

Efficiency and Envy-freeness in Fair Division of Indivisible Goods: Logical Representation and Complexity

Sylvain Bouveret

ONERA Centre de Toulouse.

2, avenue Édouard Belin, BP74025.

31055 Toulouse cedex 4, FRANCE.

SYLVAIN.BOUVERET@ONERA.FR

Jérôme Lang

IRIT-CNRS. 118, route de Narbonne.

31062 Toulouse cedex, FRANCE.

LANG@IRIT.FR

Abstract

We consider the problem of allocating fairly a set of indivisible goods among agents from the point of view of compact representation and computational complexity. We start by assuming that agents have dichotomous preferences expressed by propositional formulae. We express efficiency and envy-freeness in a logical setting, which reveals unexpected connections to nonmonotonic reasoning. Then we identify the complexity of determining whether there exists an efficient and envy-free allocation, for several notions of efficiency, when preferences are represented in a succinct way (as well as restrictions of this problem). We first study the problem under the assumption that preferences are dichotomous, and then in the general case.

1. Introduction

Allocating goods to agents is an important issue that has been considered from different perspectives in economics (especially social choice theory) and in computer science (especially Artificial Intelligence and Operations Research), and arises in various real-world settings: auctions, divorce settlements, electronic spectrum and frequency allocation, airport traffic management, or the fair and efficient exploitation of Earth Observation Satellites (see the survey from Chevaleyre, Dunne, Endriss, Lang, Lemaître, Maudet, Padget, Phelps, Rodríguez-Aguilar, & Sousa, 2006, for a detailed description). The general issue also covers a huge variety of allocation problems, depending on the following parameters (see again the work from Chevaleyre et al., 2006, for a detailed taxonomy):

- the nature of the resources to be allocated (are they divisible or not? single-unit or multi-unit?);
- the nature of the preferences of the agents (are they numerical or simply ordinal? can there be preferential dependencies between goods?);
- the nature of the permitted allocations (can goods be shared among several agents? do all goods have to be allocated? can allocations be accompanied by side payments?);
- the evaluation of the quality of an allocation (Pareto-efficiency, utilitarian or egalitarian social welfare etc.);

- the nature of the process that leads to the allocation (centralized or decentralized).

For instance, standard combinatorial auctions (Cramton, Shoham, & Steinberg, 2005) typically correspond to indivisible goods (possible multi-unit), numerical preferences with possible dependencies between goods, monetary payments, maximization of the total value of sold items, and a centralized computation.

In this paper we focus on *fair division of indivisible goods without money transfers*¹. Fair division makes a prominent use of *ex-post* fairness criteria such as *equity* and *envy-freeness*, and on this point departs from auctions, that rather focus on other kinds of fairness (as well as on efficiency of the procedure), such as truthful mechanisms, or fairness of the procedure itself (Brams & Taylor, 1996; Young, 1995). A key concept in the literature on fair division is *envy-freeness*: an allocation is envy-free if and only if each agent likes her share at least as much as the share of any other agent. Ensuring envy-freeness is considered a desirable property; however, envy-freeness alone does not suffice as a criterion for finding satisfactory allocations (this is especially obvious if it is not compulsory to allocate all goods: in this case, not allocating anything to anyone results in an envy-free allocation, yet totally unsatisfactory), therefore it has to be paired with some efficiency criterion, such as Pareto optimality or maximum social welfare. However, it is known that for any reasonable notion of efficiency, there are profiles for which no efficient and envy-free allocation exists² (Brams, Edelman, & Fishburn, 2003). This is even trivial if every good must be assigned to someone: in this case, there are profiles for which not even an envy-free allocation exists. Another well-known notion of fairness (that we do not consider in this paper, except for one result) is Rawlsian egalitarianism, which says that an allocation is equitable if it maximizes the satisfaction of the least satisfied agent. Unlike envy-freeness, egalitarianism requires interpersonal comparability of preferences.

Whereas social choice theory has developed an important literature on fair division, computational issues have rarely been considered. On the other hand, artificial intelligence has studied these issues extensively, but until now has focused mainly on combinatorial auctions and related classical utilitarian problems. Combinatorial auctions, aiming at maximizing the auctioneer’s revenue (consisting in the sum of the prices paid by all agents), are only a specific form of allocation process, namely a pure utilitarian form with money transfers in which considerations such as equity and fairness are not relevant. The literature on combinatorial auctions and related problems has been investigating issues such as compact representation (so as to allow agents to express their bids in a concise way – see the work from Nisan, 2005, for an overview) as well as computational complexity, algorithms, tractable subclasses and approximation. Complexity issues for negotiation (where agents exchange goods by means of deals) have also been studied (*e.g.* Dunne, Wooldridge, & Laurence, 2005; Chevaleyre, Endriss, Estivie, & Maudet, 2004).

The above discussion reveals the existence of a gap (summarized in Table 1): compact representation and complexity issues for fair division have received very little attention until now, apart of the recent work (Lipton, Markakis, Mossel, & Saberi, 2004) which

1. Note that the possibility of money transfers reintroduces divisibility to some extent, considering money as a particular – and divisible – good.
 2. Consider for example the situation where we have a single item and two agents who both want it: in that case, any allocation is either efficient but not envy-free (when the item is given to one of the two agents), or envy-free but not efficient (if the item is not allocated to anyone).

studies approximation schemes for envy-freeness. The need for compact representation arises from the following dilemma, formulated by several social choice theorists: either (a) allow agents to express any possible preference relation on the set of all subsets of items, and end up with an exponentially large representation (which actually happens for example³ in the work from Herreiner & Puppe, 2002); or (b) severely restrict the set of expressible preferences, typically by assuming additive independence between items, and then design procedures where agents express preferences between single items, these preferences being then extended to sets of items by assuming additivity, thus giving up the possibility of expressing preferential dependencies such as complementarity and substitutability effects among items; this is the path followed by Brams et al. (2003), Brams and King (2005) and Demko and Hill (1998). Yet, as we advocate in this paper, conciliating conciseness and expressivity is possible, by means of *compact representation*.

	AXIOMATIC STUDY	COMPUTATIONAL STUDY
AUCTIONS (and related problems)	economics	computer science (especially AI)
FAIR DIVISION	economics (especially social choice)	?

Table 1: Computational issues in fair division.

As in most works on fair allocation of indivisible items, we focus on the joint search for envy-freeness and efficiency. The impossibility to guarantee the existence of an efficient envy-free allocation implies that determining whether there exists such an allocation is a crucial task, since a positive answer leads to choose such an allocation whereas a negative answer calls for a relaxation of one of the criteria, which has been investigated in few papers (Lipton et al., 2004; Chevaleyre, Endriss, & Maudet, 2007b).

We consider the problem of determining whether there exists an efficient and envy-free allocation from the point of view of compact representation and computational complexity. First, since in most cases agents have preferential dependencies (or synergies) between goods, we raise the issue of how a fair division problem of indivisible goods should be expressed. We focus first in the simple case where agents have *dichotomous preferences*, that is, they simply express a partition between satisfactory and unsatisfactory shares. The interest of such a restriction is that in spite of the expressivity loss it imposes, it will be shown to be no less complex than the general case, while being much simpler to expose. Dichotomous preferences have been considered before in social choice contexts, such as by Bogomolnaia, Moulin, and Stong (2005) in a fair division context, and of course in *approval voting* (Brams & Fishburn, 1978), where every voter specifies a dichotomous preference on the set of candidates. The most natural representation of a dichotomous preference (with preferential dependencies between formulae – otherwise the problem is trivial) is by a single

3. Quoting from the work of Brams et al. (2003): “Herreiner and Puppe (...) assume that each person has a linear preference order in 2^B . This allows for complementarity and substitutability effects among items (...). In view of interdependencies that may beset subset evaluations (...), the procedures of Herreiner and Puppe offer a creative way of dealing with subset preference. On the other hand, the sheer number of subsets (more than a million when $n = 20$) and their presumption of clear preference between subsets, could detract from the practicability of their procedures.”

propositional formula, where variables correspond to goods. Expressing envy-freeness and efficiency within this logical representation reveals unexpected connections to nonmonotonic reasoning; this issue will be addressed in Section 3.

The following Sections are devoted to a detailed complexity study of the following problem: “given some fair division problem, does there exist an efficient end envy-free allocation?”. The latter problem will be studied for different notions of efficiency and various restrictions. We start (in Sections 3 and 4) by assuming that preferences are dichotomous, and we identify the complexity of the existence of an envy-free and Pareto-efficient allocation, which turns out to be Σ_2^P -complete. Then we consider several restrictions of the latter problem, namely, (a) fixing the number of agents to two; (b) forcing all agents to have identical preferences; (c) restricting the syntax of the propositional formulae expressing the preferences of the agents. Then we study variations of the problem obtained by replacing Pareto-efficiency by other notions of efficiency, namely: (a) asking for complete allocations (such that every good is allocated to an agent); (b) requiring that a maximum number of agents be satisfied. In Section 5 we then consider the more general problem obtained by removing the assumption that preferences are dichotomous. As the problem then becomes very dependent on the choice of a particular language for compact preference representation, we pick a particular one (weighted propositional formulae) that extends in a simple way the pure propositional representation considered in Section 3, and identify the complexity of the existence of envy-free and efficient allocations, again for several notions of efficiency. Finally, in Section 6 we sum up our contributions and discuss related work and further issues.

2. Background

In this section we provide some basic concepts and definitions that we will use all along the paper.

2.1 Fair Division Problems

Definition 1 (Fair division problem) *A fair division problem⁴ is a tuple $\mathcal{P} = \langle I, X, \mathcal{R} \rangle$ where*

- $I = \{1, \dots, N\}$ is a set of agents;
- $X = \{x_1, \dots, x_p\}$ is a set of indivisible goods;
- $\mathcal{R} = \langle R_1, \dots, R_N \rangle$ is a preference profile, where each R_i is a reflexive, transitive and complete relation on 2^X .

R_i is the preference relation of agent i . AR_iB is alternatively denoted by $R_i(A, B)$ or by $A \succeq_i B$; we write $A \succ_i B$ (strict preference) for ($A \succeq_i B$ and not $B \succeq_i A$) and $A \sim_i B$ (indifference) for ($A \succeq_i B$ and $B \succeq_i A$).

In addition, R_i is said to be *monotonic* if and only if for all A, B , $A \subseteq B \subseteq X$ implies $B \succeq_i A$. $\mathcal{R} = \langle R_1, \dots, R_N \rangle$ is monotonic if and only if R_i is monotonic for every i .

4. In the following, we will use indifferently the terms fair division and resource allocation.

Definition 2 (Allocation) An allocation for $\mathcal{P} = \langle I, X, \mathcal{R} \rangle$ is a mapping $\pi : I \rightarrow 2^X$ such that for all i and $j \neq i$, $\pi(i) \cap \pi(j) = \emptyset$. If for every $x \in X$ there exists an i such that $x \in \pi(i)$ then π is a complete allocation.

In other words, it is not possible in our framework to give the same good to different agents at the same time, but it is possible to throw away some goods. In this paper, we will focus especially on two desirable properties of the allocations: Pareto-efficiency and envy-freeness.

Definition 3 Let π, π' two allocations. π dominates π' if and only if (a) for all i , $\pi(i) \succeq_i \pi'(i)$ and (b) there exists an i such that $\pi(i) \succ_i \pi'(i)$. π is **(Pareto-) efficient** if and only if there is no π' such that π' dominates π .

Definition 4 An allocation π is **envy-free** if and only if $\pi(i) \succeq_i \pi(j)$ holds for all i and all $j \neq i$.

2.2 Propositional Logic

Let V be a finite set of *propositional variables*. L_V is the propositional language generated from V , the usual connectives \neg, \wedge and \vee and the Boolean constants \top and \perp in the usual way⁵.

An *interpretation* M for L_V is an element of 2^V , *i.e.*, a truth assignment to symbols: for all $x \in V$, $x \in M$ (resp. $x \notin M$) means that M assigns x to true (resp. to false). $\text{Mod}(\varphi) = \{M \in 2^V \mid M \models \varphi\}$ is the set of all models of φ (the satisfaction relation \models is defined as usual, as well as satisfiability and logical consequence).

A *literal* is a formula of L_V of the form x or of the form $\neg x$, where $x \in V$. A formula φ is under *negative normal form* (or NNF) if and only if any negation symbol in φ appears only in literals. Any formula can be turned in polynomial time into an equivalent NNF formula. For instance, $a \wedge \neg(b \wedge c)$ is not under NNF but is equivalent to the NNF formula $a \wedge (\neg b \vee \neg c)$.

A formula is *positive* if it contains no occurrence of the negation symbol. For instance, $a \wedge (b \vee \neg c)$ and $a \vee (\neg a \wedge b)$ are not positive, whereas $a \wedge (b \vee c)$ and $(a \wedge c) \vee (a \wedge b)$ are. \top and \perp are considered positive as well.

Let $\varphi \in L_V$. $\text{Var}(\varphi) \subseteq V$ is the set of propositional variables appearing in φ . For instance, $\text{Var}((a \wedge c) \vee (a \wedge b)) = \{a, b, c\}$ and $\text{Var}(\top) = \emptyset$.

Lastly, if $S = \{\varphi_1, \dots, \varphi_n\}$ is a finite set of formulae then $\bigwedge S = \varphi_1 \wedge \dots \wedge \varphi_n$ is the conjunction of all formulae of S , and $\bigvee S = \varphi_1 \vee \dots \vee \varphi_n$ is the disjunction of all formulae of S .

2.3 Computational Complexity

In this paper we will refer to some complexity classes located in the polynomial hierarchy. BH_2 (also referred to as DP) is the class of all languages of the form $L_1 \cap L_2$ where L_1 is in NP and L_2 in coNP . $\Delta_2^p = \text{P}^{\text{NP}}$ is the class of all languages recognizable by a deterministic Turing machine working in polynomial time using NP oracles. Likewise, $\Sigma_2^p = \text{NP}^{\text{NP}}$. $\Theta_2^p =$

5. Note that connectives \rightarrow and \leftrightarrow are not allowed; this is important for the definition of positive formulae (to come).

$\Delta_2^p[\mathcal{O}(\log n)]$ is the subclass of Δ_2^p of problems that only need a logarithmic number of oracles. See for instance the book by Papadimitriou (1994) for further details.

3. Fair Division Problems with Dichotomous Preferences: Logical Representation

We start by considering in full detail the case where preferences are dichotomous.

Definition 5 R_i is dichotomous if and only if there exists a subset $Good_i$ of 2^X such that for all $A, B \subseteq X$, $A \succeq_i B$ if and only if $A \in Good_i$ or $B \notin Good_i$. $\mathcal{R} = \langle R_1, \dots, R_N \rangle$ is dichotomous if and only if every R_i is dichotomous.

There is an obvious way of representing dichotomous preferences compactly, namely by a propositional formula φ_i (for each agent i) of the language L_X (a propositional symbol v_x for each good x) such that $Mod(\varphi) = Good_i$. Formally:

Definition 6 Let R_i be a dichotomous preference on 2^X , with $Good_i$ its associated subset of 2^X , and φ_i a propositional formula on the propositional language L_X . We say that φ_i represents R_i if and only if $Mod(\varphi_i) = Good_i$.

Clearly, for any dichotomous preference R_i there is a formula φ_i representing R_i : $\varphi_i = \bigvee_{A \in Good_i} \left(\bigwedge_{x \in A} v_x \wedge \bigwedge_{x \notin A} \neg v_x \right)$. Furthermore, this formula is unique up to logical equivalence.

Example 1 $X = \{a, b, c\}$ and $Good_i = \{\{a, b\}, \{b, c\}\}$. Note that R_i is not monotonic. Then $\varphi_i = (a \wedge b \wedge \neg c) \vee (\neg a \wedge b \wedge c)$ represents R_i . So does $\varphi'_i = b \wedge ((a \wedge \neg c) \vee (\neg a \wedge c))$, which is logically equivalent to φ .

An easy but yet useful result:

Proposition 1 Let R_i be a dichotomous preference on 2^X . The following statements are equivalent:

1. R_i is monotonic;
2. $Good_i$ is upward closed, that is, $A \in Good_i$ and $B \supseteq A$ imply $B \in Good_i$.
3. R_i is representable by a positive propositional formula.

Proof (1) \Rightarrow (2). Suppose R_i is monotonic, and let $A \in Good_i$ and $B \supseteq A$. Then we must have $B \succeq_i A$, and therefore $B \in Good_i$ (since $A \in Good_i$).

(2) \Rightarrow (3). Suppose that $Good_i$ is upward closed, and consider the set $\min_{\subseteq}(Good_i)$ of all inclusion-minimal sets in $Good_i$. Then the formula $\varphi_i = \bigvee_{A \in \min_{\subseteq}(Good_i)} \left(\bigwedge_{x \in A} v_x \right)$ represents R_i for the following reasons. For each $B \in Good_i$, there is a set $A \in \min_{\subseteq}(Good_i)$ such that $A \subseteq B$. Thus the corresponding conjunction in φ_i is satisfied, then φ_i is satisfied. Conversely, for each set $B \notin Good_i$, there is no $A \subseteq \min_{\subseteq}(Good_i)$ such that $A \subseteq B$. Therefore, none of the terms in φ_i is satisfied, then φ_i is not satisfied. Moreover, φ_i is clearly a positive propositional formula.

(3) \Rightarrow (1). Suppose that R_i is representable by a positive propositional formula φ_i , and let A and B be two sets of items such that $A \subseteq B$. If $A \notin Good_i$ then clearly $B \succeq_i A$. If $A \in Good_i$, then $A \in Mod(\varphi_i)$. Since φ_i is positive, then $B \in Mod(\varphi_i)$ also. Therefore $B \in Good_i$, and finally $B \succeq_i A$. \square

From now on, we assume that allocation problems \mathcal{P} are represented in propositional form, namely, instead of I , X and \mathcal{R} we only specify $\langle \varphi_1, \dots, \varphi_N \rangle$. I and X are obviously determined from $\langle \varphi_1, \dots, \varphi_N \rangle$. In the following, we will also write the propositional variables corresponding to the items x instead of v_x for short, since it is often unambiguous.

Let $\mathcal{P} = \langle \varphi_1, \dots, \varphi_N \rangle$ be an allocation problem with dichotomous preferences; then for each $i \leq N$, we rewrite φ_i into φ_i^* obtained from φ_i by replacing every variable x by the new symbol x_i . For instance, if $\varphi_1 = a \wedge (b \vee c)$ and $\varphi_2 = a \wedge d$ then $\varphi_1^* = a_1 \wedge (b_1 \vee c_1)$ and $\varphi_2^* = a_2 \wedge d_2$.

Example 2 Consider the following allocation problem: $\varphi_1 = a \vee (b \wedge c)$, $\varphi_2 = a$ and $\varphi_3 = a \vee b$ (therefore $I = \{1, 2, 3\}$ and $X = \{a, b, c\}$). All the formulae are positive, therefore preferences are monotonic. For instance, Good_1 is the set composed of all supersets of $\{a\}$ and of all supersets of $\{b, c\}$. Then we have

$$\begin{aligned}\varphi_1^* &= a_1 \vee (b_1 \wedge c_1); \\ \varphi_2^* &= a_2; \\ \varphi_3^* &= a_3 \vee b_3\end{aligned}$$

For all $i \leq N$, let $X_i = \{x_i, x \in X\}$. An allocation for a standard allocation problem \mathcal{P} corresponds to a model of $V = X_1 \cup \dots \cup X_N$ satisfying *at most* one x_i for each $x \in X$. In other terms, there is a bijective mapping between the set of possible allocations and the models of the following formula: $\Gamma_{\mathcal{P}} = \bigwedge_{x \in X} \bigwedge_{i \neq j} \neg(x_i \wedge x_j)$.

