

Efficiency and envy-freeness in fair division of indivisible goods: logical representation and complexity

Sylvain Bouveret

IRIT (UPS-CNRS) and

ONERA

31055 Toulouse Cedex 4, France

sylvain.bouveret@cert.fr

Jérôme Lang

IRIT

Université Paul Sabatier - CNRS

31062 Toulouse Cedex, France

lang@irit.fr

Abstract

We study fair division of indivisible goods among agents from the point of view of compact representation and computational complexity. We identify the complexity of several problems, including that of deciding whether there exists an efficient and envy-free allocation when preferences are represented in a succinct way. We also draw connections to nonmonotonic reasoning.

1 Introduction

Allocation of goods among agents has been considered from different perspectives in social choice theory and AI. In this paper we focus on *fair division of indivisible goods without money transfers*. Fair division makes a prominent use of fairness criteria such as *equity* and *envy-freeness*, and on this point totally depart from auctions, where only *efficiency* is relevant (and moreover a specific form of efficiency, since the criterion to be maximized is the total revenue of the auctioneer). Envy-freeness is a key concept in the literature on fair division: 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 as crucial; however, envy-freeness alone does not suffice as a criterion for finding satisfactory allocations, therefore it has to be paired with some efficiency criterion, such as Pareto optimality. However, it is known that for any reasonable notion of efficiency, there are profiles for which no efficient and envy-free allocation exists (see [Brams *et al.*, 2000])¹.

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 problems, investigating issues such as compact representation as well as complexity and algorithms. Complexity issues for negotiation (where agents exchange goods by means of deals) have also been studied (e.g. [Dunne *et al.*, 2005; Chevaleyre *et al.*, 2004]). See also [Bouveret *et al.*, 2005] for a preliminary complexity study of fair division problems.

¹This 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.

The above discussion reveals the existence of a gap: compact representation and complexity issues for fair division have received very little attention until now, apart of the recent work [Lipton *et al.*, 2004] which 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 (such as in [Herreiner and 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, 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.*, 2000] and [Demko and Hill, 1998]. Yet, as we advocate in this paper, conciliating conciseness and expressivity is possible, by means of *compact representation*.

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. We consider this problem from the point of view of compact representation and computational complexity. 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. The most natural representation of a dichotomous preference is by a single propositional formula, where variables correspond to goods. Expressing envy-freeness and efficiency within this logical representation reveals unexpected connections to nonmonotonic reasoning. We identify the complexity of the key problem of the existence of an envy-free and Pareto-efficient allocation, which turns out to be Σ_2^P -complete; we also identify the complexity of several other problems obtained either by imposing some restrictions on the latter or by replacing Pareto-efficiency by other criteria. In Section 5 we extend this result to the case of non-dichotomous, compactly represented preferences.

2 Background

2.1 Fair division problems

Definition 1

A **fair division problem** 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 note $A \succ_i B$ (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 *monotonous* 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 monotonous if and only if R_i is monotonous for every i .

Definition 2

- 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 a i such that $x \in \pi(i)$ then π is a complete allocation.
- 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 .
- 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 way². 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 $\neg x$, 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 formulas then $\bigwedge S = \varphi_1 \wedge \dots \wedge \varphi_n$ is the conjunction of all formulas of S .

²Note that connectives \rightarrow and \leftrightarrow are not allowed; this is important for the definition of positive formulas (to come).

2.3 Computational complexity

In this paper we will refer to some complexity classes located in the polynomial hierarchy. We assume the reader to be familiar with the classes NP and coNP. 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. $\Delta_2^P = \text{PNP}$ is the class of all languages recognizable by a deterministic Turing machine working in polynomial time using NP oracles. Likewise, $\Sigma_2^P = \text{NPNP}$. $\Theta_2^P = \Delta_2^P[\mathcal{O}(\log n)]$ is the subclass of Δ_2^P of problems that only need a logarithmic number of oracles. See for instance [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 3 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 $A \in \text{Good}_i$ or $B \notin \text{Good}_i$. $\mathcal{R} = \langle R_1, \dots, R_N \rangle$ 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 symbols for each good) such that $\text{Mod}(\varphi) = \text{Good}_i$. Formally:

Definition 4 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 $\text{Mod}(\varphi_i) = \text{Good}_i$.

