

Fair Division under Ordinal Preferences: Computing Envy-Free Allocations of Indivisible Goods

Sylvain Bouveret¹ and Ulle Endriss² and Jérôme Lang³

Abstract. We study the problem of fairly dividing a set of goods amongst a group of agents, when those agents have preferences that are ordinal relations over alternative bundles of goods (rather than utility functions) and when our knowledge of those preferences is incomplete. The incompleteness of the preferences stems from the fact that each agent reports their preferences by means of an expression of bounded size in a compact preference representation language. Specifically, we assume that each agent only provides a ranking of individual goods (rather than of bundles). In this context, we consider the algorithmic problem of deciding whether there exists an allocation that is possibly (or necessarily) envy-free, given the incomplete preference information available, if in addition some mild economic efficiency criteria need to be satisfied. We provide simple characterisations, giving rise to simple algorithms, for some instances of the problem, and computational complexity results, establishing the intractability of the problem, for others.

1 INTRODUCTION

The problem of fairly dividing a set of goods amongst a group of agents has recently started to receive increased attention in the AI literature [6, 10, 15, and others]. The study of the computational aspects of fair division, in particular, finds a natural home in AI; and fair division is immediately relevant to a range of applications in multiagent systems and electronic commerce.

To define an instance of a fair division problem, we need to specify the *type of goods* we want to divide, the nature of the *preferences* that individual agents hold, and the kind of *fairness criterion* we want to apply when searching for a solution. In this paper, we are concerned with *indivisible goods* that cannot be shared: each item needs to be allocated to (at most) one agent in its entirety. This choice renders fair division a combinatorial optimisation problem.

Regarding preferences, most work in fair division has made the assumption that the preferences of individual agents can be modelled as utility (or valuation) functions, mapping bundles of goods to a suitable numerical scale. This assumption is technically convenient, and it is clearly appropriate in the context of applications with a universal currency, rendering preferences interpersonally comparable. On the other hand, from a cognitive point of view, assuming such cardinal preferences may be questionable, as it requires an agent to be able to attach a number to every conceivable state of the world. In this paper, we make instead the (much weaker, and arguably more realistic) assumption that agents have *ordinal preferences*, and for the sake of simplicity we assume that these preferences are strict orders (which

is a common assumption in fair division and voting). That is, each agent i is equipped with a preference relation \succ_i : $A \succ_i B$ expresses that agent i prefers the set of items A over the set of items B .

The third parameter is the criterion used to define what makes an allocation “fair”. Restricting attention to ordinal preferences rules out some criteria. For instance, the Rawlsian (or egalitarian) approach to fairness ties social welfare to the welfare of the worst-off agent [16], which presupposes that interpersonal comparison of preferences is possible. Instead, we focus on the important criterion of *envy-freeness* [13]. An allocation is envy-free if each agent likes the bundle she received at least as much as any of the bundles received by others. Besides envy-freeness, a secondary criterion we shall be working with is *Pareto efficiency*, which also only requires ordinal preferences. An allocation is Pareto efficient if there is no other allocation making some agents better and no agent worse off.

A challenging aspect of devising methods for fair division with indivisible goods is its combinatorial nature [9]: the space of possible bundles grows exponentially in the number of goods. If there are 20 goods, each agent would, in principle, have to rank over one million bundles. This leads to the following dilemma: either we allow agents to express any possible preference relation on the set of all subsets of items, and end up with an exponentially large representation, as in the *descending demand procedure* of Herreiner and Puppe [14], which, while of great theoretical interest, is computationally infeasible as soon as the number of goods is more than a few units; or we restrict the range of preferences that agents may express. The latter is the path followed by Brams and King [8] and Brams et al. [7], who address the problem using the following approach: Elicit the preferences \triangleright_i of each agent i over *single* goods (the assumption is that this is a strict linear order) and induce an (incomplete) preference order \succ_i over bundles as follows: for two bundles A and B , infer $A \succ_i B$ if there exists an injective mapping $f : (B \setminus A) \rightarrow (A \setminus B)$ such that $f(a) \triangleright_i a$ for any $a \in B \setminus A$. That is, \succ_i ranks A above B if a (not necessarily proper) subset of A *pairwise dominates* B , i.e., if A is definitely preferred to B given the limited information (provided in the form of \triangleright_i) available—under reasonable assumptions on how to “lift” preferences from single goods to bundles.⁴ From a “computational” perspective, we might say that Brams and coauthors [7, 8] are using \triangleright_i as a *compact representation* of \succ_i . In fact, their approach coincides precisely with a simple fragment of the language of *conditional importance networks* (CI-nets), a compact graphical representation language for modelling ordinal preference relations that are monotonic [5]. The fragment in question are the so-called (ex-

¹ Onera Toulouse, France, email: sylvain.bouveret@onera.fr

² ILLC, Uni. of Amsterdam, The Netherlands, email: ulle.endriss@uva.nl

³ Lamsade, Uni. Paris-Dauphine, France, email: lang@lamsade.dauphine.fr

⁴ The problem of lifting preferences over items to sets of items has been studied in depth in social choice theory [3]. Indeed, pairwise dominance is closely related to the axiom of “(weak) preference dominance” put forward by Sen in the context of work on formalising freedom of choice [17].

haustive) *SCI-nets*, which we will define in Section 2.2.