If allocation are required to be complete, then $\Gamma_{\mathcal{P}}$ is replaced by $\Gamma_{\mathcal{P}}^C = \Gamma_{\mathcal{P}} \wedge \bigwedge_{x \in X} (x_1 \vee \dots \vee x_n)$. The rest is unchanged.

Let $V = \{x_i \mid i = 1, \dots, N, x \in X\}$. Any interpretation M of $\text{Mod}(\Gamma_{\mathcal{P}})$ is such that it is never the case that x_i and x_j are simultaneously true for $i \neq j$, therefore we can map $M \in \text{Mod}(\Gamma_{\mathcal{P}})$ to an allocation simply defined by $\pi(i) = \{x \mid M \models x_i\}$. This mapping is obviously bijective; we will write $F(\pi)$ the model of $\text{Mod}(\Gamma_{\mathcal{P}})$ corresponding to allocation π , and of course $F^{-1}(M)$ the allocation corresponding to an interpretation M of $\text{Mod}(\Gamma_{\mathcal{P}})$.

Example 2 (continued) For the allocation problem of Example 2, we have:

$$\begin{aligned}\Gamma_{\mathcal{P}} &= \neg(a_1 \wedge a_2) \wedge \neg(a_1 \wedge a_3) \wedge \neg(a_2 \wedge a_3) \\ &\quad \wedge \neg(b_1 \wedge b_2) \wedge \neg(b_1 \wedge b_3) \wedge \neg(b_2 \wedge b_3) \\ &\quad \wedge \neg(c_1 \wedge c_2) \wedge \neg(c_1 \wedge c_3) \wedge \neg(c_2 \wedge c_3)\end{aligned}$$

The interpretation M such that M only sets a_1 , b_3 , and c_1 to true is clearly a model of $\Gamma_{\mathcal{P}}$. It corresponds to the allocation $F^{-1}(M) = \pi$, with $\pi(1) = \{a, c\}$, $\pi(2) = \emptyset$ and $\pi(3) = \{c\}$.

3.1 Envy-freeness

We now show how the search for an envy-free allocation can be mapped to a model finding problem. Let $\varphi_{j|i}^*$ be the formula obtained from φ_i^* by substituting every symbol x_i in φ_i^* by x_j : for instance, if $\varphi_1^* = a_1 \wedge (b_1 \vee c_1)$ then $\varphi_{2|1}^* = a_2 \wedge (b_2 \vee c_2)$. Notice that obviously, we have $\varphi_{i|i}^* = \varphi_i^*$.

We first give the following lemma, which is easy but yet useful:

Lemma 1 For all i, j , $\pi(j) \in \text{Good}_i$ if and only if $F(\pi) \models \varphi_{j|i}^*$.

In particular, when $i = j$ we have $\pi(i) \in \text{Good}_i$ if and only if $F(\pi) \models \varphi_i^*$.

Proof By definition of F , $\pi(j) \in \text{Good}_i$ if and only if $\{x \mid F(\pi) \models x_j\} \in \text{Good}_i$, that is, $\{x \mid F(\pi) \models x_j\} \models \varphi_i$. The latter relation is equivalent to $\{x_i \mid F(\pi) \models x_j\} \models \varphi_i^*$, and finally to $\{x_j \mid F(\pi) \models x_j\} \models \varphi_{j|i}^*$ by definition of $\varphi_{j|i}^*$, from which we can deduce the result. \square

Using this lemma, we can now map the envy-freeness property to the satisfiability of a logical formula:

Proposition 2 Let $\mathcal{P} = \langle \varphi_1, \dots, \varphi_N \rangle$ be an allocation problem with dichotomous preferences under propositional form, and the formulae $\varphi_{j|i}^*$ and mapping F as defined above. Let

$$\Lambda_{\mathcal{P}} = \bigwedge_{i=1, \dots, N} \left[\varphi_i^* \vee \left(\bigwedge_{j \neq i} \neg \varphi_{j|i}^* \right) \right]$$

Then π is envy-free if and only if $F(\pi) \models \Lambda_{\mathcal{P}}$.

Proof Let π be an allocation. π is not envy-free if and only if there is a pair (i, j) , with $i \neq j$ such that $\pi(j) \succ_i \pi(i)$, that is $\pi(j) \in \text{Good}_i$ and $\pi(i) \notin \text{Good}_i$, which is in turn equivalent to $F(\pi) \models \varphi_{j|i}^*$ and $F(\pi) \not\models \varphi_i^*$ by lemma 1. Therefore π is envy-free if and only if $F(\pi) \models \Lambda_{\mathcal{P}}$. \square

The intuitive meaning of this result is that an allocation π is envy-free if and only if, for every agent i , either i is satisfied by her share $\pi(i)$, or she is not but she envies no one, that is, for every j , she would not be satisfied by j 's share either.

The search for envy-free allocations can thus be reduced to a model finding problem: $\{F^{-1}(M) \mid M \models \Gamma_{\mathcal{P}} \wedge \Lambda_{\mathcal{P}}\}$ is the set of envy-free allocations for \mathcal{P} . Note that, importantly, $\Gamma_{\mathcal{P}} \wedge \Lambda_{\mathcal{P}}$ has a polynomial size (precisely, quadratic) in the size of the input data.

The problem of existence of an envy-free allocation without any other property is not very interesting, because such an allocation always exists : it suffices to consider the allocation that gives an empty share to everyone. However,

- the problem of deciding whether there exists an envy-free allocation satisfying some property expressible by a polynomial size formula (*e.g.* completeness) can be reduced to a satisfiability problem;
- the problem of finding (resp. counting) all envy-free allocations comes down to the problem of finding (resp. counting) all models of $\Lambda_{\mathcal{P}} \wedge \Gamma_{\mathcal{P}}$.

Example 2 (continued) For the allocation problem of Example 2, we have:

$$\begin{aligned} \Lambda_{\mathcal{P}} = & ((a_1 \vee (b_1 \wedge c_1)) \vee (\neg(a_2 \vee (b_2 \wedge c_2)) \wedge \neg(a_3 \vee (b_3 \wedge c_3)))) \\ & \wedge (a_2 \vee (\neg a_1 \wedge \neg a_3)) \\ & \wedge ((a_3 \vee b_3) \vee (\neg(a_1 \vee b_1) \wedge \neg(a_2 \vee b_2))) \end{aligned}$$

$$\text{Mod}(\Gamma_{\mathcal{P}} \wedge \Lambda_{\mathcal{P}}) = \{\{c_1\}, \{c_1, b_3\}, \{c_2, b_3\}, \{c_2\}, \{b_3\}, \{c_3\}, \emptyset\}.$$

There are therefore 7 envy-free allocations, namely $(c, -, -)$, $(c, -, b)$, $(-, c, b)$, $(-, c, -)$, $(-, -, b)$, $(-, -, c)$ and $(-, -, -)$. Note that none of them is complete.

3.2 Efficient Allocations

Similarly as for the envy-freeness property, Pareto-efficiency can be expressed as a logical property. The logical expression of this property requires the definition of a maximal β -consistent subset of a set of formulae.

Definition 7 Let $\Delta = \{\alpha_1, \dots, \alpha_m\}$ be a set of formulae and β be a formula. $S \subseteq \Delta$ is a maximal β -consistent subset of Δ if and only if (a) $\bigwedge S \wedge \beta$ is consistent and (b) there is no S' such that $S \subset S' \subseteq \Delta$ and $\bigwedge S' \wedge \beta$ is consistent. Let $\text{MaxCons}(\Delta, \beta)$ be the set of all maximal β -consistent subsets of Δ . Moreover, we will write $\text{MaxCons}(\Delta)$ the set of all maximal-consistent subsets of Δ , that is, the set $\text{MaxCons}(\Delta, \top)$.

Proposition 3 Let $\mathcal{P} = \langle \varphi_1, \dots, \varphi_N \rangle$ an allocation problem. Let $\Phi_{\mathcal{P}} = \{\varphi_1^*, \dots, \varphi_N^*\}$. Then π is Pareto-efficient for \mathcal{P} if and only if $\{\varphi_i^* \mid F(\pi) \models \varphi_i^*\}$ is a maximal $\Gamma_{\mathcal{P}}$ -consistent subset of $\Phi_{\mathcal{P}}$.

Proof Let π be an allocation. Let $\text{Sat}(\pi)$ be the set of agents satisfied by π , that is, by Lemma 1, the set $\{i \mid F(\pi) \models \varphi_i^*\}$. For all $i \in \text{Sat}(\pi)$, $F(\pi) \models \varphi_i^*$ by definition of $\text{Sat}(\pi)$, and $F(\pi) \models \Gamma_{\mathcal{P}}$ by definition of $F(\pi)$. Therefore $F(\pi) \models \Gamma_{\mathcal{P}} \wedge \bigwedge \{\varphi_i^* \mid i \in \text{Sat}(\pi)\} = \Gamma_{\mathcal{P}} \wedge \bigwedge \{\varphi_i^* \mid F(\pi) \models \varphi_i^*\}$. Therefore $\bigwedge \{\varphi_i^* \mid F(\pi) \models \varphi_i^*\} \wedge \Gamma_{\mathcal{P}}$ is consistent.

By definition, π is Pareto-dominated if and only if there is an allocation π' such that $\text{Sat}(\pi') \supseteq \text{Sat}(\pi)$. Therefore, if π is Pareto-dominated there is a consistent subset $S \subseteq \Phi_{\mathcal{P}}$ (corresponding to $\{\varphi_i^* \mid F(\pi') \models \varphi_i^*\}$) such that $\{\varphi_i^* \mid F(\pi) \models \varphi_i^*\} \subset S$. Moreover, since π' is an allocation, $\bigwedge S \wedge \Gamma_{\mathcal{P}}$ is consistent.

Conversely, let $S \subseteq \Phi_{\mathcal{P}}$ be such that $\{\varphi_i^* \mid F(\pi) \models \varphi_i^*\} \subset S$ and $\bigwedge S \wedge \Gamma_{\mathcal{P}}$ is consistent, and M be a model of $\bigwedge S \wedge \Gamma_{\mathcal{P}}$. By definition, $\pi' = F^{-1}(M)$ is a well-defined allocation, and $\pi'(i) \in \text{Good}_i$ for all $\varphi_i^* \in S$. Since $\{\varphi_i^* \mid F(\pi) \models \varphi_i^*\} \subset S$, we have $\text{Sat}(\pi') \supseteq \text{Sat}(\pi)$. Therefore π is Pareto-dominated. \square

This simple result suggests that efficient allocations can be computed from the logical expression $\Phi_{\mathcal{P}}$ of the problem, namely, by computing the maximal $\Gamma_{\mathcal{P}}$ -consistent subsets of $\Phi_{\mathcal{P}}$; call them $\{S_1, \dots, S_q\}$. Then for each S_i , let $M_i = \text{Mod}(\bigwedge S_i \wedge \Gamma_{\mathcal{P}})$ and let $M = \bigcup_{i=1}^q M_i$. Then $F^{-1}(M)$ is the set of all efficient allocations for $\Phi_{\mathcal{P}}$. Note that there are in general exponentially many maximal $\Gamma_{\mathcal{P}}$ -consistent subsets of $\Phi_{\mathcal{P}}$ (and therefore exponentially many efficient allocations). This can be tempered by (a) there are many practical cases where the number of maximal consistent subsets is small; (b) it is generally not asked to look for *all* efficient allocations; if we look for just one, then this can be done by computing one maximal $\Gamma_{\mathcal{P}}$ -consistent subset of $\Phi_{\mathcal{P}}$.

Example 2 (continued) The maximal $\Gamma_{\mathcal{P}}$ -consistent subsets of $\Phi_{\mathcal{P}}$ are $S_1 = \{\varphi_1^*, \varphi_2^*\}$, $S_2 = \{\varphi_1^*, \varphi_3^*\}$ and $S_3 = \{\varphi_2^*, \varphi_3^*\}$. $\bigwedge S_1 \wedge \Gamma_{\mathcal{P}}$ has only one model: $\{b_1, c_1, a_2\}$. $\bigwedge S_2 \wedge \Gamma_{\mathcal{P}}$ has two models: $\{a_1, b_3\}$ and $\{b_1, c_1, a_3\}$. $\bigwedge S_3 \wedge \Gamma_{\mathcal{P}}$ has one model: $\{a_2, b_3\}$. Therefore the four efficient allocations for \mathcal{P} are $(bc, a, -)$, $(a, -, b)$, $(bc, -, a)$ and $(-, a, b)$. None of them is envy-free.

3.3 Efficient and Envy-free Allocations

We are now in position of putting things together. Since envy-free allocations correspond to the models of $\Lambda_{\mathcal{P}}$ and efficient allocations to the models of maximal $\Gamma_{\mathcal{P}}$ -consistent subsets of $\Phi_{\mathcal{P}}$, the existence of an efficient and envy-free (EEF) allocation is equivalent to the following condition: *there exists a maximal $\Gamma_{\mathcal{P}}$ -consistent subset S of $\Phi_{\mathcal{P}}$ such that $\bigwedge S \wedge$*

$\Gamma_{\mathcal{P}} \wedge \Lambda_{\mathcal{P}}$ is consistent. In this case, the models of the latter formula are the EEF allocations. Interestingly, this is an instance of a well-known problem in nonmonotonic reasoning:

Definition 8 A supernormal default theory⁶ is a pair $D = \langle \beta, \Delta \rangle$ with $\Delta = \{\alpha_1, \dots, \alpha_m\}$, where $\alpha_1, \dots, \alpha_m$ and β are propositional formulae. A propositional formula ψ is a skeptical consequence of D , denoted by $D \sim^{\forall} \psi$, if and only if for all $S \in \text{MaxCons}(\Delta, \beta)$ we have $\bigwedge S \wedge \beta \models \psi$.

Proposition 4 Let $\mathcal{P} = \langle \varphi_1, \dots, \varphi_N \rangle$ be a fair division problem. Let $D_{\mathcal{P}} = \langle \Gamma_{\mathcal{P}}, \Phi_{\mathcal{P}} \rangle$. Then there exists an efficient and envy-free allocation for \mathcal{P} if and only if $D_{\mathcal{P}} \not\sim^{\forall} \neg \Lambda_{\mathcal{P}}$.

Proof Let $\mathcal{P} = \langle \varphi_1, \dots, \varphi_N \rangle$ be a fair division problem. Let π be a Pareto-efficient and envy-free allocation, and $S = \{\varphi_i^* \mid F(\pi) \models \varphi_i^*\}$. Then $F(\pi) \models \Lambda_{\mathcal{P}}$ by Proposition 2. We also have $F(\pi) \models \bigwedge S$ by definition of S , and then $F(\pi) \models \bigwedge S \wedge \Lambda_{\mathcal{P}}$, which proves that $\bigwedge S \wedge \Lambda_{\mathcal{P}}$ is consistent, or, in other terms, $\bigwedge S \not\models \neg \Lambda_{\mathcal{P}}$. Moreover, S is a maximal $\Gamma_{\mathcal{P}}$ -consistent subset of $\Phi_{\mathcal{P}}$ by Proposition 3. Thus $S \in \text{MaxCons}(\Phi_{\mathcal{P}}, \Gamma_{\mathcal{P}})$, and $\bigwedge S \wedge \Gamma_{\mathcal{P}} \not\models \neg \Lambda_{\mathcal{P}}$, which implies $\langle \Gamma_{\mathcal{P}}, \Phi_{\mathcal{P}} \rangle \not\sim^{\forall} \neg \Lambda_{\mathcal{P}}$ by definition of \sim^{\forall} .

Conversely, suppose that $\langle \Gamma_{\mathcal{P}}, \Phi_{\mathcal{P}} \rangle \not\sim^{\forall} \neg \Lambda_{\mathcal{P}}$. Then there is a set $S \in \text{MaxCons}(\Phi_{\mathcal{P}}, \Gamma_{\mathcal{P}})$ such that $\bigwedge S \wedge \Gamma_{\mathcal{P}} \wedge \Lambda_{\mathcal{P}}$ has a model M . By Proposition 3, $F^{-1}(M)$ is a Pareto-efficient allocation (since M is a model of $\bigwedge S \wedge \Gamma_{\mathcal{P}}$), and by Proposition 2, $F^{-1}(M)$ is envy-free (since M is a model of $\Lambda_{\mathcal{P}}$). \square

This somewhat unexpected connection to nonmonotonic reasoning has several implications. First, EEF allocations correspond to the models of $\bigwedge S \wedge \Gamma_{\mathcal{P}} \wedge \Lambda_{\mathcal{P}}$ for $S \in \text{MaxCons}(\Phi_{\mathcal{P}}, \Gamma_{\mathcal{P}})$; however, $\text{MaxCons}(\Phi_{\mathcal{P}}, \Gamma_{\mathcal{P}})$ may be exponentially large, which argues for avoiding to start computing efficient allocations and then filtering out those that are not envy-free, but rather compute EEF allocations in a single step, using default reasoning algorithms. Thus, fair division may benefit from computational work in default logic and connected domains such as belief revision and answer set programming (Baral, 2003; Gebser, Liu, Namasivayam, Neumann, Schaub, & Truszczyński, 2007). Moreover, alternative criteria for selecting extensions in default reasoning (such as cardinality, weights or priorities) correspond to alternative efficiency criteria in allocation problems.

4. Allocation Problems with Dichotomous Preferences: Complexity

We will study in this section the complexity of the problem of the existence of an EEF allocation and some of its restrictions, in the case where preferences are dichotomous, for several notions of efficiency.

4.1 Complexity of the General Problem

It is known that skeptical inference is Π_2^P -complete (Gottlob, 1992); now, after Proposition 4, the problem of the existence of an EEF allocation can be reduced to the complement of a skeptical inference problem, which immediately tells that it is in Σ_2^P . Less obviously, we now show that it is complete for this class, even if preferences are required to be monotonic. To prove hardness, we will use the following restriction of the skeptical inference problem:

Problem 1: RESTRICTED SKEPTICAL INFERENCE (RSI)

INSTANCE: A set of propositional formulae $\Delta = \{\alpha_1, \dots, \alpha_n\}$.

QUESTION: do all maximal-consistent subsets of Δ contain α_1 ?

6. ‘‘Supernormal’’ defaults are also called ‘‘normal defaults without prerequisites’’ (e.g. Reiter, 1980).

