Distributed Fair Allocation of Indivisible Goods☆

Yann Chevaleyre[a], Ulle Endriss[b], Nicolas Maudet[c]

[a]*LIPN, Université Paris-Nord, France*
[b]*ILLC, University of Amsterdam, The Netherlands*
[c]*LIP6, Université Pierre et Marie Curie, France*

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 will 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 representing a social network that constrains which agents can interact with which other agents, and we 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 overall envy.

Keywords: multiagent systems, multiagent resource allocation, fair division, negotiation, social networks

1. Introduction

The problem of fairly dividing a number of goods between several agents has been studied in a variety of 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 [3, 4, 5]. As for indivisible goods, we can distinguish *assignment problems* [6], 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 [7, 8, 9]. 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 of money (such as *quasi-linearity*, for instance). Finally, there are many different ways in which to interpret the term *fairness* itself [10]. 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. We focus on two fairness criteria: first, we look for solutions that are *proportionally fair* (i.e., where the utility of each of the n agents is at least $1/n$th of the value she ascribes to the full set of goods) and, second, we look for *envy-free* solutions (where no agent envies any of the other agents). (These criteria will get formally defined in Sections 3 and 4.)

While fair division is a problem originating in Economics and Political Science, with a number of important contributions by mathematicians, it eventually also started to attract the attention of researchers in Artificial Intelligence, Multiagent Systems, and Theoretical Computer Science [11, 12, 13, 14, 15, 16, 9, 5]. The reason for this trend is twofold: first, concepts from fair division are immediately relevant to these disciplines (e.g., finding acceptable agreements in a multiagent system) and, second, 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 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 [3], 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 [17], a procedure for finding equitable allocations of indivisible goods.

The kind of distributed approach followed here has been studied by several authors [18, 13, 19, 20, 21]. 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 [22], which in principle could be employed to find an optimal allocation of indivisible goods, may be considered problematic. The same argument has also been made by Netzer et al. [23], who proposed a framework for finding fair allocations of goods based on distributed constraint optimisation. But even when we give up central control of the optimisation algorithm, as system designers, we still would want to try to set up the system in such way that it guarantees certain desirable properties, although now 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 (which we might think of as representing a social network) 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 measures 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 economically *efficient* allocation, maximising the sum of individual utilities. Section 3 studies the case of proportionally fair outcomes. We show how to adapt the classical Knaster procedure to

2

our setting and obtain a distributed procedure that guarantees convergence to a proportionally fair 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 in the case of envy-freeness arrive at a particularly satisfying definition as well as an attractive convergence result. We also argue that proportionality and efficiency are not naturally generalised to this setting. Section 6 proposes a framework for assessing the degree of envy in a society. It also establishes 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 holds for all the concrete measures of envy we define, and indeed for a much larger class of such measures, defined by a simple axiom. 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 [13, 18].

2.1. Agents, goods, allocations, valuations

Let $\mathcal{N} = \{1, \ldots, n\}$ be a finite set of *agents* and let $\mathcal{G}$, with $m = |\mathcal{G}|$, be a finite set of indivisible *goods*. An *allocation* $A : \mathcal{N} \to 2^{\mathcal{G}}$ is a partitioning of the items in $\mathcal{G}$ amongst the agents in $\mathcal{N}$, i.e., each good is owned by exactly one agent (here $2^{\mathcal{G}}$ denotes the powerset of $\mathcal{G}$). 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, denoted by A_0, has been chosen at random, 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* $v_i : 2^{\mathcal{G}} \to \mathbb{R}$, mapping bundles of goods to the reals. Valuation functions are private information. However, we assume that when asked for the valuation of a given bundle of goods S, any agent i will truthfully reveal the value $v_i(S)$ she assigns to S. To simplify presentation, throughout this paper, we will make the assumption that all valuations v_i are *normalised*, in the sense that $v_i(\emptyset) = 0$, and *non-negative*, meaning that $v_i(S) \geqslant 0$ 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 $v(S_1 \cup S_2) \geqslant v(S_1) + v(S_2) - v(S_1 \cap S_2)$. Finally, a valuation function is called *single-minded* if there exist a bundle $S^* \subseteq \mathcal{G}$ and a positive number c such that $v(S) = c$ in case $S^* \subseteq S$ and $v(S) = 0$ otherwise. Thus, in a single-minded agent is only interested in a single bundle of goods and assigns a fixed value to that bundle (and to all larger bundles including that specific bundle as a subset).

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 after the deal. 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 $\delta = (A, A')$ is denoted as $\mathcal{N}^{\delta} = \{i \in \mathcal{N} \mid A(i) \neq A'(i)\}$. (In Section 5 we will 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*: $p : \mathcal{N} \to \mathbb{R}$, 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 she *receives* money. We associate each allocation A that is reached in a sequence of deals with

	$\{g_1\}$	$\{g_2\}$	$\{g_3\}$	$\{g_4\}$	$\{g_1, g_2, g_3, g_4\}$
Agent 1	5	5	0	0	10
Agent 2	4	4	4	4	18
Agent 3	0	0	5	5	10

Table 1: Valuation functions for an "almost modular" 3-agent/4-good allocation problem

a function $\pi : \mathcal{N} \to \mathbb{R}$ 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, π) of an allocation A and a *payment balance* π. 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* $u_i : 2^{\mathcal{G}} \times \mathbb{R} \to \mathbb{R}$ 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, π), while $u_i(A(j), \pi(j))$ is the utility that i would experience if she 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 measure 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, 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*.

Example 1. *Let us consider an allocation problem with three agents,* $\mathcal{N} = \{1, 2, 3\}$, *and four goods,* $\mathcal{G} = \{g_1, g_2, g_3, g_4\}$. *Table 1 defines the valuation functions* v_i *by indicating for each agent how much she values the four individual goods and how much she values the full bundle. The values of bundles not shown are assumed to always be equal to the sum of the values of the items in them. Thus,* v_1 *and* v_3 *are modular, and all three valuation functions are supermodular.*

Maximal social welfare for this allocation problem is 20, *namely when agent 1 holds the bundle* $\{g_1, g_2\}$, *agent 2 gets nothing, and agent 3 holds* $\{g_3, g_4\}$. *Let us denote this optimal allocation by* A^*. *Suppose now we are in allocation* A, *in which all four items are held by agent 2. Social welfare for* A *is* 18. *If our agents are only willing to negotiate 1-deals, then we will remain stuck in this suboptimal allocation and never reach* A^*. *To see this, check that agent 2 giving any single item to one of the other two agents would result in a drop of her utility from* 18 *to* 12, *thus requiring a price strictly above* 6, *which neither one of the other two agents would be willing to pay. We will return to this example throughout this paper.*

A central result in distributed negotiation, due to Sandholm [18], shows that the problem we observed in Example 1 does not persist if we do not impose any structural constraints on deals:

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

4

That is, if agents continue to agree on IR deals, they will eventually reach a state from which no further IR deal is possible, and that state will be efficient. 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 [13]:

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 Endriss et al. [13] for a proof):

Lemma 3 (IR deals and social welfare). *A deal $\delta = (A, A')$ is IR if and only if $sw(A) < sw(A')$.*

We will frequently rely on this fact in the sequel, mostly without explicit reference to Lemma 3. Note that Lemma 3 also entails that we can always reach an efficient allocation by means of just a single IR deal. Importantly, this does not diminish the importance of Theorem 1. Finding that particular IR deal to reach the optimum in a single step requires full coordination between all agents, while Theorem 1 guarantees convergence also in the absence of such global coordination and only requires agents to make deals locally that benefit themselves then and there.