We will model agent preferences using SCI-nets. Each SCI-net induces an incomplete preference order over bundles, with the intended interpretation that the agent's true preference order is some complete order that is consistent with the known incomplete order. This requires a nonstandard approach to defining fairness criteria. Here, again, we follow Brams and King [8] and Brams et al. [7] and define an allocation as being *possibly* envy-free if it is envy-free for *some* set of complete preferences that are consistent with the known incomplete preferences; and we say an allocation is *necessarily* envy-free if it is envy-free under *all* possible completions. We define possible and necessary Pareto efficiency accordingly.

The main question we study in this paper is then: *Given partially specified agent preferences, modelled in terms of SCI-nets, does there exist an allocation that is possibly (necessarily) envy-free?* As the allocation that simply disposes of all goods (*i.e.*, that does not assign any goods to the agents) is always both possibly and necessarily envy-free, to be interesting, this question needs to be asked under some efficiency requirements. In particular, we will ask whether there exists such allocations that are *complete* (*i.e.*, that allocate every item to some agent) or possibly (necessarily) Pareto efficient.

Some of our results are positive: we are able to provide simple characterisations of situations in which an allocation of the desired kind exists, and these characterisations immediately suggest an algorithm for computing such an allocation. Other results are negative: deciding existence of an allocation of the desired kind (and thus also computing such an allocation) often turns out to be intractable.

The remainder of the paper is organised as follows. In Section 2 we define the model of fair division we shall be working with. In particular, this includes the language used to specify agent preferences and several fairness and efficiency criteria. In Section 3 we give the main results of this paper; namely, we show that while it is easy to compute possibly envy-free allocations that are also complete or possibly Pareto efficient, requiring necessary envy-freeness makes the problem NP-hard. The concluding Section 4 includes a short discussion of related work. (For lack of space, some proofs are only sketched.)

2 THE MODEL

Let $\mathcal{A} = \{1, \dots, n\}$ be a finite set of *agents* and $\mathcal{G} = \{x_1, \dots, x_m\}$ be a finite set of *goods* ($n \geq 2$ and $m \geq 1$). An *allocation* $\pi : \mathcal{A} \rightarrow 2^{\mathcal{G}}$ is a mapping from agents to sets of goods such that $\pi(i) \cap \pi(j) = \emptyset$ for any two distinct agents $ij \in \mathcal{A}$; thus, goods are *indivisible*. An allocation π with $\pi(1) \cup \dots \cup \pi(n) = \mathcal{G}$ is called *complete*.

In this section, we define criteria for identifying fair (or efficient) allocations of goods. These criteria will be defined in terms of the preferences of the individual agents over the bundles they receive.

2.1 Basic terminology and notation

A *strict partial order* is a binary relation that is irreflexive and transitive. A *linear order* is a strict partial order that is complete (*i.e.*, $X \succ Y$ or $Y \succ X$ whenever $X \neq Y$). A binary relation \succ on $2^{\mathcal{G}}$ is *monotonic* if $X \supset Y$ implies $X \succ Y$. If \succ (or \triangleright) is a binary relation, then \succeq (or \trianglerighteq) represents the reflexive closure of that relation (*i.e.*, $X \succeq Y$ if and only if $X \succ Y$ or $X = Y$). Given two binary relations R and R' on $2^{\mathcal{G}}$, we say that R' *refines* R if $R \subseteq R'$.

2.2 Preferences: SCI-nets

The preference relation of each agent $i \in \mathcal{A}$ is assumed to be a linear order $\succ_i^?$ over the bundles (subsets of \mathcal{G}) she might receive. However,

as argued above, eliciting $\succ_i^?$ entirely would be infeasible; so we do not assume that $\succ_i^?$ is fully known to us (or even to the agents themselves). Instead, for each agent i we are given a strict partial order \succ_i representing our partial knowledge of $\succ_i^?$, and the true preference of i is *some* complete refinement of \succ_i . The strict partial orders \succ_i are generated from expressions of a suitable preference representation language. In this paper, we focus on the language of SCI-nets, *i.e.*, precondition-free CI-nets in which all compared sets are singletons [5]. We now introduce SCI-nets;⁵ for full CI-nets see [5].

Definition 1 (SCI-nets) An SCI-net \mathcal{N} on \mathcal{G} is a linear order on \mathcal{G} , denoted by $\triangleright_{\mathcal{N}}$ (or simply \triangleright , when the context is clear). A strict partial order \succ on $2^{\mathcal{G}}$ complies with \mathcal{N} , if (i) \succ is monotonic and (ii) $S \cup \{x\} \succ S \cup \{y\}$ for any \mathfrak{y} such that $x \triangleright_{\mathcal{N}} y$ and any $S \subseteq \mathcal{G} \setminus \{x, y\}$. The preference relation $\succ_{\mathcal{N}}$ induced by \mathcal{N} is the smallest strict partial order that complies with \mathcal{N} .




