Efficient Reallocation
under Additive and Responsive Preferences

Haris Aziz

UNSW Sydney and Data61, Sydney, Australia

Péter Biró

Hungarian Academy of Sciences, Budapest, Hungary

Jérôme Lang, Julien Lesca, Jérôme Monnot

LAMSADE, Universite Paris-Dauphine, Paris, France

Abstract

Reallocating resources to get mutually beneficial outcomes is a fundamental problem in various multi-agent settings. While finding an arbitrary Pareto optimal allocation is generally easy, checking whether a particular allocation is Pareto optimal can be much more difficult. This problem is equivalent to checking that the allocated objects cannot be reallocated in such a way that at least one agent prefers her new allocation to her old one, and no agent prefers her old allocation to her new one. We consider the problem for two related types of preference relations over sets of objects. In the first part of the paper we focus on the setting in which agents express additive cardinal utilities over objects. We present computational hardness results as well as polynomial-time algorithms for testing Pareto optimality under different restrictions such as two utility values or lexicographic utilities. In the second part of the paper we assume that agents express only their (ordinal) preferences over individual objects, and that their underlying preferences are additively separable. In this setting, we present characterizations and polynomial-time algorithms for possible and necessary Pareto optimality.

Keywords: Fair Division, Resource Allocation, Pareto optimality

1. Introduction

Reallocation of resources to achieve mutually better outcomes is a central concern in multi-agent settings. A desirable way to achieve 'better' outcomes

---

is to obtain a Pareto improvement in which each agent is at least as happy and at least one agent is strictly happier. Pareto improvements are desirable for two fundamental reasons: they result in strictly more welfare for almost any reasonable notion of welfare (such as utilitarian social welfare, or egalitarian 'leximin' social welfare). Secondly, they satisfy the minimal requirement of individual rationality in the sense that no agent is worse off after the trade. If a series of Pareto improvements results in a Pareto optimal outcome, that is even better because there exists no other outcome which each agent weakly prefers and at least one agent strictly prefers.

We consider the setting in which agents are initially endowed with objects and have additive preferences for the objects. In the absence of endowments, achieving a Pareto optimal assignment is easy in almost all cases: if we know the cardinal utilities that agents have for the objects, we can simply assign every object to the agent who values it the most; if we have only ordinal information, and assume that an agent always strictly prefers a superset of objects to a subset, then assigning all objects to the same agent leads to a Pareto optimal allocation.[1] On the other hand, in the presence of endowments, finding a Pareto optimal assignment that respects individual rationality is more challenging. The problem is closely related to the problem of testing Pareto optimality of the initial assignment: a certificate of a lack of Pareto optimality is an assignment that respects individual rationality and is a Pareto improvement. In fact, if testing Pareto optimality is computationally intractable, then finding an individually rational and Pareto optimal assignment is computationally intractable as well. In view of this, we focus on the problem of testing Pareto optimality. In all cases where we are able to test it efficiently, we also present algorithms to compute individually rational and Pareto optimal assignments.

We will consider two related settings: one where agents express cardinal preferences, and one where they express ordinal preferences. In the cardinal case, we will assume that not only they express cardinal preferences but *additive* preferences, that is, they express a value for every single object, and the value they give to a set of objects is taken to be the sum of the values they give to each of its elements. We make here a (weak) monotonicity assumption: every value is non-negative.

In the ordinal case, we assume the agents' preferences are additively separable [24]: each of them is a weak order (over subsets of objects) compatible with an additive utility function. This leads to assuming that agents have *responsive* preferences: an agent prefers to exchange an object o against an object o' in a bundle of objects S (containing o) if she prefers to exchange o against o' in any other bundle of objects containing o. Because, as we do not have enough information to test Pareto optimality if some agents may be indifferent

---

[1]The only remaining case is where we have ordinal preferences and some agent may be indifferent between receiving an object and not receiving it. In this case, we do not have enough information to say whether an allocation is Pareto optimal or not, nor to find an arbitrary Pareto optimal allocation, independently of computational resources.

between having an object or not, in this setting we will have to assume strict monotonicity.

Contributions. In the cardinal setting, we first relate the problem of computing an individually rational and Pareto optimal assignment to the more basic problem of testing Pareto optimality of a given assignment. We show that for an unbounded number of agents, testing Pareto optimality is strongly coNP-complete even if the assignment assigns at most two objects per agent. We then identify some natural tractable cases. In particular, we present a pseudo-polynomial-time algorithm for the problem when the number of agents is constant. We characterize Pareto optimality under lexicographic utilities and we show that Pareto optimality can be tested in linear time. For dichotomous utilities in which utilities can take values α or β, we present a characterization of Pareto optimal assignments which also yields a polynomial-time algorithm to test Pareto optimality.

In the ordinal setting, we consider two versions of Pareto optimality: *possible Pareto optimality* and *necessary Pareto optimality*. Intuitively, an assignment is possibly Pareto-optimal if there is a collection of additive utility functions (one for each agent) compatible with the agents ordinal preferences over single objects, for which the assignment is Pareto-optimal; and an assignment is necessarily Pareto-optimal if it is Pareto-optimal for any such collection of additive utility functions. For both properties, we present characterizations that lead to polynomial-time algorithms for testing the property for a given assignment.

Related Work. The setting in which agents express additive cardinal utilities and a welfare maximizing or fair assignment is computed is a very well-studied problem in computer science. Although computing a utilitarian welfare maximizing assignment is easy, the problem of maximizing egalitarian welfare is NP-hard [19].

Algorithmic aspects of Pareto optimality have received attention in discrete allocation of indivisible goods [1], randomized allocation of indivisible goods [9, 11], two-sided matching [16, 22, 28], coalition formation under ordinal preferences [5], committee elections [8], and probabilistic voting [4]. Since we are interested in Pareto improvements, our paper is also related to housing markets with endowments [7, 23, 25, 29, 30]. Recently, Damamme et al. [17] examined restricted Pareto optimality under ordinal preferences. When each agent can get at most one object, there is a well-known characterization of Pareto optimal assignments as not admitting a trading cycle. The result implies that Pareto optimality can be tested in linear time (see e.g., [1, 6]).

De Keijzer et al. [18] studied the complexity of deciding whether there exists a Pareto optimal and envy-free assignment when agents have additive utilities. They also showed that testing Pareto optimality under additive utilities is coNP-complete. We show that this result holds even if each agent initially owns two objects.

Cechlárová et al. [15] proved that Pareto optimality of an assignment under lexicographic utilities can be tested in polynomial time. In this paper, we present

a simple characterization of Pareto optimality under lexicographic utilities that leads to a linear-time algorithm to test Pareto optimality.

Bouveret et al. [12] have considered necessary Pareto optimality as Pareto optimality for linear orders consistent with responsive set extension,[2] and have presented some computational results when necessary Pareto optimality is considered *in conjunction* with other fairness properties.

Reallocating resources to improve fairness has also been studied before [21].

2. Preliminaries

We consider a set of agents $N = \{1, \ldots, n\}$ and a set of objects $O = \{o_1, \ldots, o_m\}$. An assignment $p = (p(1), \ldots, p(n))$ is a partition of O into n subsets, where $p(i)$ is the bundle assigned to agent i. We denote by $\mathscr{X}$ the set of all possible assignments. Informally and in examples, we will present an allocation p in the form $(p(1)|\cdots|p(n))$.

In Section 3, we assume that each agent expresses a cardinal utility function u_i over O. We assume that an agent's value for an object is a non-negative rational, i.e, $u_i(o) \in \mathbb{Q}$ for all $i \in N$ and $o \in O$. We also assume additivity: u_i is lifted from O to 2^O by: $u_i(O') = \sum_{o \in O'} u_i(o)$ for each $i \in N$ and $O' \subseteq O$ (and we write $u_i(o)$ instead of $u_i(\{o\})$). The vector $u = (u_1, \ldots, u_n)$ consisting of the utility functions of the agents is referred to as a *utility function profile*.

In Section 4, we assume that each agent i expresses only her complete, transitive and reflexive preferences $\succsim_i$ over O. Agents may be indifferent among objects. Let $\sim_i$ and $\succ_i$ denote the symmetric and anti-symmetric part of $\succsim_i$, respectively. We denote by $E_i^1, \ldots, E_i^{m_i}$ the m_i equivalence classes of an agent $i \in N$. Those classes partition O into m_i sets of objects such that agent i is indifferent between two objects belonging to the same class, and strictly prefers an object of E_i^k to an object of E_i^l whenever $k < l$. The preference profile $\succsim = (\succsim_1, \ldots, \succsim_n)$ specifies for each agent i her preference relation $\succsim_i$ over single objects. We say that a utility function u_i over O is *consistent* with $\succsim_i$ if $u_i(o) \geq u_i(o')$ iff $o \succsim_i o'$. The set of all utility functions consistent with $\succsim_i$ will be denoted by $\mathscr{U}(\succsim_i)$ We will denote by $\mathscr{U}(\succsim)$ the set of all utility function profiles $u = (u_1, \ldots, u_n)$ such that $u_i \in \mathscr{U}(\succsim_i)$ for each $i \in N$. The way an agent's ordinal preference relation is lifted from single objects to sets of objects will be made precise in Section 4.

An assignment $p \in \mathscr{X}$ is said to be *individually rational* for an initial endowment $e \in \mathscr{X}$ if $u_i(p(i)) \geq u_i(e(i))$ holds for every agent i. An assignment $p \in \mathscr{X}$ is said to be *Pareto dominated* by another $q \in \mathscr{X}$ if (i) for every agent $i \in N$, $u_i(q(i)) \geq u_i(p(i))$ holds, (ii) for at least one agent $i \in N$, $u_i(q(i)) > u_i(p(i))$ holds. An assignment is *Pareto optimal* if it is not Pareto dominated by another assignment. Finally, whenever cardinal utilities are considered, the *(utilitarian) social welfare* of an assignment p is defined as $SW(p) = \sum_{i \in N} u_i(p(i))$.

---

[2]Brams et al. [14] used the term *Pareto ensuring* for Pareto optimality for all completions of the responsive set extension.