Clearly, for any dichotomous preference R_i there is a formula φ_i representing R_i ; furthermore, this formula is unique up to logical equivalence.

Example 1 $X = \{a, b, c\}$ and $\text{Good}_i = \{\{a, b\}, \{b, c\}\}$. Note that R_i is not monotonous. Then $\varphi_i = (a \wedge b \wedge \neg c) \vee (\neg a \wedge b \wedge c)$ represents R_i .

An easy but yet useful result (whose proof is omitted):

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

1. R_i is monotonous;
2. Good_i is upward closed, that is, $A \in \text{Good}_i$ and $B \supseteq A$ imply $B \in \text{Good}_i$.
3. R_i is representable by a positive propositional formula.

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 $\langle \varphi_1, \dots, \varphi_N \rangle$.

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 $\varphi_2^* = a_2 \wedge d_2$.

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 $F(\pi) = M$ simply defined by $\pi(i) = \{x \mid M \models x_i\}$. This mapping is obviously bijective, and we denote $F^{-1}(M)$ the allocation corresponding to an interpretation M of $\text{Mod}(\Gamma_{\mathcal{P}})$.

3.1 Envy-freeness

We now show how the search for an envy-free allocation can be mapped to a satisfiability 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 $\varphi_{2|1}^* = a_2 \wedge (b_2 \vee c_2)$. (Obviously, $\varphi_{i|i}^* = \varphi_i^*$.)

Proposition 2 Let $\mathcal{P} = \langle \varphi_1, \dots, \varphi_N \rangle$ be an allocation problem with dichotomous preferences under propositional form, and the formulas $\varphi_{j|i}^*$ 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 $F(\pi) \models \Lambda_{\mathcal{P}}$.

The proof is simple, so we omit it. The search for envy-free allocations can thus be reduced to a satisfiability 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.

Example 2

$\varphi_1 = a \vee (b \wedge c)$; $\varphi_2 = a$; $\varphi_3 = a \vee b$.
 $\Lambda_{\mathcal{P}} = ((a_1 \vee (b_1 \wedge c_1)) \vee (\neg(a_2 \vee (b_2 \wedge c_2)) \wedge \neg(a_3 \vee (b_3 \wedge c_3))))$
 $\wedge (a_2 \vee (\neg a_1 \wedge \neg a_3)) \wedge ((a_3 \vee b_3) \vee (\neg(a_1 \vee b_1) \wedge \neg(a_2 \vee b_2)))$;
 $\text{Mod}(\Gamma_{\mathcal{P}} \wedge \Lambda_{\mathcal{P}}) = \{\{c_1\}, \{c_1, b_3\}, \{c_2, b_3\}, \{c_2\}, \{b_3\}, \{c_3\}, \emptyset\}$.
 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

Definition 5 Let $\Delta = \{\alpha_1, \dots, \alpha_m\}$ a set of formulae and β a formula. $S \subseteq \Delta$ is a maximal β -consistent subset of Δ iff (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 .

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 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 $\Phi_{\mathcal{P}}$.

This simple result, whose proof is omitted, suggests that efficient allocations can be computed from the logical expression Φ of the problem, namely, by computing the maximal $\Gamma_{\mathcal{P}}$ -consistent subsets of Φ ; 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 \mathcal{P} ³.