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. [25] 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 3, the social surplus is positive if and only if δ is IR. The LUPF divides this amount equally amongst the *participating* agents $\mathcal{N}^\delta$; the GUPF divides it equally amongst *all* agents $\mathcal{N}$:

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

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

Importantly, note that the social surplus $sw(A') - sw(A)$ can always be computed from the valuation functions of the participating agents alone:

$$sw(A') - sw(A) \;=\; \sum_{i \in \mathcal{N}^\delta} [v_i(A') - v_i(A)]$$

Thus, the LUPF is a fully local payment function: it only requires us to consider input from the participating agents and only the payment balances of these agents will be affected. The GUPF is also local as far as the information required is concerned, except that we need to know the quantity n before we can compute it. On the other hand, the GUPF is global in the sense that payment balances of all agents are affected. Having said this, agents not participating in the deal can only ever receive money, but never have to make a payment themselves.

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, π) is called proportional if $u_i(A, \pi) \geqslant v_i(\mathcal{G})/n$ for all agents $i \in \mathcal{N}$.*

Proportionality is the central fairness criterion applied in the cake-cutting literature, which addresses the problem of fairly dividing a single divisible good [3, 4, 5]. Before we proceed, let us briefly pause and consider under what circumstances proportionality is a reasonable fairness concept also in our context. Importantly, in almost all of the cake-cutting literature, agents are assumed to have *additive* preferences, which is an assumption we do not always want to make. As a minimum requirement for Definition 3 to be meaningful, we need to make the assumption that valuation functions are *monotonic*: $S_1 \subseteq S_2$ implies $v_i(S_1) \leqslant v_i(S_2)$. 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 $1/n$th of the value of the full set of goods would often be considered more than the value of $1/n$th 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 $1/n$th 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 not withstanding, 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. We will be able to obtain a positive result for arbitrary domains, but the reader should keep in mind that this result is at its most attractive when applied to supermodular rather than to submodular domains.

To define a distributed procedure that can guarantee proportional outcomes we will take inspiration from the classical fair division procedure proposed by Bronisław Knaster in the 1940s [26, 3, 27].[1] 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 [27]. 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 her *excess* $ex_i(A^*) = v_i(A^*) - v_i(\mathcal{G})/n$ over her fair share. The *total excess* is defined as $Ex(A^*) = \sum_i ex_i(A^*)$.

(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.

Note that the individual excess of an agent could be a negative number, as could be the payment she has to make (in which case she would receive money). Clearly, the payments of all agents sum up to zero, so are feasible. While an agent's individual excess ex_i could be negative, the total excess $Ex(A^*)$ 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 $sw(A^*) \geqslant v_i(\mathcal{G})$ 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.

---

[1]In an early paper, Knaster [28] 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 [26] later discussed a procedure for indivisible goods using this concept and attributed the idea for this procedure to Knaster.

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 π_0 and p such that at each stage (A, π) 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 π_0 as follows:

$$\pi_0(i) = ex_i(A_0) - Ex(A_0)/n \tag{1}$$

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, π). This forces the following payments for a deal $\delta = (A, A')$, with π being the payment balance for allocation A and π' 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 4 (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, π) 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^*, π^*) 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). $\square$

Thus, convergence to a proportional outcome is guaranteed. 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 5 (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 4. $\square$

$\pi(1)$	$\pi(2)$	$\pi(3)$	allocation	sw	$p(1)$	$p(2)$	$p(3)$	u_1	u_2	u_3
n/a	n/a	n/a	$\langle\{g_4\},\{g_2,g_3\},\{g_1\}\rangle$	$0+8+0$	$-16/9$	$32/9$	$-16/9$	$16/9$	$40/9$	$16/9$
$-16/9$	$32/9$	$-16/9$	$\langle\{g_2\},\{g_3,g_4\},\{g_1\}\rangle$	$5+8+0$	$30/9$	$-15/9$	$-15/9$	$\mathbf{31/9}$	$\mathbf{55/9}$	$\mathbf{31/9}$
$14/9$	$17/9$	$-31/9$	$\langle\{g_2\},\emptyset,\{g_1,g_3,g_4\}\rangle$	$5+0+10$	$-6/9$	$-78/9$	$84/9$	$\mathbf{37/9}$	$\mathbf{61/9}$	$\mathbf{37/9}$
$8/9$	$-61/9$	$53/9$	$\langle\{g_1,g_2\},\emptyset,\{g_3,g_4\}\rangle$	$10+0+10$	$30/9$	$-15/9$	$-15/9$	$\mathbf{52/9}$	$\mathbf{76/9}$	$\mathbf{52/9}$

Table 2: Sequence of states using the Knaster payment scheme

Example 2. *Recall Example 1, with three agents with valuation functions as specified in Table 1. The only efficient allocation here is A^*, which assigns $\{g_1,g_2\}$ to agent 1 and $\{g_3,g_4\}$ to agent 3 and which has a social welfare of 20. We may write this more compactly as $A^* = \langle\{g_1,g_2\},\emptyset,\{g_3,g_4\}\rangle$. Suppose the initial allocation is $A_0 = \langle\{g_4\},\{g_2,g_3\},\{g_1\}\rangle$, giving utility 0, 8, and 0 to agents 1, 2, and 3, respectively.*

Let us illustrate convergence to a proportional allocation with Knaster payments. First, we compute the excess $ex_i(A_0)$ for each agent i:

- $ex_1(A_0) = 0 - 10/3 = -10/3$

- $ex_2(A_0) = 8 - 18/3 = 6/3$

- $ex_3(A_0) = 0 - 10/3 = -10/3$

The total excess $Ex(A_0)$ thus is $6/3 - 10/3 - 10/3 = -14/3$. The initial payments therefore must be:

- $\pi_0(1) = -10/3 + 14/9 = -16/9$

- $\pi_0(2) = 6/3 + 14/9 = 32/9$

- $\pi_0(3) = -10/3 + 14/9 = -16/9$

It can readily be checked that, indeed, $\sum_i \pi_0(i) = 0$. At this stage, the state (A_0, π_0) is not proportional: for instance, $u_1(A_0(1), \pi_0(1)) = 0 + 16/9 < 10/3$. Now suppose that, starting from this state, our agents negotiate the sequence of states shown in Table 2. Each row corresponds to a state. The first three columns represent the payment balances before the state in question. Then the allocation and the social welfare are shown, followed by the new payments to be made in that state, in line with the Knaster payment scheme. Finally, the utility of each agent is shown. These utilities are indicated in boldface in case an agent has reached her "proportionality threshold" (which is $10/3$ for agents 1 and 3, and $18/3$ for agent 2). Observe that, for this example, all three agents surpass this threshold already after the second deal, i.e., efficiency is not always required to achieve proportionality.

We continue 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 [27], 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 $\alpha_i \geqslant 0$ 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 4 and 5 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 [27] we have $\alpha_i = \frac{v_i(\mathcal{G})}{\sum_j v_j(\mathcal{G})}$. The two payment schemes coincide in case we have $v_i(\mathcal{G}) = 1$ for all $i \in \mathcal{N}$.

8

3.4. Path length

In the previous section we have seen that convergence to an efficient and proportional state can always be guaranteed by our scheme. But how many deals do we require to reach a social optimum? First, observe that by Lemma 3, it is always possible to reach the optimum using (at most) one deal, i.e., 1 is a (tight) upper bound on the shortest path. Much more interesting in practice, however, is the longest path, i.e., the longest sequence of IR deals that agents may end up agreeing on before convergence. For convergence to an efficient allocation by means of IR deals, $n^m - 1$, with n being the number of agents and m being the number of goods, is known to be a (tight) upper bound on the longest path, i.e., there exist profiles of valuation functions and initial allocations such that it is possible to visit *all* allocations before convergence [29]. Inspection of the proof of that result shows that it does not depend on the payment scheme used. Hence, $n^m - 1$ is also an upper bound on the number of deals for scenarios covered by Theorem 4—at least if we do not put restrictions on the nature of our agents' valuation functions.

If we do require valuation functions to be normalised and monotonic (and possibly also supermodular, which we have argued to be reasonable in the context of studies of proportionality), then the construction used in previous work [29] does not apply any longer. We now want to show that the result itself nevertheless carries over.