Figure 1. Preference relation induced by SCI-net $a \triangleright b \triangleright c \triangleright d$. Dotted arcs are obtained by monotonicity; arcs obtained by transitivity are omitted.

As discussed earlier, $\succ_{\mathcal{N}}$ is the partial order we obtain when we lift the order $\triangleright_{\mathcal{N}}$ on \mathcal{G} to an order on $2^{\mathcal{G}}$ by invoking the principles of monotonicity and pairwise dominance, as proposed by Brams and coauthors [7, 8]. We can give yet another characterisation of $\succ_{\mathcal{N}}$, in terms of a utility function: Given SCI-net \mathcal{N} and $A \subseteq \mathcal{G}$, for every $k \leq |A|$ we denote with $A_{(k)}^{\mathcal{N}}$ the k -most important element of A ; *i.e.*, if $x \in A$ and $\#\{y \in A \mid y \succeq_{\mathcal{N}} x\} = k$ then $A_{(k)}^{\mathcal{N}} = x$. Given a vector $w = (w_1, \dots, w_m) \in (\mathbb{R}^+)^m$ inducing the additive utility function $u_w : 2^{\mathcal{G}} \rightarrow \mathbb{R}$ with $u_w(A) = \sum_{x_i \in A} w_i$, and SCI-net $\mathcal{N} = x_{(1)} \triangleright \dots \triangleright x_{(m)}$ (for some permutation θ of $\{1, \dots, m\}$), we say that w and \mathcal{N} are *compatible* if $w_{(1)} > \dots > w_{(m)}$.

Proposition 1 (Dominance) Given an SCI-net \mathcal{N} and bundles $\mathcal{A} \subseteq \mathcal{G}$, the following statements are equivalent:

- (1) $A \succ_{\mathcal{N}} B$
- (2) There exists an injective mapping $f : (B \setminus A) \rightarrow (A \setminus B)$ such that $f(a) \triangleright_{\mathcal{N}} a$ for any $a \in B \setminus A$.
- (3) There exists an injective mapping $g : B \rightarrow A$ such that $g(a) \triangleright_{\mathcal{N}} a$ for all $a \in B$ and $g(a) \triangleright_{\mathcal{N}} a$ for some $a \in B$.
- (4) Either $A \supset B$, or the following three conditions are satisfied:
 - $|A| \geq |B|$;
 - for every $k \leq |B|$, $A_{(k)}^{\mathcal{N}} \triangleright_{\mathcal{N}} B_{(k)}^{\mathcal{N}}$;
 - there exists a $k \leq |B|$ such that $A_{(k)}^{\mathcal{N}} \triangleright_{\mathcal{N}} B_{(k)}^{\mathcal{N}}$.
- (5) For any w compatible with \mathcal{N} we have $u_w(A) > u_w(B)$.

The proof is simple; we omit it due to space constraints.

2.3 Criteria: envy-freeness and efficiency

For the fair division problems we study, each agent $i \in \mathcal{A}$ provides an SCI-net \mathcal{N}_i . This gives rise to a *profile* of strict partial orders ($\succ_{\mathcal{N}_1}, \dots, \succ_{\mathcal{N}_n}$). For any such profile (whether it has been induced by SCI-nets or not), we can ask whether it admits a fair solution.

As our agents are only expressing *incomplete* preferences, the standard notions of envy-freeness and efficiency need to be adapted. For any solution concept, we may say that it is *possibly* satisfied (if

⁵ What we call “SCI-nets” here were called “exhaustive SCI-nets” in [5].

some refinement of the preference profile to a profile of linear orders satisfies it) or that it is *necessarily* satisfied (if *all* such refinements do). The following definitions are a synthesis of those introduced by Brams and King [8] and Brams et al. [7].⁶ While the results reported in the sequel apply to scenarios where each agent expresses her preferences in terms of an SCI-net, we state these definitions independently from the preference representation language in use.

Definition 2 (Modes of envy-freeness) *Given a profile of strict partial orders $(\succ_1, \dots, \succ_n)$ on 2^G , an allocation π is called*

- (i) *possibly envy-free (PEF) if for every $i \in \mathcal{A}$ there exists a linear order $\succ_i^?$ refining \succ_i such that $\pi(i) \succ_i^? \pi(j)$ for all $j \in \mathcal{A}$;⁷ and*
- (ii) *necessarily envy-free (NEF) if for every $i \in \mathcal{A}$ and every linear order $\succ_i^?$ refining \succ_i we have $\pi(i) \succ_i^? \pi(j)$ for all $j \in \mathcal{A}$.*

Next we establish alternative characterisations of PEF and NEF allocations, which are more “computation-friendly”.