4

Example 1. *Let $n = 3$, $m = 5$, and the utilities of agents be represented as follows.*

	o_1	o_2	o_3	o_4	o_5
1	16	8	4	2	1
2	10	4	3	3	4
3	6	1	3	6	3

We assume that the initial assignment is $p = (o_2o_4|o_1|o_3o_5)$ in which $p(1) = \{o_2, o_4\}$, $p(2) = \{o_1\}$, and $p(3) = \{o_3, o_5\}$. We have $u_1(p(1)) = 12$, $u_2(p(2)) = 10$ and $u_3(p(3)) = 6$. We see that p is not Pareto optimal: indeed, $p' = (o_1|o_2o_3o_5|o_4)$ gives respective utilities 16, 11 and 6 to agents 1, 2 and 3. On the other hand, p' cannot be improved and is Pareto optimal. We can see this by noticing first that agent 1 cannot do better without object o_1. Let us give o_1 to 1; now, the only way for 2 to do better (without o_1) is to have all four remaining objects, in which case 3 gets nothing.

3. Additive utilities

In this section we assume that each agent expresses a cardinal utility function u_i over O, where $u_i(o) \geq 0$ for all $i \in N$ and $o \in O$.

3.1. Testing Pareto optimality: hard cases

We will consider Pareto optimality and individual rationality with respect to additive utilities. The following lemma shows that the computation of an individually rational and Pareto improving assignment is at least as hard as the problem of deciding whether a given assignment is Pareto optimal:

Lemma 1. *If there exists a polynomial-time algorithm to compute a Pareto optimal and individually rational assignment, then there exists a polynomial-time algorithm to test Pareto optimality.*

Proof. We assume that there is a polynomial-time algorithm A to compute an individually rational and Pareto optimal assignment. Consider an assignment p for which Pareto optimality needs to be tested. We can use A to compute an assignment q which is individually rational for the initial endowment p and Pareto optimal. By individual rationality $u_i(q(i)) \geq u_i(p(i))$ for all $i \in N$. If $u_i(q(i)) = u_i(p(i))$ for all $i \in N$, then p is Pareto optimal simply because q is Pareto optimal. However, if there exists $i \in N$ such that $u_i(q(i)) > u_i(p(i))$, it means that p is *not* Pareto optimal. $\square$

For Example 1, assume first that the initial assignment is p. Then computing a Pareto optimal and individually rational assignment will return p' (note that for this example there is no other individually rational and Pareto optimal assignment; of course, in general there may be several). Since p' Pareto dominates p, we conclude that p is not Pareto optimal. Now, assume that the initial assignment is p'. Then computing a Pareto optimal and individually rational assignment will return p', therefore, p' is Pareto optimal.

A Pareto optimal assignment can be computed trivially by giving each object to the agent who values it the most (if there are several different agents who value a particular object most, is it sufficient to give this object to any one of them).[3] Bouveret and Lang [13, Proposition 6] proved that a problem concerning coalitional manipulation in sequential allocation is NP-complete. The result can be reinterpreted as follows.

Theorem 1. *Under additive preferences, testing Pareto optimality of a given assignment is weakly coNP-complete, even for $n = 2$ even if the induced ordinal preferences over individual objects are the same.*

Proof. We write the proof for the sake of completeness. Testing Pareto optimality is in coNP since one can test that an assignment is Pareto dominated by another one in polynomial time. The reduction is from PARTITION. An instance of PARTITION is described by a set of t elements $E = \{e_1, \ldots, e_t\}$, and integer weight $w(e_j)$ for each $e_j \in E$ such that $\sum_{e_i \in E} w(e_i) = 2M$. The question is to decide if there exists a balanced partition of E i.e., $S \subseteq E$ such that $\sum_{e_i \in S} w(e_i) = \sum_{e_i \in E \setminus S} w(e_i) = M$. The reduction relies on a set of $t + 1$ objects $\{g^+, g_1, \ldots, g_t\}$, and two agents $\{1, 2\}$. The utility values for agent 1 are $u_1(g^+) = M$ and $u_1(g_i) = w(e_i)$ for all $i \in \{1, \ldots t\}$. The utility values for agent 2 are $u_2(g^+) = M + \varepsilon$, with $0 < \varepsilon < 1$, and $u_2(g_i) = w(e_i)$ for all $i \in \{1, \ldots, t\}$. Then, one can easily check that the assignment in which agent 1 gets g^+ and agent 2 gets all g_j objects is not Pareto optimal if and only if there is a balanced partition of E. □

Corollary 1. *Computing an individually rational and Pareto optimal assignment is weakly NP-hard for $n = 2$.*

For an unbounded number of agents, testing Pareto optimality of a given assignment is strongly coNP-complete [18]. We show that the problem remains strongly coNP-complete, for an unbounded number of agents, even if each agent initially owns exactly two objects.

Theorem 2. *Testing Pareto optimality of a given assignment is strongly coNP-complete, even if each agent initially owns two objects.*

Proof. The reduction is done from 2-NUMERICAL MATCHING WITH TARGET SUMS (2NMTS in short). The input of 2NMTS is a sequence $a_1, \ldots, a_k$ of k positive integers such that $\sum_{i=1}^{k} a_i = k(k+1)$, where $1 \le a_i \le 2k-1$ for $i = 1, \ldots, k$, and $a_1 \le a_2 \ldots \le a_k$. We want to decide if there are two permutations π and θ of the integers $\{1, \ldots, k\}$ such that $\pi(i) + \theta(i) = a_i$ for $i = 1, \ldots, k$. 2NMTS is known to be strongly NP-complete [31].

The reduction from an instance of 2NMTS is as follows. There are $3k + 1$ agents $N = L \cup C \cup R \cup \{d\}$ where $L = \{\ell_1, \ldots, \ell_k\}$, $R = \{r_1, \ldots, r_k\}$ and

---

[3]Note that this particular assignment also maximizes social welfare.

$C = \{c_1, \ldots, c_k\}$, and $6k + 2$ objects $O = F \cup G \cup H \cup \{o\}$ where $F = \{f_i^L, f_i^R : i = 1, \ldots, k\}$, $G = \{g_i^L, g_i^R : i = 1, \ldots, k\} \cup \{g^C\}$, and $H = \{h_i^{CL}, h_i^{CR} : i = 1, \ldots, k\}$. Let ε be a positive value strictly lower than $1/2$. All utility values are equal to ε except for the ones summarized in the following table, where agt#1 is the agent which receives the object in the initial assignment and $u_{agt\#1}$ is her utility for it, and where agt(s)#2 lists the other agents having a utility different from ε for the object and $u_{agt(s)\#2}$ corresponds to their utility for it:

object	agt#1	$u_{\text{agt}\#1}$	agt(s)#2	$u_{\text{agt}(s)\#2}$
h_i^{CL}	c_i	a_i	ℓ_i	$1 + \varepsilon$
h_i^{CR}	c_i	$3k$	r_i	$1 - \varepsilon$
f_i^L	ℓ_i	1	c_j with $a_j \geq i + 1$	i
f_i^R	r_i	1	c_j with $a_j \geq i + 1$	$3k + i$
g_i^R	r_i	3	r_{i+1} if $i < k$	$3 + \varepsilon$
			d if $i = k$	$3 + \varepsilon$
g_i^L	ℓ_i	3	ℓ_{i-1} if $i > 1$	$3 - \varepsilon$
			r_1 if $i = 1$	$3 + \varepsilon$
g^C	d	3	ℓ_k	$3 - \varepsilon$
o	d	1		

The initial assignment gives the following utilities to agents: $u_d(\{g^C, o\}) = 4$, $u_{c_i}(\{h_i^{CL}, h_i^{CR}\}) = 3k + a_i$, $u_{\ell_i}(\{f_i^L, g_i^L\}) = 4$ and $u_{r_i}(\{f_i^R, g_i^R\}) = 4$ for $i = 1 \ldots k$. Intuitively, the way that the objects of $\{f_i^L\}_{i \in \{1, \ldots, k\}}$ and $\{f_i^R\}_{i \in \{1, \ldots, k\}}$ will be reassigned to the agents of C represents permutations π and θ, respectively. More precisely, in order to keep the same utility for their bundle, each agent c_i will receive two objects f_j^L and f_ℓ^R such that $j + \ell = a_i$. The corresponding interpretation in the instance of 2NMTS is that $\pi(i)$ and $\theta(i)$ will be set to j and ℓ, respectively. The other agents and objects are gadgets to enforce the allocation of the objects of F to the agents of C as described above.

Clearly, this instance is constructed within polynomial time and each agent has initially two objects. We claim that there is a Pareto improvement of the initial assignment if and only if $\{a_i : i = 1 \ldots k\}$ is a yes-instance of 2NMTS.

Assume first that there exist π and θ such that $\pi(i) + \theta(i) = a_i$ for $i = 1 \ldots k$ i.e., $\{a_i : i = 1 \ldots k\}$ is a yes-instance of 2NMTS. Note that this implies for any $i = 1 \ldots k$ that

$$\pi(i) + 1 \leq a_i \quad \text{and} \quad \theta(i) + 1 \leq a_i \tag{1}$$

because $\pi(i) \geq 1$ and $\theta(i) \geq 1$. Then consider the following assignment:

- $\{h_i^{CL}, g_{i+1}^L\}$ (resp. $\{h_k^{CL}, g^C\}$) is assigned to ℓ_i with $i < k$ (resp. to ℓ_k) with utility 4.

- $\{h_i^{CR}, g_{i-1}^R\}$ (resp. $\{h_1^{CR}, g_1^L\}$) is assigned to r_i with $i > 1$ (resp. to r_1) with utility 4.