Lemma 6 (Distinct welfare). *We can find a valuation function for each agent that is normalised, non-negative, and modular such that any two allocations have distinct utilitarian social welfare.*

Proof. We first define the valuation function v_1 of the first agent in a way that ensures that she assigns a distinct value to every possible bundle. Let $\mathcal{G} = \{g_1, \ldots, g_m\}$ and define:

$$v_1(S) \quad = \quad \sum_{g \in S} \nu(g) \quad \text{with } \nu(g_k) = 2^{k-1}$$

That is, v_1 maps bundles to the numbers from 0 to $2^m - 1$ (the reader may find it helpful to think of bundles as representing binary numbers with m digits). Note that v_1 is normalised, non-negative, and modular. Now define the valuation functions of the remaining agents as follows:

$$v_i(S) \quad = \quad v_1(S) \cdot (2^m)^{i-1}$$

That is, each agent's valuation function is operating "at its own scale". Now, if we are given the utilitarian social welfare of an allocation A as a binary number, then we know that that number will have a 1 in position $m \cdot (i-1) + k$ (counting from the right) if and only if agent i has obtained good g_k under A. Hence, there is a bijection between allocations and social welfare values, i.e., any two distinct allocations must have distinct social welfare, as claimed. $\square$

Observe that any valuation function that is normalised, non-negative, and modular is also monotonic and supermodular. Hence, there exist normalised, monotonic and supermodular valuation functions such that no two allocations have the same utilitarian social welfare. This means, that there exists a sequence of IR deals that visits all n^m allocations (each of the m goods may be given to any of the n agents), and thus $n^m - 1$ is a tight upper bound on the length of the longest path for scenarios covered by Theorem 4, also when valuations are subject to the restrictions we have argued to be natural for studies of proportional fairness.

This bound applies in case we are interested in finding an allocation that is both proportional and efficient. Can we do better if we only require proportionality? This might seem a possibility, given the fact that, if agents negotiate IR deals using the Knaster payment scheme, then efficiency is only achieved at the very last deal, but proportionality will often get achieved already much earlier, namely as soon as the total excess $Ex(A)$ of the current allocation A is non-negative. Yet, the answer to our question is still negative:

Proposition 7 (Path length to a proportional state). *We can find a valuation function for each agent and an initial allocation such that there exists a sequence of IR deals with Knaster payments that will achieve proportionality only after exactly $n^m - 1$ deals.*

Proof. Take a scenario where the agents have the valuation functions used in the proof of Lemma 6, except that $v_n(\mathcal{G})$ is set to such a high value that the only way of achieving a proportional state is to give all goods to agent n (because no payments the other agents would be prepared to make could be large enough to give utility $v_n(\mathcal{G})/n$ to agent n). Then all allocations have distinct social welfare, i.e., we can construct a sequence of deals visiting all of them, and proportionality would only get achieved at the very end of the sequence. □

Finally, given Lemma 6, known results for convergence by means of 1-deals in case all valuation functions are additive (i.e., the kind of scenario covered by Theorem 5) now also carry over immediately [29]. Concretely, this means that m is a tight upper bound for the shortest path (as we need at most one deal per item if we only choose optimal deals) and $m \cdot (n-1)$ is a tight upper bound for the longest path (as each item may, in the worst case, pass through every single agent).

4. Envy-freeness and globally uniform payments

Besides proportionality, another important fairness criterion is *envy-freeness* [30, 3, 11, 10]. 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 [14]. 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 [7].

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

Definition 4 (Envy-freeness)**.** *A negotiation state (A, π) is called envy-free if $u_i(A(i), \pi(i)) \geqslant u_i(A(j), \pi(j))$ 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. [6]. When agents can obtain (and express preferences over) bundles consisting of several goods, there is in general no such guarantee, i.e., as demonstrated by Tadenuma [8], for some allocation problems there simply exists no state that is both efficient and envy-free. However, the assumption of quasi-linearity of money is very important here: if we require quasi-linearity, as we do in this paper, then the existence of EEF states can be proven by an argument similar to the one used by Alkan et al. [6], appealing to the Duality Theorem [31], as demonstrated by Beviá [7]. When all valuation functions are *supermodular*, we can give a much simpler argument for the existence of EEF states, which does not rely on the Duality Theorem. As we will see later on, 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 π^* can be arranged such that the state (A^*, π^*) is EEF. Define $\pi^*(i)$ for each agent i:

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

10

First, note that π^* 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^*, π^*). As A^* is efficient, giving both $A^*(i)$ and $A^*(j)$ to i will not increase social welfare any further:

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

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

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

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

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 [32]. 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. Convergence in supermodular domains

We are now going to prove a result on the reachability of EEF states by means of distributed negotiation. Our result applies to supermodular domains only. The payment scheme it employs consists of the GUPF in combination with the following initial payments:

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

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 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.

Theorem 8 (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, π) and every agent i, provided that agents only negotiate IR deals using the GUPF:

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

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, π_0). Now let $\delta = (A, A')$ and assume (2) 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:

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

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 π^* be the associated payment balance. Equation (2) also applies to the state (A^*, π^*). But as we have already seen in Section 4.1, the efficiency of A^* together with Equation (2) applied to (A^*, π^*) implies that (A^*, π^*) must be an EEF state. $\square$

$\pi(1)$	$\pi(2)$	$\pi(3)$	allocation	sw	$p(1)$	$p(2)$	$p(3)$	u_1	u_2	u_3
n/a	n/a	n/a	$\langle\{g_4\},\{g_2,g_3\},\{g_1\}\rangle$	$0+8+0$	$-8/3$	$16/3$	$-8/3$	$8/3$	$8/3$	$8/3$
$-8/3$	$16/3$	$-8/3$	$\langle\{g_1,g_2\},\emptyset,\{g_3,g_4\}\rangle$	$10+0+10$	$18/3$	$-36/3$	$18/3$	$20/3$	$20/3$	$20/3$

Table 3: Convergence to an EEF state using equitability payments and the GUPF

Example 3. *Observe that the final state attained in Example 2 is not envy-free. This is so, because, for instance, $u_1(A^*(1),\pi^K(1)) = 52/9$, while $u_1(A^*(2),\pi^K(2)) = 76/9$: agent 1 envies agent 2, even though the latter does not own any goods, because of the amount of money she receives through the Knaster payments (here, π^K denotes the payment balances according to the Knaster scheme for the final state, corresponding to the efficient allocation $A^* = \langle\{g_1,g_2\},\emptyset,\{g_3,g_4\}\rangle$). To illustrate the difference between the Knaster payment scheme and the payment scheme of Theorem 8, let us consider the simple example shown in Table 3, involving the initial allocation $A_0 = \langle\{g_4\},\{g_2,g_3\},\{g_1\}\rangle$ and a single IR deal leading to the efficient allocation A^*. Note that the initial state is not envy-free, not even after the initial equitability payments, as for instance agent 2 envies agent 1, given that $u_2(A_0(1),\pi_0(1)) = 4 + 8/3 = 20/3 > 8/3 = 8 - 16/3 = u_2(A_0(2),\pi_0(2))$.*

As argued in Section 3.4, Lemma 6 shows that the bounds on the lengths of paths to convergence established in previous work on negotiating efficient allocation for arbitrary valuation functions [29] apply also to the types of scenarios covered by Theorem 8: 1 is a (tight) upper bound on the shortest path and $n^m - 1$ is a (tight) upper bound on the longest path.

Theorem 8 is a surprising result. As pointed out elsewhere [13], 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 8 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 8 in place.

Theorem 9 (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 8, which is characterised by Inequality (2), the envy-freeness condition $u_i(A(i),\pi(i)) \geqslant u_i(A(j),\pi(j))$ is equivalent to the following inequality (which must hold for any two agents i and j):

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

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

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

12

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

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

But now, in a scenario with two agents with these two valuation functions, 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 $v_1(\{g_1\}) + v_1(\{g_2\}) - d/3$. 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 her own share at least as much as anyone else would. $\qquad\square$

Theorem 9 is reminiscent of so-called *maximality theorems* in multiagent resource allocation [20], which are results that show that certain classes of valuation functions that can guarantee convergence under a certain negotiation protocol are maximal in the sense of no strictly larger class being able to still provide the same guarantee. Observe that the proof of Theorem 9 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 9 offers an interesting characterisation of a domain of valuation functions guaranteeing convergence to states without envy.