Proposition 5 *The problem RSI is Π_2^p -complete.*

Proof Membership to Π_2^p comes easily from the fact that the problem RSI is a restriction of the skeptical inference problem, where the formula φ to infer is simply α_1 . Hardness comes from the fact that from any instance $\langle \Delta, \beta, \psi \rangle$ of the skeptical inference problem, we have $\langle \Delta, \beta \rangle \vdash^\forall \psi$ if and only if $\langle \{\psi\} \cup \Delta, \beta \rangle \vdash^\forall \psi$, if and only if $\langle \{\psi \wedge \beta\} \cup \{\alpha_1 \wedge \beta, \dots, \alpha_n \wedge \beta\}, \top \rangle \vdash^\forall \psi$, if and only if all maximal β -consistent subsets of $\{\psi, \alpha_1, \dots, \alpha_n\}$ contain ψ , which is an instance of RSI. \square

Proposition 6 *The problem EEF EXISTENCE of determining whether there exists an efficient and envy-free allocation for a given problem \mathcal{P} with monotonic, dichotomous preferences under logical form is Σ_2^p -complete.*

We show hardness by the following reduction from $\overline{\text{RSI}}$ (the complement problem of RSI, that is, is there one maximal-consistent subset of Δ that does not contain α_1 ?) to EEF EXISTENCE. Given any finite set Δ of propositional formulae, let $V_\Delta = \text{Var}(\Delta)$ be the set of propositional symbols appearing in Δ , and let $\mathcal{P}(\Delta) = \langle I, X, \Phi_{\mathcal{P}(\Delta)} \rangle$ be the following instance of EEF EXISTENCE:

1. $I = \{1, 2, \dots, n + 3\}$;
2. $X = \{v^i | v \in V_\Delta, i \in 1 \dots n\} \cup \{\bar{v}^i | v \in V_\Delta, i \in 1 \dots n\} \cup \{x^i | i \in 1 \dots n + 1\} \cup \{y\}$;
3. for each $i = 1, \dots, n$, let β_i be obtained from α_i by the following sequence of operations:
 - (i) put α_i into NNF form (let α'_i be the result);
 - (ii) for every $v \in V_\Delta$, replace, in α'_i , each (positive) occurrence of v by v^i and each occurrence of $\neg v$ by \bar{v}^i ; let β_i be the formula obtained. Then:
 - for $i = 1, \dots, n$, $\varphi_i = \beta_i \wedge x^i$,
 - $\varphi_{n+1} = \left(\left(\bigwedge_{v \in V_\Delta} \left(\bigwedge_{i=1}^n v^i \right) \vee \left(\bigwedge_{i=1}^n \bar{v}^i \right) \right) \wedge x^{n+1} \right) \vee y$,
 - $\varphi_{n+2} = y$,
 - $\varphi_{n+3} = \varphi_1$.

We will now prove Proposition 6 using several lemmas.

Lemma 2 *An allocation π for \mathcal{P} is said to be regular if and only if for all $i \leq n + 3$, $\pi(i) \subseteq \sigma(i)$, where*

- for all $i \leq n$, $\sigma(i) = \bigcup_{v \in V_\Delta} \{v^i, \bar{v}^i\} \cup \{x^i\}$;
- $\sigma(n + 1) = \bigcup_{v \in V_\Delta, i=1, \dots, n} \{v^i, \bar{v}^i\} \cup \{x^{n+1}, y\}$;
- $\sigma(n + 2) = \{y\}$.
- $\sigma(n + 3) = \sigma(1)$.

Given some allocation π , let now π_R be defined by $\pi_R(i) = \pi(i) \cap \sigma(i)$. Then

1. π_R is regular;
2. π is efficient if and only if π_R is efficient;

3. if π is envy-free then π_R is envy-free.

Proof (1) is obvious. For all i , the goods outside $\sigma(i)$ do not have any influence on the satisfaction of i (since they do not appear in φ_i), therefore $\pi_R(i) \sim_i \pi(i)$, from which (2) follows. The formulae φ_i being positive, the preference relations \succeq_i are monotonic, therefore $\pi(j) \succeq_i \pi_R(j)$ holds for all i, j . Now, if π is envy-free then for all i, j we have $\pi(i) \succeq_i \pi(j)$, therefore $\pi_R(i) \sim_i \pi(i) \succeq_i \pi(j) \succeq_i \pi_R(j)$ and therefore π_R is envy-free, from which (3) follows. \square

Lemma 3 *If π is regular then*

1. 1 can only envy $n + 3$;
2. $n + 3$ can only envy 1;
3. $2, \dots, n$ envy no one;
4. $n + 1$ can only envy $n + 2$;
5. $n + 2$ can only envy $n + 1$;

Proof First, note that for any $i, j \neq i$, i envies j if and only if $\pi(i) \models \neg\varphi_i$ and $\pi(j) \models \varphi_i$.

1. Let $i = 1$ and $j \in \{2, \dots, n, n + 2\}$. If 1 envies j , then $x^1 \in \pi(j)$. π being regular, $x^1 \notin \pi(j)$, therefore i cannot envy j .
2. Since $\varphi_{n+3} = \varphi_1$, the same holds for agent $n + 3$.
3. Let $i \in \{2, \dots, n\}$ and $j \neq i$. If i envies j then $\pi(j) \models \beta_i \wedge x^i$, which is impossible because $x^i \notin \pi(j)$, due to the regularity of π .
4. Assume $n + 1$ envies j for $j \in \{1, \dots, n, n + 3\}$. Then $\pi(j) \models \varphi_{n+1}$. Since $\pi(j) \models y$ is impossible (because π is regular), we have $\pi(j) \models \left(\bigwedge_{v \in V_\Delta} \left(\bigwedge_{i=1}^n v^i \right) \vee \left(\bigwedge_{i=1}^n \bar{v}^i \right) \right) \wedge x^{n+1}$, thus $\pi(j) \models x^{n+1}$ which is impossible as well, since π is regular.
5. Let $i = n + 2$ and $j \in \{1, \dots, n, n + 3\}$. If i envies j then $\pi(j) \models y$, which is impossible because π is regular. \square

Lemma 4 *Let π be a regular allocation satisfying $n + 1$ and $n + 2$ and leaving 1 and $n + 3$ unsatisfied. Let $M(\pi)$ be the interpretation on V_Δ obtained from π by: for all $v \in V_\Delta$, $M(\pi) \models v$ (i.e., $v \in M(\pi)$) if $n + 1$ receives $\bar{v}^1, \dots, \bar{v}^n$, and $M(\pi) \models \neg v$ otherwise, i.e., if $n + 1$ receives v^1, \dots, v^n . Then π is efficient and envy-free only if $M(\pi) \not\models \alpha_1$.*

Proof Let π be a regular allocation satisfying $n + 1$ and $n + 2$. Since π satisfies $n + 2$, $y \in \pi(n + 2)$. Now, π satisfies $n + 1$ without giving her y , therefore, for any v , $n + 1$ receives either all the v^i 's or all the \bar{v}^i 's. This shows that our definition of $M(\pi)$ is well-founded.

Now assume that π is efficient and envy-free, and suppose that $M(\pi) \models \alpha_1$. One of the agents 1 and $n + 3$ is satisfiable without spoiling any agent $j \notin \{1, n + 3\}$, by giving her $\{x^1\} \cup \bigcup \{\bar{v}^i \mid M(\pi) \models \neg v^i\} \cup \bigcup \{v^i \mid M(\pi) \models v^i\}$. Then since π is efficient it must satisfy at least one of 1 and $n + 3$. It cannot satisfy them simultaneously (because of x^1). Thus only one among 1 and $n + 3$ is satisfied by π , the other one envying her. It contradicts the envy-freeness of π , thus proving that $M(\pi) \not\models \alpha_1$. \square

Lemma 5 *For each interpretation M over V_Δ , let us define $\pi_M : I \rightarrow 2^X$ by:*

- for each $i \in 1, \dots, n$, $\pi_M(i) = \{v^i \mid M \models v\} \cup \{\bar{v}^i \mid M \models \neg v\} \cup \{x^i\}$;
- $\pi_M(n + 1) = \{x^{n+1}\} \cup \{\bar{v}^i \mid M \models v, i = 1, \dots, n\} \cup \{v^i \mid M \models \neg v, i = 1, \dots, n\}$;
- $\pi_M(n + 2) = \{y\}$

- $\pi_M(n+3) = \emptyset$.

Then:

1. π_M is a well-defined and regular allocation satisfying $n+1$ and $n+2$;
2. $M(\pi_M) = M$ ($M(\pi_M)$ is obtained from π_M as in Lemma 4).
3. for any $i \in 1, \dots, n$, π_M satisfies i if and only if $M \models \alpha_i$.
4. π_M is efficient if and only if M satisfies a maximal consistent subset of Δ .

Proof 1. One can easily check that π_M does not give the same good to more than one individual, and that π_M can only give to each agent i a set of items that are in $\sigma(i)$. Therefore it is a well-defined and regular allocation. This allocation obviously satisfies $n+1$ and $n+2$.

2. If $M \models v$ then $\pi_M(n+1)$ contains $\{\bar{v}^i \mid i = 1, \dots, n\}$ and therefore $M(\pi_M) \models v$. The case $M \models \neg v$ is similar.

3. Let $i \in 1, \dots, n$. Since π_M gives x^i to i , π_M satisfies i if and only if $F(\pi_M(i)) \models \beta_i$, which is equivalent to $M \models \alpha_i$.

4. From point 3, $\{i \mid \pi_M \text{ satisfies } i\} = \{i \mid M \models \alpha_i\} \cup \{n+1, n+2\}$ (obviously, $n+3$ is not satisfied). Now, since preferences are dichotomous, an allocation π is efficient if and only if the set of individuals it satisfies is maximal with respect to inclusion. Therefore, π_M is efficient if and only if M satisfies a maximal consistent subset of Δ . \square

Lemma 6 *Let π be a regular and efficient allocation satisfying $n+1$ and $n+2$. Then $M(\pi)$ satisfies a maximal consistent subset of Δ .*

Proof π is regular and satisfies $n+1$ and $n+2$, therefore obviously $\pi(n+2) = \{y\}$ and by lemma 4 $M(\pi)$ is well-defined. We will now consider the allocation $\pi_{M(\pi)}$, defined from π like in the previous lemmas. We have $\pi_{M(\pi)}(n+1) = \{x^{n+1}\} \cup \{v^i \mid M(\pi) \models \neg v\} \cup \{\bar{v}^i \mid M(\pi) \models v\} = \{x^{n+1}\} \cup \{v^i \mid \{v^1, \dots, v^n\} \subset \pi(n+1)\} \cup \{\bar{v}^i \mid \{\bar{v}^1, \dots, \bar{v}^n\} \subset \pi(n+1)\}$. Since $n+1$ is satisfied by π , $\pi(n+1)$ must contain $\{x^{n+1}\}$, from what we assert that $\pi_{M(\pi)}(n+1) \subseteq \pi(n+1)$.

Let $i \in \{2, \dots, n\}$. Since π is regular, we have $\pi(i) \subset \sigma(i)$. Since $\pi_{M(\pi)}$ is a complete allocation by definition, and regular by lemma 5, $\sigma(i) \subset \pi_{M(\pi)}(i) \cup \pi_{M(\pi)}(n+1)$. Since $\pi_{M(\pi)}(n+1) \subseteq \pi(n+1)$ we have $\sigma(i) \subset \pi_{M(\pi)}(i) \cup \pi(n+1)$, and thus $\pi(i) \subset \pi_{M(\pi)}(i) \cup \pi(n+1)$. π being an allocation, we have of course $\pi(i) \cap \pi(n+1) = \emptyset$, from what follows $\pi(i) \subseteq \pi_{M(\pi)}(i)$.

π being regular, we have $\pi(1) \cup \pi(n+3) \subseteq \sigma(1) \cup \sigma(n+3)$. Since $\sigma(1) = \sigma(n+3)$, the latter inclusion comes down to $\pi(1) \cup \pi(n+3) \subseteq \sigma(1)$. Now, we have $\sigma(1) \subset \pi_{M(\pi)}(1) \cup \pi_{M(\pi)}(n+1) \cup \pi_{M(\pi)}(n+3)$ for similar reasons as for $i \in \{2, \dots, n\}$, which, together with $\pi_{M(\pi)}(n+1) \subseteq \pi(n+1)$ and $\pi_{M(\pi)}(n+3) = \emptyset$, comes down to $\sigma(1) \subset \pi_{M(\pi)}(1) \cup \pi(n+1)$, and, with $\pi(1) \cap \pi(n+1) = \emptyset$, we deduce the inclusion $\pi(1) \cup \pi(n+3) \subseteq \pi_{M(\pi)}(1)$.

We can now prove that $\pi_{M(\pi)}$ is efficient. Since preferences are monotonic, all individuals but $n+3$ satisfied by π are satisfied by π' as well (since $\forall i \neq n+3, \pi(i) \subseteq \pi_{M(\pi)}(i)$).

- If $n+3$ was not satisfied by π , then immediately we deduce that $\pi_{M(\pi)}$ is efficient.
- If $n+3$ was satisfied by π , suppose that $\pi_{M(\pi)}$ is not efficient. In that case, there is an allocation π' such that $\forall i, \pi_{M(\pi)}$ satisfies i implies π' satisfies i and there is a particular $j \neq 1$ such that π' satisfies j and $\pi_{M(\pi)}$ does not satisfy j . Clearly, π' satisfies 1 (since $\pi_{M(\pi)}$ does), thus $j \neq n+3$ (satisfying simultaneously 1 and $n+3$ is impossible). Consider the allocation π'' deduced from π' just from swapping the shares of 1 and $n+3$. We have, for all $i \in \{2, \dots, n+2\}$, π satisfies i implies $\pi_{M(\pi)}$ satisfies i implies π' satisfies i implies π'' satisfies i . We also have π satisfies $n+3$ and does not satisfy 1, which is the same for π'' . Thus π satisfies i implies π'' satisfies i for all i . Moreover, π'' satisfies $j \in \{2, \dots, n+2\}$ (the same j as above) while $\pi_{M(\pi)}$ does not, and therefore neither does π . It proves that π is Pareto-dominated, which is contradictory with the hypotheses.

Therefore $\pi_{M(\pi)}$ is efficient, from which we conclude, together with lemma 5 (point 4), that $M(\pi)$ satisfies a maximal consistent subset of Δ . \square

Lemma 7 *Any envy-free and efficient allocation for $\mathcal{P}(\Delta)$ satisfies $n + 1$ and $n + 2$, and leaves 1 and $n + 3$ unsatisfied.*

Proof Suppose that π does not satisfy $n + 1$; then $y \notin \pi(n + 1)$; now, if $y \in \pi(n + 2)$ then $n + 1$ envies $n + 2$; if $y \notin \pi(n + 2)$ then π is not efficient because giving y to $n + 2$ would satisfy $n + 2$ and thus lead to a better allocation than π .

Now, suppose π does not satisfy $n + 2$, that is, $y \notin \pi(n + 2)$; if $y \in \pi(n + 1)$ then $n + 2$ envies $n + 1$; if $y \notin \pi(n + 1)$ then again, π is not efficient because giving y to $n + 2$ would satisfy $n + 2$ and thus lead to a better allocation than π .

Concerning agents 1 and $n + 3$, one can notice that since they have identical preferences, any envy-free allocation must either satisfy them both, or leave them both unsatisfied. Since they cannot be simultaneously satisfied (because of x^1), any envy-free allocation leaves them both unsatisfied. \square

Lemma 8 *If there exists an EEF allocation, then there exists a maximal consistent subset of Δ that does not contain α_1 .*

Proof Let π be an efficient and envy-free allocation. By Lemma 2, π_R is regular, efficient and envy-free. By Lemma 7, π_R satisfies $n + 1$ and $n + 2$ and leaves 1 and $n + 3$ unsatisfied. Then by Lemma 6, $M(\pi_R)$ satisfies a maximal consistent subset of Δ , and by Lemma 4, $M(\pi_R) \not\models \alpha_1$. Therefore $\{\alpha_i \in \Delta \mid M(\pi_R) \models \alpha_i\}$ is a maximal consistent subset of Δ and does not contain α_1 . \square

Lemma 9 *If there exists a maximal consistent subset of Δ that does not contain α_1 then there exists an EEF allocation.*

Proof Assume that there exists a maximal consistent subset S of Δ that does not contain α_1 , and let M be a model of $\bigwedge S$. By point 4 of Lemma 5, π_M is efficient.

By point 1 of Lemma 5, π_M is regular; then by Lemma 3, π_M is envy-free if and only if (i) 1 does not envy $n + 3$, (ii) $n + 3$ does not envy 1 (iii) $n + 1$ does not envy $n + 2$ and (iv) $n + 2$ does not envy $n + 1$. By definition of π_M , π_M does not satisfy $n + 3$, hence (i) holds. By point 3 of Lemma 5, $M \not\models \alpha_1$ implies that π_M does not satisfy 1, therefore (ii) holds as well. And finally, by point 1 of lemma 4, $n + 1$ and $n + 2$ are satisfied by π_M , thus (iii) and (iv) also hold. Therefore π_M is envy-free. \square

We are now in position of putting things together and proving Proposition 6:

Proof (Proposition 6) From Lemmas 8 and 9, the existence of a maximal consistent subset of Δ that does not contain α_1 and the existence of an efficient and envy-free allocation for $\mathcal{P}(\Delta)$ are equivalent. Clearly, $\mathcal{P}(\Delta)$ is computed in polynomial time. Therefore, \mathcal{P} is a polynomial reduction from $\overline{\text{RSI}}$ to EEF EXISTENCE, which shows that the latter problem is Σ_2^P -hard, and therefore Σ_2^P -complete. \square

As a corollary, this Σ_2^P -completeness result holds for general (not necessarily monotonic) dichotomous preferences:

Corollary 1 *EEF EXISTENCE for general, dichotomous preference under logical form is Σ_2^P -complete.*

4.2 Restrictions on the Language

As a consequence of this high complexity, it is worth studying restrictions and variants of the latter problem for which complexity may fall down. We will investigate three kind of intuitive restrictions of the EEF EXISTENCE problem, defined by:

- fixing the number of agents, and especially, restricting the problem to the case where there are only 2 agents;

- forcing the agents to have identical preferences;
- restricting the syntax of the agents' goals, by limiting their expression to some subclasses of propositional formulae (e.g. clauses, cubes, ...).

Contrary to the general EEF EXISTENCE problem, the complexity of these restrictions may be sensitive to whether the preferences are monotonic or not.

4.2.1 IDENTICAL PREFERENCES

We start by considering *identical* dichotomous preference profiles, that is, all agents have the same preference, i.e. the same formula φ .

Proposition 7 EEF EXISTENCE *with N identical dichotomous, monotonic preferences is NP-complete. This result holds even for a fixed number of agents $N \geq 2$.*

Proof If preferences are identical, an envy-free allocation satisfies either *all agents* or *none*; now, if preferences are monotonic, it is always possible to satisfy at least one agent (by giving her all the items). Therefore, an allocation is EEF if and only if it satisfies all the agents. Clearly, it can be checked in polynomial time that a given allocation satisfies all agents, hence membership to NP.

Hardness comes from a simple reduction from SET SPLITTING:

Problem 2: SET SPLITTING

INSTANCE: A collection $C = \{X_1, \dots, X_n\}$ of subsets of a finite set S .

QUESTION: Is there a partition $\langle S_1, S_2 \rangle$ of S such that no subset in C is entirely contained in either S_1 or S_2 ?

Given an instance $\langle C, S \rangle$ of SET SPLITTING, let $\mathcal{P}(C, S)$ be the following EEF EXISTENCE instance:

Agents: 2 agents,

Objects: one object $x(a)$ per element $a \in S$,

Preferences: $\varphi_1 = \varphi_2 = \bigwedge_{X_i \in C} \bigvee_{a \in X_i} x(a)$ (and as usual $\varphi_k^* = \bigwedge_{X_i \in C} \bigvee_{a \in X_i} x_k(a)$): each agent wants at least one object from each set.

It is easy to see that if there is a set splitting $\langle S_1, S_2 \rangle$ of $\langle C, S \rangle$, it is possible to find an allocation that satisfy the two agents, by giving them respectively $x(S_1)$ and $x(S_2)$. Conversely, suppose that there is an efficient and envy-free allocation π , then this allocation must satisfy the two agents. Let $\langle S_1, S_2 \rangle = \langle x^{-1}(\pi(1)), x^{-1}(\pi(2)) \rangle$. Suppose that there is an $X_i \in C$ such that $X_i \subset S_1$ or $X_i \subset S_2$ (say e.g. $X_i \subset S_1$). Then $\bigvee_{a \in X_i} x_2(a)$ is false, thus making φ_2^* false, which is contradictory with the initial hypothesis. Therefore $\langle S_1, S_2 \rangle$ is a set splitting of $\langle C, S \rangle$.

This is clearly a polynomial-time reduction, hence NP-hardness of EEF EXISTENCE with 2 agents having identical dichotomous and monotonic preferences. \square

Unlike Proposition 6, Proposition 7 is sensitive to whether preferences are required to be monotonic or not.

Proposition 8 EEF EXISTENCE *with N identical dichotomous preferences is coBH_2 -complete. This result holds even for a fixed number of agents $N \geq 2$.*

Proof If preferences are identical, an envy-free allocation satisfies either *all agents* or *none*. Let φ be the formula representing one agent's preferences (of course φ is identical for all the agents). If φ is satisfiable then it is possible to satisfy at least one agent. In that case, an allocation π is EEF if and only if π satisfies all the agents. If φ is not satisfiable, then every allocation is EEF. Therefore, there exists an EEF allocation if and only if $\Gamma \wedge \varphi_1^* \dots \dots \varphi_N^*$ is satisfiable or φ is unsatisfiable. This shows membership to coBH_2 .

Hardness comes from a simple reduction from SAT-OR-UNSAT. Let $I = \langle \varphi, \psi \rangle$ be a pair of propositional formulae, assumed w.l.o.g. to have no variables in common; we map I to the following allocation problem:

Agents: 2 agents;
Objects: 2 objects v and v' per propositional variable v appearing in φ , one object w per propositional variable w appearing in ψ , and one object y ;
Preferences: $\varphi_1 = \varphi_2 = \varphi \vee \varphi' \vee (y \wedge \psi)$ where φ' denotes the formula φ where each variable v has been replaced by v' .

1. Suppose that φ is not satisfiable, but ψ is (corresponding to a negative instance of SAT-OR-UNSAT). Then it is possible to satisfy at least one agent by giving her y and the objects v corresponding to the variables assigned to true in the model of ψ . However it is not possible to satisfy simultaneously the second agent because φ is unsatisfiable (and by the way so is φ'), and the first agent has already taken y . Therefore in that case there is no EEF allocation.
2. Suppose now that φ is satisfiable or ψ is unsatisfiable (corresponding to a positive instance of SAT-OR-UNSAT). There are two cases:
 - φ is satisfiable. In that case, no matter whether ψ is satisfiable or not, it is possible to satisfy the two agents by satisfying simultaneously φ for the first one and φ' for the second one. Consequently there is an EEF allocation.
 - φ and ψ are both unsatisfiable (recall that the case φ unsatisfiable and ψ satisfiable is covered by point 1). In that case it is clearly impossible to satisfy any agent. In that case, the empty allocation is efficient and envy-free.

Therefore there is an EEF allocation if and only if φ is satisfiable or ψ is unsatisfiable, which proves the proposition. \square

4.2.2 TWO AGENTS

Note that for the two previous results, the hardness result holds for any *fixed* number of agents (≥ 2). Things are different with Proposition 6, for which hardness does not hold when N is fixed. Namely, we have the following results.

Proposition 9 EEF EXISTENCE for two agents with monotonic dichotomous preferences is NP-complete.

Proof Hardness is a corollary of Proposition 7. Membership is obtained as follows. Let $\langle \varphi_1, \varphi_2 \rangle$ be the preference profile, where φ_1, φ_2 are both positive. The formulae Γ, Λ and the formulae φ_i^* are defined as earlier (see section 3), as well as $F(\pi)$ for all allocation π . Now π is efficient if and only if either (a) it satisfies both agents, or (b) it satisfies only one agent, and $\Gamma \wedge \varphi_1^* \wedge \varphi_2^*$ is unsatisfiable, or (c) it is impossible to satisfy even one agent, i.e., both φ_1 and φ_2 are unsatisfiable. (c) is impossible because φ_1, φ_2 are positive. Now, π is envy-free if and only if $M_\pi \models \Lambda$. Therefore, π is EEF if and only if (a) $M_\pi \models \Gamma \wedge \varphi_1^* \wedge \varphi_2^*$ or (b) $M_\pi \models \Gamma \wedge (\varphi_1^* \vee \varphi_2^*) \wedge \Lambda$. Thus, there exists an EEF allocation if and only if $(\Gamma \wedge \varphi_1^* \wedge \varphi_2^*) \vee (\Gamma \wedge (\varphi_1^* \vee \varphi_2^*) \wedge \Lambda)$ is satisfiable, hence membership to NP. \square

Proposition 10 EEF EXISTENCE for 2 agents with dichotomous preferences is coBH₂-complete.

Proof Membership comes from the following reduction to SAT-OR-UNSAT. Consider an instance \mathcal{P} of the EEF problem with 2 agents having respectively the preferences φ_1 and φ_2 . We translate this instance to an instance $\langle \psi, \psi' \rangle$ of SAT-OR-UNSAT defined as follows (Γ is defined as usual): $\psi = (\Gamma \wedge \varphi_1^* \wedge \varphi_2^*) \vee (\varphi_1 \wedge \neg \varphi_2) \vee (\neg \varphi_1 \wedge \varphi_2)$ and $\psi' = (\varphi_1 \vee \varphi_2)$. We will prove that ψ is satisfiable or ψ' is unsatisfiable if and only if there is an EEF allocation for \mathcal{P} .

1. Suppose that ψ is not satisfiable, but ψ' is. Since ψ is not satisfiable, no valid allocation can satisfy both agents (because $\Gamma \wedge \varphi_1^* \wedge \varphi_2^*$ is not satisfiable). Since ψ' is satisfiable, it is possible to satisfy at least one agent (because $\varphi_1 \vee \varphi_2$ is satisfiable). We can then deduce that every efficient allocation satisfies exactly one agent. Since $(\varphi_1 \wedge \neg \varphi_2) \vee (\neg \varphi_1 \wedge \varphi_2)$ is not satisfiable, $Mod(\varphi_1) = Mod(\varphi_2)$ (in other words, φ_1 and φ_2 are logically equivalent). Let π be an allocation

satisfying the agent 1 (the case is similar with agent 2), $F(\pi) \models \varphi_1^*$, and thus $F(\pi) \models \varphi_{2|1}^*$. Since $F(\pi) \not\models \varphi_2^*$ (because it is impossible to satisfy both agents), $F(\pi) \not\models \Lambda$ and thus π is not envy-free. Hence every efficient allocation raises envy: there is no EEF allocation.

2. Suppose now that ψ is satisfiable or ψ' is not satisfiable.
 - ψ' is not satisfiable. This case is easy because it is not possible to satisfy even one agent. Therefore every allocation is efficient and envy free.
 - φ and ψ are both satisfiable. There are two cases:
 - $\Gamma \wedge \varphi_1 \wedge \varphi_2$ is satisfiable. In this case, there is an allocation, corresponding to the model of $\Gamma \wedge \varphi_1 \wedge \varphi_2$, that satisfies both agents. This allocation is clearly EEF.
 - $\varphi_1 \wedge \neg \varphi_2$ is satisfiable but $\Gamma \wedge \varphi_1 \wedge \varphi_2$ is not satisfiable (the case with $\varphi_2 \wedge \neg \varphi_1$ is similar). In that case it is not possible to satisfy both agents. However, since ψ is satisfiable, it is possible to satisfy at least one, and, like in point 1, every efficient allocation satisfies exactly one agent. Since $\varphi_1 \wedge \neg \varphi_2$ is satisfiable, there is a model of φ_1 which is not a model of φ_2 . The allocation corresponding to this model is such that the agent 1 is satisfied and not the agent 2, but the latter agent cannot envy the first one.

It finally proves the correctness of the reduction, which is clearly polynomial.

Hardness comes directly from Proposition 8. □

4.2.3 RESTRICTION ON THE PROPOSITIONAL LANGUAGE

In the previous results, we made no specific assumptions on the formulae expressing the agents' preferences, except (sometimes) monotonicity. However, if we restrict the set of possible propositional formulae, it can decrease the complexity of the EEF EXISTENCE problem. We investigate first two natural restrictions on the propositional language: in the first case we restrict the formulae to the set of clauses (disjunctions of literals), in the second case the preferences are expressed using cubes only (conjunctions of literals). These restrictions match two different kinds of real-world problems.

- The case, where the agents' preferences are represented by clauses, corresponds to the kind of problems where each agent only wants one single object of a certain class. One can consider for example a set of patients all waiting for a kidney transplant. Each patient only needs one kidney, but several ones may be compatible.
- The case, where the agents' preferences are represented by cubes, corresponds to the kind of problems where each agent needs a single bundle of objects. It is typical from the kind of problems where the agents build the object they want from a set of basic material (or virtual) components: the set of objects stand for the basic components, and the cube of one agent stands for the complete device she wants.

Making one of these two assumptions actually decreases the complexity of the EEF EXISTENCE problem: it even renders it tractable in the case of clauses of objects.

Proposition 11 *EEF EXISTENCE for agents having dichotomous preferences restricted to clauses of objects can be solved in polynomial time.*

Proof We will first make two additional assumptions and prove that the complexity of the problem does not decrease under these two assumptions. (1) We suppose first that the agents' preferences are monotonic. If one agent has non-monotonic preferences, it means that there is one negative literal in her clause. Giving her an empty share will satisfy her without spoiling another agent, thus she can be safely removed from the problem. (2) We also suppose that each agent wants at least one object. If one

agent has an empty clause as her goal, it means that no matter what she gets, she cannot be satisfied. Thus she can also be safely removed from the problem. In the rest of the proof, we will only consider problems that verify the assumptions (1) and (2).

The proof is based on the following result: when the agents' preferences are disjunction of objects under assumptions (1) and (2), an allocation is Pareto-efficient and envy-free if and only if it satisfies every agent. The implication \Leftarrow is immediate. To prove the implication \Rightarrow we need to notice that an agent is satisfied if and only if she receives at least one object of her clause. Now take an allocation π such that there is an agent i which is not satisfied by π . Then either it is possible to satisfy her without spoiling another agent, because one object of her clause is given to an agent that does not want it, or that is already satisfied by another agent: in that case π is not Pareto-efficient. Or it is not possible to do so because all the objects of agent i 's disjunction have been given to some other agents that truly want them: in that case, the agent i envies these other agents, and π is thus not envy-free.

Therefore, finding a Pareto-efficient and envy-free allocation here comes down to finding an allocation that gives to each agent one object that she wants. Thus, an instance \mathcal{P} of the EEF EXISTENCE problem can be reduced to a maximal matching problem, in a bipartite graph $G_{\mathcal{P}}$ where there is one node per agent on one side, and one node per object on the other side, and where there is an edge between an agent-node i and an object-node o if and only if o is in the agent i 's clause. It can be easily checked that there is a Pareto-efficient and envy-free allocation if and only if there is a matching of size n in $G_{\mathcal{P}}$. The latter problem can be solved in time $O(nm)$ (Ford & Fulkerson, 1962), where m is the size of the biggest disjunction. \square

Now we investigate the case where the agents' preferences are cubes of objects. Not very surprisingly, this case is harder than the previous one, but remains in NP.

Proposition 12 *EEF EXISTENCE for agents having dichotomous preferences restricted to cubes of objects is NP-complete. The same result holds if we require the preferences to be monotonic.*

Proof The proof is organized as follows. We will first prove membership to NP without any assumption on the monotonicity of the preferences. Then we will show hardness in the monotonic case.

We first introduce some additional notations. In the following, we will denote by $Obj^+(i)$ (resp. $Obj^-(i)$) the set of objects appearing as positive (resp. negative) literals in agent i 's cube. Let π be an allocation. π will be said *minimally regular* if $\forall i$, either $\pi(i) = Obj^+(i)$ or $\pi(i) = \emptyset$. For a given allocation π , we will denote by π_{MR} its corresponding minimally regular allocation, that is, the allocation such that for all i , $\pi_{MR}(i) = \emptyset$ if $Obj^+(i) \not\subseteq \pi(i)$, and $\pi_{MR}(i) = Obj^+(i)$ if $Obj^+(i) \subseteq \pi(i)$. We will also write $Sat(\pi)$ the set of agents satisfied by π : $Sat(\pi_{MR}) = \{i \mid \pi_{MR}(i) = Obj^+(i)\}$, and $Allocated(\pi) = \bigcup_{i \in I} \pi(i)$ (the set of objects allocated to an agent).

We have the following result:

Lemma 10 *Let π be an allocation. We have:*

- π_{MR} is minimally regular;
- $Sat(\pi) \subseteq Sat(\pi_{MR})$;
- If π is Pareto-efficient, then π_{MR} is also Pareto-efficient.

Proof • For all agent i , $\pi_{MR}(i) = \emptyset$ or $\pi_{MR}(i) = Obj^+(i)$ by definition of π_{MR} . Therefore π_{MR} is minimally regular.

- Let π be an allocation, and let i be an agent. If i is satisfied by π , then $Obj^+(i) \subseteq \pi(i)$ and $Obj^-(i) \cap \pi(i) = \emptyset$. By definition of π_{MR} , we have $\pi_{MR}(i) = Obj^+(i)$, and thus we still have $Obj^+(i) \subseteq \pi(i)$ and $Obj^-(i) \cap \pi(i) = \emptyset$, therefore the agent i is still satisfied by π_{MR} . It proves that $Sat(\pi) \subseteq Sat(\pi_{MR})$.
- Suppose that π is Pareto-efficient, and suppose that there is a π' that Pareto-dominates π_{MR} . Then we have $Sat(\pi) \subseteq Sat(\pi_{MR}) \subsetneq Sat(\pi')$, which contradicts the fact that π is Pareto-efficient. It proves the third point. \square

Lemma 11 *A minimally regular allocation π_{MR} is Pareto-efficient if and only if there is no i such that (a) $i \notin \text{Sat}(\pi_{MR})$ and (b) $\text{Obj}^+(i) \cap \text{Allocated}(\pi_{MR}) = \emptyset$.*

Proof Let π_{MR} be a minimally regular allocation, and suppose that there is some i such that $i \notin \text{Sat}(\pi_{MR})$ and $\text{Obj}^+(i) \subseteq X \setminus \text{Allocated}(\pi_{MR})$. Then the allocation π' such that $\forall j \neq i$ $\pi'(j) = \pi_{MR}(j)$ and $\pi'(i) = \text{Obj}^+(i)$ is well-defined (since $\text{Obj}^+(i)$ is among the set of unallocated objects for π_{MR}), and Pareto-dominates π_{MR} , since all the agents satisfied by π_{MR} are also satisfied by π' , and i is now satisfied by π' whereas she was not by π_{MR} .

Conversely, suppose that π_{MR} is not Pareto-efficient, and let π' be a Pareto-efficient allocation that Pareto-dominates π_{MR} . Then by Lemma 10, π'_{MR} is Pareto-efficient, and also Pareto-dominates π_{MR} . For all $i \in \text{Sat}(\pi_{MR})$, $\pi_{MR}(i) = \pi'_{MR}(i) = \text{Obj}^+(i)$ since these two allocations are both minimally regular, and each agent satisfied by π_{MR} is also satisfied by π'_{MR} . Moreover, there is a $j \notin \text{Sat}(\pi_{MR})$ such that $\pi'_{MR}(j) = \text{Obj}^+(j)$. Since $\text{Allocated}(\pi_{MR}) = \bigcup_{i \in \text{Sat}(\pi_{MR})} \text{Obj}^+(i)$ and since $\pi'_{MR}(j) \in X \setminus \bigcup_{i \in \text{Sat}(\pi_{MR})} \text{Obj}^+(i)$, we have $\text{Obj}^+(j) \in X \setminus \text{Allocated}(\pi_{MR})$, which finally proves the lemma. \square

The two lemmas provide a procedure to check if a given allocation π is Pareto-efficient: first compute π_{MR} (which can be done in polynomial time). By Lemma 10, $\text{Sat}(\pi) \subseteq \text{Sat}(\pi_{MR})$. If the inclusion is strict (that is, $\text{Sat}(\pi) \subsetneq \text{Sat}(\pi_{MR})$), then obviously π is not Pareto-efficient since π_{MR} Pareto-dominates it. Otherwise, checking if π is Pareto-efficient comes down to checking if π_{MR} is Pareto-efficient, which comes down, according to Lemma 11, to n set inclusion tests. Now, checking if π is envy-free is still polynomial. Hence the problem is in NP.

We will now prove the hardness of the problem by focusing on monotonic preferences (that is, such that $\text{Obj}^-(i) = \emptyset$ for all i). For the hardness proof, we will need a few additional lemmas.

Lemma 12 *Let π be an allocation and suppose that the agents have monotonic preferences. If π is envy-free then π_{MR} is envy-free.*

Proof Let π be an allocation. By Lemma 10, $\text{Sat}(\pi) \subseteq \text{Sat}(\pi_{MR})$. Conversely, let $i \in \text{Sat}(\pi_{MR})$. We have $\pi_{MR}(i) \subseteq \pi(i)$, which proves that $i \in \text{Sat}(\pi)$, because we deal with monotonic preferences. Now suppose that i envies j in π_{MR} . Then i is not satisfied by π_{MR} , and thus neither satisfied by π . Since $\pi_{MR}(j) \subseteq \pi(j)$ and agent i 's preferences are monotonic, i will still envy j in π . Thus if π is envy-free, so is π_{MR} . \square

An important corollary of this lemma is that when we deal with monotonic cubes, the existence of a Pareto-efficient and envy-free allocation is equivalent to the existence of a minimally regular Pareto-efficient and envy-free allocation. Therefore we can restrict our existence problem to the minimally regular allocations.

Lemma 13 *Let i and j be two different agents (and still suppose that the agents have monotonic preferences). Then (there exists a minimally regular allocation π_{MR} such that i envies j) if and only if $\text{Obj}^+(i) \subseteq \text{Obj}^+(j)$.*

Proof Let π_{MR} be a minimally regular allocation, and suppose that i envies j . Then obviously i is not satisfied but j is; hence $\pi_{MR}(j) = \text{Obj}^+(j)$. Since i envies j , we thus have directly $\text{Obj}^+(i) \subseteq \text{Obj}^+(j)$.

Conversely, suppose that $\text{Obj}^+(i) \subseteq \text{Obj}^+(j)$. Then the allocation π_{MR} that gives $\text{Obj}^+(j)$ to j and nothing to the other agents is clearly minimally regular, and is also obviously such that i envies j . \square

Now we introduce the NP-complete problem that we will use to prove NP-hardness (in the monotonic case):

Problem 3: EXACT COVER BY 3-SETS (Karp, 1972)

INSTANCE: A set S of size $3q$, and a collection $C = \langle S_1, \dots, S_{|C|} \rangle$ of 3-element subsets of S

QUESTION: Does C contain an exact cover for S , i.e. a sub-collection $C' \subseteq C$ such that every element of S occur in exactly one member of C' ?

Given an instance $\langle S, C = \langle S_1, \dots, S_{|C|} \rangle \rangle$ of EXACT COVER BY 3-SETS (we will assume w.l.o.g. that the S_i are all different), let $\mathcal{P}(S, C)$ be the following EEF INSTANCE:

- Agents:** a set of $|C| + 2|S|$ agents $I = I_1 \cup I_2$, with $I_1 = \{1, \dots, |C|\}$ and $I_2 = \{|C| + 1, \dots, |C| + 2|S|\}$,
- Objects:** a set of $2|S|$ items $X = X_1 \cup X_2$, with $X = \{x_1, \dots, x_{|S|}\}$ and $X' = \{x'_1, \dots, x'_{|S|}\}$, each pair (x_i, x'_i) corresponding to a different element a_i of S ,
- Preferences:** for each agent $i \in I_1$, $\varphi_i = \bigwedge_{a_j \in S_i} x_j$, and for each $k \in \{1, \dots, |S|\}$, $\varphi_{|C|+2k-1} = \varphi_{|C|+2k} = x_k \wedge x'_k$.

In other words, the first $|C|$ agents' preferences correspond to the sets in the collection C , and the last $2|S|$ agents are gathered by pairs, each member of the same pair having the same preferences as the other member.

Since all the S_i are different and of size 3, for all $i \neq j$, $S_i \not\subseteq S_j$, and thus $Obj^+(i) \not\subseteq Obj^+(j)$. By definition of the preferences, we also have $Obj^+(i) \not\subseteq Obj^+(j)$ for each $(i, j) \in I_1 \times I_2$ and for each $(i, j) \in I_2 \times I_1$ as well. Hence, by Lemma 13, the only potential source of envy in such an instance comes from an agent in I_2 envying her partner. Since it is impossible to satisfy the two agents of the same pair at the same time, an allocation is envy-free if and only if it does not satisfy any agent from I_2 .

By Lemma 11, a minimally regular allocation π_{MR} is Pareto-efficient if and only if there is no i such that $i \notin Sat(\pi_{MR})$ and $Obj^+(i) \subseteq X \setminus Allocated(\pi_{MR})$. Therefore a minimally regular allocation π_{MR} is Pareto-efficient and envy-free if and only if there is no $k \in \{1, \dots, |S|\}$ such that $\pi_{MR}(|C| + 2k - 1) = \{x_k, x'_k\}$ or $\pi_{MR}(|C| + 2k) = \{x_{k'}, x'_{k'}\}$, and there is no $k' \in \{1, \dots, |S|\}$ such that $\{x_{k'}, x'_{k'}\} \subset X \setminus Allocated(\pi_{MR})$ (this last condition comes down to $x_{k'} \notin Allocated(\pi_{MR})$, since $x_{k'}$ and $x'_{k'}$ must be allocated together in a minimally regular allocation). Finally, π_{MR} is Pareto-efficient and envy-free if and only if $\forall k \in \{1, \dots, |S|\}$, there is an $i \in I_1$ such that $x_k \in \pi_{MR}(i)$, that is, if and only if $\bigcup_{i \in I_1} \pi_{MR}(i) = \bigcup_{a_j \in S} \{x_j\}$.

Let π_{MR} be a minimally regular allocation. Then we can define the sub-collection $g(\pi_{MR})$ by $g(\pi_{MR}) = \{S_i \in C \mid \pi_{MR}(i) = Obj^+(i)\}$. The mapping g clearly defines a bijection between the set of non-overlapping sub-collections and the set of minimally regular allocations, and one can notice that $\bigcup_{i \in I_1} \pi_{MR}(i) = \bigcup_{S_j \in g(\pi_{MR})} \bigcup_{a_k \in S_j} \{x_k\}$.

Let $C' \subseteq C$ is an exact cover for S . Then $g^{-1}(C')$ exists and is a valid minimally regular allocation. We also have $\bigcup_{i \in I_1} g^{-1}(C')(i) = \bigcup_{S_j \in C'} \bigcup_{a_k \in S_j} \{x_k\} = \bigcup_{a_j \in S} \{x_j\}$ because C' is a cover. Therefore $g^{-1}(C')$ is Pareto-efficient and envy-free by the previous result.

Conversely, suppose that there is a minimally regular Pareto-efficient and envy-free allocation π_{MR} . Then $g(\pi_{MR})$ is a non-overlapping sub-collection of C , and is such that $\bigcup_{S_i \in g(\pi_{MR})} \bigcup_{a_j \in S_i} \{a_j\} = \bigcup_{i \in I_1} \bigcup_{x_j \in \pi_{MR}(i)} \{a_j\} = \bigcup_{x_j \in \{\pi_{MR}(i) \mid i \in I_1\}} \{a_j\}$. π_{MR} being Pareto-efficient and envy-free, we have $\bigcup_{i \in I_1} \pi_{MR}(i) = \bigcup_{a_j \in S} \{x_j\}$, and hence $\bigcup_{x_j \in \{\pi_{MR}(i) \mid i \in I_1\}} \{a_j\} = \bigcup_{x_j \in \{x_j \mid a_j \in S\}} \{a_j\} = S$. It proves that $g(\pi_{MR})$ is an exact cover for S .

The reduction is clearly polynomial; hence the NP-hardness. \square

The previous proof (and especially Lemma 13) sheds the light on the hard case of the EEF EXISTENCE problem with conjunctive preferences. In an instance of this problem, the only possible source of envy comes from $Obj^+(i) \subseteq Obj^+(j)$ when we deal with minimally regular allocations. In that case, we cannot satisfy j without raising envy from i . Now if $Obj^+(i) \subsetneq Obj^+(j)$, one can remove j from the instance, because if she is satisfied, then i will envy her (note that this is not true for non-monotonic preferences, because one can give an object to an agent that does not want it to prevent envy from another agent, and thus we cannot restrict our problem to minimally regular allocations).

If there is no pair (i, j) , $i \neq j$, such that $Obj^+(i) = Obj^+(j)$, then one can remove every agent i such that there is a $j \neq i$ such that $Obj^+(j) \subsetneq Obj^+(i)$. After that, it is easy to see that every minimally regular allocation is envy-free. Since there is at least one Pareto-efficient minimally regular allocation, it guarantees the existence of a Pareto-efficient and envy-free allocation in that case, may the preferences be monotonic or not. More formally:

Proposition 13 *There always exists an efficient and envy-free allocation for an instance of the EEF EXISTENCE problem with agents having dichotomous preferences restricted to cubes of objects when the following condition holds:*

$$\forall (i, j) \in I^2, i \neq j, (\varphi_i = \varphi_j) \Rightarrow (\exists k \text{ such that } k \neq i, k \neq j \text{ and } Obj^+(k) \subsetneq Obj^+(i)). \quad (1)$$

Of course, there is no equivalence between condition 1 and the existence of a Pareto-efficient and envy-free allocation⁷, because it may happen that, given two agents i and j having the same preferences, the satisfaction of one of these two agents is prevented by another agent k such that $Obj^+(i) \cap Obj^+(k) \neq \emptyset$, but $Obj^+(k) \not\subseteq Obj^+(i)$. This is the hard case: when two agents i and j have identical preferences, but no agent k is such that $Obj^+(k) \subsetneq Obj^+(i)$, it may however be possible to prevent i and j from being satisfied by satisfying another agent k' such that $Obj^+(i) \cap Obj^+(k') \neq \emptyset$, as the following example shows: $\varphi_1 = \varphi_2 = x_1 \wedge x_2$, and $\varphi_3 = x_2 \wedge x_3$. Satisfying the agent 2 only leads to an efficient and envy-free allocation, whereas the condition 1 does not hold.

Proof (Proposition 13) In the following we will denote by I_1 the set of all the agents whose preferences are “inclusion-minimal”, that is, $I_1 = \{i \mid \nexists j \text{ such that } Obj^+(j) \subset Obj^+(i)\}$. We will denote by I_2 the other agents: $I_2 = I \setminus I_1$.

Here is a simple procedure for finding a Pareto-efficient and envy-free allocation: greedily select a maximal set S of agents from I_1 , such that each agent $i \in S$ receives $Obj^+(i)$ (until it becomes impossible to select another unsatisfied agent in I_1).

The allocation π resulting from this procedure is minimally regular, and by Lemma 13 it is clearly envy-free (by definition of I_1). Moreover, suppose there is an $i \notin Sat(\pi)$ such that $Obj^+(i) \subseteq X \setminus Allocated(\pi)$. Then $i \notin I_1$, since if it was the case, the procedure would have selected her and therefore she should be satisfied. We have also $i \notin I_2$, because if it was the case, then there is a $j \in I_1$ such that $Obj^+(j) \subset Obj^+(i)$, and therefore $Obj^+(j) \subseteq X \setminus Allocated(\pi)$, which is impossible for the same reasons as above. Therefore, π is also Pareto-efficient by Lemma 11. \square

After having investigated two natural restrictions on the propositional language used for dichotomous preferences, we introduce a more general result, based on the fact that the hardness result of Proposition 6 is clearly linked to the NP-completeness of the SAT problem. Now what happens if we restrict the expression of preferences to a certain class \mathcal{C} such that $SAT(\mathcal{C})$ can be solved in polynomial time? In the general case, where no additional assumption is made on \mathcal{C} we cannot say anything more about the complexity of the EEF EXISTENCE problem than about the complexity of the general problem. However, if \mathcal{C} is also closed by conjunction, the complexity falls down to NP:

Proposition 14 *Let \mathcal{C} be a class of propositional formulae closed by conjunction and such that $SAT(\mathcal{C})$ is in P. Then the EEF EXISTENCE for agents having dichotomous preferences expressed with formulae from class \mathcal{C} is in NP.*

Proof Membership to NP comes from the fact that, after having non-deterministically guessed an allocation π , checking if it is envy-free and Pareto-efficient can be done in polynomial time. Given an allocation, checking if it is envy-free can be done in time $O(nm)$ (where m is the length of the biggest CNF), just by checking, for each unsatisfied agent, if it would have been satisfied with the share of another agent. Given the set $Sat(\pi)$ of the agents satisfied by π , checking the Pareto-efficiency of π comes down to check $\forall i \in I \setminus Sat(\pi)$ if $\bigwedge_{j \in Sat(\pi)} \varphi_j \wedge \varphi_i$ is unsatisfiable. This can be done by making a linear number of calls to a $SAT(\mathcal{C})$ oracle, since all the preferences are in \mathcal{C} , and since this class is closed by conjunction. It thus proves that the EEF EXISTENCE with formulae in \mathcal{C} is in NP. \square

7. Otherwise, the Proposition 12 would be false, or we would have proved that $P = NP$.

As a corollary, for all classes of propositional formulae such that $\text{SAT}(\mathcal{C})$ is polynomial, and that contain cubes, the EEF EXISTENCE problem is NP-complete. This applies for example to the class of 2-CNF formulae or to the class of Horn clauses.

4.3 Alternative Efficiency Criteria

The main reason for the high complexity of the EEF EXISTENCE problem is that the Pareto-efficiency of an allocation is hard to check. As a consequence, complexity can decrease if we choose an alternative notion of efficiency. We investigate here two alternative efficiency criteria: completeness of the allocation, and maximal number of satisfied agents.

First, we weaken the Pareto-efficiency by only requiring the allocations to be *complete*. Unsurprisingly, this makes the complexity fall down to NP.

Proposition 15 *The problem of deciding whether a complete envy-free allocation for agents with dichotomous preferences exists is NP-complete, even for 2 agents with identical preferences.*

Proof Since checking that an allocation is complete can be done in polynomial time, membership to NP is straightforward.

We will prove hardness by a reduction from the SAT problem. Let φ be a propositional formula. We create the following instance of the resource allocation problem : we map each propositional variable of φ to a different object and we add another object y ; the two agents have the same preferences, represented by the formula $\varphi \vee y$. Obviously, every complete allocation satisfies at least one agent (the one who receives y). If φ is satisfiable, it is possible to satisfy the other agent as well with a share that corresponds to a model of φ : thus in this case there exists a complete and envy-free allocation. Conversely, suppose that there exists a complete and envy-free allocation. Then one of the two agents must be satisfied thanks to φ (since y cannot be given to both agents), hence proving that φ is satisfiable. \square

Notice that the hardness proof above is no longer valid if we require preferences to be monotonic. As an anonymous referee pointed out, it can be proved by a reduction from EXACT-COVER-BY-3-SETS that the result holds if we require monotonicity, while relaxing the restriction on the number of agents. However, we do not know if NP-hardness holds for *two* agents with monotonic preferences (we conjecture it does).

Secondly, we can think of looking for cardinality-maximal subsets of satisfied agents, instead of inclusion-maximal subsets like Pareto-efficiency does.

Proposition 16 *The problem of deciding whether an envy-free allocation satisfying a maximal number of agents with monotonic dichotomous preferences exists is Θ_2^P -complete.*

Proof Checking whether there exists an envy-free allocation satisfying at least k agents is in NP; therefore, the maximal number of agents who can be satisfied simultaneously can be computed by dichotomy within $\log n$ NP oracles. It then suffices, after this step is done, to guess an allocation and check that it is envy-free and satisfies a maximal number of agents, adding just one more NP oracle. Hence the membership to Θ_2^P .

Hardness is obtained by a reduction from the following problem⁸:

Problem 4: MAX-INDEX-SAT_{odd} (Wagner, 1990)

INSTANCE: A sequence of propositional formulae $\langle \chi_1, \dots, \chi_n \rangle$ such that $(\chi_i \text{ is unsatisfiable}) \Rightarrow (\chi_{i+1} \text{ is unsatisfiable.})$

QUESTION: Is the maximum index i such that χ_i is satisfiable an odd number ?

8. This problem is referred to several times in the literature, but it does not seem to have a name.

First notice that the complexity of the latter problem does not decrease under the following assumptions:

- n is even (if it is not, we can just add the \perp formula at the end of the sequence);
- the sets of propositional variables of the formulae χ_i are pairwise-disjoint (if two formulae χ_i and χ_{i+1} share some variables, just transform each variable v of say χ_i into a copy v' : it will not change the (un)satisfiability of χ_i but the set of propositional variables of χ_i and χ_{i+1} will now be disjoint).

Let $\langle \chi_1, \dots, \chi_n \rangle$ be an instance of MAX-INDEX-SAT_{odd} with the two additional latter assumptions, and let V_i denote the set of propositional variables appearing in χ_i . We translate this instance to the following instance $\mathcal{P}(\chi_1, \dots, \chi_n)$:

- Agents:** $2n$ agents : $I_{1,2} \cup I_{3,4} \cup \dots \cup I_{n-1,n}$, where the group $I_{2i-1,2i}$ contains the four agents $\{4i-3, 4i-2, 4i-1, 4i\}$.
- Objects:** we create for each $v \in V_i$ (for each $i \in \{1, \dots, n\}$) four objects $x_v, \bar{x}_v, y_v, \bar{y}_v$, and we add n dummy objects d_k ($k \in \{1, \dots, n\}$);
- Preferences:** for each group $I_{2i-1,2i}$ ($i \in \{1, \dots, n/2\}$), the preferences of the agents are:
- $\varphi_{4i-3} = \varphi_{4i-2} = (\chi'_{2i-1} \wedge d_{2i-1}) \vee (\chi'_{2i} \wedge d_{2i})$,
 - $\varphi_{4i-1} = \bigwedge_{v \in V_{2i-1} \cup V_{2i}} x_v \vee \bar{x}_v$,
 - $\varphi_{4i} = \bigwedge_{v \in V_{2i-1} \cup V_{2i}} y_v \vee \bar{y}_v$,
- where χ'_k is the formula χ_k where v has been replaced by $x_v \wedge y_v$, and $\neg v$ has been replaced by $\bar{x}_v \wedge \bar{y}_v$.

The proof of the proposition is primarily based on the fact that the problem can be split into $n/2$ subproblems concerning the agents from $I_{2i-1,2i}$:

Lemma 14 *We denote by \mathcal{P}_i the restriction of $\mathcal{P}(\chi_1, \dots, \chi_n)$ to the set of agents $I_{2i-1,2i}$ and to the objects they want. An allocation π will be said to be splittable if $\forall i \neq j, \pi(I_{2i-1,2i}) \cap \pi(I_{2j-1,2j}) = \emptyset$. The restriction of a splittable allocation π to $\pi(I_{2i-1,2i})$ will be written $\pi|_i$.*

There exists an envy-free allocation satisfying a maximal number of agents for $\mathcal{P}(\chi_1, \dots, \chi_n)$ if and only if there exists a splittable allocation π such that $\forall i \in \{1, \dots, n/2\}$, $\pi|_i$ is envy-free and satisfies a maximal number of agents for \mathcal{P}_i .

Proof First, we restrict our attention to *regular* allocations, where regular means, like in Lemma 2, that an allocation gives an object to an agent only if she wants it. This is safe for the same reasons as in Lemma 2: the existence of an envy-free allocation satisfying a maximal number of agents is equivalent to the existence of a regular envy-free allocation satisfying a maximal number of agents. Since the sets V_i are pairwise-disjoint, two different subproblems \mathcal{P}_i and \mathcal{P}_j do not share any object, and therefore any regular allocation is also splittable.

Let π be a regular allocation. Suppose that there is an allocation π'_i for the problem \mathcal{P}_i that satisfies more agents than $\pi|_i$. Then the splittable allocation made from the sub-allocations $\pi|_j$ for $j \neq i$ and π'_i is valid, regular, and satisfies more agents than π . Conversely, suppose that there is a regular allocation π' that satisfies more agents than π . Then there is at least one i such that strictly more agents from $I_{2i-1,2i}$ are satisfied by π'_i than by $\pi|_i$. It proves that any regular allocation π satisfies a maximal number of agents if and only if for each i , $\pi|_i$ satisfies a maximal number of agents.

Suppose now that π is envy-free. Then obviously all the $\pi|_i$ are as well. Conversely, suppose that all the $\pi|_i$ are envy-free. Then π is envy-free, because (1) no agent can envy another agent from the same group, because the $\pi|_i$ are envy-free, and (2) no agent from a group i can envy an agent from another group j , since $\pi(I_{2i-1,2i}) \cap \pi(I_{2j-1,2j}) = \emptyset$. \square

For each interpretation Int_k of V_k , we define the following sets of objects:

- $f(Int_k) = \{x_v \mid Int_k \models v\} \cup \{\bar{x}_v \mid Int_k \not\models v\}$;
- $g(Int_k) = \{y_v \mid Int_k \models v\} \cup \{\bar{y}_v \mid Int_k \not\models v\}$;
- $\bar{f}(Int_k) = \{x_v \mid Int_k \not\models v\} \cup \{\bar{x}_v \mid Int_k \models v\}$;

- $\bar{g}(Int_k) = \{y_v \mid Int_k \not\models v\} \cup \{\bar{y}_v \mid Int_k \models v\}$.

Moreover, given two interpretations Int_{2i-1} and Int_{2i} of respectively V_{2i-1} and V_{2i} , we will write $\pi_{Int_{2i-1}, Int_{2i}}$ the following allocation of \mathcal{P}_i :

- $\pi_{Int_{2i-1}, Int_{2i}}(4i-3) = f(Int_{2i-1}) \cup g(Int_{2i-1}) \cup \{d_{2i-1}\}$;
- $\pi_{Int_{2i-1}, Int_{2i}}(4i-2) = f(Int_{2i}) \cup g(Int_{2i}) \cup \{d_{2i}\}$;
- $\pi_{Int_{2i-1}, Int_{2i}}(4i-1) = \bar{f}(Int_{2i-1}) \cup \bar{f}(Int_{2i})$;
- $\pi_{Int_{2i-1}, Int_{2i}}(4i) = \bar{g}(Int_{2i-1}) \cup \bar{g}(Int_{2i})$.

Lemma 15 *Let Int_{2i-1} and Int_{2i} be two respective interpretations of V_{2i-1} and V_{2i} .*

- $\pi_{Int_{2i-1}, Int_{2i}}$ satisfies both agents $4i-1$ and $4i$;
- $\pi_{Int_{2i-1}, Int_{2i}}$ satisfies $4i-3$ if and only if $Int_{2i-1} \models \chi_{2i-1}$, and $\pi_{Int_{2i-1}, Int_{2i}}$ satisfies $4i-2$ if and only if $Int_{2i} \models \chi_{2i}$;

Proof Let Int_{2i-1} and Int_{2i} be two respective interpretations of V_{2i-1} and V_{2i} .

- By definition, $\bar{f}(Int_k)$ contains x_v or \bar{x}_v for each $v \in V_k$, thus $\pi_{Int_{2i-1}, Int_{2i}}(4i-1)$ contains x_v or \bar{x}_v for each $v \in V_{2i-1} \cup V_{2i}$. Therefore the agent $4i-1$ is satisfied by $\pi_{Int_{2i-1}, Int_{2i}}(4i-1)$. The same reasoning holds for the agent $4i$.
- By definition, χ_{2i-1} is satisfied by Int_{2i-1} if and only if χ'_{2i-1} is satisfied by the interpretation defined by setting to true all the x_v and y_v (resp. all the \bar{x}_v and \bar{y}_v) such that $Int_{2i-1} \models v$ (resp. $Int_{2i-1} \not\models v$). Thus, if $Int_{2i-1} \models \chi_{2i-1}$, $\pi_{Int_{2i-1}, Int_{2i}}$ satisfied χ'_{2i-1} . Since it also satisfies d_{2i-1} , then $4i-3$ is satisfied by $\pi_{Int_{2i-1}, Int_{2i}}$. Conversely, if $4i-3$ is satisfied by $\pi_{Int_{2i-1}, Int_{2i}}$, then obviously χ'_{2i-1} must be satisfied by $\pi_{Int_{2i-1}, Int_{2i}}$ (because $4i-3$ does not receive d_{2i}), which proves that χ_{2i-1} is satisfied by Int_{2i-1} . The same result holds for χ_{2i} and the agent $4i-2$. \square

Lemma 16 *Consider the restricted problem \mathcal{P}_i .*

- If neither χ_{2i-1} nor χ_{2i} is satisfiable, then for any interpretations Int_{2i-1} and Int_{2i} of V_{2i-1} and V_{2i} respectively, $\pi_{Int_{2i-1}, Int_{2i}}$ is envy-free and satisfies a maximal number of agents.
- If only χ_{2i-1} is satisfiable, then if M_{2i-1} is a model of χ_{2i-1} , $\pi_{M_{2i-1}, Int_{2i}}$ satisfies a maximal number of agents. Moreover, there is no envy-free allocation satisfying a maximal number of agents in this case.
- If both χ_{2i-1} and χ_{2i} are satisfiable, then if M_{2i-1} and M_{2i} are respective models of χ_{2i-1} and χ_{2i} , $\pi_{M_{2i-1}, M_{2i}}$ satisfies a maximal number of agents and is envy-free.

Proof Suppose that neither χ_{2i-1} nor χ_{2i} is satisfiable. Then any allocation π_i satisfying $4i-3$ (resp. $4i-2$) must be such that there is at least one $v \in V_{2i-1} \cup V_{2i}$ such that $\{x_v, \bar{x}_v, y_v, \bar{y}_v\} \subset \pi_i(4i-3)$ (resp. $\pi_i(4i-2)$), because otherwise one could deduce a model of χ_{2i-1} or χ_{2i} from $\pi_i(4i-3)$ (resp. $\pi_i(4i-2)$). Thus any of the agents $4i$ and $4i-1$ can be satisfied in this case: the maximal number of agents it is possible to satisfy is 2. Since every allocation of the form $\pi_{Int_{2i-1}, Int_{2i}}$ satisfies both agents $4i-1$ and $4i$, it satisfies a maximal number of agents in this case. It is also obviously envy-free, since neither d_{2i} nor d_{2i-1} are in the shares of agents $4i-1$ and $4i$, and thus the 2 other agents cannot envy them.

Suppose that only χ_{2i-1} is satisfiable. Then any allocation satisfying both agents $4i-3$ and $4i-2$ must satisfy χ'_{2i-1} for one of these two agents, and χ'_{2i} for the other one (because of d_{2i} and d_{2i-1}). Since χ_{2i} is not satisfiable, in that case neither $4i-1$ nor $4i$ can be satisfied by π_i , for the same reasons as above. From that we can deduce that it is not possible to satisfy the 4 agents. Neither is it possible to satisfy 3 agents with both $4i-3$ and $4i-2$ satisfied. Now consider the allocation $\pi_{M_{2i-1}, Int_{2i}}$, M_{2i-1} being a model of χ_{2i-1} . By Lemma 15, $\pi_{M_{2i-1}, Int_{2i}}$ satisfies 3 agents: $4i-3$, $4i-1$ and $4i$. This allocation is not envy-free, but no allocation satisfying a maximal number of agents can be in this case (because either $4i-3$ or $4i-2$ remains unsatisfied in such an allocation, envying her partner).

Lastly, suppose that both χ_{2i-1} and χ_{2i} are satisfiable, and let M_{2i-1} and M_{2i} be their models. Then by Lemma 15, $\pi_{M_{2i-1}, M_{2i}}$ satisfies the 4 agents, thus satisfying a maximal number of agents and obviously being envy-free. \square

We can now conclude the proof. By Lemma 14, there exists an envy-free allocation satisfying a maximal number of agents for $\mathcal{P}(\chi_1, \dots, \chi_n)$ if and only if there exists a splittable allocation π such that $\forall i \in \{1, \dots, n/2\}$, $\pi|_i$ is envy-free and satisfies a maximal number of agents for $\mathcal{P}|_i$. By Lemma 15, there exists an envy-free allocation $\pi|_i$ satisfying a maximal number of agents for $\mathcal{P}|_i$ if and only if either none of the two formulae χ_{2i-1} and χ_{2i} are satisfiable, or both are. Now suppose that the maximum index j such that χ_j is satisfiable is an odd number (say $2i - 1$). In that case, there is not any envy-free allocation satisfying a maximal number of agents for $\mathcal{P}|_i$ since χ_{2i-1} is satisfiable but χ_{2i} is not. Conversely, suppose that the maximum index j such that χ_j is satisfiable is an even number (say $2i$). In that case, there is an envy-free allocation satisfying a maximal number of agents for all $\mathcal{P}|_k$, since for each $\mathcal{P}|_k$ either the two formulae χ_{2k-1} and χ_{2k} are satisfiable (if $k \leq i$), or none of them is (if $k > i$).

We thus have a reduction from the latter problem to the problem of existence of an envy-free allocation satisfying a maximal number of agents. It proves the Θ_2^P -completeness. \square

5. Non-dichotomous Preferences

We will now consider the case where preferences are no longer dichotomous.

5.1 General Logical Preferences

Again, since an explicit description of preferences is exponentially large, the need for a compact description thereof is clear. Many languages exist for succinct representation of preference. We limit our investigation to the following class of languages:

Definition 9 (Compact language under logical form) *Let \mathcal{L} be a language representing a set of preference relations over a set of alternatives 2^X . \mathcal{L} will be said to be a compact language under logical form if and only if :*

- (a) *it is able to express any dichotomous preference as compactly as the previous language introduced, that is, the language expressing dichotomous preferences in propositional form can be polynomially reduced to \mathcal{L} ;*
- (b) *comparing two sets of goods can be done in polynomial time.*

These two previous conditions are in practise not very restricting, and are met by many languages for succinct representation of representation. See for instance the paper by Lang (2004) for a survey of logical languages for compact preference representation. Note that several widely studied representation languages, such as CP-nets and other graphical languages, are not under logical form, because they fail to represent preferences expressed by logical formulas within polynomial space⁹. Interestingly, Proposition 6 extends to any compact representation language under logical form:

Corollary 2 *EEF EXISTENCE with monotonic compact preference under logical form is Σ_2^P -complete.*

Proof The EEF EXISTENCE problem can be solved using the following algorithm:

1. non-deterministically guess an allocation π ;
2. check that it is envy-free;
3. check that it is Pareto-efficient.

9. A natural question is the complexity of the fair division problems when preferences are expressed in these languages. This is left for further study.

By condition (b), step two can be done in polynomial time, since it requires at most a quadratic number of polynomial oracles. By condition (b) also, the problem of checking whether a given allocation is Pareto-efficient is in co-NP. Therefore, the previous non-deterministic algorithm uses 1 NP oracle, and runs in polynomial-time. Hence membership to Σ_2^P .

Hardness is a corollary of Proposition 6 together with condition (a). \square

5.2 Numerical Preferences under Logical Form

For the latter result preferences do not have to be numerical since Pareto-efficiency and envy-freeness are purely ordinal notions. Now, if preferences are numerical, which implies the possibility of intercomparing and aggregating preferences of several agents, then, besides Pareto-efficiency, we may consider efficiency based on social welfare functions. We consider here only the two most classical way of aggregating a collection of utility functions into a social welfare function:

Definition 10 (Classical utilitarianism and egalitarianism) *Given a collection of individual utility functions $\langle u_1, \dots, u_n \rangle$, with for each i , $u_i : 2^X \rightarrow \mathbb{Z}$:*

- *the classical utilitarian social welfare function is the function defined by $sw^* : \pi \mapsto \sum_i u_i(\pi(i))$;*
- *the egalitarian social welfare function is the function defined by $sw^{(e)} : \pi \mapsto \min_i u_i(\pi(i))$;*

Maximizing the egalitarian social welfare function is often viewed as an alternative criterion of fairness, encoding the Rawlsian egalitarian point of view (Rawls, 1971). However, as we will see in Proposition 17, egalitarianism (as well as classical utilitarianism) is not always compatible with envy-freeness. The link between these two alternative points of view about fairness is deeply investigated by Brams and King (2005).

Since we do not deal anymore with ordinal (or dichotomous) preferences, we have to define precisely what we will mean by compact representation of numerical preferences. Here we will pick as the basic numerical language one of the most simple compact languages, consisting in associating numerical weights to propositional formulae – see e.g. (Chevalere, Endriss, & Lang, 2006), or (Jeong & Shoham, 2005) in the context of coalitional games:

Definition 11 (Weighted propositional language) *Given a set of goods X , the weighted propositional language associated to X is the set of all possible subsets of $L_X \times \mathbb{Z}$.*

Given a set of weighted propositional formulae $\varphi = \{\langle \psi_1, w_1 \rangle, \dots, \langle \psi_r, w_r \rangle\}$, the utility function associated is:

$$u_\varphi : \begin{array}{l} 2^X \\ S \end{array} \rightarrow \mathbb{Z} \quad , \quad \text{with } \sigma_k = \begin{cases} 1 & \text{if } S \models \psi_k \\ 0 & \text{otherwise.} \end{cases}$$

$$S \mapsto \sum_{k=1}^r w_k \times \sigma_k$$

Using this language, preferences are monotonic if all formulae are positive and all weights are positive.

Now, we can define our notion of compact numerical language:

Definition 12 (Compact numerical language under logical form) *Let \mathcal{L} be a language representing a set of utility functions over a set of alternatives 2^X . \mathcal{L} will be said to be a compact numerical language under logical form if and only if :*

- (a) *it is able to express any utility function as compactly as the weighted propositional language, that is, the weighted propositional language can be polynomially reduced to \mathcal{L} ;*
- (b) *computing the utility of one set of goods can be done in polynomial time.*

Of course, since a compact numerical language under logical form is also a compact language under logical form, the complexity result from Corollary 2 still holds. However, it appears that the complexity of the problem of deciding whether an efficient and envy-free allocation exists decreases when the Pareto-efficiency is replaced by a weaker notion: the maximization of one of the two latter social welfare functions.

Proposition 17 *Given a collection of utility functions on $2^{\mathcal{R}}$ given in a compact numerical language under logical form:*

- *the problem of deciding whether there exists an envy-free allocation among those that maximize utilitarian social welfare is Δ_2^p -complete, even if $N = 2$, and even if the agents have identical preferences.*
- *the problem of deciding whether there exists an envy-free allocation among those that maximize egalitarian social welfare is Δ_2^p -complete, even if $N = 2$.*

Proof For both results, membership comes easily from the fact that the maximum value of social welfare can be computed by dichotomy over the set of all possible social welfare values; there are exponentially many, therefore we need a polynomial number of NP oracles to do this; after this step is done, it suffices to guess an allocation and check that it is envy-free and that it maximizes social welfare, adding just one more NP oracle.

Hardness is obtained in both utilitarian and egalitarian cases by a simple reduction to an instance of the fair division problem with preferences expressed in the weighted propositional language from the following problem:

Problem 5: MAX-SAT-ASG_{even} (Wagner, 1987)

INSTANCE: A propositional formula χ in Conjunctive Normal Form, over a set of propositional variables $V = \{v_1, \dots, v_n\}$, and a weight function w over interpretations $I : V \rightarrow \{0, 1\}$, defined by $w(I) \stackrel{\text{def}}{=} \sum_i I(v_i) \times 2^{i-1}$.

QUESTION: Is $\max_M \text{model of } \chi w(M)$ an even number (in other words, is v_1 falsified in the model of maximal weight) ?

We will suppose that the formula χ has at least on model M such that $M \not\models v_1$. It does not change the complexity, because if v_1 is verified in every model of χ , clearly the answer to the problem MAX-SAT-ASG_{even} is no: as a consequence, every instance $\langle \chi, V \rangle$ without any assumption on χ can be solved by first checking if $\neg v_1 \wedge \chi$ is unsatisfiable (which is a coNP-complete problem), and then, if not, by solving the UNSAT-OR-MAX-SAT-ASG_{even} problem on an instance that has at least one model falsifying v_1 .

Utilitarian social welfare: From an instance $\langle \chi, V \rangle$ of MAX-SAT-ASG_{even}, we create the following instance $\mathcal{P}(\chi, V)$:

- Agents:** 2 agents;
- Objects:** for each literal v_i of χ , we create two objects x_i and x'_i , except for v_1 , for which only one object x_1 is created, and we add two objects y and y' ;
- Preferences:** the agents 1 and 2 have identical preferences, and both ask for $\langle (\psi \wedge y) \vee (\psi' \wedge y'), 2^{n+1} \rangle, \langle x_1 \wedge y, 1 \rangle, \dots, \langle x_n \wedge y, 2^{n-1} \rangle, \langle x'_2 \wedge y', 2 \rangle, \dots, \langle x'_n \wedge y', 2^{n-1} \rangle$, with ψ the formula χ in which each symbol v_i has been replaced by x_i , and ψ' the formula χ in which each symbol v_i (except v_1 replaced by x_1) has been replaced by x'_i .

Let (M_1, M_2) be a pair of models of χ (with possibly $M_1 = M_2$) such that $M_2 \not\models v_1$. Then we define the allocation π_{M_1, M_2} by: $\pi_{M_1, M_2}(1) = \{y\} \cup \{x_i | M_1 \models v_i\}$ and $\pi_{M_1, M_2}(2) = \{y'\} \cup \{x'_i | M_2 \models v_i\}$.

The proof is based on the following lemma:

Lemma 17 *There exists an envy-free allocation among those that maximize utilitarian social welfare if and only if there are two models M_1 and M_2 of χ (possibly $M_1 = M_2$) with $M_2 \not\models v_1$, such that π_{M_1, M_2} is envy-free and maximizes utilitarian social welfare.*

Proof Let π be an allocation maximizing the utilitarian social welfare. Let M be a model of χ falsifying v_1 (our hypothesis is that there is at least one). Then $F(\pi_{M, M}(1)) \models \psi \wedge y$ and $F(\pi_{M, M}(2)) \models \psi' \wedge y'$, which proves that the individual utility of the two agents is at least 2^{n+1} . Hence there is at least one allocation whose utilitarian social welfare is greater than or equal to 2^{n+2} . Therefore, an allocation π maximizing the utilitarian social welfare must be such that $F(\pi(1)) \models y \wedge \psi$ and $F(\pi(2)) \models y' \wedge \psi'$, or vice versa. Moreover, either $x_1 \notin \pi(1)$, or $x_1 \notin \pi(2)$. Suppose that $x_1 \in \pi(2)$: swapping the shares of the agents leads to an allocation π' which is completely equivalent to π with respect to Pareto-efficiency and envy-freeness, due to the identical preferences. We can therefore assume w.l.o.g that π is such that $x_1 \notin \pi(2)$

Since $F(\pi(1)) \models \psi$, then there is a model M_1 of χ such that $\pi(1) = \{y\} \cup \{x_i | M_1 \models v_i\} \cup S_1$, where $S_1 \subseteq \{x_1, x'_2, \dots, x'_n\}$. Similarly, there is a model M_2 such that $M_2 \not\models v_1$, and $\pi(2) = \{y'\} \cup \{x'_i | i > 1, M_2 \models v_i\} \cup S_2$, where $S_2 \subseteq \{x_1, \dots, x_n\}$. Now consider the allocation π_{M_1, M_2} , that is well-defined since $M_2 \not\models v_1$. We have $u_1(\pi) = u_1(\pi_{M_1, M_2})$, since the x'_i do not satisfy any formula of the preferences of the agent 1 without y' (given to the agent 2), and $u_2(\pi) = u_2(\pi_{M_1, M_2})$ for the same reasons. In other terms, π_{M_1, M_2} gives the same utility as π to both agents. π_{M_1, M_2} is thus envy-free and maximizes the utilitarian social welfare. \square

By Lemma 17, we can thus restrict our problem to the allocations of the form π_{M_1, M_2} . We have, for each M_1 and M_2 as defined earlier, $u_1(\pi_{M_1, M_2}) = 2^{n+1} + w(M_1)$ and $u_2(\pi_{M_1, M_2}) = 2^{n+1} + w(M_2)$; thus $sw^*(\pi_{M_1, M_2}) = 2^{n+2} + w(M_1) + w(M_2)$. We have: $\operatorname{argmax}_{\pi_{M_1, M_2}} sw^*(\pi_{M_1, M_2}) = \pi_{\operatorname{argmax}_{(M_1, M_2)} \{w(M_1) + w(M_2) | M_1 \not\models v_1 \text{ or } M_2 \not\models v_1\}}$. Given the symmetry of the problem, we can assume that only M_2 has to satisfy $M_2 \not\models v_1$, thus the latter allocation becomes: $\pi_{M_{opt}, \operatorname{argmax}_{M_2} \{w(M_2) | M_2 \not\models v_1\}}$, where M_{opt} is the model of χ of maximal weight.

Suppose $M_{opt} \not\models v_1$, then the allocation maximizing the utilitarian social welfare is $\pi_{M_{opt}, M_{opt}}$ and it is clearly envy-free, because both agents have the same utility. Now suppose that $M_{opt} \models v_1$, then the allocation maximizing the utilitarian social welfare is $\pi_{M_{opt}, M_{opt}'}$, where M_{opt}' is the model of χ of maximal weight that assigns v_1 to false. We have $w(M_{opt}') < w(M_{opt})$, thus $u_1(\pi_{M_{opt}, M_{opt}'}) > u_2(\pi_{M_{opt}, M_{opt}'})$, hence this allocation is not envy-free.

The latter reduction is clearly polynomial (recall that the weights 2^{n+1} can be encoded using a linear space). This proves the proposition for the utilitarian case.

Egalitarian social welfare: From an instance $\langle \chi, V \rangle$ of MAX-SAT-ASG_{even}, we create the following instance $\mathcal{P}(\chi, V)$:

Agents: 2 agents;

Objects: for each literal v_i of χ , we create two objects x_i and x'_i , and we add two objects y and y' ;

Preferences: the preferences of the agent 1 are $\langle x_1, 1 \rangle, \dots, \langle x_n, 2^{n-1} \rangle, \langle \psi \wedge y, 2^n \rangle$, and the preferences of the agent 2 are $\langle y \vee y', 2^{2n} \rangle, \langle x_1, 1 \rangle$, with ψ the formula χ in which each symbol v_i has been replaced by x_i .

Every allocation that maximizes egalitarian social welfare must give at least a utility of 2^{2n} to agent 2. In that case, the egalitarian social welfare will be given by the utility of agent 1, because this utility cannot be greater than 2^{2n} . Therefore, maximizing social welfare comes down to maximizing the utility of agent 1, or in other words to give her the items corresponding to the model of χ of maximum weight. If v_1 is set to true in this model, then x_1 is given to agent 1, and since y is also given to agent 1, agent 2 could get a strictly higher utility with agent 1's share. Therefore this allocation is not envy-free. If v_1 is set to false in the latter model, then x_1 is not given to agent 1, and thus can be given to agent 2, producing an envy-free allocation. Since the reduction is polynomial, it proves hardness for the egalitarian case. \square

We can notice here that the combination of envy-freeness with a numerical criterion such as classical utilitarianism or egalitarianism induces a complexity gap, since, as stated by Bouveret, Fargier, Lang, and Lemaître (2005), the complexity of the problems of maximizing the classical utilitarian or egalitarian collective utility functions, with agents having weighted logical preferences, are “only” NP-complete.

From the previous proof in the utilitarian case, we can notice that the hardness result still holds if we require the allocation to be Pareto-efficient instead of maximizing the utilitarian social welfare. It suggests that in the case of a language extending the weighted propositional formulae, the EEF EXISTENCE problem with identical preferences is much harder than in the case of agents having identical dichotomous preferences. We have actually the following result:

Proposition 18 *Given a collection of N identical utility functions on $2^{\mathcal{R}}$ given in compact numerical language under logical form the problem of deciding whether a Pareto-efficient and envy-free allocation exists is Δ_2^p -complete, even if $N = 2$, even if the preferences are monotonic.*

Proof Since the preferences are identical, any envy-free allocation satisfies all the agents equally. Thus, a Pareto-efficient and envy-free allocation, if there is one, is an allocation that gives everyone a utility of M , maximal among the set of allocations that satisfy everyone equally. This value M can be computed using a polynomial number of NP oracles (like in the previous proof). Now that we have this value of M , checking if there is a Pareto-efficient and envy-free allocation comes down to check if there is no allocation giving at least M to all the agents, and at least $M + 1$ to at least one agent, which is a problem in coNP, and hence adds only one call to an NP oracle.

For the hardness proof, one may notice that the same reduction as the one used in the utilitarian case of the proof of Proposition 17 works in this case, because an allocation is Pareto-efficient and envy-free in this particular problem if and only if it is envy-free and maximizes the utilitarian social welfare. \square

5.3 Additive Numerical Preferences

A last case that we consider here is the case of *additive* numerical preferences. Additive numerical preferences is a degenerate case of weighted logical preferences, where all the formulae are single positive literals. In other words, the preferences of an agent i are given by a set Δ_i of pairs $\langle x_k, w_k \rangle$, where x_k is an object and w_k is the weight (possibly 0) associated with the object. The utility function associated with such preferences is thus the following:

$$\begin{aligned} u : 2^X &\rightarrow \mathbb{Z} \\ S &\mapsto \sum_{x_k \in S} w_k. \end{aligned}$$

Notice that the agents’ preferences are monotonic if and only if all the numbers w_k are positive.

This preference representation language is the most natural one when dealing with resource allocation problems; however, it is unable to express compactly any kind of dependencies (superadditivity or subadditivity) between the objects. In particular, it does not extend dichotomous preferences. Hence the previous hardness results do not extend to additive preferences. However, since we are still able to compare two alternatives in polynomial time, membership to Σ_2^p is guaranteed.

Our intuition is that this problem is as hard as the EEF EXISTENCE problem with dichotomous preferences:

Conjecture 1 EEF EXISTENCE *with additive numerical preferences is Σ_2^P -complete, even if the preferences are monotonic.*

All we know about this problem is that it is NP-hard (this is implied by Proposition 20 presented later) and in Σ_2^P , but its precise complexity remains open. However, things become much easier if we only require the allocation to be *complete*, instead of Pareto-efficient. This case has already been investigated by Lipton et al. (2004), and we have the following result:

Proposition 19 (Lipton et al., 2004) *The problem of deciding whether there exists a complete envy-free allocation for agents having additive preferences is NP-complete, even if their preferences are monotonic.*

Other restrictions of the EEF EXISTENCE problem with additive preferences are worth being studied. First, we study as in the dichotomous case the restriction to identical additive preferences:

Proposition 20 EEF EXISTENCE *with N identical additive numerical preferences is NP-complete, for any fixed $N \geq 2$. The same result holds if we require the preferences to be monotonic.*

Proof Membership is easy to prove. Since all the preferences are identical (we write $\langle u(x_1), \dots, u(x_p) \rangle$ the utility vector associated to the set of object), an allocation is Pareto-efficient if and only if it gives each object x_j such that $u(x_j) > 0$ to one agent, and trashes each object x_j such that $u(x_j) \leq 0$. Moreover, an allocation is envy-free if and only if all the agents have the same utility. These two latter properties can be checked in polynomial-time, hence membership to NP.

Hardness comes from a reduction from PARTITION:

Problem 6: PARTITION

INSTANCE: A finite set S and a size $s(a) \in \mathbb{N}$ for each $a \in S$.

QUESTION: Is there a subset $S' \subseteq S$ such that $\sum_{a \in S'} s(a) = \sum_{a \in S \setminus S'} s(a)$?

From a given instance $\langle S, s \rangle$ of the PARTITION problem, we create the following instance $\mathcal{P}(S, s)$ of the EEF EXISTENCE problem:

Agents: 2 agents;

Objects: to each $a \in S$, we associate an object x_a ;

Preferences: The two agents' preferences are identical and defined by the size of the elements in the initial set: $u(x(a)) = s(a)$.

There is a Pareto-efficient and envy-free allocation for $\mathcal{P}(S, s)$ if and only if there is an allocation π such that $\sum_{x \in \pi(1)} u(x) = \sum_{x \in \pi(2)} u(x)$, that is, the PARTITION problem returns true. The reduction can be clearly done in polynomial time, which proves the proposition. \square

Now another interesting case is the case where the preferences are not necessary identical, but where all the atomic utilities $u_i(x_j)$ are either 0 or 1. In other words, an agent either wants an object or does not want it, and each agent wants to maximize the number of desired objects she gets.

Proposition 21 EEF EXISTENCE *with additive 0–1-preferences (i.e. $\forall i, j, u_i(x_j) \in \{0, 1\}$) is NP-complete.*

Proof Pareto-efficiency is easy to check in this case. We can first safely remove the objects that do not appear anywhere in the preferences. Afterwards, an allocation is Pareto-efficient if and only if each object x_j is given to an agent i such that $u_i(x_j) = 1$, for the following reasons. (\Rightarrow) Let π be a Pareto-efficient allocation, and suppose that there is an x_j which is not given to any agent, or which is given to an agent such that $u_i(x_j) = 0$. Let k be an agent such that $u_k(x_j) = 1$ (there is one since we previously trashed the undesired objects). Then giving x_j to the agent k increases k 's utility while the others' utilities remains the same. Thus π is Pareto-dominated. (\Leftarrow) Let π be an allocation such that each object x_j is given to an agent i such that $u_i(x_j) = 1$, and suppose π is Pareto-dominated by an allocation π' . Then $\sum_{i \in I} u_i(\pi'(i)) = \sum_{i \in I} \sum_{x_j \in \pi'(i)} u_i(x_j) > \sum_{i \in I} u_i(\pi(i)) = \sum_{i \in I} \sum_{x_j \in \pi(i)} u_i(x_j) = p$. Therefore there is at least one $u_i(x_j) > 1$, which is not possible due to our restriction to 0–1-preferences.

It thus give an simple way to check Pareto-efficiency, by just checking if the sum of utilities is equal to the number of objects p that are desired by at least one agent. As usual, envy-freeness can be verified in polynomial time; therefore EEF EXISTENCE with additive 0–1-preferences is in NP.

Hardness can be proved by a polynomial reduction from EXACT COVER BY 3-SETS (problem 3). Given an instance $\langle S, C = \langle S_1, \dots, S_{|C|} \rangle \rangle$ of EXACT COVER BY 3-SETS, create the following EEF EXISTENCE instance $\mathcal{P}(C, S)$ (we will suppose that the elements of S will be written a_i , with $i \in \{1, \dots, |S|\}$):

- Agents:** a set of $3|C|$ agents gathered by triples $\{3i - 2, 3i - 1, 3i\}$;
Objects: a set of $|S| + 3|C|$ items $X = M \cup D$ (M for “main”, and D for “dummy”), with $M = \{m_1, \dots, m_{|S|}\}$, and $D = \bigcup_{i \in \{1, \dots, |C|\}, j \in \{1, 2, 3\}} \{d_{i,j}\}$;
Preferences: the agents from $\{3i - 2, 3i - 1, 3i\}$ all desire the same set of objects $\bigcup_{a_k \in S_i} \{m_k\} \cup \{d_{i,1}, d_{i,2}, d_{i,3}\}$ (the three objects corresponding to S_i plus the three dummy objects $d_{i,j}$).

If there is an exact cover C' for the instance $\langle C, S \rangle$, then we will consider the following allocation: each agent from the triple $\{3i - 2, 3i - 1, 3i\}$ gets respectively $d_{i,1}$, $d_{i,2}$, and $d_{i,3}$, and if $S_i \in C'$, each one of these three agents gets one of the three objects m_k corresponding to the elements of the set S_i . This allocation is admissible and Pareto-efficient (because all the objects are allocated). It is also envy-free, for the following reasons:

- The agents from the same triple cannot envy each other, because they are equally satisfied.
- An agent k_1 cannot envy any agent k_2 from another triple, because the only objects k_1 could envy from k_2 's share are the m_i . k_2 having at most one m_i , and k_1 having a utility of at least one, k_1 cannot envy k_2 .

The rest of the proof is based on the following result: if an allocation π is Pareto-efficient and envy-free for $\mathcal{P}(C, S)$ then we must have $\pi(3i - 2) \cup \pi(3i - 1) \cup \pi(3i) = \bigcup_{a_k \in S_i} \{m_k\} \cup \{d_{i,1}, d_{i,2}, d_{i,3}\}$ or $\pi(3i) \cup \pi(3i - 1) \cup \pi(3i) = \{d_{i,1}, d_{i,2}, d_{i,3}\}$. This is easy to show. Since the agents from the triple $\{3i - 2, 3i - 1, 3i\}$ are the only ones to desire the objects $d_{i,k}$, these three objects must be given to these three agents, for the allocation to be efficient. Since these three agents have the same preferences, the allocation has to satisfy them equally in order to be envy-free. Thus the number of objects allocated to the three agents must be divisible by 3, which gives only two possible numbers, 3 and 6, and hence only two possible allocations.

Suppose that there is a Pareto-efficient and envy-free allocation π for $\mathcal{P}(C, S)$. Consider the sub-collection $C' = \langle S_1, \dots, S_{|C'|} \rangle$ made of the triples S_i from the collection C such that $\pi(3i - 2) \cup \pi(3i - 1) \cup \pi(3i) = \bigcup_{a_k \in S_i} \{m_k\} \cup \{d_{i,1}, d_{i,2}, d_{i,3}\}$. Then we have the following results.

- The S_i are pairwise disjoint. Suppose that there is a pair (i, j) such that $i \neq j$ and there is an element a_k belonging to both S_i and S_j . Then m_k is allocated to two different agents: one member of triple $\{3i - 2, 3i - 1, 3i\}$, and one member of triple $\{3j + 1, 3j + 2, 3j + 3\}$, which is impossible.
- $\bigcup_{i \in \{1, \dots, |C'|\}} S_i = S$. Let a_k be an element of S . Since π is Pareto-efficient, then m_k must be allocated to one agent that wants it (say this agent belongs to the triple $\{3j + 1, 3j + 2, 3j + 3\}$), unless no one wants it, which occurs if $\bigcup_{i \in \{1, \dots, |C'|\}} S_i \neq S$. Then, by the previous result, all the objects from $\bigcup_{a_l \in S_j} \{m_l\}$ must be allocated to this triple. Consequently, $S_j \in C'$. Since $a_k \in S_j$, a_k belongs to at least one set from the collection C' .

Therefore, C' is an exact cover of S , which finally proves the proposition. \square

We can see with Proposition 21 and Conjecture 1 that there is a huge complexity gap between the problem where we allow the weights to be freely given and the problem where we require the weights to be 0 or 1 (at least if the conjecture is true). The natural question it raises is to know if this complexity fall is specific to the 0–1-preferences, or if it occurs for any fixed upper bound of the weights.

Conjecture 2 *The complexity of the EEF EXISTENCE problem with additive 0–1–...– k preferences for $k \geq 2$ fixed is as hard as the general problem with unbounded additive preferences.*

The precise complexity of this problem remains an open problem, but as it is stated in the conjecture, our intuition is that it is as hard as the EEF EXISTENCE problem with unbounded additive preferences.

Another natural problem is raised by the Proposition 21: what is the complexity of the EEF EXISTENCE problem with *stratified* 0–1-preferences? By “stratified 0–1-preferences”, we mean that the preferences are given by a set of pairs $\langle x_k, p \rangle$, where x_k is an object and p is a priority level. Comparing two sets of objects comes down to compare lexicographically the vectors where the component at index i is the number of objects of priority i in the share of the agent. Notice that this problem is not an instance of the EEF EXISTENCE problem with additive preferences, nor an instance of the EEF EXISTENCE problem with logical numerical preferences under logical form. However it is easy to see that it remains in Σ_2^p , but its precise complexity remains unknown.

Finally, we investigate the case where the number of objects is less than the number of agents. One could think intuitively that the problem is trivial in this case. However, it is not always the case, as we will see. To begin with, the following results shows that the problem is easy for monotonic preferences:

Proposition 22 *Let \mathcal{P} be an allocation problem with n agents having additive monotonic numerical preferences and all wanting at least one object, and p objects that are desired by at least one agent.*

- *If $p < n$, then there is no Pareto-efficient and envy-free allocation.*
- *If $p = n$, the problem of deciding whether there exists an efficient and envy-free allocation for agents having monotonic additive preferences is in \mathbf{P} .*

Proof • Each object being desired by at least one agent, every Pareto-efficient allocation is complete. If the number of objects p is strictly lower than the number of agents N , then at least one agent i is unsatisfied. Consequently, there is an agent j that obtains an object wanted by i , hence creating envy. Thus no Pareto-efficient allocation can be envy-free.

- Since there are as many objects as agents, each agent should receive one object that she values the most (that is, such that $u_i(\{x\})$ is maximal), for an allocation π to be Pareto-efficient and envy-free. Indeed, if an agent i receives an object that is not her preferred one, it means that another agent j receives it (otherwise the allocation would not be Pareto-efficient), creating envy from i . Therefore, checking the existence of a Pareto-efficient and envy-free allocation comes down in this case to checking if it is possible to give to every agent one of her best-valued objects. It comes down to checking if there is a perfect matching in the bipartite graph made with one node per agent on the one side, and one node per object on the other side, and connecting an

agent i to an object x if and only if x is among the best-valued objects in agent i 's preferences. Such a perfect matching can be computed in polynomial time, hence the result. \square

Interestingly, the latter result does not hold at all if we allow non-monotonic preferences. In that case, the complexity increases up to the complexity of the general EEF EXISTENCE problem with additive preferences.

Proposition 23 *The EEF EXISTENCE problem with additive numerical preferences and such that the number of objects is less than the number of agents has the same complexity as the EEF EXISTENCE problem with additive numerical preferences, but no assumption on the number of objects.*

Proof Let us consider an instance $\langle I, X, \langle \dots, u_i(x_j), \dots \rangle$ of the EEF EXISTENCE problem with additive numerical preferences, with N agents and p objects ($p > N$). We create the following instance $\mathcal{P}(\langle I, X, \langle \dots, u_i(x_j), \dots \rangle)$:

Agents: $p + 3$ agents (the number of agents is not important, it just has to be greater than the number of objects and of the initial number of agents);

Objects: the p initial objects x_i plus two dummy ones d_1 and d_2 ;

Preferences: the preferences of the N first agents are the same as in $\langle I, X, \langle \dots, u_i(x_j), \dots \rangle$; the preferences of the $(N + 1)^{\text{st}}$ agent are $u_{N+1}(\{d_1\}) = u_{N+1}(\{d_2\}) = 1$ and $u_{N+1}(x_j) = 0$ for the other items x_j ; and the preferences of the remaining agents are $u_{N+1}(\{d_1\}) = 1$, $u_{N+1}(\{d_2\}) = -2$ and $u_{N+1}(x_j) = 0$ for the remaining objects.