³Note that there are in general exponentially many maximal $\Gamma_{\mathcal{P}}$ -consistent subsets of Φ (and therefore exponentially many efficient

Example 2 (cont'd) The maximal $\Gamma_{\mathcal{P}}$ -consistent subsets of Φ 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 corresponds 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 6 A supernormal default theory⁴ 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 formulas. A propositional formula ψ is a skeptical consequence of D , denoted by $D \vdash^{\forall} \psi$, if and only if for all $S \in \text{MaxCons}(\Delta, \beta)$ we have $\bigwedge S \wedge \beta \models \psi$.

Proposition 4 Let $\mathcal{P} = \langle \varphi_1, \dots, \varphi_N \rangle$ 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 $D \not\vdash^{\forall} \neg \Lambda_{\mathcal{P}}$.

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 connex domains such as belief revision and answer set programming. 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

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

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 Φ .

⁴“Supernormal” defaults are also called “normal defaults without prerequisites” (e.g., [Reiter, 1980]).

is in Σ_2^P . Less obviously, we now show that it is complete for this class, even if preferences are required to be monotonous.

Let us first note that skeptical inference remains Π_2^P -complete under these two restrictions (to which we refer as RSI, for RESTRICTED SKEPTICAL INFERENCE): (a) $\psi = \varphi_1$; (b) $n \geq 2$. (Here is the justification: $\Delta \sim \psi$ if and only if $\Delta \cup \{\psi, \psi\} \sim \psi$.) Equivalently, RSI is the problem of deciding whether, given $\Delta = \langle \alpha_1, \dots, \alpha_n \rangle$ with $n \geq 2$, all maximal consistent subsets of Δ contain α_1 .

Proposition 5 *The problem EEF EXISTENCE of determining whether there exists an efficient and envy-free allocation for a given problem \mathcal{P} with monotonous, 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) to EEF EXISTENCE. Given any finite set Δ of propositional formulae, let $V_\Delta = \text{Var}(\Delta)$ the set of propositional symbols appearing in Δ , and let $\mathcal{P} = H(\Delta)$ the following instance of EEF EXISTENCE:

(1) $I = \{1, 2, \dots, n+2\}$;

(2) $X = \{v^i | v \in V_\Delta, i \in 1..n\} \cup \{\bar{v}^i | v \in V_\Delta, i \in 1..n\} \cup \{x^i | i \in 1..n\} \cup \{y\}$;

(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); (b) 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

- $\varphi_1 = \beta_1 \vee x^1$;
- for $i = 2, \dots, n$, $\varphi_i = \beta_i \wedge x^i$;
- $\varphi_{n+1} = \left(\left(\bigwedge_{v \in \text{Var}(\Delta)} (\bigwedge_{i=1}^n v^i) \vee (\bigwedge_{i=1}^n \bar{v}^i) \right) \wedge x^1 \right) \vee y$;
- $\varphi_{n+2} = y$.

Lemma 1 *An allocation π for \mathcal{P} is said to be regular if and only if for all $i \neq n$, $\pi(i) \subseteq \sigma(i)$, where*

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

Let now π_R 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 α_i), therefore $\pi_R(i) \sim_i \pi(i)$, from which (2) follows. The formulas α_i being positive, the preference relations \succeq_i are monotonous, 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 π_R is envy-free, from which (3) follows. ■

Lemma 2 *If π is regular then*

1. 1 can only envy $n+1$;

2. $2, \dots, n$ envy noone;
3. $n+1$ can only envy $n+2$;
4. $n+2$ 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 $j \in \{2, \dots, n, n+2\}$. Assume 1 envies j . Then $\pi(j) \models \varphi_1$. π being regular, $x^1 \notin \pi(j)$, therefore $\pi(j) \models \beta_1$. Now, since π is regular, $\pi(j)$ does not contain any v^i nor any \bar{v}^i ; now, β_1 can only be made true by variables v^i or \bar{v}^i (which cannot be the case here) unless it is a tautology. Now, if β_1 is a tautology, then 1 is satisfied by π and cannot envy j , a contradiction.

2. 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 .