Vice versa, as we will see next, if we fix the supermodularity condition, then the payment scheme of Theorem 8 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 *path-independent* if there exists a function $f : \mathcal{N} \times \mathbb{R}^n \to \mathbb{R}$ such that for any deal $\delta = (A, A')$ the payment $p(i)$ of any agent i is equal to $f(i, \langle \Delta_1, \ldots, \Delta_n \rangle)$, where $\Delta_j = v_j(A') - v_j(A)$. That is, a payment function is path-independent 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 path-independent if the associated payment function is path-independent.

Theorem 10 (Payment schemes). *No path-independent payment scheme Π 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 Π 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 8. 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 π^*. We need to show that, amongst the *path-independent* payment schemes, the scheme of Theorem 8 is the only one with this property.

Path-independent schemes do not allow us to make payments dependent on where in the sequence of deals we currently are. In particular, knowing only the differences in individual valuations between two allocations A and A' is not sufficient information to determine whether A' is efficient, i.e., whether the process will terminate once A' is reached. Therefore, $\pi(i) = v_i(A) - sw(A)/n$ must hold after *every* deal. But this *forces* initial payments to be exactly as in Theorem 8 (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. $\qquad\square$

In summary, Theorems 9 and 10 together show that our convergence result, Theorem 8, 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.

In related work, there have been some proposals for procedures for finding EEF states in the social choice literature [6, 33, 34], albeit with little or no attention to computational issues. The work of Haake et al. [34] 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, 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. [34] 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 8, is not an NP-hard problem [35].

4.3. Convergence in modular domains

In modular domains, we can strengthen Theorem 8 and even guarantee convergence to an EEF state by means of 1-deals (over one item at a time):

Theorem 11 (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 8, 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. □

5. Fair division on social networks

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 becomes unrealistic. The pervasiveness of applications exhibiting underlying graph-like structures, like for instance small-world 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 will assume that not every agent is able to "see" all of the other agents and formulate an appropriate generalisation of our model.[2] A *negotiation topology* (or *social network*) 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. The visibility relation E is symmetric (the graph is undirected); this is important to be able to define negotiation along graphs in a meaningful manner.

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 G. (Recall that 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*.

Definition 5 (Clique-deals). *A deal $\delta = (A, A')$ is called a clique-deal if the set of involved agents $\mathcal{N}^\delta$ is a clique of the negotiation topology G.*

---

[2]De Weerdt et al. [36] explore a similar model in the context of task allocation problems.

We will be particularly interested in deals that are both clique-deals and IR. Note that we will soon 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 will not consider this possibility here.)

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

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 natural 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 an agent 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 6 (GEF states). *A state (A, π) is called graph-envy-free (GEF) with respect to the graph $G = (\mathcal{N}, E)$ if $u_i(A(i), \pi(i)) \geqslant u_i(A(j), \pi(j))$ for all agents $(i, j) \in E$.*

As we have seen in Section 4, our guarantee for 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 7 (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 $i \notin C$.*
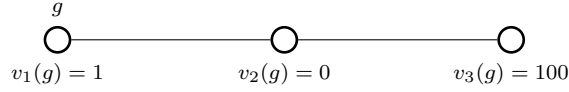
$$g$$

$$v_1(g) = 1 \qquad v_2(g) = 0 \qquad v_3(g) = 100$$

Figure 1: Low global efficiency despite high clique-wise efficiency

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

Definition 8 (Clique-wise efficiency)**.** *An allocation A is called clique-wise efficient if $sw(A) \geqslant sw(A')$ 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. In particular, as the following example demonstrates, an allocation that is clique-wise efficient can be arbitrarily bad in terms of global efficiency. Consider Figure 1, showing a negotiation topology between three agents, together with the valuation each of these three agents assigns to the good g. Now suppose that in allocation A, item g is assigned to agent 1, i.e., social welfare is 1. The only clique-variant of A is the allocation where g is given to agent 2, which would reduce social welfare to 0. Thus, A is clique-wise efficient. On the other hand, allocation A', giving g to agent 3 instead, would increase social welfare by a factor of 100.

The reader might doubt the significance of this example and point out that the chosen valuation functions are extremely different from each other, or that a line topology might be a particularly problematic case to handle. But in fact, this kind of situation can occur also under much more favourable circumstances, as the following result shows:

Proposition 12 (Global inefficiency of clique-wise efficiency)**.** *If the negotiation topology is not complete, then the social welfare of a clique-wise efficient allocation can be arbitrarily far from the optimum. This holds even when all agents share the same single-minded valuation function.*

Proof. Let $G = (\mathcal{N}, E)$ be a graph that is missing at least one edge. W.l.o.g., assume that agent 1 is not connected to agent 2. Let the set of goods be $\{g_1, g_2\}$. Assume all agents share the same single-minded valuation function v with $v(\{g_1, g_2\}) = 1$ and $v(S) = 0$ for all bundles $S \neq \{g_1, g_2\}$. Consider the allocation A in which agent 1 owns g_1 and agent 2 owns g_2. In allocation A, if no payments have been made in the past, the utility of every agent is equal to 0, and thus so is social welfare. Nevertheless, A is clique-wise efficient, because any clique of G must be a subset of either $\{1, 3, 4, \ldots, n\}$ or $\{2, 3, 4, \ldots, n\}$, i.e., no clique collectively owns both g_1 and g_2. On the other hand, any allocation where one agent owns both items has a social welfare of 1. Thus, the ratio between the social welfare of efficient allocations and the social welfare of clique-wise efficient allocations cannot be bounded by a constant. $\square$

This result may call into question the usefulness of Definition 8. While our definition of envy with respect to a graph is very natural and reaching GEF states seems indeed desirable, it is questionable whether it is at all possible to relativise the standard notion of efficiency with respect to a negotiation topology in a meaningful manner. Our interest in the purely technical definition of clique-wise efficiency given above 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 13 (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 3, $\delta = (A, A')$ must be IR (besides being a clique-deal). This contradicts our assumption of A being a terminal allocation. $\square$

$\pi(1)$	$\pi(2)$	$\pi(3)$	allocation	sw	$p(1)$	$p(2)$	$p(3)$	u_1	u_2	u_3
n/a	n/a	n/a	$\langle\{g_2,g_3\},\{g_1,g_4\},\emptyset\rangle$	$5+8+0$	$2/3$	$11/3$	$-13/3$	$13/3$	$13/3$	$13/3$
$2/3$	$11/3$	$-13/3$	$\langle\{g_1,g_2\},\{g_3,g_4\},\emptyset\rangle$	$10+8+0$	$10/3$	$-5/3$	$-5/3$	$18/3$	$18/3$	$18/3$

Table 4: Convergence via deals between neighbouring agents 1 and 2

Lemma 13 generalises Theorem 1. The latter corresponds to the case of a fully connected graph. While Lemma 13 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. And even if it does, as Proposition 12 shows, global efficiency may still be arbitrarily bad. Nevertheless, as we will see next, clique-wise efficiency is a sufficiently strong notion to serve as a basis for negotiating envy-free states.

5.3. Convergence

We now prove a convergence theorem for GEF states, which extends Theorem 8 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 14 (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 8 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, π) reached during negotiation and every agent $i \in \mathcal{N}$. By Lemma 13, 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^*, π^*) 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^*, π^*). $\qquad\square$

Theorem 14 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.