If there is an efficient and envy free allocation π for $\langle I, X, \langle \dots, u_i(x_j), \dots \rangle$ then it can be easily checked that the allocation that gives the same items to the N first agents of $\mathcal{P}(\langle I, X, \langle \dots, u_i(x_j), \dots \rangle)$, $\{d_1, d_2\}$ to the $(N + 1)^{\text{st}}$ agent, and nothing to the remaining ones is efficient and envy-free. Conversely, any Pareto-efficient and envy-free allocation for $\mathcal{P}(\langle I, X, \langle \dots, u_i(x_j), \dots \rangle)$ yields a Pareto-efficient and envy-free allocation for $\langle I, X, \langle \dots, u_i(x_j), \dots \rangle$ by restricting it to the N first agents and all the objects but the two dummy ones. \square

6. Related Work and Discussion

As already argued in the Introduction, computational studies in resource allocation either concern divisible goods, or focus on classical utilitarianism such as in combinatorial auctions. Existing work on fair division of indivisible goods, on the other hand, is mainly axiomatic, and its computational aspects have been neglected, except in a few papers that we are mentioning below.

Our results have of course a lot in common with complexity results for combinatorial auctions. After all, the structure of the problems are, to some extent, similar: items are indivisible, allocations are preemptive¹⁰, and each agent has preferences over sets of items expressed in some compact representation language. Logical bidding languages have also been designed (Boutilier & Hoos, 2001). However, the complexity results completely differ: the standard decision problem for combinatorial auctions is NP-complete (Rothkopf, Pekeč, & Harstad, 1998) while the decision problems considered here are typically located at the second level of the polynomial hierarchy, even in the degenerate case where preferences are dichotomous. This can be explained by the fact that combinatorial auctions care only about efficiency, not about envy-freeness. Requiring both criteria together (efficiency and

10. Here, “preemptive” means that an object cannot be allocated to more than one agent. This assumption can be absent from some problems implying for example virtual objects, such as software licences.

envy-freeness) makes things much more difficult: while, under the usual assumption that preferences are monotonic, efficiency is a monotonic property (allocating more goods never makes an allocation less efficient), envy-freeness is not (allocating more goods to an agent may generate envy)¹¹. This is the reason why there may not exist an EEF allocation, and this is also the source of the high complexity of the problem.

Moreover, due to the failure of monotonicity for envy-freeness, searching for an EEF allocation cannot simply be formulated as the maximization or minimization of a simple criterion, which is problematic when designing local search algorithms or approximation schemes. A few authors (Lipton et al., 2004; Chevaleyre, Endriss, Estivie, & Maudet, 2007a; Brams, Jones, & Klamler, 2007) have suggested to relax the envy-freeness criterion and make it a gradual notion, by defining a measure on envy-freeness. Lipton et al. (2004) assume that the input consists of numerical utility functions over sets of goods; the degree to which agent i envies agent j in allocation π is either defined as the envy difference $d_{i,j}(\pi) = \max(0, u_i(\pi(j)) - u_i(\pi(i)))$ or the envy ratio $r_{ij}(\pi) = \max\left(1, \frac{u_i(\pi(j))}{u_i(\pi(i))}\right)$. In both cases, the global degree of envy is then the maximum degree of envy between any pair of players. Alternative definitions of the degree of envy in a society have been proposed by Chevaleyre et al. (2007a), such as the number of envious agents, the number of pairs (i, j) such as i envies j , or the sum of all local degrees of envy, which is relatively similar to the measure of envy considered by Brams et al. (2007), based on the maximum number of agents that any single agent may envy. Chevaleyre et al. (2007b) suggest a radically different way of relaxing envy-freeness: the society comes along with an undirected graph (which reflects acquaintance between agents), and an allocation is envy-free if and only if no agent envies an agent to whom she is connected.

Then, Lipton et al. (2004) focus on the search for a complete allocation with minimum envy; moreover, in the case of additive utilities, they provide approximation schemes. Another work on approximation algorithms for the fair allocation of indivisible goods is the one by Asadpour and Saberi (2007), who assume from the beginning that utilities are linear and consider the problem of finding a maximally equitable allocation, that is, an allocation maximizing the utility of the least satisfied agent (cf. the problem considered in our Proposition 17); they do not consider envy-freeness at all¹².

All the approaches we have considered so far (including ours) assume that the allocation is computed in a centralized way by a neutral, objective agent. In other contexts, this centralized approach is not possible or not realistic, and the allocation is obtained in a decentralized way, by successive negotiations between groups of agents. Such an approach has been initiated by Sandholm (1998), who studies convergence properties on the allocation depending on the structural restrictions made on exchanges of goods that may occur. It has been further investigated by Dunne (2005) and Dunne et al. (2005), who study the computational complexity of these negotiation problems, and by Endriss, Maudet, Sadri, and Toni (2006). In these approaches, the optimality criterion is the classical utilitarian

11. Note that it is not antimonotonic either: allocating less goods to an agent may generate envy as well.

12. Note that alternative ways of relaxing the search for EEF allocations exist. One may for instance keep envy-freeness as a hard requirement and relax efficiency, or allow for relaxing both and look for an allocation that shows a good trade-off between efficiency and envy-freeness (considering then the problem as a two-criteria optimization problem).

social welfare. Envy-freeness is considered in a distributed setting by Chevaleyre et al. (2007a, 2007b).

So far, by “computational issues” we referred to the design and the study of algorithms to be run on computers so as to find an allocation, and not to designing and studying protocols that query agents interactively so as to gather enough information for a solution to be determined. These “procedural issues”, although extensively studied in the literature of fair division of *divisible* goods (see below), as well as in voting (Conitzer & Sandholm, 2005) have rarely been considered for allocating indivisible goods, with the notable exception of Herreiner and Puppe (2002), who study the properties of interactive protocols where two agents enumerate their preferred bundles one by one, until an allocation is found. As mentioned in the Introduction, the drawback of such protocols is that they are exponentially long, and henceforth infeasible as soon as the number objects is more than a few units.


Beyond these few works on the computational aspects of fair division of *indivisible* goods, there are much more that consider the computational and procedural aspects of fair division of *divisible goods* (or, at least, assume that at least one good is divisible — e.g., money). The literature on the subject is vast and the techniques are quite far from those used for allocating indivisible goods (see in particular the literature on cake-cutting algorithms, e.g. Brams & Taylor, 1996; Robertson & Webb, 1998) so we do not find it relevant to give here a detailed bibliography on this subject. The interested reader can refer to the book by Brams (2008) that covers fair resource allocation in both the indivisible and the divisible case.

7. Conclusion

We have studied several computational aspects the search for efficient and envy-free allocations in fair division problems with indivisible goods. Our results of Section 3 show that in the case of dichotomous preferences, the search for such allocations can be reduced to the search for preferred models in prerequisite-free default logic. Such a connection was rather unexpected, and it implies that the practical search for EEF allocations can be done using existing algorithms for default logic. However this search is likely to be very time-consuming, due to our complexity results: indeed, our extensive study of the complexity of deciding whether an efficient and envy-free allocation exists, under various restrictions (dichotomous preferences or not, two agents or more, all agents having identical preferences or not, monotonic preferences or not) and for various notions of efficiency (Pareto-efficiency, completeness, maximum number of satisfied agents, maximum social welfare — classical utilitarian or egalitarian), seems to show that the problem is intrinsically very difficult, since it lies at the second level of the polynomial hierarchy, *even if preferences are dichotomous and monotonic*. This implies that designing fast algorithms for solving the problem in the general case is out of reach. We may first focus on those restrictions for which the problem is at most NP-complete. Unfortunately, these restrictions (agents with identical preferences; only two agents; purely conjunctive, purely disjunctive or 2-CNF preferences; search for complete allocations without any other efficiency requirement; additive 0-1 preferences) are very compelling and imply a huge loss of generality.

The complexity results introduced in this paper are summed-up in Figure 1 and Table 2.

Several issues for further research remain to be explored.



★ Proof not included in this paper.

- 1 A problem whose complexity has been proved in this paper (the mapping between the numbers and the problems is specified in table 2).
- 17 A problem whose complexity was already known in the literature.
- 22? A problem whose complexity remains unknown.
- The intersection of the problems corresponding to its outgoing edges.
- i → j Problem *i* is included in problem *j*. Arcs that can be obtained by transitivity are omitted.

Figure 1: The different problems and their complexity classes and inclusion relations.

	Efficiency	number of agents	preferences	monotonicity	comp.
Dichotomous preferences					
1, 1'	Pareto-eff.	not fixed	any	yes (1) or no (1')	Σ_2^P -c.
2	Pareto-eff.	not fixed or fixed with $N \geq 2$	identical	yes	NP-c.
3	Pareto-eff.	not fixed or fixed with $N \geq 2$	identical	no	coBH ₂ -c.
4	Pareto-eff.	2 agents	any	yes	NP-c.
5	Pareto-eff.	2 agents	any	no	coBH ₂ -c.
6, 6', 6''	complete all.	not fixed (6, 6', 6'') or fixed with $N \geq 2$ (6, 6')	ident. or not	yes (6) or no (6', 6'')	NP-c.
7, 7'	max nb of ag.	not fixed	any	yes (7) or no (7')	Θ_2^P -c.
8, 8'	Pareto-eff.	any	disjunctions	yes (8) or no (8')	P
9, 9'	Pareto-eff.	any	conjunctions	yes (9) or no (9')	NP-c.
10, 10'	Pareto-eff.	any	conj. with condition 1	yes (10) or no (10')	$O(1)$
11, 11'	Pareto-eff.	any	\mathcal{C} st $\text{SAT}(\mathcal{C}) \in \mathbf{P}$ and closed by \wedge	yes (11) or no (11')	NP-c.
Non-dichotomous preferences					
12, 12'	Pareto-eff.	not fixed	numerical	yes (12) or no (12')	Σ_2^P -c.
13	utilitarian sw	not fixed or fixed with $N \geq 2$	numerical	no	Δ_2^P -c.
14	egalitarian sw	not fixed or fixed with $N \geq 2$	numerical	no	Δ_2^P -c.
15, 15'	Pareto-eff.	not fixed or fixed with $N \geq 2$	numerical, identical	yes (15) or no (15')	Δ_2^P -c.
16	Pareto-eff.	not fixed	additive	no	Σ_2^P -c. ?
17, 17'	complete all.	not fixed	additive	yes (17) or no (17')	NP-c.
18, 18'	Pareto-eff.	not fixed or fixed with $N \geq 2$	additive ident.	yes (17) or no (17')	NP-c.
19	Pareto-eff.	any	additive 0-1	yes	NP-c.
20	Pareto-eff.	> Nb of objects	additive	yes	$O(1)$
21	Pareto-eff.	= Nb of objects	additive	yes	P
22	Pareto-eff.	\geq Nb of objects	additive	no	Σ_2^P -c. ?

Table 2: The set of resource allocation problems studied in this paper. Their complexity classes are represented in figure 1.

First, knowing that there is no efficient and envy-free allocation for a given problem is not very helpful in practice when an allocation has to be found anyway. The solution then consists in defining functions that return an allocation in all cases, even when envy-freeness and efficiency cannot be jointly met. A way of addressing this issue consists in defining suitable relaxations of the problem, such as: (a) using a measure of envy instead of seeing envy-freeness as a strict criterion (as suggested by Lipton et al., 2004; Chevaleyre et al., 2007b); (b) make envy-freeness a relative notion, for instance by introducing an acquaintance graph between agents (Chevaleyre et al., 2007b); or (c) keeping envy-freeness as a strict criterion and relaxing efficiency. In all cases, new problems arise, the complexity of which has to be identified.

Second, our results are mostly negative, since most of the interesting problems we studied are NP-hard (and often even worse), therefore, it would be worth pursuing work and design practical algorithms for these problems. Most likely this would require coming up with appropriate optimization criteria and then either (a) giving polynomial algorithms that can approximate the desired objective (Lipton et al., 2004) or (b) implementing and experimenting local search algorithms with relevant heuristics.

Third, throughout our paper it is assumed that everyone’s preferences are completely known. In reality, presumably agents need to report their preferences, which introduces the issue of strategic misreporting (manipulation). One direction for future research would be to investigate how to prevent this, that is, mechanism design aspects.

Acknowledgments

We thank Michel Lemaître for stimulating discussions about fair division and compact representation; Thibault Gajdos for stimulating discussions about envy-freeness and for pointing to us some relevant papers; Steven Brams, for giving us some feedback on an earlier version of the paper and pointing to us some relevant references; and the participants of the AgentLink *Technical Forum Group on Multiagent Resource Allocation*. This work has been partly supported by the project ANR-05-BLAN-0384 “Preference Handling and Aggregation in Combinatorial Domains”, funded by *Agence Nationale de la Recherche*.

References

- Asadpour, A., & Saberi, A. (2007). An approximation algorithm for max-min fair allocation of indivisible goods. Tech. rep., Department of Management Science and Engineering, Stanford University, Stanford.
- Baral, C. (2003). *Knowledge Representation, Reasoning and Declarative Problem Solving*. Cambridge University Press.
- Bogomolnaia, A., Moulin, H., & Stong, R. (2005). Collective choice under dichotomous preferences. *Journal of Economic Theory*, 122, 165–184.
- Boutilier, C., & Hoos, H. H. (2001). Bidding languages for combinatorial auctions. In *Proc. of the 17th International Joint Conference on Artificial Intelligence (IJCAI-01)*, pp. 1211–1217, Seattle, Washington, USA.

- Bouveret, S., Fargier, H., Lang, J., & Lemaître, M. (2005). Allocation of indivisible goods: a general model and some complexity results. In Dignum, F., Dignum, V., Koenig, S., Kraus, S., Singh, M. P., & Wooldridge, M. (Eds.), *Proceedings of AAMAS'05*, Utrecht, The Netherlands. ACM Press.
- Brams, S., & Fishburn, P. (1978). Approval voting. *American Political Science Review*, 72(3), 831–857.
- Brams, S. J. (2008). *Mathematics and Democracy: Designing Better Voting and Fair-Division Procedures*. Princeton University Press.
- Brams, S. J., Edelman, P. H., & Fishburn, P. C. (2003). Fair division of indivisible items. *Theory and Decision*, 55(2), 147–180.
- Brams, S. J., Jones, M. A., & Klamler, C. (2007). Divide-and-conquer: A proportional, minimal-envy cake-cutting procedure. Tech. rep., NYU Department of Politics.
- Brams, S. J., & King, D. L. (2005). Efficient fair division: Help the worst off or avoid envy?. *Rationality and Society*, 17(4), 387–421.
- Brams, S. J., & Taylor, A. (1996). *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge Univ. Press.
- Chevalleyre, Y., Dunne, P. E., Endriss, U., Lang, J., Lemaître, M., Maudet, N., Padget, J., Phelps, S., Rodríguez-Aguilar, J. A., & Sousa, P. (2006). Issues in multiagent resource allocation. *Informatica*, 30, 3–31. Survey paper.
- Chevalleyre, Y., Endriss, U., Estivie, S., & Maudet, N. (2004). Multiagent resource allocation with k -additive utility functions. In *Proc. DIMACS-LAMSADE Workshop on Computer Science and Decision Theory*, Vol. 3 of *Annales du LAMSADE*, pp. 83–100.
- Chevalleyre, Y., Endriss, U., Estivie, S., & Maudet, N. (2007a). Reaching envy-free states in distributed negotiation settings. In Veloso, M. (Ed.), *Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI-2007)*, pp. 1239–1244.
- Chevalleyre, Y., Endriss, U., & Lang, J. (2006). Expressive power of weighted propositional formulas for cardinal preference modelling. In *Proceedings of the 10th International Conference on Principles of Knowledge Representation and Reasoning (KR)*, pp. 145–152, Lake District, UK. AAAI Press.
- Chevalleyre, Y., Endriss, U., & Maudet, N. (2007b). Allocating goods on a graph to eliminate envy. In *Proceedings of AAAI-07*, pp. 700–705.
- Conitzer, V., & Sandholm, T. (2005). Communication complexity of common voting rules. In *Proceedings of the EC-05*.
- Cramton, P., Shoham, Y., & Steinberg, R. (Eds.). (2005). *Combinatorial Auctions*. MIT Press.
- Demko, S., & Hill, T. P. (1998). Equitable distribution of indivisible items. *Mathematical Social Sciences*, 16, 145–158.
- Dunne, P. E. (2005). Extremal behaviour in multiagent contract negotiation. *Journal of Artificial Intelligence Research*, 23, 41–78.

- Dunne, P. E., Wooldridge, M., & Laurence, M. (2005). The complexity of contract negotiation. *Artificial Intelligence*, 164(1-2), 23–46.
- Endriss, U., Maudet, N., Sadri, F., & Toni, F. (2006). Negotiating socially optimal allocations of resources. *Journal of Artificial Intelligence Research*, 25, 315–348.
- Ford, L. R., & Fulkerson, D. R. (1962). *Flows in Networks*. Princeton University Press.
- Gebser, M., Liu, L., Namasivayam, G., Neumann, A., Schaub, T., & Truszczyński, M. (2007). The first answer set programming system competition. In Baral, C., Brewka, G., & Schlipf, J. (Eds.), *Proceedings of the Ninth International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR'07)*, Vol. 4483 of *Lecture Notes in Artificial Intelligence*, pp. 3–17. Springer-Verlag.
- Gottlob, G. (1992). Complexity results for nonmonotonic logics. *Journal of Logic and Computation*, 2, 397–425.
- Herreiner, D. K., & Puppe, C. (2002). A simple procedure for finding equitable allocations of indivisible goods. *Social Choice and Welfare*, 19, 415–430.
- Ieong, S., & Shoham, Y. (2005). Marginal contribution nets: a compact representation scheme for coalitional games. In *Proceedings of EC'05*.
- Karp, R. M. (1972). Reducibility among combinatorial problems. In Miller, R. E., & Thatcher, J. W. (Eds.), *Complexity of Computer Computations*, pp. 85–103, New York. Plenum Press.
- Lang, J. (2004). Logical preference representation and combinatorial vote. *Annals of Mathematics and Artificial Intelligence*, 42(1), 37–71.
- Lipton, R. J., Markakis, E., Mossel, E., & Saberi, A. (2004). On approximately fair allocations of indivisible goods. In *EC '04: Proceedings of the 5th ACM conference on Electronic commerce*, pp. 125–131, New York, NY, USA. ACM Press.
- Nisan, N. (2005). *Bidding Languages for Combinatorial Auctions*. MIT Press.
- Papadimitriou, C. H. (1994). *Computational complexity*. Addison–Wesley.
- Rawls, J. (1971). *A Theory of Justice*. Harvard University Press, Cambridge, Mass.
- Reiter, R. (1980). A logic for default reasoning. *Artificial Intelligence*, 13, 81–132.
- Robertson, J., & Webb, W. (1998). *Cake-Cutting Algorithms*. A.K. Peters.
- Rothkopf, M., Pekeč, A., & Harstad, R. (1998). Computationally manageable combinatorial auctions. *Management Science*, 8(44), 1131–1147.
- Sandholm, T. (1998). Contract types for satisficing task allocation: I theoretical results. In *Proc. AAAI Spring Symposium: Satisficing Models*.
- Wagner, K. W. (1987). More complicated questions about maxima and minima, and some closures of NP. *Theoretical Computer Science*, 51, 53–80.
- Wagner, K. W. (1990). Bounded query classes. *SIAM Journal on Computing*, 19(5), 833–846.
- Young, H. P. (1995). *Equity in Theory and Practice*. Princeton University Press.