can be guaranteed for the outcome of a distributed negotiation process, under specific assumptions. But this is certainly not the final word on the matter. The payment scheme used in these theorems introduces a non-local element, in the sense that it redistributes the social surplus over the whole society. One may ask how a local payment function, such as the LUPF, would fare in comparison. Also, one could object to the requirement that initial payments have to be made. The restriction to supermodular valuation functions in Theorems 5 and 9 is another limitation. Lastly, our convergence theorems do not say how envy evolves over the course of negotiation. Because negotiation can become very long in practice, it is possible that the process will have to 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 be able to address such questions, more than the mere classification of a negotiation state as being either envy-free or not is needed. We require a way to measure the *degree of envy* in a society. In this section, we shall propose a systematic approach to defining metrics for assessing the degree of envy. We propose to analyse the degree of envy of a society as a three-level aggregation process, starting with envy between two agents, over envy of a single agent towards everyone else, to eventually provide a definition for the degree of envy of a society. Towards the end of the section we then prove a theorem on the computational complexity of the problem of finding an IR deal that reduces envy according to such a metric.

### 6.1 Envy between two individual agents

How much does agent  $i$  envy agent  $j$  in state  $(A, \pi)$ ? Any such metric will be based on the difference in utility that  $i$  assigns to their own lot and that of  $j$ :

$$e(i, j) = u_i(A(j), \pi(j)) - u_i(A(i), \pi(i))$$

We now use  $e(i, j)$  to define different metrics for the degree of envy between individual agents:<sup>3</sup>

$$\begin{aligned} e^{id}(i, j) &= e(i, j) \\ e^{pos}(i, j) &= \max\{e(i, j), 0\} \\ e^{sgn}(i, j) &= \text{sgn}(e(i, j)) \\ e^{0-1}(i, j) &= \max\{\text{sgn}(e(i, j)), 0\} \end{aligned}$$

The last of these ( $e^{0-1}$ ) only allows us to specify whether or not  $i$  envies  $j$ . In most cases this is all that is needed. Sometimes we may also want to be able to express that  $i$  may experience “negative envy” towards  $j$  ( $e^{sgn}$ ). This can be useful if we want agents to be able to compensate for envy. The first two metrics are useful if we also want to say *how much*  $i$  envies (or does not envy)  $j$ .

### 6.2 Degree of envy of a single agent

How envious is a single agent  $i$  in state  $(A, \pi)$ ? This notion considers the agent in relation to many other agents. That is, we need to aggregate the individual envies that  $i$  experiences *vis-à-vis* each of

---

<sup>3</sup>Here  $\text{sgn} : \mathbb{R} \rightarrow \{-1, 0, 1\}$  is the function mapping positive numbers to 1, negative numbers to  $-1$ , and 0 to 0.

the other agents. The most natural aggregation operators are *max* and summation, but other options are possible as well.

$$\begin{aligned} e^{sum,op}(i) &= \sum_{j \in \mathcal{N}} e^{op}(i,j) \\ e^{max,op}(i) &= \max_{j \in \mathcal{N}} e^{op}(i,j) \end{aligned}$$

Here *op* may stand for any of  $\{id, pos, sgn, 0-1\}$ . For instance,  $e^{max,pos}(i)$  measures how much  $i$  envies the agent it envies the most, while  $e^{sum,0-1}(i)$  counts the number of agents envied by  $i$ .

### 6.3 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 [24], there are a multitude of different options available for doing this. Here again, just as for the case of the degree of envy of a given agent, we only list two options that appear particularly appropriate in our context, namely *max* and summation.

$$\begin{aligned} e^{sum,op_1,op_2} &= \sum_{i \in \mathcal{N}} e^{op_1,op_2}(i) \\ e^{max,op_1,op_2} &= \max_{i \in \mathcal{N}} e^{op_1,op_2}(i) \end{aligned}$$

Using the *max*-operator means focusing on the most envious agent of the society (whatever operator was chosen to measure that), while the summation (or, equivalently, averaging) provides a more global picture of the situation. The former has an egalitarian flavour, while the latter is utilitarian in nature.