Example 4. *We consider once more the allocation problem of Table 1, but now the agents are arranged on a line, with agent 1 in the middle position, i.e., $E = \{(1,2),(2,1),(1,3),(3,1)\}$. Suppose the initial allocation is $A_0 = \langle\{g_2,g_3\},\{g_1,g_4\},\emptyset\rangle$. Table 4 shows the allocations visited during a sequence consisting of a single IR clique-deal, which culminates in a clique-wise efficient state (A', π'), in which the final payment balances of the agents are as follows: $\pi'(1) = 12/3$, $\pi'(2) = 6/3$, and $\pi'(3) = -18/3$. The final utility of each agent is $18/3$. It can be checked that the following statements hold.*

- *Agent 2 does not envy agent 1, as $u_2(A'(1), \pi'(1)) = 8 - 12/3 = 12/3 < 18/3$*

- *Agent 3 does not envy agent 1, as $u_3(A'(1), \pi'(1)) = 0 - 12/3 = -12/3 < 18/3$*

- *Agent 1 does not envy agent 2 nor agent 3, as both:*

  - *$u_1(A'(2), \pi'(2)) = 0 + 6/3 = 6/3 < 18/3$*

$$- u_1(A'(3), \pi'(3)) = 0 + 18/3 = 18/3$$

Thus, no agent envies any of her neighbours. However, note that if agent 3 could see agent 2, envy would occur as $u_3(A'(2), \pi'(2)) = 10 - 6/3 = 24/3 > 18/3$.

A critical point in Theorem 14 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 again there are three agents on a line, 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 defined by the social network we belong to, 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 5 and 11, if all valuation functions are modular, then we can strengthen Theorem 14 and prove convergence by means of IR 1-deals between connected agents using our by now familiar technique.

Theorem 15 (GEF 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 between pairs of connected agents using the GUPF will eventually result in a GEF state.*

5.4. Path length

Just as we did for Theorems 4 and 8, also for Theorem 14 we may ask how many deals are required before the system converges. Now the bounds on the longest (as well as the shortest) path will depend on the negotiation topology. At one extreme, if the graph is fully connected then Theorem 14 reduces to Theorem 8 and we obtain the same bounds as before. In particular, $n^m - 1$ will always be an upper bound on the longest path. At the other extreme, if no two agents are connected to each other, then no deals are possible, and the upper bound reduces to 0. In general, it is clear that the sparser the graph, the shorter the longest path to convergence we can construct. The question now arises whether there is a class of graphs that is relatively sparse (but still natural and attractive for applications) that would result in a significant reduction of the length of the longest possible path. The answer to this question, unfortunately, is negative. As we will see next, even if agents are arranged on a line (a highly simplistic negotiation topology), there may still be a number of IR deals that is exponential in the number of goods.

Proposition 16 (Path length for line topology). *Suppose the negotiation topology $G = (\mathcal{N}, E)$ is a line: $E = \{(i, j) \in \mathcal{N}^2 \mid |i - j| = 1\}$. Then we can find a valuation function for each agent such that there exists a sequence of IR clique-deals of length $(2^m - 1) \cdot (n - 1)$.*

Proof. Suppose the agents have the valuation functions defined in the proof of Lemma 6 and initially allocate all goods to agent 1. Then agents 1 and 2 can go through all 2^m allocations that divide all goods between the two of them and end up in the allocation where agent 2 owns all items (that is, there are $2^m - 1$ IR deals between them). We can do the same for every one of the $n - 1$ pairs $(i, i+1)$ of connected agents and thus end up with a path of the claimed length of $(2^m - 1) \cdot (n - 1)$ deals. □

Note that our construction is independent of the payment scheme used. In particular, it applies in case the payment scheme of Theorem 14 is in operation.

Interestingly, the upper bound of the *shortest* path may increase as the graph connecting the agents becomes sparser, simply because this means that some transactions cannot be executed as a single deal anymore. For the scenario described in the proof of Proposition 16 the shortest path would be $n - 1$ (rather than 1, the bound for fully connected graphs): we can reach the social optimum by implementing one deal per pair of neighbouring 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 of 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 8 and 14 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 will propose a systematic approach to defining measures 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. We also formulate a number of fundamental axioms that any reasonable measure of envy, not only the concrete measures proposed here, should satisfy. 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 any such measure.

6.1. Envy between two individual agents

How much does agent i envy agent j in state (A, π)? Any measure of envy between i and j will be based on $\Delta_{i,j}(A, \pi) = u_i(A(j), \pi(j)) - u_i(A(i), \pi(i))$, the difference in utility that i assigns to their own lot and that of j. If this number is negative or equal to 0, then agent i does not envy agent j, nor does she in case she cannot see agent j in the social network. This gives rise to a matrix of dimension $n \times n$, with entries $e_{i,j}$:

$$e_{i,j} = \begin{cases} \Delta_{i,j}(A, \pi) & \text{if } \Delta_{i,j}(A, \pi) > 0 \text{ and } (i, j) \in E \\ 0 & \text{otherwise} \end{cases}$$

We now use $e_{i,j}$ to define two different measures for the degree of envy between individual agents:

$$e^{raw}(i, j) = e_{i,j} \qquad e^{bool}(i, j) = \begin{cases} 1 & \text{if } e_{i,j} > 0 \\ 0 & \text{otherwise} \end{cases}$$

The second measure (e^{bool}) only allows us to specify whether or not i envies j. In many cases this is all that is needed. The first measure is useful if we also want to say *how much* i envies j. In principle one could also consider the notion of "negative envy" [1], i.e., one could consider treating the case of $\Delta_{i,j}(A, \pi) < 0$ differently from the case where that difference is exactly 0 or where i cannot see j, but we will not pursue this route here.

6.2. 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 to all the other agents. That is, we need to aggregate the individual envies that i experiences *vis-à-vis* each of the other agents. The most natural aggregation operators are *max* and summation, but other options are possible as well.

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

Here *op* may stand for either *raw* or *bool*. For instance, $e^{max,raw}(i)$ measures how much i envies the agent she envies the most, while $e^{sum,bool}(i)$ 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.

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

Using the *max*-operator means focusing on the most envious agent of the society (whatever operator was chosen to measure that), while 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 measures for measuring envy. Several of these have been proposed in the literature before:

- Feldman and Kirman [37] 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 ($e^{sum,sum,bool}$). They also define two alternative measures for the degree of envy of a single agent, one of which corresponds to our $e^{sum,raw}(i)$ and the other of which takes negative envies into account (so is not covered by our system). They then propose to measure the degree of envy of a society by means of a weighted sum of either one of their measures for individual envy. Thus, their discussion includes, in particular, the measure $e^{sum,sum,raw}$.

- Lipton et al. [12] work with two envy measures. The first of these is the "maximum envy between any pair of players" ($e^{max,max,raw}$). Their second envy measure is the maximum *envy ratio*, defined as $\max_{i,j \in \mathcal{N}} \{\frac{u_i(A(j),\pi(j))}{u_i(A(i),\pi(i))}, 1\}$, which is not covered by our system. Caragiannis et al. [15] also work with $e^{max,max,raw}$ and Netzer et al. [23] use both $e^{max,max,raw}$ and $e^{sum,max,raw}$.

- Brams et al. [38] refer to two measures of envy: the "number of envies of all players" ($e^{sum,sum,bool}$) and the "maximum number of envies of an individual player" ($e^{max,sum,bool}$).

- Lindner and Rothe [39] also work with $e^{sum,sum,bool}$, at least implicitly—explicitly they count the number of ordered pairs of agents such that the first agent does *not* envy the second.

Thus, our simple system covers many of the measures of envy proposed in the literature. Another obvious choice for measuring envy is the number of envious agents in a society ($e^{sum,max,bool}$). Finally, $e^{max,max,bool}$ 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.

As noted above, one of the measures proposed by Lipton et al. [12] does not fit our system, as that measure is grounded in utility ratios rather than utility differences. Such ratio-based measures are useful

when establishing approximability results (as envy-freeness then corresponds to a value of 1 rather than 0) and have been used by several authors [12, 40, 23]. In related work, Nguyen and Rothe [40] have adapted an early version of our three-stage system [1] to define a family of ratio-based measures of envy. The basic idea of measuring degrees of envies has also been adapted to the context of matching [41].