Proposition 2 (PEF and NEF allocations) *Given $(\succ_1, \dots, \succ_n)$,*

- *π is NEF if and only if for all ij , we have $\pi(i) \succ_i \pi(j)$;*
- *π is PEF if and only if for all ij , we have $\pi(j) \not\succeq_i \pi(i)$.*

Proof. The first point is obvious: π is NEF iff for every i and j , and every $\succ_i^?$ refining \succ_i we have $\pi(i) \succ_i^? \pi(j)$, i.e., iff $\pi(i) \succ_i \pi(j)$ holds for every i, j . For the second point, suppose $\pi(j) \succ_i \pi(i)$ for some ij ; then $\pi(j) \succ_i^? \pi(i)$ holds for any refinement $\succ_i^?$ of \succ_i , which implies that π is not PEF. The converse direction is less immediate, because the condition C_i : “for all j , $\pi(j) \not\succeq_i \pi(i)$ ” only guarantees that for every i and every $j \neq i$ there exists a refinement $\succ_i^?$ of \succ_i such that $\pi(i) \succ_i^? \pi(j)$. Assume that C_i holds and let the relation R_i be defined by $R_i = [\succ_i \cup \{(\pi(i), B) \mid B \neq \pi(i) \text{ and } B \not\succeq_i \pi(i)\}]$. We show that R_i is acyclic. First, suppose there is an X such that $\mathcal{R} \succ_i X$. Then by definition of R_i , $X \succ_i X$ ($X \neq \pi(i)$ by definition of R_i), which cannot be the case since \succ_i is a well-defined strict order. Suppose now that there exists an irreducible cycle X_1, \dots, X_q of length at least 2 such that $X_1 R_i X_2 \dots R_i X_q R_i X_{q+1} = X_1$, and $X_j \neq X_k$ for every $1 \leq j \neq k \leq q$. From the definition of R_i , for every $k \leq q$ we have either $X_k \succ_i X_{k+1}$ or $(X_k = \pi(i) \text{ and } X_{k+1} \not\succeq_i \pi(i))$. Because \succ_i is acyclic, there is at least one k such that $X_k = \pi(i)$. Because the cycle is irreducible, there is at most one k such that $X_k = \pi(i)$. Therefore, there is exactly one k such that $X_k = \pi(i)$; without loss of generality, let $k = 1$. We have (a) $X_2 \succ_i \pi(i)$ and (b) for every $j \neq 1$, $X_j \succ_i X_{j+1}$, that is, $X_1 = \pi(i) R_i X_2 \succ_i X_3 \succ_i \dots \succ_i X_q \succ_i X_1 = \pi(i)$. Because \succ_i is transitive, $X_2 \succ_i X_3 \succ_i \dots \succ_i X_q \succ_i \pi(i)$ implies $X_2 \succ_i \pi(i)$, which contradicts (a). Therefore, R_i is acyclic, and its transitive closure $R_i^?$ is a strict partial order. Take $\succ_i^?$ to be any linear order refining $R_i^?$. Because R_i contains \succ_i , $\succ_i^?$ refines \succ_i ; and for every j , because $\pi(j) \not\succeq_i \pi(i)$, by construction of R_i we have that $\pi(i) R_i \pi(j)$, therefore also $\pi(i) \succ_i^? \pi(j)$. \square

⁶ Brams and coauthors [7, 8] use a different terminology: our necessarily (resp. possibly) envy-free allocations correspond to their allocations that are not envy-possible (resp. that are not envy-ensuring), and our necessarily (resp. possibly) Pareto efficient allocations correspond to their Pareto-ensuring (resp. Pareto-possible) allocations. We believe that applying the standard modalities of “necessary” and “possible” to basic fairness and efficiency criteria is the most systematic way of defining these notions.

⁷ The usual definition of envy-freeness only requires that each agent should be *at least as happy* with her share as with the share of anyone else, i.e., that $\pi(i) \succeq_i^? \pi(j)$ holds for all $i, j \in \mathcal{A}$. Here, $\pi(i) \succeq_i^? \pi(j)$ and $\pi(j) \succ_i^? \pi(i)$ are equivalent, because $\pi(i) \succeq_i^? \pi(j)$ is equivalent to $\pi(i) \succ_i^? \pi(j)$ or $\pi(i) = \pi(j)$, and of course we have $\pi(i) \neq \pi(j)$.

Example 1 *Let $m = 5$, $n = 2$, $\mathcal{N}_1 = a \triangleright b \triangleright c \triangleright d$ and $\mathcal{N}_2 = d \triangleright c \triangleright b \triangleright a$. Consider the allocation π defined by $\pi(1) = \{ad\}$ and $\pi(2) = \{bc\}$. We have $\{bc\} \not\succeq_1 \{ad\}$ and $\{ad\} \not\succeq_2 \{bc\}$, therefore π is PEF. However, π is not NEF, but the allocation π' such that $\pi'(1) = \{\emptyset\}$ and $\pi'(2) = \{ad\}$ is NEF (hence also PEF).*

Recall that for a profile of linear orders $(\succ_1^?, \dots, \succ_n^?)$ on 2^G , an allocation π' is said to *Pareto-dominate* another allocation π if $\pi'(i) \succeq_i^? \pi(i)$ for all $i \in \mathcal{A}$ and $\pi'(j) \succ_j^? \pi(j)$ for some $j \in \mathcal{A}$.