- $\{f_{\pi(i)}^R, f_{\theta(i)}^L\}$ is assigned to c_i. By using (1), it is easy to check that the utility of agent c_i is $3k + \pi(i) + \theta(i) = 3k + a_i$.
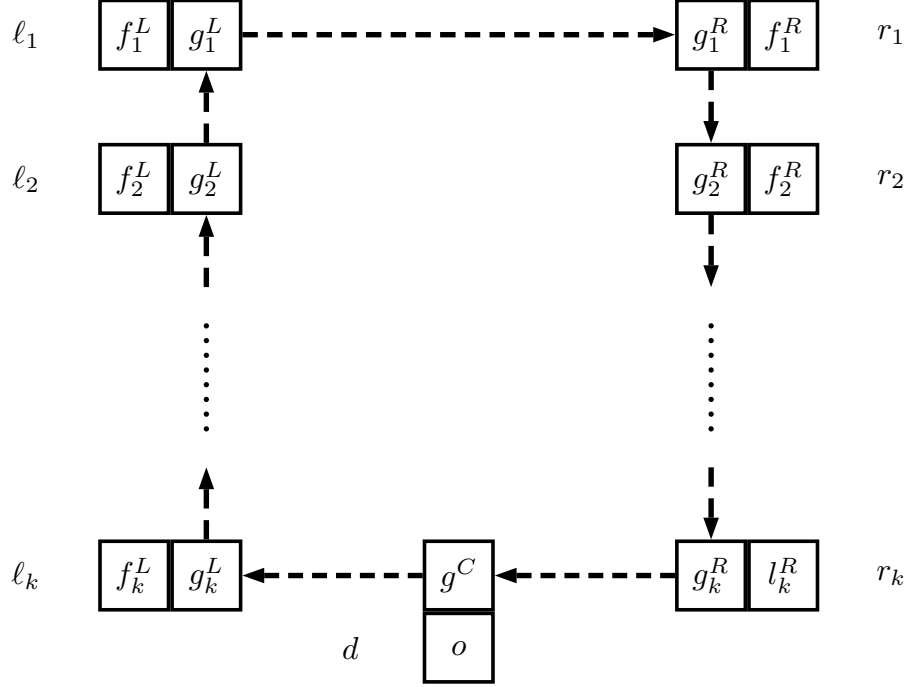
Figure 1: Initial assignment for agents of $L \cup R \cup \{d\}$.

- $\{o, g_k^R\}$ is assigned to d with utility $4 + \varepsilon$.

This allocation is clearly a Pareto improvement of the initial allocation.

Assume now that $\{a_i : i = 1 \ldots k\}$ is a no-instance of 2NMTS. By contradiction, assume that there exists a Pareto improvement p of the initial assignment. Note first that every agent should receive in p at least two objects. Indeed, there is no object which provides a utility greater than $3 + \varepsilon$ to every agent of $L \cup R \cup \{d\}$, and any of those agents receives a utility of 4 in the initial assignment. Furthermore, any object f_i^R provides a utility of at most $3k + i$ to agent c_j, which is strictly lower than her utility of $3k + a_j$ in the initial assignment because $a_j \geq i + 1$ (otherwise c_j would get utility 0 from f_i^R). Since the number of objects is twice the number of agents, we can conclude that p assigns exactly 2 objects to every agent.

Let us focus first on the objects of G. We are going to show in the next two paragraphs that there is only two possible allocations for these objects in order to obtain an allocation p which Pareto dominates the initial assignment. These objects are the only ones that can provide a utility of at least $3 - \varepsilon$ to the agents of $L \cup R \cup \{d\}$. Any other object provides a utility of at most $1 + \varepsilon$ to the agents of $L \cup R \cup \{d\}$. Therefore, each agent in $L \cup R \cup \{d\}$ should receive exactly one object from G to achieve a utility of at least 4 since $|L \cup R \cup \{d\}| = |G| = 2k + 1$.

8

Figure 1 illustrates the initial assignment for the agents of $L \cup R \cup \{d\}$. In this figure, a dashed-line arrow from an object of G means that the agent pointed by the arrow is the only one that has a utility value different from ε for the object. Figure 1 illustrates the fact that the objects of G could be allocated in only two different manners in p to be a Pareto improvement of the initial endowment: either (i) every object of G is assigned to the same agent as in the initial assignment, or (ii) every object of G is assigned to the agent pointed at by the corresponding arrow in Figure 1.

First, we consider case (i) where all objects of G are assigned in p exactly as in the initial assignment. To achieve a utility of at least 4, every agent r_i should receive object f_i^R to complete her bundle of two objects. This implies that these objects cannot be assigned to agent c_i, with $i = 1 \dots k$, in order to ensure that they get a utility of at least $3k + a_i$. Therefore, each agent c_i should receive object h_i^{CR} with utility $3k$. Furthermore no agent c_i can receive an object f_j^L to complete her bundle of two objects because this object would provide her a utility of at most $a_i - 1$. Hence, each agent c_i should receive object h_i^{CL}. All in all, p should be exactly the same assignment as the initial assignment, which contradicts our assumption that p Pareto dominates this initial assignment.

We now consider case (ii) where each object of G should be assigned in p to the agent pointed at by the corresponding dotted arrow in Figure 1. Note that this is the only remaining case for a Pareto improvement of the initial assignment to exist. We examine now how the objects should be allocated in p for Pareto dominance to hold. To achieve a utility of at least 4, every agent ℓ_i should receive object h_i^{CL} to complete her bundle of two objects. If agent c_i does not receive at least one object f_j^R such that $a_i \geq j + 1$, then the maximal utility achievable by c_i would be $3k + a_i - 1$, which would be strictly lower than her utility in the initial assignment. Hence, each agent c_i should receive exactly one object f_j^R such that $a_i \geq j + 1$. This implies that no object f_i^R can be assigned to agent r_i. Hence, to achieve a utility of at least 4, every agent r_i should receive object h_i^{CR} to complete her bundle of two objects. Then object o should be assigned to agent d to complete her bundle of two objects. Finally it remains to assign to each agent c_i object f_j^L such that $a_i \geq j + 1$.

Now we focus on the pair of objects assigned to each agent c_i in p. Note that these two objects belong to F. We know that the total amount of utility provided by the objects of F to the agents of C should be exactly equal to $3k^2 + k(k+1)$. Furthermore, every agent c_i should receive an allocation of at least $3k + a_i$ of this total amount of utility. Since $\sum_{i=1}^{k}(3k + a_i) = 3k^2 + k(k+1)$, every agent c_i should receive two objects f_j^L and $f_{j'}^R$ such that $u_{c_i}(\{f_j^L, f_{j'}^R\}) = 3k + a_i$. Let π and θ be the two permutations of $\{1, \dots, k\}$ such that for any $i = 1 \dots k$, objects $f_{\pi(i)}^L$ and $f_{\theta(i)}^R$ are assigned in p to agent c_i. These two permutations are such that for any $i = 1 \dots k$, $\pi(i) + \theta(i) = a_i$. This leads to a contradiction with our assumption that $\{a_i : i = 1 \dots k\}$ is a no-instance. $\qquad \square$

Note that Theorem 2 is tight according to the number of objects initially owned by each agent because if initially each agent has exactly one object in assignment p, then our problem is well-known to be solvable in linear time [1].

9

Our final result shows that even if agents' utilities for objects are bounded to three values then the complexity of testing Pareto optimality remains strongly coNP-complete.

Theorem 3. *Testing Pareto optimality of a given assignment is strongly coNP-complete, even if there are only three utility values.*

Proof. The reduction is from (3,B2)-SAT [26] which is a restriction of 3SAT where each literal appears exactly twice in the clauses, and therefore, each variable appears four times. An instance of (3,B2)-SAT is composed by set $X = \{x_1, ..., x_s\}$ of s variables and set $C = \{c_1, ..., c_t\}$ of t clauses. Each clause in C contains exactly 3 literals. The decision problem is to define whether there exists a truth assignment of X such that all clauses in C are true.

The reduction from an instance of (3,B2)-SAT is as follows. Let q denote the initial endowment. For each variable x_i, we create two variable-agents V_i and W_i with $q(V_i) = \{v_i^1, v_i^2\}$ and $q(W_i) = \{w_i^1, w_i^2\}$. Intuitively, variable-agents V_i and W_i represent the truth value assigned to variable which can be either true or false, respectively. Object h_i is a gadget which ensures that variable x_i is considered either true or false, but not both. Both V_i and W_i have utility 1 for each good in their respective initial endowment, and both of them have utility t for h_i.

Furthermore, for each clause c_j we create clause-agent Z_j with $q(Z_j) = \{z_j\}$. Agent Z_j has utility 1 for z_j, and she has utility 1 for any good corresponding to a literal of c_j i.e., she has utility 1 for v_i^1 and v_i^2 (w_i^1 and w_i^2, respectively) if x_i ($\neg x_i$, respectively) is a literal appearing in c_j. Finally, we create two additional agents A and H with $q(A) = \{a\}$ and $q(H) = \{h_1, ..., h_t\}$. Agent A has utility t for a, and she has utility 1 for each object z_j, with $j = 1..t$. Agent H has utility 1 for each good in $q(H)$, and she has utility t for a. All the other utility values are equal to $\varepsilon = \frac{1}{6t}$. Intuitively, in any assignment which Pareto dominates the initial assignment, we will see that each clause-agent should receive an object corresponding to one of its literal, and agent A should receive all the objects initially owned by clause-agents. In other words, this gadget will enforce all clauses to be true.

The following table summarizes the utility values which differ from ε, where agt#1 is the agent which receives the object in the initial assignment and $u_{\text{agt\#1}}$ is her utility for it, and where agt(s)#2 lists the other agents having a utility different from ε for the object and $u_{\text{agt(s)\#2}}$ corresponds to their utility for it:

object	agt#1	$u_{\text{agt\#1}}$	agt(s)#2	$u_{\text{agt(s)\#2}}$
v_i^ℓ	V_i	1	Z_j where c_j contains x_j	1
w_i^ℓ	W_i	1	Z_j where c_j contains $\neg x_j$	1
z_j	Z_j	1	A	1
h_i	H	1	V_i	t
			W_i	t
a	A	t	H	t

We claim that the instance of (3,B2)-SAT is a yes-instance if and only if q is not Pareto optimal.