Our framework allows us to define a range of different metrics for measuring envy. Several of these have been proposed before in the literature:

- Feldman and Kirman [16] make several proposals for how to measure envy. Their first suggestion is to count the number of pairs of agents where the first agent envies the second ( $e^{sum,sum,0-1}$ ). They also define two alternative measures for the degree of envy of a single agent, corresponding to our  $e^{sum,id}$  and  $e^{sum,pos}$ , and propose to measure the degree of envy of a society by means of a weighted sum of either one of these two measures. That is, their discussion includes, in particular,  $e^{sum,sum,id}$  and  $e^{sum,sum,pos}$ .
- Lipton et al. [23] work with two envy measures. The first of these is the “maximum envy between any pair of players” ( $e^{max,max,pos}$ ). Their second envy measure is the maximum *envy ratio*, defined as  $\max_{i,j \in \mathcal{N}} \left\{ \frac{u_i(A(j),\pi(j))}{u_i(A(i),\pi(i))}, 1 \right\}$ , which is not covered by our system. Caragiannis et al. [6] also work with  $e^{max,max,pos}$ .
- Brams et al. [4] refer to two measures of envy: the “number of envies of all players” ( $e^{sum,sum,0-1}$ ) and the “maximum number of envies of an individual player” ( $e^{max,sum,0-1}$ ).

Another obvious choice for measuring envy is the number of envious agents in a society ( $e^{sum,max,0-1}$ ). Finally,  $e^{max,max,0-1}$  should be considered a degenerate measure of envy, as it only allows us to distinguish between situations that are envy-free and those that are not.

## 6.4 Complexity of reducing envy

As we have seen (in Section 4.2), the objectives of achieving envy-freeness and of maintaining individual rationality during negotiation are not entirely compatible. Without initial payments, there can be negotiation states such that no IR deal would allow us to eliminate envy. So one may ask whether a given negotiation state admits a follow-up deal that would not only be IR but that also reduces envy in the system (or, as a special case, that completely eliminates envy). Next, we want to analyse the computational complexity of this problem.

Our result will apply to any measure of envy  $envy(., .)$  mapping negotiation states to the rationals that satisfies the following three conditions:

- (1)  $envy(A, \pi) \leq 0$  if (though not necessarily only if)  $(A, \pi)$  is envy-free.
- (2)  $envy(., .)$  is monotonically non-decreasing in  $e(i, j)$  for any  $i, j \in \mathcal{N}$ .
- (3)  $envy(., .)$  is polynomial-time computable.

In particular, all of the measures of envy defined above satisfy the first two of these conditions.<sup>4</sup> The third condition is satisfied as long as the problem of evaluating a valuation function over a given set of goods is a polynomial problem. This is the case for most reasonable languages for representing valuation functions.<sup>5</sup>

We now define the decision problem we want to analyse. For the purposes of stating the problem, let a *negotiation problem* be specified by a negotiation topology  $G = (\mathcal{N}, E)$ ; a set of goods  $\mathcal{G}$ ; and a profile of valuations  $\langle v_1, \dots, v_n \rangle$ .

---

WELFARE IMPROVEMENT WITH ENVY REDUCTION (WIER)

---

**Instance:** negotiation problem; state  $(A, \pi)$ ;  $K \in \mathbb{Q}$ ;  
measure of envy  $envy(., .)$

**Question:** Is there an alternative state  $(A', \pi')$  satisfying

- $\delta = (A, A')$  is a clique-deal;
  - $v_i(A') - v_i(A) > \pi'(i) - \pi(i)$  for all  $i \in \mathcal{N}$ ;
  - $envy(A, \pi) - envy(A', \pi') \geq K$ ?
- 

That is, WIER asks whether, for a given negotiation state, there exists a clique-deal that is IR (second condition) and for which payments can be arranged such that the degree of envy reduces at least by an amount of  $K$  (third condition). Observe that, if we set  $K = envy(A, \pi)$ , then WIER asks whether there exists an IR clique-deal that eliminates envy entirely (independently of the measure of envy chosen).