Definition 3 (Modes of dominance) *Given a profile of strict partial orders $(\succ_1, \dots, \succ_n)$ on 2^G and two allocations π and π' ,*

- (i) *π' possibly Pareto-dominates π if π' Pareto-dominates π for some profile of linear orders $(\succ_1^?, \dots, \succ_n^?)$ refining $(\succ_1, \dots, \succ_n)$.*
- (ii) *π' necessarily Pareto-dominates π if π' Pareto-dominates π for all profiles of linear orders $(\succ_1^?, \dots, \succ_n^?)$ refining $(\succ_1, \dots, \succ_n)$.*

We get characterisations of possible and necessary Pareto dominance that are similar as those of Proposition 2.

Proposition 3 (Pareto dominance) *Given $(\succ_1, \dots, \succ_n)$,*

- *π' necessarily Pareto-dominates π if and only if (a) for all i , we have $\pi'(i) \succeq_i \pi(i)$ and (b) for some i , we have $\pi'(i) \succ_i \pi(i)$;*
- *π' possibly Pareto-dominates π if and only if (c) for all i , we have $\pi(i) \not\succeq_i \pi'(i)$ and (d) for some i , we have $\pi(i) \not\succeq_i \pi'(i)$.*

Proof. For the first point: (a) and (b) together clearly imply that π' necessarily dominates π . Conversely, assume π' necessarily dominates π . Then, by definition, π' Pareto-dominates π for all profiles of linear orders refining the partial orders. Exchanging the position of the two universal quantifiers immediately gives (a). Now, suppose that there is no i such that $\pi'(i) \succ_i \pi(i)$. Then for each i there is at least one refinement $\succ_i^?$ such that $\pi(i) \succeq_i^? \pi'(i)$. Let $P^? = (\succ_1^?, \dots, \succ_n^?)$. $P^?$ refines $(\succ_1, \dots, \succ_n)$, and for $P^?$, π' does not Pareto dominate π , which contradicts the initial assumption, and we are done. The proof for the second point is similar. \square

Definition 4 (Modes of efficiency) *Given a profile of strict partial orders $(\succ_1, \dots, \succ_n)$ on 2^G , an allocation π is called*

- (i) *possibly Pareto efficient (PPE) if there exists no allocation π' that necessarily Pareto-dominates π ; and*
- (ii) *necessarily Pareto efficient (NPE) if there exists no allocation π' that possibly Pareto-dominates π .*

Above concepts naturally extend to the case where preferences are modelled using a representation language, such as SCI-nets. For example, given a profile of SCI-nets $(\mathcal{N}_1, \dots, \mathcal{N}_n)$, an allocation π is PEF if π is PEF for the profile $(\succ_{\mathcal{N}_1}, \dots, \succ_{\mathcal{N}_n})$.

3 COMPUTING ENVY-FREE ALLOCATIONS

In this section, we consider the problem of checking whether, for a given profile of SCI-nets, there exists an allocation that is (possibly or necessarily) envy-free, and that also satisfies a secondary efficiency requirement (in particular completeness).

3.1 Possible envy-freeness

We first ask whether a given profile of SCI-nets permits an allocation that is both PEF and complete. It turns out that there is a very simple characterisation of those profiles that do: all that matters is the number of distinct goods that are ranked at the top by one of the agents (in

relation to the number of agents and goods). As will become clear in the proof of this result, the algorithm for computing a complete PEF allocation is also very simple.

Proposition 4 (PEF: general case) *If n agents express their preferences over m goods using SCI-nets and k distinct goods are top-ranked by some agent, then there exists a complete PEF allocation if and only if $m \geq 2n - k$.*

Proof. First, suppose there are $m \geq 2n - k$ goods. Executing the following protocol will result in a PEF allocation of $2n - k$ of those goods: (1) Go through the agents in ascending order, ask them to pick their top-ranked item if it is still available and ask them leave the room if they were able to pick it. (2) Go through the remaining $n - k$ agents in ascending order and ask them to claim their most preferred item from those still available. (3) Go through the remaining agents in descending order and ask them to claim their most preferred item from those still available. The resulting allocation is PEF, because for no agent the bundle of (one or two) goods(s) she obtained is pairwise dominated by any of the other bundles: she either is one of the k agents who received their top-ranked item or she was able to pick her second item before any of the agents preceding her in the first round were allowed to pick *their* second item. The remaining goods (if any) can be allocated to any of the agents; the resulting allocation remains PEF and is furthermore complete.