Assume that there exists truth assignment φ of X such that all clauses in C are true. We show that we can construct from φ an assignment p that Pareto dominates q. For each variable x_i that is true according to φ, $p(V_i) = \{h_i\}$ and $p(W_i) = \{w_i^1, w_i^2\}$. For each variable x_i that is false according to φ, $p(V_i) = \{v_i^1, v_i^2\}$ and $p(W_i) = \{h_i\}$. Note that in both cases, $u_{V_i}(p(V_i)) \geq u_{V_i}(q(V_i)) = 2$ and $u_{W_i}(p(W_i)) \geq u_{W_i}(q(W_i)) = 2$ hold, and at least one of these inequalities is strict. For each clause c_j and for each literal x_i ($\neg x_i$, respectively) appearing in c_j, if x_i is true (false, respectively) according to φ then v_i^ℓ (w_i^ℓ, respectively) is assigned to Z_j in p, where ℓ equals 1 if the first occurrence of x_i ($\neg x_i$, respectively) is in c_j and ℓ equals 2 otherwise. Since there is at least one literal true according to φ for each clause in C, at least one object is assigned to each Z_j in p and $u_{Z_j}(p(Z_j)) \geq u_{Z_j}(q(Z_j)) = 1$ holds. Finally, $p(H) = \{a\}$ and $p(A) = \{z_1, \ldots, z_t\}$. Note that both $u_A(p(A)) = u_A(q(A)) = t$ and $u_H(p(H)) = u_H(q(H)) = t$ hold. Therefore, p Pareto dominates q.

Assume now that there exists an assignment p that Pareto dominates q. Note first that $p(V_i) = q(V_i)$ and $p(W_i) = q(W_i)$ cannot hold for every $i = 1..t$. Indeed, otherwise it is easy to check that each agent should receive in p the same bundle as in q to ensure that each of them receives a bundle at least as good as her bundle in p, leading to a contradiction with p Pareto dominating q. This means that for some $k \in \{1, \ldots, t\}$, either V_k or W_k receive a bundle different from q. In both cases, this bundle should contain h_k to ensure that the utility of V_k or W_k does not decrease. But this implies that h_k cannot be assigned to H in p. To ensure $u_H(p(H)) \geq u_H(q(H)) = t$, agent H should receive a. Since a is not assigned to A, $\{z_1, \ldots, z_t\} \subseteq p(A)$ should hold to ensure $u_A(p(A)) \geq u_A(q(A)) = t$. Hence, no agent Z_j receives object z_j. To achieve a utility greater than or equal to $q(Z_j) = 1$, each agent Z_j should receive in p either *(i)* good v_i^ℓ corresponding to one of its positive literals, or *(ii)* good w_i^ℓ corresponding to one of its negative literals. In case *(i)*, v_i^ℓ is not assigned to V_i, and $u_{V_i}(p(V_i)) \geq u_{V_i}(q(V_i))$ implies $h_i \in p(V_i)$. Therefore, h_i is not assigned to W_i, and $p(W_i) = q(W_i) = \{w_i^1, w_i^2\}$ is necessary to ensure $u_{W_i}(p(W_i)) \geq u_{W_i}(q(W_i))$. This means that no agent $Z_{j'}$ receives object w_i^1 or w_i^2. In case *(ii)*, the same type of reasoning leads to conclude that if agent Z_j receives w_i^ℓ then no agent $Z_{j'}$ receives object v_i^1 or v_i^2. Therefore, we can construct from p a truth assignment of X that satisfies all clauses in C by setting the truth value of variable x_i to true if object v_i^1 or v_i^2 is assigned to some agent Z_j, and to false otherwise. In other words, the instance of (3,B2)-SAT is a yes-instance. □

3.2. Complexity of testing Pareto optimality: tractable cases

We now identify conditions under which the problem of computing individually rational and Pareto optimal assignments is polynomial-time solvable.

3.2.1. Constant number of agents and small utilities

We first focus on the case of a constant number of agents and small integer utility weights.

Lemma 2. *If there is a constant number of agents and the utilities are all integers, then the set of all vectors of utilities that correspond to an assignment can be computed in pseudo-polynomial-time.*

Proof. Consider the following algorithm (by 0^k we denote $0, \ldots, 0$ with k occurrences of 0). The symbol $\leftarrow$ denotes assignment.

```
1: L ← {(0^n)};
2: for j = 1 to m do
3:     L' ← {l + (0^{i-1}, u_i(o_j), 0^{n-i}) | i ∈ N; l ∈ L}
4:     L ← L'
5: end for
6: return  L
```

Let W be the maximal social welfare that is achievable; then, at any step of the algorithm, the number of vectors in L cannot exceed $(W + 1)^n$. Hence the algorithm runs in $O(W^n \cdot n \cdot m)$. Now, $W \leq \sum_{i,j} u_i(o_j)$, and since n is constant, the algorithm runs in pseudo-polynomial-time.

We can prove by induction on k that a vector of utilities $l = (v_1, \ldots, v_n)$ can be achieved by assigning objects $o_1, \ldots, o_k$ to the agents if and only if l belongs to L after objects $o_1, \ldots, o_k$ have been considered. This is obviously true at the start of the algorithm, when no object at all has been considered. Now, suppose the induction assumption is true for k. If l belongs to L after iteration k, then l' belongs to L after iteration $k + 1$ if l' is obtained from l by adding $u_i(o_k)$ to the utility of some agent i, that is, if $l = (v_1, \ldots, v_n)$ can be achieved by assigning objects $o_1, \ldots, o_{k+1}$. □

Theorem 4. *If there is a constant number of agents and the utilities are all integers, then there exists a pseudo-polynomial-time algorithm to compute a Pareto optimal and individually rational assignment.*

Proof. We apply the algorithm of Lemma 2, but in addition we keep track, for each $l \in L$, of a partial assignment that supports it: every time we add $l + (0^{i-1}, u_i(o_j), 0^{n-i})$ to L', the corresponding partial assignment is obtained from the partial assignment corresponding to l, and then mapping o_j to i. If several partial assignments correspond to the same utility vector, we keep an arbitrary one. At the end, we obtain the list L of all feasible utility vectors, together with, for each of them, one corresponding assignment. For each of them, check whether there is at least one l' in L that Pareto dominates it, which takes at most $O(|L|^2)$, which is polynomially large if L is encoded in unary. The assignments that correspond to the remaining vectors are Pareto optimal.[4] □

---

[4]Note that it is generally not the case that we get *all* Pareto optimal assignments: if there

3.2.2. Lexicographic Utilities

We say that utilities are *lexicographic* if for each agent $i \in N$ and each object $o \in O$, $u_i(o) > \sum_{o' \prec_i o} u_i(o')$ with the convention $\sum_{o \in \emptyset} u_i(o) = 0$, which implies that $u_i(o) > 0$ for each o. Note that agent 1 in Example 1 has lexicographic utilities. So would be an agent with utilities $(8, 8, 3, 3, 1)$.

In order to test the Pareto optimality of an assignment p, we construct a directed graph $G(p) = (V(p), E(p))$, called the envy graph of p [1]: The set of vertices $V(p)$ contains one vertex per object belonging to O. Furthermore, for any vertex of $V(p)$ associated to object o, the set of edges $E(p)$ contains one edge (o, o') for any object o' belonging to $O \setminus \{o\}$ such that $o' \succsim_i o$, where i is the agent who receives object o in p.

Example 2. *Let $n = 3$, $m = 5$, and the following ordinal information about preferences corresponding to the lexicographic utilities. (Due to the lexicographicity assumption, this ordinal information is enough to know completely an agent's preference relation over all sets of objects.)*

$$1 : o_1 \succ o_2 \succ o_3 \succ o_4 \succ o_5$$
$$2 : o_1 \succ o_2 \sim o_5 \succ o_3 \sim o_4$$
$$3 : o_1 \sim o_4 \succ o_3 \sim o_5 \succ o_2$$

Let $p = (o_2 o_4 | o_1 | o_3 o_5)$ be the initial assignment. Then, the corresponding envy graph $G(p)$ is depicted in Figure 2. In Figure 2, dashed-line edges represent indifferences (when $o' \sim o$) and solid-line edges represent strict preferences (when $o' \succ o$).

It follows from [16] that Pareto optimality of an assignment under lexicographic utilities can be tested in polynomial time. We provide a simple characterization of a Pareto optimal assignment under lexicographic utilities. The characterization we present also provides an interesting connection with the possible Pareto optimality that we consider in the next section.

Theorem 5. *An assignment p is not Pareto optimal with respect to lexicographic utilities if and only if there exists a cycle in $G(p)$ which contains at least one edge corresponding to a strict preference.*

Proof. Assume that there exists a cycle C that contains at least one edge corresponding to a strict preference. Then, the exchange of objects along the cycle by agents owning these objects corresponds to a Pareto improvement.

Assume now that p is not Pareto optimal and let q_1 be an assignment that Pareto dominates p. For at least one agent i, $q_1(i) \succ_i p(i)$. Therefore, there exists at least one object o_1 in $q_1(i) \setminus p(i)$. Let i_1 be the owner of o_1 in p. Since preferences are lexicographic, in compensation of the loss of o_1, agent i_1 must receive an object o_2 in q_1 which is at least as good as o_1 according to her own preferences. Let i_2 be the owner of o_2 in p. Agent i_2 must also receive an

---

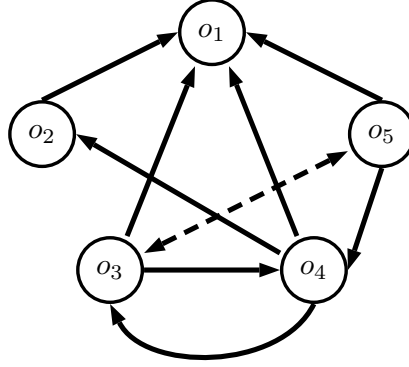are several assignments corresponding to the same utility vector, then we'll obtain only one.