6.4. Axioms for envy measures

Next, we want to formulate some minimal requirements on any reasonable measure of envy, which will allow us to state results regarding the reduction of envy in generic terms.

Let $envy : (2^{\mathcal{G}})^{\mathcal{N}} \times \mathbb{R}^{\mathcal{N}} \to \mathbb{R}$ be a measure of envy, i.e., a function mapping states to the non-negative reals, that satisfies $envy(A, \pi) \geqslant 0$ for every state (A, π). Every state (A, π) gives rise to a matrix $(e_{i,j})_{i,j \in \mathcal{N}}$ of local envies. Our most fundamental requirement on $envy(.,.)$ is that this function must be computable from the values recorded in the matrix alone. We call any measure of envy that satisfies this requirement a *proper* measure of envy. Thus, we may think of a proper measure of envy as a function $envy : \mathbb{R}^{n \times n} \to \mathbb{R}$ mapping envy matrices to the reals.[3]

Beyond the basic requirement of $envy(.,.)$ being proper, we propose four natural axioms that any reasonable measure of envy should satisfy:

- **Discernibility:** It must be the case that $envy((e_{ij})_{i,j \in \mathcal{N}}) = 0$ if and only if $e_{i,j} = 0$ for all agents $i, j \in \mathcal{N}$. Thus, the degree of envy of a state should be zero if and only if it is an envy-free state. This excludes situations where states that are "almost" envy-free are not discernible from the truly envy-free states by using the measure of envy in question.

- **Dominance:** If $(e_{i,j})_{i,j \in \mathcal{N}}$ and $(e'_{i,j})_{i,j \in \mathcal{N}}$ are envy matrices satisfying $e_{i,j} \geqslant e'_{i,j}$ for all agents $i, j \in \mathcal{N}$, then we must have $envy((e_{i,j})_{i,j \in \mathcal{N}}) \geqslant envy((e'_{i,j})_{i,j \in \mathcal{N}})$. Thus, if all local envies either increase or stay the same, then overall envy cannot go down.

- **Source Anonymity:** If $(e_{i,j})_{i,j \in \mathcal{N}}$ and $(e'_{i,j})_{i,j \in \mathcal{N}}$ are envy matrices such that there exists a permutation $\sigma : \mathcal{N} \to \mathcal{N}$ for which $e_{i,j} = e'_{\sigma(i),j}$ for all agents $i, j \in \mathcal{N}$, then we must have $envy((e_{i,j})_{i,j \in \mathcal{N}}) = envy((e'_{i,j})_{i,j \in \mathcal{N}})$. In other words, $envy(.,.)$ is indifferent under permutations of the rows of the matrix of local envies. Thus, every individual agent's envy is equally important.

- **Target Anonymity:** If $(e_{i,j})_{i,j \in \mathcal{N}}$ and $(e'_{i,j})_{i,j \in \mathcal{N}}$ are envy matrices such that there exist an agent $i^* \in \mathcal{N}$ and a permutation $\sigma : \mathcal{N} \to \mathcal{N}$ for which $e_{i^*,j} = e'_{i^*,\sigma(j)}$ for all $j \in \mathcal{N}$ and $e_{i,j} = e'_{i,j}$ for all $i, j \in \mathcal{N}$ with $i \neq i^*$, then we must have $envy((e_{i,j})_{i,j \in \mathcal{N}}) = envy((e'_{i,j})_{i,j \in \mathcal{N}})$. In other words, $envy(.,.)$ is indifferent under permutations of the elements of a single row of the envy matrix. Thus, we do not care about *who* an individual agent envies, but only about how many agents she envies and to what degree she does so.

A weaker form of target anonymity would merely require indifference under permutations of the columns of the matrix of local envies. Such a weaker axiom would allow us to distinguish, say, a situation where $n - 1$ agents all envy the nth agent from a situation where the same $n - 1$ agents all envy a different agent, which may be useful in some circumstances.

Source and target anonymity together mean that we can think of the input to the function $envy$ as a multiset of multisets of local envies, rather than as a matrix of envies. Note that we do not require $envy(.,.)$ to be computable from the single multiset of all local envies—that would be an even stronger requirement that, for instance, would be violated by the measure of envy counting the number of envious agents in a society ($e^{sum,max,bool}$).

It is readily verified that all of the concrete measures of envy for a society defined earlier meet our axiomatic requirements:

Proposition 17 (Envy measures). *Let $op_1, op_2 \in \{max, sum\}$ and let $op_3 \in \{raw, bool\}$. Then e^{op_1, op_2, op_3} is a proper measure of envy that satisfies discernibility, dominance, target anonymity, and source anonymity.*

---

[3]Note that we do not necessarily require that $envy(.,.)$ can be computed in terms of the degrees of envy of single agents, even though this is the case for all our concrete measures of envy defined earlier.

6.5. Complexity of reducing envy

The objectives of achieving envy-freeness and of maintaining individual rationality during negotiation are not entirely compatible. This is why we have required initial payments for our convergence results. Indeed, as the following example shows, without initial payments, there can be negotiation states such that no IR deal would allow us to eliminate envy. Consider a scenario with two agents and one good:

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

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.

So it is a relevant question whether a given negotiation state admits a deal that would not only be IR but that also reduces envy in the system (or, as a special case, that completely eliminates envy). As a final contribution, we now want to analyse the computational complexity of this problem, for arbitrary measures of envy of the form $envy(.,.)$. 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 $\langle v_1, \ldots, v_n \rangle$.

WELFARE IMPROVEMENT WITH ENVY REDUCTION (WIER)

Instance:	negotiation problem, state (A, π), measure $envy(.,.)$
Question:	Is there an alternative state (A', π') such that:

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

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 is reduced (third condition).

WIER is related to the WELFARE IMPROVEMENT problem (WI) studied in the literature [11], which is identical to WIER, except that it does not impose any conditions on envy. WI is known to be NP-complete for a number of languages for representing valuation functions (i.e., the input to the problem). Despite this similarity, we are unaware of any simple reduction from WI to WIER, which could be used to prove NP-completeness of WIER.[4] Instead, below we will provide a reduction from the fundamental CLIQUE problem. Clearly, there is no chance for WIER being computable in polynomial time unless $envy(.,.)$ is computable in polynomial time, so we will require this. Regarding the language used to represent valuation functions, we will use the simplest of all such languages, namely the language for expressing single-minded valuations [42]. Specifying a valuation function in this setting amounts to fixing one bundle and its value. Finally, our result applies to *every* proper envy measure that satisfies the discernibility axiom, rather than just to the specific measures defined earlier.

Theorem 18 (Complexity of WIER)**.** *WIER is NP-complete for single-minded agents for every proper measure of envy that is polynomial-time computable and that satisfies the discernibility axiom.*

Proof. NP-membership follows from the fact that the conditions of WIER can be verified in polynomial time for any proposed solution state.

To establish NP-hardness we give a polynomial-time reduction from the NP-hard CLIQUE problem [43], defined as follows: Given an undirected graph $G = (V, E)$ and a positive integer $K \leqslant |V|$, is it the case that G contains a clique of size K or larger? W.l.o.g., assume $K \geqslant 2$.

---

[4]In previous work we have given such a reduction [2], but that reduction involves a negotiation problem with valuation functions that are not monotonic, which somewhat limits the interestingness of that result.

Auctioneer x^a
– valuation: $v^a(\mathcal{R}) = K - 1$
– holds all goods in $\mathcal{R}$
– has no money

Real agent x_i^r
– valuation: $v_i^r(\{r_i\}) = 1$
– holds no goods
– has amount of $K - \epsilon$

Dummy agent x_i^d
– valuation: $v_i^d(\{d_i\}) = \Omega$
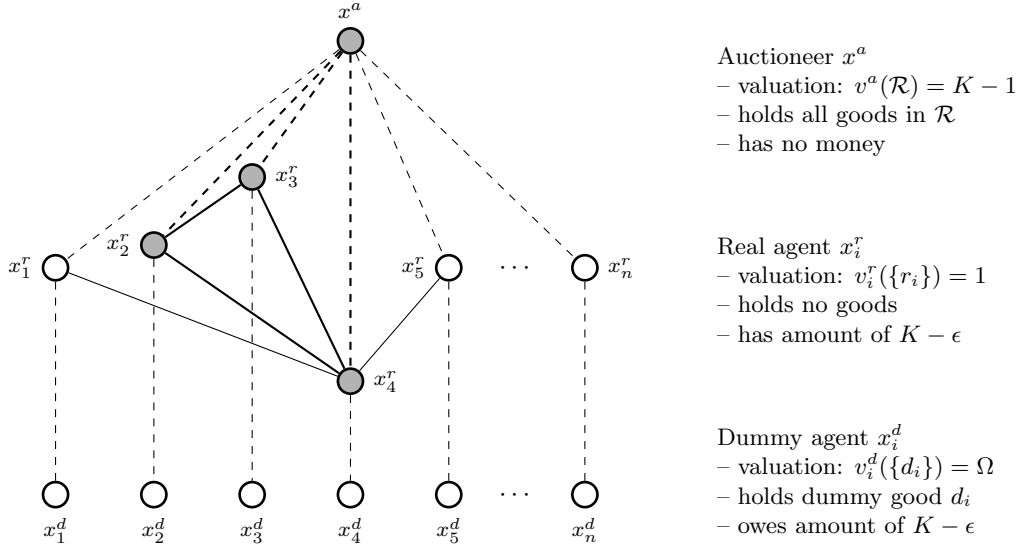– holds dummy good d_i
– owes amount of $K - \epsilon$

Figure 2: Reduction from Clique to WIER (initial state)

Given a Clique instance with $|V| = n$, we build the following instance of WIER. Fix some $\Omega > 2K$ and some positive $\epsilon < 1$. There are $2n$ goods in the system, $\mathcal{G} = \mathcal{R} \cup \mathcal{D}$ with $\mathcal{R} = \{r_1, \ldots, r_n\}$ (the "real" goods) and $\mathcal{D} = \{d_1, \ldots, d_n\}$ (the "dummy" goods). There furthermore are $2n+1$ agents with single-minded valuations, meaning that they each value just a single bundle of goods (and all their supersets). We indicate this single bundle and its value for each agent:

- a single *auctioneer agent* x^a, with $v^a(\mathcal{R}) = K - 1$;

- one *real agent* x_i^r, with $v_i^r(\{r_i\}) = 1$, for all $i \in \{1, \ldots, n\}$;

- one *dummy agent* x_i^d, with $v_i^d(\{d_i\}) = \Omega$, for all $i \in \{1, \ldots, n\}$.

The real agents are connected amongst themselves exactly as in the given Clique instance. Apart from that, the agents are connected via a graph as indicated in Figure 2. Thus, the auctioneer is connected to every real agent and each real agent x_i^r is connected to the corresponding dummy agent x_i^d. Now consider a state in which the auctioneer holds all the goods in $\mathcal{R}$ and each dummy agent x_i^d holds the dummy item d_i she desires. Furthermore, each real agent has a positive payment balance of $K - \epsilon$, while each dummy agent owes an amount of $K - \epsilon$. These payment balances are feasible, even if we do not wish to permit initial payments: in a previous state, the real agents may have held the dummy items and sold them to the dummy agents for a price of $K - \epsilon$ a piece.[5]

Let us take note of the envies experienced in this state. The auctioneer envies each of the real agents: her own utility is $K - 1$, but each real agent has a monetary amount of $K - \epsilon$. The real agents and the dummy agents do not experience any envy. In particular, no real agent envies the auctioneer, because $K - \epsilon$ (the amount of money she owns) is greater than 1 (the utility she ascribes to the bundle held by the auctioneer).

Now suppose there is a clique of size at least K in the original graph. We will show how to construct an envy-reducing IR clique-deal in this case. The clique in question must be part of some *maximal* clique (which could be that clique itself). Let $K' \geqslant K$ be the size of that maximal clique. Let us call the real agents corresponding to the vertices in the maximal clique the *clique agents* and let us call the remaining real agents the *non-clique agents*. Consider the following deal, starting from the initial state described earlier and shown in Figure 2:

---
[5]This, the justification of the payment balances in this initial state, has been the only reason for introducing dummy agents and dummy goods.

- Each clique agent x_i^r obtains item r_i from the auctioneer and pays her an amount of $(K - \epsilon)/K'$ in return, i.e., the auctioneer obtains a total amount of $K - \epsilon$.

- The remaining real goods are also given to the clique agents, in such a way that no non-clique agent x_i^r can see the item r_i she desires. This is possible due to the maximality of the clique (for every non-clique agent there exists a clique agent she is not connected to).

This deal is an IR clique-deal: all the agents involved are fully connected to each other, and social welfare increases from $n \cdot \Omega + (K - 1)$ to $n \cdot \Omega + K'$. This deal is also envy-reducing. Indeed, this deal reduces envy to zero (which, by the discernibility axiom, makes it an envy-reducing deal). To see this, first note that nobody envies anyone for goods anymore (nobody can see another agent holding a good she desires herself), so we only need to compare every agent's utility with the monetary balances of the agents she can see:

- Each dummy agent has utility $\Omega - (K - \epsilon) > K$, while no agent possesses an amount of money exceeding K.

- The auctioneer has an amount of $K - \epsilon$ (which is also her utility), while no agent she can see either holds a bundle she desires or owns an amount of money exceeding $K - \epsilon$. The same is true for every non-clique agent.

- Each clique agent x_i^r has utility $1 + (K - \epsilon) - (K - \epsilon)/K'$, composed of the value she ascribes to r_i, the amount of money she held before the deal, and the payment she made to the auctioneer. But this sum is strictly greater than $K - \epsilon$, the maximal amount of money held by any of the agents she can see.

This concludes the first direction of the proof: if there is a clique of size K, then there exists an envy-reducing IR clique-deal. For the other direction, suppose the largest clique in the given graph is strictly smaller than K. We will see that then there exists no IR clique-deal (and thus certainly no such deal that would also be envy-reducing). So consider all maximal cliques of agents that might conduct a deal for the state shown in Figure 2:

- No group of real agents could raise enough money to reimburse the auctioneer for giving up some or all of the real goods.

- No real agent would have an interest in buying a dummy good from a dummy agent.

Hence, if there is no clique of size K, then no deal is possible. This concludes the proof. $\square$

7. Conclusion

In this paper we have addressed the problem of designing mechanisms to fairly allocate indivisible goods to a set of agents. While previously this question has been mostly investigated in the context of assignment problems, in which agents can receive at most one such good [6], 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 for other proposals in which payments ensuring fairness are arranged for a given efficient allocation [33, 34], the approach advocated here is *distributed:* agents locally implement mutually beneficial reallocations of goods (furthering the efficiency of the allocation), while the payment rule updates the payments after each deal in a way that ensures fairness at the end of the process. That is, in this approach the quests for efficiency and fairness are intimately interleaved.

The first contribution of this paper has been 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, however, 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 [29]. 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 4 demonstrates that the classical Knaster procedure can be given a distributed flavour and still achieve proportional outcomes. Theorem 8 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 9) or different types of payment schemes (Theorem 10).

Having established these results, we turned our attention to a situation of great practical relevance, namely societies of agents where the information available to agents is constrained by a social network. Interestingly, we saw that, while the notion of proportionality was difficult to characterise in a meaningful manner, envy can naturally be restricted to agents that can be "seen" from an agent's position in the network. Such topological constraints can be very damaging as far as efficiency is concerned, which suggests that it could potentially jeopardise also the possibility of reaching envy-free states. Theorem 14 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 graph, an interesting lead for future research would be to relax this assumption. Typically, the negotiation graph would be a subgraph of the envy graph, so convergence to fair outcomes will be more difficult to guarantee.

For all of our convergence results we have also provided 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 5, 11, and 15).