Second, suppose there are $m < 2n - k$ goods. Then, by the pigeon hole principle, there must be at least one agent i who receives an item that is not her top-ranked item \mathfrak{h} and no further items beyond that. But then i will necessarily envy the agent who does receive \mathfrak{h} ; thus, the allocation cannot be PEF. \square

Example 2 *Let $m = 6, n = 4, \mathcal{N}_1 = a \triangleright b \triangleright c \triangleright d \triangleright e \triangleright f, \mathcal{N}_2 = a \triangleright d \triangleright b \triangleright c \triangleright e \triangleright f, \mathcal{N}_3 = b \triangleright a \triangleright c \triangleright d \triangleright f \triangleright e$ and $\mathcal{N}_4 = b \triangleright a \triangleright c \triangleright e \triangleright f \triangleright d$. We have $k = 2$ and $m \geq 2n - k$. Therefore, the algorithm returns a complete PEF allocation, namely, if we consider the agents in the order $1 > 2 > 3 > 4$: $\pi(1) = \{a\}$; $\pi(2) = \{d\}$; $\pi(3) = \{b\}$; $\pi(4) = \{e\}$. However, if f were unavailable, there would not be any complete PEF allocation.*

It is possible to show that Proposition 4 remains true if we require allocations to be PPE rather than just complete:

Proposition 5 (PPE-PEF: general case) *If n agents express their preferences over m goods using SCI-nets and k distinct goods are top-ranked by some agent, then there exists a PPE-PEF allocation if and only if $m \geq 2n - k$.*

Proof. First, any PPE allocation is complete; therefore, if there exists a PPE-PEF allocation, there also exists a complete PEF allocation. Conversely, if we refine the protocol given in the proof of Proposition 4 by allowing the last agent in round three to take all the remaining items at the end, then that protocol returns an allocation that is the product of *sincere choices* [8] by the agents for the sequence $1, 2, \dots, m, 1, \dots, 1$. By Proposition 1 of Brams and King [8], any such allocation is PPE. \square

The complexity of determining whether there exists an NPE-PEF allocation is still an open problem.

3.2 Necessary envy-freeness

Next, we turn attention to the problem of checking whether a NEF allocation exists, given a profile of SCI-nets. This is a considerably

more demanding property than possible envy-freeness. For instance, it is easy to see that a necessary precondition for the existence of a complete NEF allocation is that all agents have distinct top-ranked goods (because any agent not receiving her top-ranked good *might* envy the agent receiving it, whatever other goods the two of them may obtain). Another necessary precondition is the following:

Lemma 6 (NEF: necessary condition) *If n agents express their preferences over m goods using SCI-nets and a complete NEF allocation does exist, then m must be a multiple of n .*

Proof. If m is not a multiple of n , then for an allocation to be complete, some agent i must receive fewer goods than another agent j . But *any* SCI-net (including that of i) is consistent with a linear order ranking any bundle of size k above any bundle of size less than k (for all k). Hence, such an allocation cannot be NEF. \square

If there are as many goods as there are agents ($m = n$), then checking whether a complete NEF allocation exists is easy: it does if and only if all agents have distinct top-ranked goods. The next most simple case in which there is a chance that a complete NEF allocation might exist is when there are twice as many goods as agents ($m = 2n$). We now show that checking whether such an allocation exists (and computing it) is intractable:

Proposition 7 (NEF: general case) *If n agents express their preferences over m goods using SCI-nets, then deciding whether there exists a complete NEF allocation is NP-complete (even if $m = 2n$).*

Proof. Membership in NP is straightforward from Proposition 2. Hardness is proved by reduction from [X3C] (exact cover by 3-sets): given a set S of size $3q$, and a collection $C = \langle C_1, \dots, C_n \rangle$ of subsets of S of size 3, does there exist a subcollection C' of C such that every element of S is present exactly once in C' ?

Without loss of generality, we have $n \geq q$. To any instance $\langle \mathcal{E} \rangle$ of [X3C] we associate the following allocation problem:

- $6n$ objects: $3n$ “dummy” objects $\{d_i^1, d_i^2, d_i^3 \mid i = 1, \dots, n\}$, $3q$ “main” objects $\{m_i \mid i = 1, \dots, 3q\}$ and $3(n - q)$ “auxiliary” objects $\{o_i \mid i = 1, \dots, 3(n - q)\}$
- $3n$ agents $\{c_i, c'_i, c''_i \mid i = 1, \dots, n\}$. c_i, c'_i and c''_i are called agents of type i and if $C_i = \{\mathfrak{h}\}$, their preferences are expressed by the following SCI-nets:

$$\begin{aligned} c_i: & d_i^1 \triangleright d_i^2 \triangleright d_i^3 \triangleright m_j \triangleright m_k \triangleright m_l \triangleright o_1 \triangleright o_2 \triangleright o_3 \triangleright \dots \triangleright \\ & o_{3(n-q)-2} \triangleright o_{3(n-q)-1} \triangleright o_{3(n-q)} \triangleright D \triangleright M; \\ c'_i: & d_i^2 \triangleright d_i^3 \triangleright d_i^1 \triangleright m_k \triangleright m_l \triangleright m_j \triangleright o_2 \triangleright o_3 \triangleright o_1 \triangleright \dots \triangleright \\ & o_{3(n-q)-1} \triangleright o_{3(n-q)} \triangleright o_{3(n-q)-2} \triangleright D \triangleright M; \\ c''_i: & d_i^3 \triangleright d_i^1 \triangleright d_i^2 \triangleright m_l \triangleright m_j \triangleright m_k \triangleright o_3 \triangleright o_1 \triangleright o_2 \triangleright \dots \triangleright \\ & o_{3(n-q)} \triangleright o_{3(n-q)-2} \triangleright o_{3(n-q)-1} \triangleright D \triangleright M; \end{aligned}$$

where D (resp. M) means “all other dummy (resp. main) objects in any arbitrary order”. m_j, m_k and m_l will be called “first-level objects” for c_i, c'_i and c''_i .