Figure 2: Envy graph $G(p)$ for assignment p in Example 2.

object o_3 in q_1 which is at least as good as o_2 according to her own preferences. More generally, we can find a sequence of agents $i_1, i_2, \ldots, i_k$ and a sequence of objects $o_1, o_2, \ldots, o_{k+1}$ such that each agent i_j receives object o_{j+1} in q_1 in exchange of the loss of object o_j, and o_{j+1} is at least as good as o_j according to the preference of agent i_j. Since O is finite, there must exist k and k' such that the sequence $o_k \rightarrow o_{k+1} \rightarrow \ldots \rightarrow o_{k'}$ forms a cycle, i.e., $o_k = o_{k'}$. If $\not\exists l \in [k, k'-1]$ such that $o_{l+1} \succ_{i_l} o_l$ then we consider assignment q_2 derived from q_1 by reassigning every object o_{l+1}, with $l \in [k, k'-1]$, to agent i_l. It is obvious that assignment q_2 is at least as good as q_1 for all agents. Hence, q_2 Pareto dominates p. By following the same reasoning as above, we can state that there exists a sequence of objects $o_k \rightarrow o_{k+1} \rightarrow \ldots \rightarrow o_{k'}$ such that $o_k = o_{k'}$ and for every $l \in [k, k'-1]$, o_{l+1} is assigned to agent i_l in q_2 to compensate the loss of o_l assigned to him in p (obviously with $o_{l+1} \succsim_{i_l} o_l$). Once again, if there is no $l \in [k, k'-1]$ such that $o_{l+1} \succ_{i_l} o_l$ then we consider assignment q_3 derived from q_2 by reassigning every object o_{l+1}, with $l \in [k, k'-1]$, to agent i_l, and so on. Since for any $s > 1$ we have $\sum_{i \in N} |q_{s-1}(i) \cap p(i)| > \sum_{i \in N} |q_s(i) \cap p(i)|$, there must exist a finite value t where $\exists l \in [k, k'-1]$ such that $o_{l+1} \succ_{i_l} o_l$ for the cycle $o_k \rightarrow o_{k+1} \rightarrow \ldots \rightarrow o_{k'}$ founded in q_t. Indeed otherwise after a finite number of steps t we would have $q_t(i) = p(i)$ for all $i \in N$, leading to a contradiction with our assumption that q_t Pareto dominates p. Therefore, we have shown that there exists a cycle $o_k \rightarrow \ldots \rightarrow o_{k'}$ in $G(p)$ with at least one edge corresponding to a strict preference. $\qquad\square$

It is clear that graph $G(p)$ can be constructed in linear time for any assignment p. Furthermore, the search of a cycle containing at least one strict preference edge in $G(p)$ can be computed in linear time by applying a graph traversal algorithm for each strict preference edge in $G(p)$. Therefore, testing if a given assignment is Pareto optimal can be done in linear time when utilities are lexicographic.

The construction of Theorem 5, and illustrated in Figure 2, gives us that

assignment p in Example 2 is Pareto dominated by $(o_2 o_3 | o_1 | o_4 o_5)$.

3.2.3. Two utility values

In this section we assume that agents use at most two utility values to evaluate objects. We say that the collection of utility functions $(u_1, \ldots, u_n)$ is *bivalued* if there exist two numbers $\alpha > \beta \geq 0$ such that for every agent i and every object o, $u_i(o) \in \{\alpha, \beta\}$.[5] This implies that for each agent i, the set of objects O is partitioned into two subsets $E_i^1 = \{o \in O, u_i(o) = \alpha\}$ and $E_i^2 = \{o \in O, u_i(o) = \beta\}$ (with possibly $E_i^1 = \emptyset$ or $E_i^2 = \emptyset$). Given an assignment q, let $q^+(i) = q(i) \cap E_i^1$, and $q^-(i) = q(i) \cap E_i^2$.

We provide a first requirement for an assignment to Pareto dominate another one:

Lemma 3. *If an assignment p is Pareto dominated by an assignment q then $|\bigcup_{i \in N} q^+(i)| > |\bigcup_{i \in N} p^+(i)|$.*

Proof. For contradiction we assume that $|\bigcup_{i \in N} q^+(i)| \leq |\bigcup_{i \in N} p^+(i)|$. In that case, $SW(q) = |\bigcup_{i \in N} q^+(i)|\alpha + |\bigcup_{i \in N} q^-(i)|\beta = |\bigcup_{i \in N} q^+(i)|(\alpha - \beta) + |O|\beta \leq |\bigcup_{i \in N} p^+(i)|(\alpha - \beta) + |O|\beta = \alpha|\bigcup_{i \in N} p^+(i)| + \beta|\bigcup_{i \in N} p^-(i)|$. Therefore, $SW(p) \geq SW(q)$ holds, which contradicts the assumption that q Pareto dominates p. $\square$

Lemma 4. *If an assignment p is not Pareto optimal then there exists an assignment q that Pareto dominates p with the two following properties: (i) $\forall i \in N, |q^+(i)| \geq |p^+(i)|$ and (ii) $\exists j \in N, |q^+(j)| > |p^+(j)|$ and $p^-(j) \neq \emptyset$.*

Proof. Assume that p is not Pareto optimal. Then there exists an assignment which Pareto dominates p. Among the set of assignments that Pareto dominate p, let q_* be the one that is the closest to p i.e., $|\bigcup_{i \in N} q_*(i) \setminus p(i)|$ is minimal.

First we note that the above assumption implies that there is no *clear winner* agent i such that $p^-(i) = \emptyset$ and $q_*^+(i) \supset p^+(i)$. Indeed, otherwise we could reallocate each object in $q_*^+(i) \setminus p^+(i)$ to its owner in p, and obtain another assignment q_{**} from q_* that also Pareto dominates p, but which is closer to p than q_*.

Lemma 3 implies that $|\bigcup_{i \in N} q_*^+(i)| > |\bigcup_{i \in N} p^+(i)|$ holds. Let o_1 be an object of $(\bigcup_{i \in N} q_*^+(i)) \setminus (\bigcup_{i \in N} p^+(i))$. Assume that object o_1 belongs to $q_*^+(i_2)$ for a given agent i_2, and to $p^-(i_1)$ for another agent i_1. If $p^-(i_2) \neq \emptyset$ then starting from p, i_1 and i_2 could exchange o_1 with an object of $p^-(i_2)$ leading to an assignment q where both (i) and (ii) hold. More precisely, agent i_1 weakly prefers $q(i_1)$ to $p(i_1)$ since she gives object o_1 (belonging to $E_{i_1}^2$) in exchange of an arbitrary object (either belonging to E_{i_1} or E_{i_2}). Furthermore, agent i_2 receives object o_1 (belonging to $E_{i_2}^1$) in exchange of an object belonging to $E_{i_2}^2$,

---

[5]Note that the result would still hold if each agent i has a different pair of values (α_i, β_i), provided that $\frac{\alpha_i}{\beta_i} = \frac{\alpha}{\beta}$ for all i. This is due to the fact that $\alpha|q^+(i)| + \beta|q^-(i)| \geq \alpha|p^+(i)| + \beta|p^-(i)| \Leftrightarrow \alpha_i|q^+(i)| + \beta_i|q^-(i)| \geq \alpha_i|p^+(i)| + \beta_i|p^-(i)|$ holds for every i since the second inequality can be obtained by multiplying the first one by $\frac{\alpha_i}{\alpha}$.

and strictly prefers $q(i_2)$ to $p(i_2)$ (note that $p^-(i_2) \neq \emptyset$ holds by hypothesis). Otherwise, if $p^-(i_2) = \emptyset$, then let $o_2 \in p^+(i_2) \setminus q_*^+(i_2)$ (which must exists since i_2 is not a clear winner). Assume that o_2 belongs to $q_*(i_3)$ for a given agent i_3. Note that o_2 must belong to $q_*^+(i_3)$, as otherwise i_2 and i_3 could exchange o_1 and o_2 in q_* and we would obtain another assignment q_{**} that still Pareto dominates p, but which is closer to p. Now, again, if $p^-(i_3) \neq \emptyset$ then starting from p we could create a Pareto dominating assignment q with properties (i) and (ii) by exchanging objects along this cycle, namely, by assigning o_1 to i_2, o_2 to i_3 and o_3 to i_1, where o_3 is an object of $p^-(i_3)$. However, if $p^-(i_3) = \emptyset$ then we continue the construction of the sequence.

The last case that we have to discuss is a possible repetition occurring in the above sequence. Suppose that for some indices $k < l$, $o_l \in q_*^+(i_k)$ for the first time in the sequence. Therefore, the agents involved in this sub-sequence exchange their top objects in q_* compared to p. But then we can construct another assignment q_{**} from q_* by reassigning these objects to their original owners in p, contradicting with our assumption of q_* being as close to p as possible. $\square$

Based on the lemma, we obtain the following characterization of Pareto optimality in the bivalued case.

Lemma 5. *An assignment p is Pareto dominated if and only if there exists an assignment q such that (i) $\forall i \in N, |q^+(i)| \geq |p^+(i)|$ and (ii) $\exists j \in N, |q^+(j)| > |p^+(j)|$ and $p^-(j) \neq \emptyset$.*

Proof. One implication has already been proved in Lemma 4. To prove the second implication we assume first that there exists q such that (i) and (ii) hold. Let j be the index described in (ii). For any $i \in N \setminus \{j\}$, let A_i be an arbitrary subset of $q^+(i)$ of size $|A_i| = |p^+(i)|$. Let A_j be an arbitrary subset of $q^+(j)$ of size $|A_j| = |p^+(j)| + 1$. Let $\bar{A} = O \setminus \bigcup_{i \in N} A_i$. Note that by definition $|\bar{A}| = |\bigcup_{i \in N} p^-(i)| - 1$ because all objects are assigned, and therefore $|\bar{A}| = |O| - |\bigcup_{i \in N} A_i| = \sum_{i \in N} |p(i)| - \sum_{i \in N} |p^+(i)| - 1 = \sum_{i \in N} |p^-(i)| - 1$.