3. Let $i = n+1$. Assume $n+1$ envies 1 then $\pi(1) \models \varphi_{n+1}$. Since $\pi(1) \models y$ is impossible (because π is regular), we have $\pi(1) \models \bigwedge_{v \in V_\Delta} (\bigwedge_{i=1}^n v^i) \vee (\bigwedge_{i=1}^n \bar{v}^i) \wedge x^1$, which implies that either $\pi(1) \models \bigwedge_{v \in V_\Delta} (\bigwedge_{i=1}^n v^i)$ or $\pi(1) \models (\bigwedge_{i=1}^n \bar{v}^i)$. Both are impossible because π is regular and $n \geq 2$. The case $j \in \{2, \dots, n\}$ is similar.

4. Let $i = n+2$ and $j \neq n$. If i envies j then $\pi(j) \models y$, which is impossible because π is regular. ■

Lemma 3 *Let π be a regular allocation satisfying $n+1$ and $n+2$. 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 envy-free iff $M(\pi) \models \alpha_1$.*

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 him 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, since π is regular, it is envy-free if and only if (a) 1 does not envy $n+1$, (b) $n+1$ does not envy $n+2$ and (c) $n+2$ does not envy $n+1$. Since $n+1$ and $n+2$ are satisfied by π , we get that π is envy-free if and only if 1 does not envy $n+1$, that is, if and only if either $\pi(1) \models \varphi_1$ or $\pi(n+1) \not\models \varphi_1$. Now, $\pi(n+1)$ contains x^1 , therefore $\pi(n+1) \models \varphi_1$, which entails that π is envy-free if and only if $\pi(1) \models \varphi_1$. This is equivalent to $\pi(1) \models \beta_1$, because 1 does not get x^1 (which is assigned to $n+1$), which in turn is equivalent to $M(\pi) \models \alpha_1$ by construction. ■

Lemma 4 *For each interpretation M over V_Δ , let us define $\pi_M : I \rightarrow 2^X$ by:*

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

Then:

1. π_M is a well-defined and regular allocation satisfying $n+1$ and $n+2$;
2. $M_{\pi_M} = M$ (M_{π_M} is obtained from π_M as in Lemma 3).
3. for any $i \in 1, \dots, n$, π_M satisfies i iff $M \models \alpha_i$.

4. π_M is efficient iff M satisfies a maximal consistent subset of Δ .

Proof:

1. π_M does not give the same good to more than one individual, therefore it is an allocation. The rest is straightforward.
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 $i \in 2, \dots, n$. Since π_M gives x^i to i , π_M satisfies i if and only if $\pi_M(i) \models \beta^i$, which is equivalent to $M \models \alpha_i$. If $i = 1$ then, since π_M does not give x^1 to 1, π_M satisfies 1 if and only if $\pi_M(1) \models \beta^1$, which is equivalent to $M \models \alpha_1$.
4. from point 3, $\{i, \pi_M \text{ satisfies } i\} = \{i, M \models \alpha_i\}$. 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 . ■

Lemma 5 *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 $\pi(n+2) = \{y\}$, $x^1 \in \pi(n+1)$, and $M(\pi)$ is well-defined. Let π' obtained from π by

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

π being regular and satisfying $n+1$ and $n+2$, we have $\pi(n+1) = \pi'(n+1)$, $\pi(n+2) = \pi'(n+2)$, and then for each i , $\pi(i) \subseteq \pi'(i)$: indeed, let $j \in \{2, \dots, n\}$ (for $n+1$ and $n+2$ this inclusion is obviously satisfied); then (a) π is regular, therefore $\pi(j) \subseteq \sigma(j)$; now, all goods of $\sigma(j)$ are either in $\pi'(j)$ or in $\pi(n+1)$ (namely: x^1 if $j = 1$ and all the v^j such that $M(\pi) \models \neg v$ and all the \bar{v}^j such that $M(\pi) \models v$); therefore, $\pi(j) \subseteq \pi'(j) \cup \pi(n+1)$, which, together with $\pi(1) \cap \pi(n+1) = \emptyset$, implies $\pi(j) \subseteq \pi'(j)$. Since preferences are monotonous, all individuals satisfied by π are satisfied by π' as well; and since π is efficient, π and π' satisfy the same set of individuals. Now, we remark that $\pi' = \pi_{M(\pi)}$. By Lemma 4, π' is efficient iff $M(\pi)$ satisfies a maximal consistent subset of Δ , from which we conclude. ■