Suppose there exists an exact cover C' of C . C' contains exactly q subsets, therefore $C \setminus C'$ contains $n - q$ subsets. Let $f : C \setminus C' \rightarrow \{1, \dots, n - q\}$ be an arbitrary bijective mapping. Define the allocation $\pi_{C'}$ as follows:

1. every agent gets her preferred dummy object d_i^j ;
2. if $C_i \in C'$ then every agent of type i gets her preferred (first-level) main object (we will call these agents “lucky” ones);
3. if $C_i \notin C'$, every (unlucky) agent of type i gets an auxiliary object: c_i gets $o_{3f(i)-2}$, c'_i gets $o_{3f(i)-1}$, and c''_i gets $o_{3f(i)}$.

Let us check that $\pi_{C'}$ is a complete allocation. Obviously, every dummy object is allocated (by point 1 above). Since C' is a cover,

every main object is allocated as first-level object for some agent (by point 2 above). Since f is a bijective mapping, every auxiliary object is allocated (by point 3 above). Every agent gets exactly 2 objects, so no object can be allocated twice and the allocation is complete.

Now, we check that $\pi_{C'}$ is NEF. Since every agent receives her top-ranked object and another one, then by Proposition 1, checking that a does not necessarily envy b comes down to checking that $\pi(a)_{(2)}^a \triangleright_a \pi(b)_{(2)}^a$ (hence comparing only the ranks of the worst objects in $\pi(a)$ and $\pi(b)$).

- For each lucky agent a , $\mathbf{r}(\pi(a)_{(2)}^a) = 4$. Each other agent gets either one main object or an auxiliary one. In both cases, the rank is obviously worse than 4, hence preventing a from possibly envying anyone else.
- The worst object received by any unlucky agent a of type i (say w.l.o.g. c_i) is her best one among the triple $\{O_{3f(i)} -$

can build a one-to-one mapping g from J to I such that for every $j \in J$, agent 2 prefers x_j to $x_{g(j)}$. This implies that π is NEF.

Reciprocally, assume there exists a $k \leq q$ such that π gives agent 1 at most $k - 1$ objects among $\{x_1, \dots, x_{2k-1}\}$. Then π gives agent 2 at least k objects among $\{x_1, \dots, x_{2k-1}\}$. This implies that for any one-to-one mapping f from $\pi(1)$ to $\pi(2)$, there is some $i \leq k$ such that $x_{f(i)} \succ_1 x_i$, therefore π is not NEF. Symmetrically, if there exists a $k \leq q$ such that π gives agent 2 at most $k - 1$ objects among her $2k - 1$ preferred objects, then π is not NEF. \square

Proof (Proposition 10). Let the preference relation of agent 1 be, w.l.o.g., $x_1 \succ_1 x_2 \succ_1 \dots \succ_1 x_{2q}$. From that SCI-net, we build the flow network shown in Figure 2 (edge labels $\frac{c}{y}$ correspond to the edge lower bound x and capacity y).

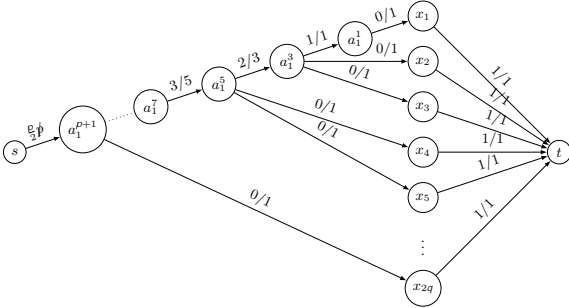


Figure 2. The flow network corresponding to one agent.

We build the same flow network for agent a_2 (nodes a_1^k are now called a_2^k) and identify, between the two networks, the nodes corresponding to the same objects, the source s , and the sink t .

We claim (but do not show due to lack of space) that there is an allocation π satisfying the condition stated in Lemma 11 if and only if there is a feasible flow of value p in the latter network.

The problem of finding a feasible flow in a network with lower bounds as well as capacities is known as the *circulation problem* and is known to be solvable in (deterministic) polynomial-time [12]. Hence the problem of deciding whether there exists a complete NEF allocation for a problem with two agents is in P. \square

4 CONCLUSION AND RELATED WORK

We have studied the problem of computing envy-free allocations of indivisible goods, when agents have ordinal preferences over bundles of goods and when we only know their preferences over single items with certainty. Building on work from the (“non-computational”) fair division literature, in particular the contributions by Brams et al. [7, 8], we have proposed a framework in which to study such questions, we have provided a number of alternative characterisations of the central concepts involved, and we have analysed the computational complexity of computing allocations of the desired kind.

We have been able to show that computing an allocation that is possibly envy-free is easy (whether paired with a requirement for completeness or possible Pareto efficiency). We have also been able to show that computing necessarily envy-free allocations is NP-hard (whatever the secondary efficiency requirement); only for problems with just two agents there is a polynomial (but non-trivial) algorithm. The complexity of finding envy-free allocation that are necessarily Pareto efficient is not fully understood at this stage. In particular, it is conceivable that deciding the existence of allocations that are both necessarily envy-free and necessarily Pareto efficient might not even be in NP; we leave the full analysis of this question to future work.

Future work should also seek to extend our results to *nonstrict* SCI-nets, where indifference between single goods is allowed. Problems that are still easy with strict SCI-nets, such as the existence of a complete PEF allocation, could conceivably become NP-complete. Intuitively, the more indifferences the agents express, the more complete the preference relations and the closer the notions of possible and necessary envy-freeness, which means that possible envy-freeness will be harder to guarantee.

Our work is part of a growing literature on computational aspects of fair division. In particular, complexity aspects of envy-freeness have been studied, for example in the works of Lipton et al. [15] and de Keijzer et al. [11], who address the problem of finding envy-free and complete (resp. Pareto efficient) allocations, when the agents have numerical additive preferences. Bouveret and Lang [6] also address the same problem, for various notions of efficiency, in a context where the agents have utilities expressed in compact form. However, none of these computational works concerns ordinal preferences, and none have considered possible or necessary satisfaction of fairness criteria. There is also a related stream of works on the *Santa Claus problem*, consisting in computing maxmin fair allocations (see e.g., Bansal and Sviridenko [2], Bezáková and Dani [4], Asadpour and Saberi [1]). These works encode fairness by an egalitarian collective utility function and do not consider envy-freeness.

Acknowledgements This research has been partially supported by the ANR project ComSoc (ANR-09-BLAN-0305).

REFERENCES

- [1] A. Asadpour and A. Saberi, ‘An approximation algorithm for max-min fair allocation of indivisible goods’, in *Proc. of STOC-2007*, (2007).
- [2] N. Bansal and M. Sviridenko, ‘The Santa Claus problem’, in *Proc. of STOC-2006*, pp. 31–40, (2006).
- [3] S. Barberà, W. Bossert, and P. Pattanaik, ‘Ranking sets of objects’, in *Handbook of Utility Theory*, volume 2, Kluwer, (2004).
- [4] I. Bezáková and V. Dani, ‘Allocating indivisible goods’, *SIGecom Exch.*, 5(3), 11–18, (2005).
- [5] S. Bouveret, U. Endriss, and J. Lang, ‘Conditional importance networks: A graphical language for representing ordinal, monotonic preferences over sets of goods’, in *Proc. of IJCAI-2009*, (2009).
- [6] S. Bouveret and J. Lang, ‘Efficiency and envy-freeness in fair division of indivisible goods: Logical representation and complexity’, *JAIR*, 32, 525–564, (2008).
- [7] S. J. Brams, P. H. Edelman, and P. C. Fishburn, ‘Fair division of indivisible items’, *Theory and Decision*, 5(2), 147–180, (2004).
- [8] S. J. Brams and D. King, ‘Efficient fair division—help the worst off or avoid envy?’, *Rationality and Society*, 17(4), 387–421, (2005).
- [9] Y. Chevaleyre, U. Endriss, J. Lang, and N. Maudet, ‘Preference handling in combinatorial domains: From AI to social choice’, *AI Magazine*, 29(4), 37–46, (2008).
- [10] Y. Chevaleyre, U. Endriss, and N. Maudet, ‘Allocating goods on a graph to eliminate envy’, in *Proc. of AAAI-2007*, (2007).
- [11] B. de Keijzer, S. Bouveret, T. Klos, and Y. Zhang, ‘On the complexity of efficiency and envy-freeness in fair division of indivisible goods with additive preferences’, in *Proc. of ADT-2009*, (2009).
- [12] J. Edmonds and R. M. Karp, ‘Theoretical improvements in algorithmic efficiency for network flow problems’, *J. ACM*, 19(2), 248–264, (1972).
- [13] D. K. Foley, ‘Resource allocation and the public sector’, *Yale Economic Essays*, 7(1), 45–98, (1967).
- [14] D. Herreiner and C. Puppe, ‘A simple procedure for finding equitable allocations of indivisible goods’, *Social Choice and Welfare*, 19(2), 415–430, (2002).
- [15] R. J. Lipton, E. Markakis, E. Mossel, and A. Saberi, ‘On approximately fair allocations of indivisible goods’, in *Proc. EC-2004*. ACM Press.
- [16] H. Moulin, *Axioms of Cooperative Decision Making*, Cambridge University Press, 1988.
- [17] A. K. Sen, ‘Welfare, preference and freedom’, *Journal of Econometrics*, 50(1–2), 15–29, (1991).