We arbitrarily partition $\bar{A}$ into n subsets $\bar{A}_1, \ldots \bar{A}_n$ such that for all $i \in N \setminus \{j\}, |\bar{A}_i| = |p^-(i)|$ and $|\bar{A}_j| = |p^-(j)| - 1$. Finally, let q_* be the assignment such that for all $i \in N, q_*(i) = A_i \cup \bar{A}_i$. By the construction of q_*, $|q_*^+(i)| \geq |p^+(i)|$ and $|q_*(i)| = |p(i)|$ hold for every $i \in N$, and $|q_*^+(j)| > |p^+(j)|$ and $|q_*(j)| = |p(j)|$ hold. Therefore, p is Pareto dominated by q_*. $\square$

If at least one object is not assigned then a trivial Pareto improvement would be to assign this object to an agent. Therefore, we can focus on the case where all objects are assigned. According to Theorem 5, a Pareto improvement can be computed by focusing on the assignment of top objects for the agents. We describe an algorithm based on maximum flow problems to obtain such an assignment. For any $i \in N$, let $G_i = (V_i, E_i)$ be a directed graph that models the search for a Pareto improvement for agent i as a flow problem. Set of vertices V_i contains one vertex per agent, one vertex per object, plus source s and sink t. To ease the notation, vertices are denoted as the agents or the objects that

they represent, and therefore, $V_i = N \cup O \cup \{s,t\}$. The set of edges E_i and their capacities are constructed as follows:

- For any $l \in N$ and $o \in O$ such that $o \in E_l^1$, there is an arrow (l, o) with capacity 1.

- For any $o \in O$, there is an arrow (o, t) with capacity 1.

- For any $l \in N \setminus \{i\}$, there is an arrow (s, l) with capacity $|p^+(l)|$, and there is an edge (s, i) with capacity $|p^+(i)| + 1$.

It is easy to show that there exists a flow of value $\sum_{l \in N} |p^+(l)| + 1$ if and only if there exists an assignment such that every agent $l \in N \setminus \{i\}$ receives at least $|p(l) \cap E_l^1|$ top objects and agent i receives $|p^+(i)| + 1$ top objects. Hence, by Lemma 5, there exists a Pareto improvement of p if and only if there exists $i \in N$ such that $p(i) \cap E_i^2 \neq \emptyset$ and there exists a flow of value $\sum_{l \in N} |p^+(l)| + 1$ in G_i. Therefore, finding a Pareto improvement can be performed in polynomial time by solving at most n maximum flow problems.

Theorem 6. *Under bivalued utilities, there exists a polynomial-time algorithm for checking Pareto optimality and finding a Pareto improvement, if any.*

Note that we can find a Pareto optimal assignment in polynomial time as well by repeatedly obtaining Pareto improvements of the next assignment. In each Pareto improvement the number of top objects increases by at least one, and therefore there is at most m Pareto improvements before reaching a Pareto optimal assignment.

Example 3. *Let $n = 3$, $m = 6$, $p = (o_1 o_4 | o_2 o_5 | o_3 o_6)$, $E_1^1 = \{o_1, o_2, o_3\}$, $E_2^1 = \{o_2\}$ and $E_3^1 = \{o_1, o_3, o_5, o_6\}$. G_1 is depicted in Figure 3. The flow of value 5 (bold-lines) gives assignment $(o_1 o_3 | o_2 o_4 | o_5 o_6)$, which Pareto dominates p.*

# agents \ utility values	≤ 2	lexicographic, ≥ 3	≥ 3
≥ 2, bounded	in P	in P	weakly NP-complete (Th. 1) pseudo-polynomial (Th. 4)
unbounded	in P (Th. 6)	in P (Th. 5)	strongly NP-complete (Ths. 2 and 3)

Table 1: Complexity of testing Pareto optimality

We conclude this section with Table 1 that summarizes our results for cardinal utilities.

3.3. Conservative Pareto optimality

In this section, we consider a related concept of conservative Pareto optimality and provide connections between testing Pareto optimality with testing
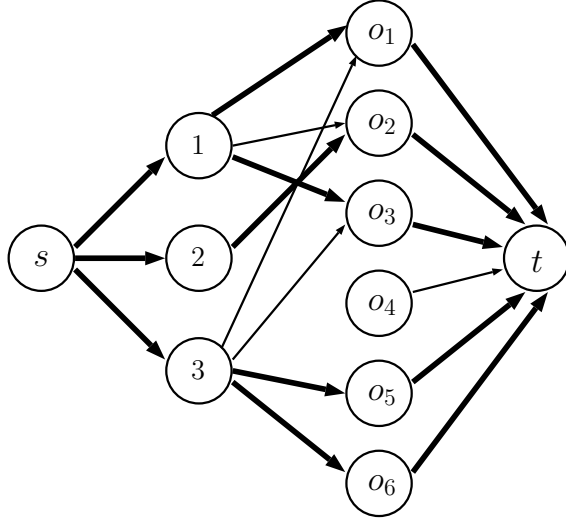
Figure 3: Flow network G_1 in Example 3.

conservative Pareto optimality. An assignment p is *conservatively Pareto optimal* if there does not exist another assignment q that Pareto dominates p and $|q(i)| = |p(i)|$ for all $i \in N$. Conservative Pareto optimality applies in several scenarios in which number of objects allocated to each agent needs to be conserved. This requirement is applicable in time-bank markets in which each agent received the same units of services as they give to their peers [2, 27]. Similarly, it applies to the Erasmus exchange program in which there is a requirement that if a university sends k students abroad, then it has to welcome exactly k from abroad [20].

Next we point out that a polynomial-time algorithm to test Pareto optimality can also be used to test conservative Pareto optimality in polynomial time.

Lemma 6. *There exists a polynomial-time algorithm to test Pareto optimality if and only if there exists a polynomial-time algorithm to test conservative Pareto optimality.*

Proof. For the left-to-right direction: for each $i \in N$ and $o \in O$, let us update the utility $u_i(o)$ to utility $u_i(o)+C$ where C is $m \times \max_{i \in N, o \in O} u_i(o)$. Note that by modifying the utility function, we ensure that an agent primarily cares about how many objects she gets and then cares about what are the particular utilities for the objects. Hence, for the modified utilities, in any Pareto improvement, no agent will get less objects than in her original allocation and hence each agent will get the same number of objects after reallocation. Hence, q Pareto dominates p with respect to the modified utilities if and only if q conservatively Pareto dominates p with respect to the original utilities. Hence, by simply modifying the utilities as specified above, a polynomial-time algorithm to test

18

Pareto optimality can be used to test conservative Pareto optimality.

For the right-to-left direction, consider an instance consisting of n agents, m objects, u the utility matrix, and p an assignment. We define the new set of objects by adding $(n-1)m$ dummy objects that each agent values at 0 (agent value the initial objects as in u). The assignment q is defined from p by giving $m - |p(i)|$ objects to agent i, for all i. Then q is conservatively Pareto optimal for the modified instance if p is Pareto optimal for the original instance. $\square$

4. Ordinal preferences

In this section, we consider the setting in which agents have additive cardinal utilities but only their ordinal preferences over objects are known by the central authority. This could be because the elicitation protocol did not ask the agents to communicate their utilities, or simply because they do not know them precisely. We will assume in this section that for any underlying additive utility $u_i(o) > 0$ for $i \in N$ and $o \in O$. In this case, one can still reason whether a given assignment is Pareto optimal with respect to some or all cardinal utilities consistent with the ordinal preferences. An assignment p is *possibly Pareto optimal* with respect to preference profile $\succsim$ if there exists $u \in \mathscr{U}(\succsim)$ such that p is Pareto optimal for u. An assignment is *necessarily Pareto optimal* with respect to preference profile $\succsim$ if for any $u \in \mathscr{U}(\succsim)$ the assignment p is Pareto optimal for u.

4.1. Possible Pareto Optimality

We first note that necessary Pareto optimality implies possible Pareto optimality. Secondly, at least one necessarily Pareto optimal assignment exists in which all the objects are given to one agent. Since the computation of a possibly or necessary Pareto optimal assignment can be performed in polynomial time, we focus on the problems of testing possible and necessary Pareto optimality.

Theorem 7. *A assignment is (1) possibly Pareto optimal if and only if (2) there exists no cycle in $G(p)$ which contains at least one edge corresponding to a strict preference if and only if (3) it is Pareto optimal under lexicographic utilities.*

Proof. We already established (2) $\Longleftrightarrow$ (3) in Theorem 5. We also know that (3) $\Longrightarrow$ (1) because if an assignment that is Pareto optimal with respect to lexicographic utilities, it is possibly Pareto optimal. It remains to be shown that (1) $\Longrightarrow$ (2). Suppose p is not Pareto optimal with respect to lexicographic utilities, then by Theorem 5, $G(p)$ admits a cycle which contains at least one edge corresponding to a strict preference. Now consider an assignment q which is a result of exchange of objects along the cycle. If o points to o' in the cycle, then the agent getting o in p now gets o' in place of o. Since the cycle contains at least one one edge corresponding to a strict preference, assignment q is a Pareto improvement over p with respect to all utilities consistent with the ordinal preferences. Hence p is not possibly Pareto optimal. $\square$

Since the characterization in Theorem 5 also applies to possible Pareto optimality, hence possible Pareto optimality can be tested in linear time. The argument in the proof above also showed that possible Pareto optimality is equivalent to Pareto optimality under lexicographic preferences.

The following example points out that a possibly Pareto optimal assignment may not be a necessarily Pareto optimal assignment.

Example 4. *Consider two agents with identical preferences* $o_1 \succ o_2 \succ o_3 \succ o_4$. *Every assignment is possibly Pareto optimal; however assignment* $p = (o_1o_4|o_2o_3)$ *is not necessarily Pareto optimal since it is not Pareto optimal for the following utilities.*

	o_1	o_2	o_3	o_4
1	10	9	8	7
2	10	3	2	1

4.2. Necessary Pareto Optimality

Next we present two characterizations of necessary Pareto optimality. The first highlights that necessary Pareto optimality is identical to the necessary Pareto optimality considered by Bouveret et al. [12].

We first define the *responsive set extension* [10], which extends preferences over objects to preferences over sets of objects. Formally, for agent $i \in N$, her preferences $\succsim_i$ over O are extended to her preferences $\succsim_i^{RS}$ over 2^O as follows: $q(i) \succsim_i^{RS} p(i)$ if and only if there exists an injection f from $p(i)$ to $q(i)$ such that for each $o \in p(i)$, $f(o) \succsim_i o$. Since $\succsim_i^{RS}$ is a partial order, we say a preference R_i is a *completion* of $\succsim_i^{RS}$ if it is a complete and transitive relation over sets of objects that is consistent with $\succsim_i^{RS}$.

We say that an assignment is *RS-efficient* if it is Pareto optimal with respect to the RS set extension relation of the agents. An assignment p is RS-efficient, if there does not exist another assignment q such that $q(i) \succsim_i^{RS} p(i)$ for all $i \in N$ and $q(i) \succ_i^{RS} p(i)$ for some $i \in N$. We say that q strictly RS-dominates p if q Pareto dominates p with respect to RS.

Theorem 8. *An assignment is necessarily Pareto optimal if and only if it is Pareto optimal under all completions of the responsive set extension.*

Proof. If an assignment is not Pareto optimal under certain additive preferences, it is by definition not Pareto optimal under this particular completion of responsive preferences.

Assume that assignment p is not Pareto optimal under some completion of the responsive set extension. Then there exists another assignment q in which (1) for all $i \in N$ $q(i) \succsim_i^{RS} p(i)$ or $((p(i) \not\succ_i^{RS} q(i)$ and $q(i) \not\succ_i^{RS} p(i)))$, and (2) for some $j \in N$, $q(j) \succ_j^{RS} p(j)$ or $(p(j) \not\succ_j^{RS} q(j)$ and $q(j) \not\succ_j^{RS} p(j))$ For both cases, if the allocations are incomparable with respect to responsive set extension, then there exists an object o such that $|q(i) \cap \{o' : o \succsim_i o\}| > |p(i) \cap \{o' : o \succsim_i o\}|$. In that case, for a small enough $\epsilon > 0$, consider a utility function u_i in which $u_i(o''') - u_i(o'') \leq \epsilon$ for all $o''', o'' \succsim_i o$ and $u_i(o) > \sum_{o' \prec_i o} u_i(o') + |O|\epsilon$. For u_i, $u_i(q(i)) > u_i(p(i))$. $\square$

For our second characterization of necessarily Pareto optimal assignments, we define a *one-for-two Pareto improvement swap* as an exchange between two agents i_j and i_k involving objects o_j^1, $o_j^2 \in p(i_j)$ and $o_k \in p(i_k)$ such that $o_k \succ_{i_j} o_j^1 \succsim_{i_j} o_j^2$.

Theorem 9. *An assignment p is necessarily Pareto optimal if and only if*

    (i) it is possibly Pareto optimal and

    (ii) it does not admit a one-for-two Pareto improvement swap.

Proof. We first show that if an assignment does not satisfy the two conditions, then it is not necessarily Pareto optimal. Possible Pareto optimality is a requirement for the assignment to be necessarily Pareto optimal. To see that the second condition is also necessary, we have to show that if p admits a one-for-two Pareto improvement swap then p is not necessarily Pareto optimal. This is because the swap could indeed be a Pareto improvement for these two agents with the following utilities: $u_{i_j}(o_k) > 2u_{i_j}(o_j^1)(\geq u_{i_j}(o_j^1) + u_{i_j}(o_j^2))$ and $u_{i_k}(o_k) < u_{i_k}(o_j^1) + u_{i_k}(o_j^2)$. These utilities are compatible with the ordinal preferences of these agents, because of the assumption $o_k \succ_{i_j} o_j^1 \succsim_{i_j} o_j^2$ (and irrespective to the ordinal preferences of i_k).

Conversely, to show that conditions (i) and (ii) are sufficient for the assignment to be necessarily Pareto optimal, suppose for a contradiction that (1) p is not necessarily Pareto optimal and (2) p does not admit a one-for-two Pareto improvement swap. We will then show that there is an assignment that strictly RS-dominates p, implying that p cannot be possibly Pareto optimal.

From (1) and Theorem 8, we have that (3) there is another assignment q and a collection of additive utility functions $u = (u_1, \ldots, u_n) \in \mathscr{U}(\succsim)$ such that q Pareto dominates p with respect to u.

Without loss of generality we may assume that each agent receives a nonempty bundle in p. Regarding **the structure** of p, first we observe that the lack of one-for-two Pareto improvement swaps implies that every agent is assigned to some (or none) of her top objects and possibly to one additional object that she ranks lower. Formally, let $T_p(i)$ denote a set of i's top objects she is assigned to in p, i.e., $T_p(i) = \{o : o \in p(i) \text{ s.t. } \nexists o' \notin p(i), o' \succ_i o\}$. Then $p(i) = T_p(i) \cup \{w_p(i)\}$, where $w_p(i)$ is either a single object or no object.

We show that $|q(i)| = |p(i)|$ must hold for every agent i. Suppose not, then there is an agent i for which $|q(i)| < |p(i)|$. By the definition of $T_p(i)$ it is straightforward that if $w_p(i) = \emptyset$ then $u_i(p(i)) = u_i(T_p(i)) > u_i(q(i))$, and if $w_p(i) \neq \emptyset$ then $u_i(p(i)) = u_i(T_p(i) \cup \{w_p(i)\}) > u_i(T_p(i)) \geq u_i(q(i))$, a contradiction. Furthermore, for every agent i, if $\{w_p(i)\} \neq \emptyset$ then for any object $o \in q(i)$ we have $o \succsim_i w_p(i)$. Otherwise, if there was an agent i with $o \in q(i)$ such that $w_p(i) \succ_i o$, then $u_i(T_p(i)) \geq u_i(q(i) \setminus \{o\})$ would imply $u_i(p(i)) = u_i(T_p(i) \cup \{w_p(i)\}) > u_i(q(i))$.

Now we construct a so-called **Pareto improvement sequence** with respect to p and q, which consists of a sequence of agents $\{i_1, i_2, \ldots i_k\}$ with possible repetitions and a set of distinct objects $\{o_1, o_2, \ldots, o_m\}$ such that

- $o_1 \in q(i_2) \setminus p(i_2)$, $o_2 \in p(i_2) \setminus q(i_2)$, and $o_1 \succsim_{i_2} o_2$;

- $o_2 \in q(i_3) \setminus p(i_3)$, $o_3 \in p(i_3) \setminus q(i_3)$, and $o_2 \succsim_{i_3} o_3$;

- $\ldots$

- $o_m \in q(i_1) \setminus p(i_1)$, $o_1 \in p(i_1) \setminus q(i_1)$, and $o_m \succsim_{i_1} o_1$.

and with strict preference for at least one agent.

The presence of the above Pareto improvement sequence would imply the existence of an assignment q' that RS-dominates p, obtained by letting the agents exchange their objects along the sequence, i.e., with $q'(i) = p(i) \cup \{o_{k-1} : i_k = i, k = 1, \ldots, m\} \setminus \{o_k : i_k = i, k = 1, \ldots, m\}$. This would contradict our assumption that p is possibly Pareto optimal.

We first define three types of agents, placed in the corresponding sets, namely X, Y, Z, and a **one-to-one mapping** π from a subset of O to itself such that for all i, if $o \in p(i) \setminus q(i)$ and $\pi(o) \in q(i) \setminus p(i)$ then i is indifferent between these two objects. In the set X we put all the agents with either no $w_p(i)$ or with $w_p(i) \in q(i)$. Each agent i in this set must be indifferent between all objects in $(p(i) \setminus q(i)) \cup (q(i) \setminus p(i))$ (i.e., these object are in a single tie in i's preference list) by the following reasons. $|p(i)| = |q(i)|$ implies $|p(i) \setminus q(i)| = |q(i) \setminus p(i)|$. By the definition of $T_p(i)$ it follows that any object in $p(i) \setminus q(i)$ is weakly preferred to any object in $q(i) \setminus p(i)$ by i. However, from (3) we have $u_i(q(i)) \geq u_i(p(i))$, which implies that $u_i(q(i) \setminus p(i)) \geq u_i(p(i) \setminus q(i))$, which can only happen if i is indifferent between any two objects in $(p(i) \setminus q(i)) \cup (q(i) \setminus p(i))$. Let π map $q(i) \setminus p(i)$ to $p(i) \setminus q(i)$ as a bijective function.

Next, let Y contain every agent i who has object $w_p(i)$ such that there is an object $o \in q(i) \setminus p(i)$ with $o \sim_i w_p(i)$. In this case i must be indifferent between all objects in $(T_p(i) \setminus (q(i) \setminus \{o\})) \cup ((q(i) \setminus \{o\}) \setminus T_p(i))$.

Indeed, $|p(i)| = |q(i)|$ implies $|T_p(i) \setminus (q(i) \setminus \{o\})| = |(q(i) \setminus \{o\}) \setminus T_p(i)|$.

By the definition of $T_p(i)$ any object in $T_p(i) \setminus (q(i) \setminus \{o\})$ is weakly preferred to any object in $(q(i) \setminus \{o\}) \setminus T_p(i)$ by i. On the other hand, $u_i(q(i)) \geq u_i(p(i))$ and $o \sim_i w_p(i)$ implies $u_i((q(i) \setminus \{o\}) \setminus T_p(i)) \geq u_i(T_p(i) \setminus (q(i) \setminus \{o\}))$, leading to the conclusion that i must be indifferent between all objects in $(T_p(i) \setminus (q(i) \setminus \{o\})) \cup ((q(i) \setminus \{o\}) \setminus T_p(i))$. Therefore we let π map o to $w_p(i)$ and $(q(i) \setminus \{o\}) \setminus T_p(i)$ to $T_p(i) \setminus q(i) \setminus \{o\}$.

Thirdly, let Z contain every agent i with object $w_p(i)$ such that for every $o \in q(i)$, $o \succ_i w_p(i)$. Note that there is at least one agent in Z, the one who gets strictly better off in q, as otherwise, if there was an object $o \in q(i)$ such that $w_p(i) \succsim_i o$, then $u_i(T_p(i)) \geq u_i(q(i) \setminus \{o\})$ would imply $u_i(p(i)) = u_i(T_p(i) \cup \{w_p(i)\}) \geq u_i(q(i))$.

Finally, we shall note that if $T_p(i)$ is empty then $|p(i)| = |q(i)| = 1$, so either i is indifferent between $p(i) = \{w_p(i)\}$ and $q(i)$, in which case i is in Y with $\pi(q(i)) = p(i)$, or i strictly prefers $q(i)$ to $p(i)$ and then i belongs to Z.

To summarize, so far we have that for any $i \in X \cup Y$ and $o \in q(i) \setminus p(i)$ we associate an object $\pi(o) \in p(i) \setminus q(i)$ such that $o \sim_i \pi(o)$. Furthermore, for any $i \in Z$ and $o \in q(i) \setminus p(i)$ we have that $o \succ_i w_p(i)$.

We build a Pareto improvement sequence as a part of a sequence involving agents $i_1, i_2, \ldots$ with corresponding objects $o_1, o_2, \ldots$ starting from any $i_1 \in Z$ with $o_1 = w_p(i)$. For every $k \geq 2$, let i_k be the agent who receives o_{k-1} in q. If $i_k \in X \cup Y$ then let $o_k = \pi(o_{k-1})$, and if $i_k \in Z$ then let $o_k = w_p(i)$. We terminate the sequence when an object is first repeated. This repetition must occur at some agent in Z, since for every agent i the objects in $q(i) \setminus p(i)$ are in a one-to-one correspondence with those in $p(i) \setminus q(i)$ by π.

Let the first repeated object be, say, $i_s = i_t \in Z$ for indices $1 \leq s < t$. We show that the sequence $i_s, \ldots, i_{t-1}$ is a Pareto improvement sequence. To see this, let us first consider an agent $i \in X \cup Y$. Whenever i appears in the sequence as $i_k \in \{i_{s+1}, \ldots, i_t\}$ she receives object $o_{k-1} \in q(i) \setminus p(i)$ and in return she gives away $\pi(o_{k-1}) = o_k \in p(i) \setminus q(i)$, where i is indifferent between o_{k-1} and o_k. Now, let $i \in Z \setminus \{i_t\}$ that appears as $i_l \in \{i_{s+1}, \ldots, i_t\}$. She receives object $o_{l-1} \in q(i) \setminus p(i)$ and in return she gives away $w_p(i) = o_k \in p(i) \setminus q(i)$, where $o_{l-1} \succ_i w_p(i)$ by the definition of Z. Since i appears in this sequence only once, it is obvious that $u_i(q(i)) > u_i(p(i))$. Finally, regarding $i = i_s = i_t \in Z$, i receives $o_{t-1} \in q(i) \setminus p(i)$ and she gives away $w_p(i) = o_s \in p(i) \setminus q(i)$, where $o_{t-1} \succ_i w_p(i)$. So we constructed a Pareto improvement sequence, and therefore p is not possibly Pareto optimal, a contradiction. $\qquad \square$

In Example 4, p is not necessarily Pareto optimal because it admits a one-for-two Pareto improvement swap: $o_2, o_3 \in p(2)$, $o_1 \in p(1)$ and $o_1 \succ_2 o_2 \succsim_2 o_3$. Example 4 also shows that although an assignment may not be necessarily Pareto optimal there may not be an assignment that Pareto dominates it for *all* utilities consistent with the ordinal preferences. The characterization above also gives us a polynomial-time algorithm to test necessary Pareto optimality.

Finally, we would like to highlight that in some natural situations, where the preferences of the agents are aligned, all necessarily Pareto optimal assignments can be seen as very unfair, as we illustrate in the following example.

Example 5. *Suppose that we have m objects and n agents, where n is much larger than m and all agents have the same preferences, e.g. they order objects according to their indices. In this situation in any necessarily Pareto efficient assignment essentially one agent must get all the $m - n$ top objects and at most one less preferred one, whilst every other agent may receive at most one object. So, if agent 1 is the lucky one and p is a necessarily Pareto optimal assignment then $p(1) = \{o_1, o_2, \ldots, o_{m-n}\}$ must hold, with the possibility of agent 1 getting a further extended set of the top objects, whilst each of the other agents may hold at most one object from the rest of the objects (including agent 1).*

5. Conclusions

We have studied, from a computational point of view, Pareto optimality in resource allocation under additive utilities and ordinal preferences. Many of our positive algorithmic results come with characterizations of Pareto optimality

that improve our understanding of the concept and may be of independent interest.

Future work includes identifying restrictions on the preferences under which Pareto optimal and individually rational reallocation can be done in a computationally efficient manner.

Acknowledgements

References

[1] Abraham, D. J., Cechlárová, K., Manlove, D., Mehlhorn, K., 2005. Pareto optimality in house allocation problems. In: Proceedings of the 16th International Symposium on Algorithms and Computation (ISAAC). Vol. 3341. pp. 1163–1175.

[2] Andersson, T., Cseh, A., Ehlers, L., Erlanson, A., 2018. Organizing time banks: Lessons from matching markets. Working Papers 2018:19, Lund University, Department of Economics.

[3] Aziz, H., Biro, P., Lang, J., Lesca, J., Monnot., J., 2016. Optimal reallocation under additive and ordinal preferences. In: Proceedings of the 15th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS). pp. 402–410.

[4] Aziz, H., Brandl, F., Brandt, F., 2015. Universal Pareto dominance and welfare for plausible utility functions. Journal of Mathematical Economics 60, 123–133.

[5] Aziz, H., Brandt, F., Harrenstein, P., 2013. Pareto optimality in coalition formation. Games and Economic Behavior 82, 562–581.

[6] Aziz, H., de Haan, R., Rastegari, B., 2017. Pareto optimal allocation under uncertain preferences. In: Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI). pp. 1472–1474.

[7] Aziz, H., de Keijzer, B., 2012. Housing markets with indifferences: a tale of two mechanisms. In: Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI). pp. 1249–1255.

[8] Aziz, H., Lang, J., Monnot, J., 2016. Computing Pareto Optimal Committees. In: Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI). pp. 60–66.

[9] Aziz, H., Mackenzie, S., Xia, L., Ye, C., 2015. Ex post efficiency of random assignments. In: Proceedings of the 14th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS). pp. 1639–1640.

[10] Barberà, S., Bossert, W., Pattanaik, P. K., 2004. Ranking sets of objects. Springer.

[11] Bogomolnaia, A., Moulin, H., 2001. A new solution to the random assignment problem. Journal of Economic Theory 100 (2), 295–328.

[12] Bouveret, S., Endriss, U., Lang, J., 2010. Fair division under ordinal preferences: Computing envy-free allocations of indivisible goods. In: Proceedings of the 19th European Conference on Artificial Intelligence (ECAI). pp. 387–392.

[13] Bouveret, S., Lang, J., 2014. Manipulating picking sequences. In: Proceedings of the 21st European Conference on Artificial Intelligence (ECAI). pp. 141–146.

[14] Brams, S. J., Edelman, P. H., Fishburn, P. C., 09 2003. Fair Division Of Indivisible Items. Theory and Decision 55 (2), 147–180.

[15] Cechlárová, K., Eirinakis, P., Fleiner, T., Magos, D., Manlove, D., Mourtos, I., Ocelákova, E., Rastegari, B., 2016. Pareto optimal matchings in many-to-many markets with ties. Theory of Computing Systems 59 (4), 700–721.

[16] Cechlárová, K., Eirinakis, P., Fleiner, T., Magos, D., Manlove, D. F., Mourtos, I., Oceľákova, E., Rastegari, B., 2015. Pareto optimal matchings in many-to-many markets with ties. In: Proceedings of the 8th International Symposium on Algorithmic Game Theory (SAGT). pp. 27–39.

[17] Damamme, A., Beynier, A., Chevaleyre, Y., Maudet, N., 2015. The power of swap deals in distributed resource allocation. In: Proceedings of the 14th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS). pp. 625–633.

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

[19] Demko, S., Hill, T. P., 1988. Equitable distribution of indivisible objects. Mathematical Social Sciences 16, 145–158.

[20] Dur, U. M., Unver, M. U., Nov. 2012. Two-Sided Matching via Balanced Exchange. Boston College Working Papers in Economics 815, Boston College Department of Economics.
URL https://ideas.repec.org/p/boc/bocoec/815.html

[21] Endriss, U., 2013. Reduction of economic inequality in combinatorial domains. In: Proceedings of the 12th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS). pp. 175–182.

[22] Erdil, A., Ergin, H., June 2015. Two-sided matching with indifferences.

[23] Fujita, E., Lesca, J., Sonoda, A., Todo, T., Yokoo, M., 2015. A complexity approach for core-selecting exchange with multiple indivisible goods under lexicographic preferences. In: Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI). pp. 907–913.

[24] Keeney, R. L., Raiffa, H., 1976. Decisions with Multiple Objectives: Preferences and Value Tradeoffs. Wiley, New York.

[25] Konishi, H., Quint, T., Wako, J., 2001. On the Shapley-Scarf economy: the case of multiple types of indivisible goods. Journal of Mathematical Economics 35 (1), 1–15.

[26] Kratochvíl, J., Savický, P., Tuza, Z., 1993. One more occurrence of variables makes satisfiability jump from trivial to np-complete. SIAM Journal on Computing 22, 203–210.

[27] Manjunath, V., Westkamp, A., 2018. Strategy-proof exchange under trichotomous preferences.

[28] Manlove, D., 2013. Algorithmics of Matching Under Preferences. World Scientific Publishing Company.

[29] Sonoda, A., Todo, T., Sun, H., Yokoo, M., 2014. Two case studies for trading multiple indivisible goods with indifferences. In: Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI). pp. 791–797.

[30] Todo, T., Sun, H., Yokoo, M., 2014. Strategyproof exchange with multiple private endowments. In: Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI). pp. 805–811.

[31] Yu, W., Hoogeveen, H., Lenstra, J. K., 2004. Minimizing makespan in a two-machine flow shop with delays and unit-time operations is NP-hard. Journal of Scheduling 7, 333–348.