---

<sup>4</sup>Envy measures that allow for compensation between positive and negative envy, such as  $e^{sum, sum, id}$ , may also signal zero envy when a state is not envy-free according to Definition 4. Whether or not such a measure of envy is of interest depends on the application at hand. We give such weak conditions here, because our complexity result does not require stronger assumptions.

<sup>5</sup>However, there are exceptions. Some of the languages used in the literature on combinatorial auctions for representing the valuations of bidders do not satisfy this criterion. For the so-called OR-language, for instance, computing the price for a given bundle given a valuation expressed in that language is NP-hard [25].

To establish the complexity of WIER, we are going to use a reduction from another decision problem:

WELFARE IMPROVEMENT (WI)
<b>Instance:</b> negotiation problem; state $(A, \pi)$
<b>Question:</b> Is there a state $(A', \pi')$ such that
<ul style="list-style-type: none"> <li>• <math>v_i(A') - v_i(A) &gt; \pi'(i) - \pi(i)</math> for all <math>i \in \mathcal{N}</math>?</li> </ul>

That is, WI asks whether a given state admits an IR deal (without any restriction to clique-deals). The name of the problem stems from the fact that any IR deal increases social welfare. WI has been shown to be NP-complete for a range of different languages for representing the valuation functions of agents [12, 7]. As our interest here is not in preference representation languages, we state our complexity result without explicit reference to the language used; it applies to any language for which WI is NP-complete.

**Theorem 10 (Complexity)** *WIER is NP-complete.*

*Proof.* NP-membership follows from the fact that the conditions of WIER can be verified in polynomial time for any proposed solution state. We are going to show NP-hardness by reduction from WI, which is known to be NP-hard even for the case of just two agents. Given an instance of WI for two agents with valuations  $\langle v_1, v_2 \rangle$ , goods  $\mathcal{G}$ , and state  $(A, \pi)$ , we build an instance of WIER (for three agents) as follows: The negotiation topology is  $G = (\{1, 2, 3\}, \{(1, 2), (2, 1), (2, 3), (3, 2)\})$  and the set of goods is  $\mathcal{G} \cup \{g_1^*, g_2^*, g_3^*\}$ . The initial state  $(A^*, \pi^*)$  is defined as follows:

$$\begin{aligned} A^*(1) &= A(1) \cup \{g_1^*\} & \pi^*(1) &= \pi(1) \\ A^*(2) &= A(2) \cup \{g_2^*\} & \pi^*(2) &= \pi(2) \\ A^*(3) &= \{g_3^*\} & \pi^*(3) &= 0 \end{aligned}$$

The valuation functions of the first two agents are given by<sup>6</sup>  $v_1^* = v_1 - 3\Omega g_2^* - 3\Omega g_3^*$  and  $v_2^* = v_2 - 3\Omega g_1^* - 3\Omega g_3^*$ , and  $v_3^*$  is defined as follows:

$$v_3^*(S) = \begin{cases} \Omega & \text{if } S = A^*(1) \text{ or } S = A^*(2) \\ -\Omega & \text{if } S = \{g_3^*\} \\ -2\Omega & \text{otherwise} \end{cases}$$

Here,  $\Omega$  is a constant, the value of which will be chosen big enough. Finally, we set  $K = \text{envy}(A^*, \pi^*)$ , i.e., we are looking for an IR clique-deal that will reduce the degree of envy to 0 (or below). This completes the description of the WIER instance. In state  $(A^*, \pi^*)$ , the utility enjoyed by agent 3 will be  $-\Omega$ , and that of agents 1 and 2 will be  $v_1(A) - \pi(1)$  and  $v_2(A) - \pi(2)$ , respectively. Now consider the following facts:

- In the initial state  $(A^*, \pi^*)$ , agent 3 envies agent 2 ( $e(3, 2) = 2\Omega - \pi^*(2)$ ), and agents 1 and 2 are envy-free.

---

<sup>6</sup>We use a shorthand to define  $v_1^*$  and  $v_2^*$ : to compute  $v_1^*(S)$  take  $v_1(S \setminus \{g_1^*, g_2^*, g_3^*\})$  and then subtract  $3\Omega$  if  $g_2^* \in S$  and again  $3\Omega$  if  $g_3^* \in S$ , and similarly for  $v_2^*$ .

- *No IR deal can involve agent 3.* Suppose agent 3 gets  $A^*(1)$  or  $A^*(2)$ , then its utility increases by  $2\Omega$ . Then, either 1 or 2 would get  $g_3^*$ , which would reduce social welfare by  $3\Omega$ . Also, if 3 gets any other bundle, social welfare will decrease by an amount of  $\Omega$ .
- *No IR deal can involve any of the goods  $\{g_1^*, g_2^*, g_3^*\}$ .* Such a deal would decrease social welfare by at least  $\Omega$ .
- *A deal is IR on the WI instance if and only if it is IR on the WIER instance.* This is a consequence of the previous two facts.
- *Any IR deal  $\delta = (A^*, A^{*'})$  will reduce the envy of agent 3.* Specifically, 3 can only envy 2 and the bundle held by 2 will become unattractive for agent 3 ( $v_3 = -2\Omega$ ). That is,  $e(3, 2)$  will go down to  $-\Omega - \pi^{*'}(2)$ . Thus, if  $\Omega$  has been chosen sufficiently big, then the envy of agent 3 will become 0 (or negative) after any IR deal.
- *Agents 1 and 2 will stay envy-free after any IR deal.* This must be so because of the goods  $g_2^*, g_3^*$  (and  $g_1^*, g_3^*$ ) they highly dislike.

The last three points imply that: if the answer to the WI problem is YES, then there is also an IR deal for the WIER instance, and any such deal is bound to bring envy down to 0 (or below). If the answer to the WI problem is NO, then there is also no IR deal for the WIER instance. Hence, if we have an algorithm for solving WIER, we can use it to solve any instance of WI by first using the translation of the input problem described above. In other words, the WIER problem is at least as hard as the WI problem.  $\square$

Our proof demonstrates that the restriction to clique-deals does not affect computational complexity for this particular kind of decision problem. The same is true for the addition of the condition that the chosen deal is required to decrease envy, but this is a lot more surprising. Intuitively, finding a deal that is not just IR, but also envy-reducing, seems a lot harder than finding an IR deal. Of course, this perceived increase in complexity could conceivably both increase or decrease the complexity of checking *whether* such a deal exists (which is the question asked by WI and WIER, respectively). Theorem 10 shows that neither is the case; the problem remains NP-complete.

## 7 Conclusion

In this paper we have addressed the problem of designing mechanisms to fairly allocate indivisible goods to a set of agents. While this question has been mostly investigated in the context of assignment problems, in which agents can receive at most one such good [1], the domain considered here is combinatorial: agents may receive more than one object and have preferences over the possible *sets* of goods that they may receive. Unlike other proposals [20, 18], the approach advocated here is *distributed*: agents are able to locally implement mutually beneficial reallocations of goods (with monetary compensation), so the quests for efficiency and fairness are intimately interleaved.

The first contribution of this paper is to provide simple conditions under which any sequence of such reallocative steps can be *guaranteed* to converge to a fair outcome. The fairness criteria

considered are proportionality and envy-freeness. As the problem is combinatorial in nature, the approach necessitates that agents can potentially agree on complex deals, and the process itself can involve a large number of steps. If we assume though that preferences are modular (additive), then negotiation can be conducted by means of very simple deals, and the length of the process will be linearly bounded [13]. Note that the proposed mechanism never requires heavy *central* computation: only some simple redistribution of the social surplus (in terms of money) is performed at the global level.

Theorem 3 demonstrates that the classical Knaster procedure can be given a distributed flavour and still achieve proportional outcomes. Theorem 5 shows that envy-free outcomes can be negotiated in a distributed manner, at least when the agents have supermodular preferences. The discussion which follows our exposition of this last result demonstrates the limitations of the approach, by showing that the result does not generalise to a wider class of preferences (Theorem 6) or a different type of payment scheme (Theorem 7).

Having established these results, we turned our attention to a situation of great practical relevance, namely societies of agents where information available is constrained by a graph structure. Interestingly, we saw that while the notion of proportionality was difficult to characterise in a meaningful manner, envy is intuitively restricted to agents that can be “seen” from an agent’s position on the network. These topological constraints can be very damaging for efficiency, which suggest that it could potentially jeopardise the possibility to reach envy-free states. Theorem 9 shows that this is not the case: a weaker (localised) notion of efficiency is sufficient to guarantee convergence to envy-free states on graphs. While this result assumes that negotiation and envy are constrained by the same topology, an interesting lead for future research could be to relax this assumption. Typically, the negotiation graph would be a subgraph of the envy graph, so convergence to fair outcomes should be more difficult to guarantee.

For all of our convergence results we have also proved corresponding results for the special case of agents with modular preferences. In these cases, convergence can be achieved by means of individually rational deals involving only a single item at a time (Theorems 4 and 8, the latter of which can easily be generalised to the case of fair division on a graph).

By construction (and following Lemma 1), the proposed mechanisms enjoy an “anytime behaviour” as far as efficiency is concerned. That is, efficiency is guaranteed to improve monotonically as negotiation progresses. The situation is less clear when it comes to envy, first of all because there is no obvious way to quantify “how much” envy occurs in a society. We have addressed this question by proposing a systematic *typology of envy measures*. Equipped with these degrees of envy, we will be in a position to appreciate how envy dynamically evolves within the society when negotiation takes place. Preliminary experiments suggest that most measures do decrease *on average*, albeit not strictly monotonically. More generally, these measures allow to go beyond a mere classification of a society as being either envy-free or not. In particular, we may observe that while some of the conditions stated in our convergence theorems may be necessary to guarantee envy-freeness in all cases, their absence would usually not be too damaging in practice. Finally, we have shown that the local decision problem faced by agents when they have to agree on a deal is not harder (in computational terms) if we require in addition to being rational that this deal also reduces envy (Theorem 10).

## References

- [1] A. Alkan, G. Demange, and D. Gale. Fair allocation of indivisible goods and criteria of justice. *Econometrica*, 59(4):1023–1039, 1991.
- [2] C. Beviá. Fair allocation in a general model with indivisible goods. *Review of Economic Design*, 3(3):195–213, 1998.
- [3] S. Bouveret and J. Lang. Efficiency and envy-freeness in fair division of indivisible goods: Logical representation and complexity. *Journal of Artificial Intelligence Research*, 32:525–564, 2008.
- [4] S. J. Brams, M. A. Jones, and C. Klamler. Divide-and-conquer: A proportional, minimal-envy cake-cutting procedure. Manuscript, February 2007.
- [5] S. J. Brams and A. D. Taylor. *Fair Division: From Cake-cutting to Dispute Resolution*. Cambridge University Press, 1996.
- [6] I. Caragiannis, C. Kaklamanis, P. Kanellopoulos, and M. Kyropoulou. On low-envy truthful allocations. In *Proc. 1st International Conference on Algorithmic Decision Theory (ADT-2009)*, volume 5783 of *LNAI*, pages 111–119. Springer-Verlag, 2009.
- [7] Y. Chevaleyre, P. E. Dunne, U. Endriss, J. Lang, M. Lemaître, N. Maudet, J. Padget, S. Phelps, J. A. Rodríguez-Aguilar, and P. Sousa. Issues in multiagent resource allocation. *Informatica*, 30:3–31, 2006.
- [8] Y. Chevaleyre, U. Endriss, and N. Maudet. Simple negotiation schemes for agents with simple preferences: Sufficiency, necessity and maximality. *Journal of Autonomous Agents and Multiagent Systems*, 2009. In press.
- [9] P. Cramton, Y. Shoham, and R. Steinberg, editors. *Combinatorial Auctions*. MIT Press, 2006.
- [10] B. de Keijzer, S. Bouveret, T. Klos, and Y. Zhang. On the complexity of efficiency and envy-freeness in fair division of indivisible goods with additive preferences. In *Proc. 1st International Conference on Algorithmic Decision Theory (ADT-2009)*, volume 5783 of *LNAI*, pages 98–110. Springer-Verlag, 2009.
- [11] M. de Weerd, Y. Zhang, and T. Klos. Distributed task allocation in social networks. In *Proc. 6th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2007)*. IFAAMAS, 2007.
- [12] P. E. Dunne, M. Wooldridge, and M. Laurence. The complexity of contract negotiation. *Artificial Intelligence*, 164(1–2):23–46, 2005.
- [13] U. Endriss and N. Maudet. On the communication complexity of multilateral trading: Extended report. *Journal of Autonomous Agents and Multiagent Systems*, 11(1):91–107, 2005.
- [14] U. Endriss, N. Maudet, F. Sadri, and F. Toni. Negotiating socially optimal allocations of resources. *Journal of Artificial Intelligence Research*, 25:315–348, 2006.



- [15] S. Estivie, Y. Chevaleyre, U. Endriss, and N. Maudet. How equitable is rational negotiation? In *Proc. 5th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2006)*, pages 866–873. ACM Press, 2006.
- [16] A. Feldman and A. Kirman. Fairness and envy. *The American Economic Review*, 64(6):995–1005, 1974.
- [17] D. Gale. *The Theory of Linear Economic Models*. University of Chicago Press, 1960.
- [18] 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.
- [19] D. Herreiner and C. Puppe. A simple procedure for finding equitable allocations of indivisible goods. *Social Choice and Welfare*, 19:415–430, 2002.
- [20] F. Klijn. An algorithm for envy-free allocations in an economy with indivisible objects and money. *Social Choice and Welfare*, 17:201–215, 2000.
- [21] B. Knaster. Sur le problème du partage pragmatique de H. Steinhaus. *Annales de la Société Polonaise de Mathématique*, 19:228–230, 1946.
- [22] H. W. Kuhn. The Hungarian Method for the assignment problem. *Naval Research Logistic Quarterly*, 2:83–97, 1955.
- [23] R. Lipton, E. Markakis, E. Mossel, and A. Saberi. On approximately fair allocations of indivisible goods. In *Proc. 5th ACM Conference on Electronic Commerce (EC-2004)*, pages 125–131. ACM Press, 2004.
- [24] H. Moulin. *Axioms of Cooperative Decision Making*. Cambridge University Press, 1988.
- [25] N. Nisan. Bidding languages for combinatorial auctions. In P. Cramton et al., editors, *Combinatorial Auctions*. MIT Press, 2006.
- [26] M. G. Raith. Fair-negotiation procedures. *Mathematical Social Sciences*, 39(3):303–322, 2000.
- [27] J. Robertson and W. Webb. *Cake-Cutting Algorithms: Be fair if you can*. A.K. Peters, 1998.
- [28] M. H. Rothkopf, A. Pekeč, and R. M. Harstad. Computationally manageable combinatorial auctions. *Management Science*, 44(8):1131–1147, 1998.
- [29] T. W. Sandholm. Contract types for satisficing task allocation: I Theoretical results. In *Proc. AAAI Spring Symposium: Satisficing Models*, 1998.
- [30] H. Steinhaus. The problem of fair division. *Econometrica*, 16(1):101–104, 1948.
- [31] K. Tadenuma. Trade-off between equity and efficiency in a general economy with indivisible goods. *Social Choice and Welfare*, 13:445–450, 1996.
- [32] W. Thomson. Fair allocation rules. RCER Working Paper 539, University of Rochester, Center for Economic Research, 2007.