By construction (and following Lemma 3), 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 standard way of quantifying "how much" envy occurs in a society. We have addressed this question by proposing a systematic *typology of envy measures* as well as a set of basic axioms that should be met by any such measure. Equipped with these envy measures, it will be possible to investigate how envy within a society evolves over time when negotiation takes place. For example, experiments reported in related work [1] suggest that envy tends to decrease, albeit not strictly monotonically. More generally, these measures of envy allow us to go beyond a mere classification of a society as being either envy-free or not. In particular, while some of the conditions stated in our convergence theorems are necessary to guarantee full envy-freeness in all cases, it would be interesting to investigate by how much envy can be decreased when not all of these conditions are met. Finally, we have shown that the local decision problem faced by agents when they have to agree on a rational deal that reduces envy is computationally intractable (Theorem 18), albeit not harder than identifying an arbitrary rational deal.

References

[1] Y. Chevaleyre, U. Endriss, S. Estivie, N. Maudet, Reaching envy-free states in distributed negotiation settings, in: Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI-2007), 2007, pp. 1239–1244.

[2] Y. Chevaleyre, U. Endriss, N. Maudet, Allocating goods on a graph to eliminate envy, in: Proceedings of the 22nd AAAI Conference on Artificial Intelligence (AAAI-2007), AAAI Press, 2007, pp. 700–705.

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

[4] J. Robertson, W. Webb, Cake-Cutting Algorithms: Be fair if you can, A.K. Peters, 1998.

[5] A. D. Procaccia, Cake cutting algorithms, in: F. Brandt, V. Conitzer, U. Endriss, J. Lang, A. D. Procaccia (Eds.), Handbook of Computational Social Choice, Cambridge University Press, 2016, Ch. 13, pp. 311–329.

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

[7] C. Beviá, Fair allocation in a general model with indivisible goods, Review of Economic Design 3 (3) (1998) 195–213.

[8] K. Tadenuma, Trade-off between equity and efficiency in a general economy with indivisible goods, Social Choice and Welfare 13 (1996) 445–450.

[9] S. Bouveret, Y. Chevaleyre, N. Maudet, Fair allocation of indivisible goods, in: F. Brandt, V. Conitzer, U. Endriss, J. Lang, A. D. Procaccia (Eds.), Handbook of Computational Social Choice, Cambridge University Press, 2016, Ch. 12, pp. 284–310.

[10] W. Thomson, Fair allocation rules, in: K. J. Arrow, A. Sen, K. Suzumura (Eds.), Handbook of Social Choice and Welfare, Vol. 2, Elsevier, 2011, Ch. 21, pp. 393–506.

[11] Y. Chevaleyre, P. E. Dunne, U. Endriss, J. Lang, M. Lemaître, N. Maudet, J. Padget, S. Phelps, J. A. Rodríguez-Aguilar, P. Sousa, Issues in multiagent resource allocation, Informatica 30 (2006) 3–31.

[12] R. Lipton, E. Markakis, E. Mossel, A. Saberi, On approximately fair allocations of indivisible goods, in: Proceedings of the 5th ACM Conference on Electronic Commerce (EC-2004), ACM Press, 2004, pp. 125–131.

[13] U. Endriss, N. Maudet, F. Sadri, F. Toni, Negotiating socially optimal allocations of resources, Journal of Artificial Intelligence Research 25 (2006) 315–348.

[14] S. Bouveret, J. Lang, Efficiency and envy-freeness in fair division of indivisible goods: Logical representation and complexity, Journal of Artificial Intelligence Research 32 (2008) 525–564.

[15] I. Caragiannis, C. Kaklamanis, P. Kanellopoulos, M. Kyropoulou, On low-envy truthful allocations, in: Proceedings of the 1st International Conference on Algorithmic Decision Theory (ADT-2009), Vol. 5783 of LNAI, Springer-Verlag, 2009, pp. 111–119.

[16] B. de Keijzer, S. Bouveret, T. Klos, Y. Zhang, On the complexity of efficiency and envy-freeness in fair division of indivisible goods with additive preferences, in: Proceedings of the 1st International Conference on Algorithmic Decision Theory (ADT-2009), Vol. 5783 of LNAI, Springer-Verlag, 2009, pp. 98–110.

[17] D. Herreiner, C. Puppe, A simple procedure for finding equitable allocations of indivisible goods, Social Choice and Welfare 19 (2) (2002) 415–430.

[18] T. W. Sandholm, Contract types for satisficing task allocation: I Theoretical results, in: Proceedings of the 1998 AAAI Spring Symposium on Satisficing Models, 1998.

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

[20] Y. Chevaleyre, U. Endriss, N. Maudet, Simple negotiation schemes for agents with simple preferences: Sufficiency, necessity and maximality, Journal of Autonomous Agents and Multiagent Systems 20 (2) (2010) 234–259.

[21] S. Airiau, U. Endriss, Multiagent resource allocation with sharable items, Journal of Autonomous Agents and Multiagent Systems 28 (6) (2014) 956–985.

[22] P. Cramton, Y. Shoham, R. Steinberg (Eds.), Combinatorial Auctions, MIT Press, 2006.

[23] A. Netzer, A. Meisels, R. Zivan, Distributed envy minimization for resource allocation, Journal of Autonomous Agents and Multiagent Systems 30 (2) (2016) 364–402.

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

[25] S. Estivie, Y. Chevaleyre, U. Endriss, N. Maudet, How equitable is rational negotiation?, in: Proceedings of the 5th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2006), ACM Press, 2006, pp. 866–873.

[26] H. Steinhaus, The problem of fair division, Econometrica 16 (1) (1948) 101–104.

[27] M. G. Raith, Fair-negotiation procedures, Mathematical Social Sciences 39 (3) (2000) 303–322.

[28] B. Knaster, Sur le probléme du partage pragmatique de H. Steinhaus, Annales de la Société Polonaise de Mathématique 19 (1946) 228–230.

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

[30] D. K. Foley, Resource allocation and the public sector, Yale Economic Essays 7 (1) (1967) 45–98.

[31] D. Gale, The Theory of Linear Economic Models, University of Chicago Press, 1960.

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

[33] F. Klijn, An algorithm for envy-free allocations in an economy with indivisible objects and money, Social Choice and Welfare 17 (2) (2000) 201–215.

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

[35] H. W. Kuhn, The Hungarian Method for the assignment problem, Naval Research Logistic Quarterly 2 (1–2) (1955) 83–97.

[36] M. de Weerdt, Y. Zhang, T. Klos, Multi-agent task allocation in social networks, Journal of Autonomous Agents and Multi-Agent Systems 25 (1) (2012) 46–86.

[37] A. Feldman, A. Kirman, Fairness and envy, The American Economic Review 64 (6) (1974) 995–1005.

[38] S. J. Brams, M. A. Jones, C. Klamler, Divide-and-conquer: A proportional, minimal-envy cake-cutting procedure, SIAM Review 53 (2) (2011) 291–307.

[39] C. Lindner, J. Rothe, Degrees of guaranteed envy-freeness in finite bounded cake-cutting protocols, in: Proceedings of the 5th International Workshop on Internet and Network Economics (WINE-2009), Vol. 5929 of LNCS, Springer-Verlag, 2009, pp. 149–159.

[40] T. T. Nguyen, J. Rothe, Minimizing envy and maximizing average Nash social welfare in the allocation of indivisible goods, Discrete Applied Mathematics 179 (2014) 54–68.

[41] H. Hosseini, K. Larson, R. Cohen, Matching with dynamic ordinal preferences, in: Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI-2015), 2015, pp. 936–943.

[42] D. Lehmann, L. I. O'Callaghan, Y. Shoham, Truth revelation in approximately efficient combinatorial auctions, Journal of the ACM 49 (5) (2002) 577–602.

[43] M. R. Garey, D. S. Johnson, Computers and Intractability: A Guide to the Theory of NP-Completeness, W. H. Freeman and Company, New York, 1979.