Lemma 6 *Any envy-free and efficient allocation for \mathcal{P} satisfies $n+1$ and $n+2$.*

Proof: Suppose π doesn't 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 $n+1$, i.e., $y \notin \pi(n+2)$; 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 . ■

Lemma 7 *If there exists an EEF allocation, then there exists a maximal consistent subset of Δ containing α_1 .*

Proof: Let π be an efficient and envy-free allocation. By Lemma 1, π_R is regular, efficient and envy-free. By Lemma 6, π_R satisfies $n+1$ and $n+2$. Then by Lemma 5, $M(\pi_R)$ satisfies a maximal consistent subset of Δ , and by Lemma 3, $M(\pi_R) \models \alpha_1$. Therefore $Sat(M(\pi_R), \Delta)$ is a maximal consistent subset of Δ and contains α_1 . ■

Lemma 8 *If there exists a maximal consistent subset of Δ containing α_1 then there exists an EEF allocation.*

Proof: Assume that there exists a maximal consistent subset S of Δ containing α_1 , and let M be a model of S . By point 4 of Lemma 4, π_M is efficient.

By point 1 of Lemma 4, π_M is regular; then by Lemma 2, π_M is envy-free if and only if (i) 1 does not envy $n+1$, (ii) $n+1$ does not envy $n+2$ and (iii) $n+2$ does not envy $n+1$. By point 1 of Lemma 4, π_M satisfies $n+1$ and $n+2$, therefore (ii) and (iii) hold. Lastly, by point 5 of Lemma 4, $M \models \alpha_1$ implies that π_M satisfies 1, therefore (i) holds as well and π_M is envy-free. ■

Proof of Proposition 5: from Lemmas 7 and 8, the existence of a maximal consistent subset of Δ containing α_1 and the existence of an efficient and envy-free allocation for $\mathcal{P} = H(\Delta)$ are equivalent. Clearly, H is computed in polynomial time. Therefore, H 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. ■

As a corollary, this Σ_2^p -completeness result holds for general (not necessarily monotonous) dichotomous preferences.

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

Proposition 6 EEF EXISTENCE with N identical dichotomous, monotonous preferences is NP-complete, for any fixed $N \geq 2$.

Due to space limitations the proofs of this result and the following ones are omitted⁵. Note that we have here a hardness result for any *fixed* number of agents (≥ 2). Things are different with Proposition 5, for which hardness does not hold when N is fixed. Namely, the following holds for $N = 2$:

Proposition 7 EEF EXISTENCE for two agents with monotonous dichotomous preferences is NP-complete.

Unlike Proposition 5, these results are sensitive to whether preferences are required to be monotonous or not.

Proposition 8 EEF EXISTENCE with N identical dichotomous preferences is coBH_2 -complete, for any fixed $N \geq 2$.

Proposition 9 EEF EXISTENCE for 2 agents with dichotomous preferences is coBH_2 -complete.

Complexity decreases as well if we weaken Pareto-efficiency by only requiring allocations to be *complete*:

⁵They can be found in the long version of the paper, accessible at <http://www.irit.fr/recherches/RPDM/persos/JeromeLang/papers/eef.pdf>.

Proposition 10 *The existence of a complete envy-free allocation for agents with monotonous, dichotomous preferences is NP-complete, even for 2 agents with identical preferences.*

Lastly, replacing Pareto-efficiency by an utilitarianistic notion of efficiency results in a complexity decrease as well: