Optimal Reallocation under Additive and Ordinal Preferences

Haris Aziz, Péter Biró, Jérôme Lang, Julien Lesca, and Jérôme Monnot

Abstract

Reallocating resources to get mutually beneficial outcomes is a fundamental problem in various multi-agent settings. 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 single objects, and that their preferences are additively separable. In this setting, we present characterizations and polynomial-time algorithms for possible and necessary 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 [1, 5, 28, 30]. Pareto improvements are desirable for two fundamental reasons: they result in strictly more welfare for any reasonable notion of welfare (such as utilitarian or leximin). 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 they also have additive preferences for the objects. In the absence of endowments, achieving a Pareto optimal assignment is easy: simply assign every object to the agent who values it the most. 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 Pareto dominance gives an assignment that respects individual rationality and is a Pareto improvement. In fact, if testing Pareto optimality is NP-hard, then finding an individually rational and Pareto optimal assignment is NP-hard 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.

Contributions 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 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 (i.e., lexicographic preferences) and we show that Pareto optimality can be tested in linear time. For dichotomous preferences 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*. 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 [2, 10, 11, 13, 19, 18, 23, 26, 27, 31, 32]. Although computing a utilitarian welfare maximizing assignment is easy, the problem of maximizing egalitarian welfare is NP-hard.

Algorithmic aspects of Pareto optimality have received attention in discrete allocation of indivisible goods, randomized allocation of indivisible goods, two-sided matching, and coalition formation under ordinal preferences [1, 5, 8, 21, 28]. Since we are interested in Pareto improvements, our paper is also related to housing markets with endowments and ordinal preferences [4, 22, 25, 33, 34]. Recently, Damamme et al. [17] examined restricted Pareto optimality under ordinal preferences.

de Keijzer et al. [18] studies 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 has two objects.

Cechlárová et al. [16] 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. [14] consider necessary Pareto optimality as Pareto optimality for all completions of the responsive set extension,[1] and present 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 [20].

2 Preliminaries

We consider the setting in which we have $N = \{1, \dots, n\}$ a set of agents, $O = \{o_1, \dots, o_m\}$ a set of objects, and the preference profile $\succsim = (\succsim_1, \dots, \succsim_n)$ specifies for each agent i 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 $E_i^1, \dots, E_i^{k_i}$ the k_i equivalence classes of an agent $i \in N$. Those classes partition O into k_i sets of objects such that agent i is indifferent between two objects belonging to the same class, and she strictly prefers an object of E_i^k to an object of E_i^l whenever $k < l$.

Each agent may additionally express a cardinal utility function u_i consistent with $\succsim_i$: $u_i(o) > u_i(o')$ iff $o \succ_i o'$ and $u_i(o) = u_i(o')$ iff $o \sim_i o'$. We will assume that each object is positively valued, i.e, $u_i(o) > 0$ for all $i \in N$ and $o \in O$. The set of all utility functions consistent with $\succsim_i$ is denoted by $\mathscr{U}(\succsim_i)$. We will denote by $\mathscr{U}(\succsim)$ the set of all utility profiles $u = (u_1, \dots, u_n)$ such that $u_i \in \mathscr{U}(\succsim_i)$ for each $i \in N$. When we consider agents' valuations according to their cardinal utilities, then we will assume additivity, that is $u_i(O') = \sum_{o \in O'} u_i(o)$ for each $i \in N$ and $O' \subseteq O$. An assignment $p = (p(1), \dots, p(n))$ 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.

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 any agent i. An assignment $p \in \mathscr{X}$ is said to be *Pareto*

[1] Brams et al. [15] used the term *Pareto ensuring* for Pareto optimality for all completions of the responsive set extension.

dominated by another $q \in \mathscr{X}$ if (i) for any 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* iff it is not Pareto dominated by another assignment. Finally, whenever cardinal utilities are considered, the *social welfare* of an assignment p is defined as $SW(p) = \sum_{i \in N} u_i(p(i))$.

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

	o_1	o_2	o_3	o_4	o_5
1	16	8	4	2	1
2	9	3	1	1	3
3	6	1	2	6	2

Since $u_1(o_1) > u_1(o_2)$, we can say that $o_1 \succ_1 o_2$. An example of an assignment is $p = (o_2o_4|o_1|o_3o_5)$ in which $p(1) = \{o_2, o_4\}$, $p(2) = \{o_1\}$, and $p(3) = \{o_3, o_5\}$.

3 Additive utilities

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

3.1 Complexity of testing Pareto optimality

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. □

A Pareto optimal assignment can be computed trivially by giving each object to the agent who values it the most. Bouveret and Lang [12] proved that a problem concerning coalitional manipulation in sequential allocation is NP-complete (Proposition 6). The result can be reinterpreted as follows.

Theorem 1. *Testing Pareto optimality of a given assignment is weakly coNP-complete for $n = 2$ and identical preferences.*

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

One may additionally require the *balanced* property, i.e., each agent gets as many objects as she initially owned. Both the theorem above and the corollary above can be extended to that case easily. If there are an unbounded number of agents, then testing Pareto optimality of a given assignment is strongly coNP-complete [18]. Next, we show that the problem remains strongly coNP-complete even if each agent receives exactly 2 objects.

Theorem 2. *Testing Pareto optimality of a given assignment is strongly coNP-complete for an unbounded number of agents even if each agent receives exactly 2 objects.*

We relegate the proof of Theorem 2 to the Appendix. Note that Theorem 2 is the best possible NP-hardness result that we can obtain according to the number of objects received by each agent because if initially each agent has exactly one object in assignment p, then our problem can be solved in linear time.

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 weights

Lemma 2. *If there is a constant number of agents, 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).

$L \leftarrow \{(0^n)\}$;
for $j = 1$ to m **do**
  $L' \leftarrow \{l + (0^{i-1}, u_i(o_j), 0^{n-i}) \mid i \in N; l \in L\}$
  $L \leftarrow L'$
end for
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 algorithms runs in pseudopolynomial 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 $k+1$ iff 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 3. *If there is a constant number of agents, 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 $O(|L|^2)$, and we recall that L is polynomially large. The assignments that correspond to the remaining vectors are Pareto optimal.[2] □

[2]Note that it is generally not the case that we get *all* Pareto optimal assignments: if there are several assignments corresponding to the same utility vector, then we'll obtain only one.

3.2.2 Lexicographic Utilities

We say that utilities are *lexicographic* if for each agent $i \in N$, $u_i(o) > \sum_{o' \prec_i o} u_i(o')$. By $q(i) \succsim_i p(i)$, we will mean $u_i(q(i)) \geq u_i(p(i))$.

In order to test the Pareto optimality of an assignment p, we construct a directed graph $G(p) = (V(p), E(p))$. The set of vertices $V(p)$ contains one vertex per object belonging to O. Furthermore, for any vertex of $V(p)$ associated to an 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 the good o in p. For example, Figure 1 illustrates such a graph for the assignment p provided by Example 2. In Figure 1, dotted edges represent indifferences (when $o' \sim o$) and plain 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 4. *An assignment p is not Pareto optimal wrt lexicographic utilities iff 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 which contains at least one edge corresponding to a strict preference. Then, the exchange of objects along the cycle by agents owning the objects corresponds to a Pareto improvement.

Assume now that p is not Pareto optimal and let q_1 be an assignment which Pareto dominates p. For at least one agent i, $q_1(i) \succ_i p(i)$. So 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 and so on. Since O is finite, there must exist k and k' such that the sequence $o_k \to o_{k+1} \to \ldots \to o_{k'}$ forms a cycle, i.e., $o_k = o_{k'}$. If $\nexists l \in [k, k'-1]$ such that $o_{l+1} \succ_{i_l} o_l$ then we consider the assignment q_2 derived from q_1 by reassigning any object o_{l+1}, with $l \in [k, k'-1]$, to agent i_l. It is obvious that this assignment q_2 is at least as good as q_1 for all the agents. So q_2 Pareto dominates p. By following the same reasoning as above, we can state that there exists a sequence of objects $o_k \to o_{k+1} \to \ldots \to o_{k'}$ such that $o_k = o_{k'}$ and for any $l \in [k, k'-1]$, 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 $\nexists l \in [k, k'-1]$ such that $o_{l+1} \succ_{i_l} o_l$ then we consider the assignment q_3 derived from q_2 by reassigning any object o_{l+1}, with $l \in [k, k'-1]$, to agent i_l... 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 such that $\exists l \in [k, k'-1]$ such that $o_{l+1} \succ_{i_l} o_l$ for the cycle $o_k \to o_{k+1} \to \ldots \to o_{k'}$ founded in q_t. Indeed otherwise after a finite number of steps t we should have $q_t(i) = p(i)$ for all $i \in N$, which leads to a contradiction with the assumption that q_t Pareto dominates p. So there exists a cycle $o_k \to \ldots \to o_{k'}$ in $G(p)$ with at least one edge corresponding to a strict preference. $\square$

It is clear that the 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 any 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.

Example 2. Let $n = 3$, $m = 5$, and the following ordinal information about preferences corresponding to the lexicographic utilities in Example 1 (as a consequence of Theorem 4, ordinal preferences are enough information to check Pareto optimality).

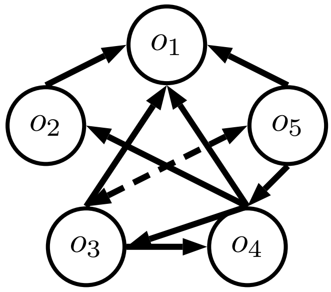

Figure 1: Graph $G(p)$ for assignment p in Example 2.

$$\begin{aligned}1 &: o_1 \succ o_2 \succ o_3 \succ o_4 \succ o_5\\ 2 &: o_1 \succ o_2 \sim o_5 \succ o_3 \sim o_4\\ 3 &: o_1 \sim o_4 \succ o_3 \sim o_5 \succ o_2\end{aligned}$$

Let $p = (o_2o_4|o_1|o_3o_5)$ be the initial assignment. The construction of Theorem 4 gives us that it is Pareto dominated by $(o_2o_3|o_1|o_4o_5)$, hence it is not Pareto optimal.

3.2.3 Two utility values

In this section we assume the agents use at most two utility values for the objects. We say that the collection of utility functions $(u_1, \ldots, u_n)$ is *bivalued* if there exist two numbers $\alpha > \beta > 0$ such that for every agent i and every object o, $u_i(o) \in \{\alpha, \beta\}$. (The result would still hold if each agent i had a different pair of values (α_i, β_i), provided that $\frac{\alpha_i}{\beta_i} = \frac{\alpha_j}{\beta_j}$ for all i, j.) This means that for every 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 $E_i^2 = \emptyset$). 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, where all objects are assigned, 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 $SW(q) = |\bigcup_{i\in N} q^+(i)|\alpha + |\bigcup_{i\in N} q^-(i)|\beta \leq |\bigcup_{i\in N} q^+(i)|(\alpha - \beta) + |O|\beta \leq |\bigcup_{i\in N} p^+(i)|(\alpha - \beta) + |O|\beta = \alpha|\bigcup_{i\in N} p^+(i)| + \beta|\bigcup_{i\in N} p^-(i)|$. So $SW(p) \geq SW(q)$, which contradicts the assumption that q Pareto dominates p. □

Lemma 4. *If an assignment p, where all objects are assigned, is not Pareto optimal then there exists an assignment q that Pareto dominates p with the following two 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 q_* which Pareto dominates p. Let q_* chosen to be as close to p as possible, namely such that $|\bigcup_{i\in N} q_*(i) \setminus p(i)|$ is minimal.

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

Lemma 3 implies that $|\bigcup_{i\in N} q_*^+(i)| > |\bigcup_{i\in N} p^+(i)|$, so there exists an object $o_1 \in (\bigcup_{i\in N} q_o^+(i)) \setminus (\bigcup_{i\in N} p^+(i))$. Let us suppose 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. 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), and let $o_2 \in q_*(i_3)$. 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 the 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. So 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. □

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

Theorem 5. *An assignment p, where all objects are assigned, is Pareto dominated iff 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) holds. Let j be as described as in (ii). For any $i \in N\setminus\{j\}$, let $A_i \subseteq q^+(i)$ such that $|A_i| = |p^+(i)|$. Let $A_j \subseteq q^+(j)$ such that $|A_j| = |p^+(j)| + 1$. 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 $|\bar{A}| = |O| - |\bigcup_{i\in N} A_i| = \sum_{i\in N} |p(i)| - \sum_{i\in N} |p^+(i)| - 1 = \sum_{i\in N} |p^-(i)| - 1$.

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

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

Proof. If at least one object is not assigned then a trivial Pareto improvement would be to assign this object to an agent. So 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 which models the search for a Pareto improvement for agent i as a flow problem. The set of vertices V_i contains one vertex per agent and per object, plus a source s and a sink t. To ease the notation, we do not discriminate between the vertices and the agents or objects that they are representing, therefore, we note $V_i = N \cup O \cup \{s, t\}$. The set of edges E_i and their capacities are constructed as follow:

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

- For any $l \in N \setminus \{i\}$ there is an edge (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 $\sum_{l \in N} |p^+(l)| + 1$ iff there exists an assignment such that any agent $l \in N \setminus \{i\}$ receives at least $|p(l) \cap E_l^1|$ top objects and agent i receives $|p^+(l)| + 1$ top objects. So by Theorem 5, there exists a Pareto improvement of p iff 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. In each Pareto improvement the number of top objects increases by at least one so there can be at most m Pareto improvements. $\square$

Note that we can find a Pareto optimal Pareto improvement in polynomial time as well: in each Pareto improvement the number of top objects increases by at least one so there can be at most m Pareto improvements.

Example 3. Let $n = 3$, $m = 6$, $E_1^1 = \{o_1, o_2, o_3\}$, $E_2^1 = \{o_2\}$, $E_3^1 = \{o_1, o_3, o_5, o_6\}$, and $p = (o_1 o_4 | o_2 o_5 | o_3 o_6)$. G_1 is depicted in Figure 2. The flow of value 5 (boldface) gives the assignment $(o_1 o_3 | o_2 o_4 | o_5 o_6)$, which Pareto-dominates p.

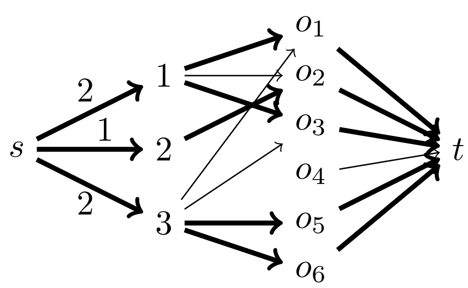

Figure 2: Flow network G_1 in Example 3.

4 Ordinal preferences

In this section, we consider the setting in which the agents have additive cardinal utilities but only their ordinal preferences over the 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 don't know them precisely. 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 $\succsim$ 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 $\succsim$ 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. We focus on the problems of testing possible and necessary Pareto optimality.

In order to characterize possible Pareto optimality, we first define *stochastic dominance (SD)* which extends ordinal preferences over objects to preferences over sets of objects (and even over fractional allocations in which agents can get fractions of items). We say that an allocation $q(i)$ *stochastically dominates* an allocation $p(i)$, denoted by $q(i) \succsim_i^{SD} p(i)$, iff $|q(i) \cap \bigcup_{j=1}^{k} E_i^j| \geq |p(i) \cap \bigcup_{j=1}^{k} E_i^j|$ for all $k \in \{1, \ldots, k_i\}$. In the case of fractional allocations, $q(i) \cap \bigcup_{j=1}^{k} E_i^j$ denotes the units of items given to i for items in $\bigcup_{j=1}^{k} E_i^j$.

The SD relation is equivalent to the *responsive set extension* [9], which also 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: $q(i) \succsim_i^{RS} p(i)$ iff 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 *SD-efficient* if it is Pareto optimal with respect to the SD relation of the agents, and *RS-efficient* if it is Pareto optimal with respect to the RS set extension relation of the agents. Under ordinal preferences, an agent i prefers one allocation $q(i)$ over another $p(i)$ with respect to responsive set extension iff she prefers it with respect to stochastic dominance [14, 7]. Thus, a (discrete) assignment is RS-efficient iff it is SD-efficient. We say that q strictly RS-dominates p if q Pareto dominates p with respect to RS.

Theorem 7. *An assignment is possibly Pareto optimal iff it is SD-efficient iff it is RS-efficient iff there exists no cycle in $G(p)$ which contains at least one edge corresponding to a strict preference.*

Proof. By the ordinal welfare theorem, a fractional assignment is possibly Pareto optimal iff it is SD-efficient (among the set of fractional assignments) [3, 6, 29]. Furthermore, a discrete assignment p that is SD-efficient among all discrete assignments is also SD-efficient among all fractional assignments because SD-efficiency of p depends on the non-existence of a cycle with a strict edge in the underlying graph $G(p)$ [3, 24]. Hence, we obtain the equivalences. □

Since the characterization in Theorem 4 also applies to RS-efficiency and 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.

We point 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 the assignment p in which agent 1 gets $\{o_1, o_4\}$ and 2 gets $\{o_2, o_3\}$ 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. [14].

Theorem 8. *An assignment is necessarily Pareto optimal iff 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 an assignment p is not Pareto optimal under some completion of the responsive set extension. Then there exists another assignment q in which 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 for some $i \in N$, $q(i) \succ_i^{RS} p(i)$ or $p(i) \not\succ_i^{RS} q(i)$ and $q(i) \not\succ_i^{RS} p(i)$. 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, 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, $u_i(q(i)) > u_i(p(i))$. □

For characterizing 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 iff*

(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 (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 $\{i_1, i_2, \dots i_k\}$ with possible repetitions and a set of distinct objects $\{o_1, o_2, \dots, o_m\}$ 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 $o_2 \succsim_{i_3} o_3$;
- $\dots$
- $o_m \in q(i_1) \setminus p(i_1)$, $o_1 \in p(i_1) \setminus q(i_1)$, 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 $q'(i) = p(i) \cup \{o_{k-1} : i_k = i, k = 1, \dots, m\} \setminus \{o_k : i_k = i, k = 1, \dots, m\}$. This would contradict our assumption that p is possibly Pareto optimal.

We first define three types of agents, and a **one-to-one mapping** π from a subset of O to itself such that 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 π can 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 $i_1, i_2, \dots$ with corresponding objects $o_1, o_2, \dots$ 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 any 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 belong to, say, $i_s = i_t \in Z$ for indices $1 \le s < t$. 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. □

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 $o_1 \succ_2 o_2 \succsim_2 o_3$. It also shows that although an assignment may not be necessarily Pareto optimal there may not be any 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.

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 other important subdomains in which Pareto optimal and individually rational reallocation can be done in a computationally efficient manner.

Acknowledgements

Haris Aziz is funded by the Australian Government through the Department of Communications and the Australian Research Council through the ICT Centre of Excellence Program. Péter Biró acknowledges support from the Hungarian Academy of Sciences under its Momentum Programme (LD-004/2010), the Hungarian Scientific Research Fund, OTKA, Grant No. K108673, and the János Bolyai Research Scholarship of the Hungarian Academy of Sciences. Jérôme Lang, Julien Lesca and Jérôme Monnot acknowledge support from the ANR project CoCoRICo-CoDec. Part of the work was conducted when Péter Biró visited Paris Dauphine and when Julien Lesca visited Corvinus University of Budapest sponsored by COST Action IC1205 on Computational Social Choice.

References

1. D. J. Abraham, K. Cechlárová, D. Manlove, and K. Mehlhorn. Pareto optimality in house allocation problems. In *Proceedings of the 16th International Symposium on Algorithms and Computation (ISAAC)*, volume 3341, pages 1163–1175, 2005.
2. A. Asadpour and A. Saberi. An approximation algorithm for max-min fair allocation of indivisible goods. *SIAM Journal on Computing*, 39(7):2970–2989, 2010.
3. S. Athanassoglou. Efficiency under a combination of ordinal and cardinal information on preferences. *Journal of Mathematical Economics*, 47(2):180–185, 2011.

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

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

6. H. Aziz, F. Brandl, and F. Brandt. Universal dominance and welfare for plausible utility functions. *Journal of Mathematical Economics*, 2015.

7. H. Aziz, S. Gaspers, S. Mackenzie, and T. Walsh. Fair assignment of indivisible objects under ordinal preferences. *Artificial Intelligence*, 227:71–92, 2015.

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

9. S. Barberà, W. Bossert, and P. K. Pattanaik. *Ranking sets of objects*. Springer, 2004.

10. I. Bezáková and V. Dani. Allocating indivisible goods. *SIGecom Exchanges*, 5(3):11–18, 2005.

11. S. Bouveret and J. Lang. Efficiency and envy-freeness in fair division of indivisible goods: logical representation and complexity. *Journal of Artificial Intelligence Research*, 32(1): 525–564, 2008.

12. S. Bouveret and J. Lang. Manipulating picking sequences. In *Proceedings of the 21st European Conference on Artificial Intelligence (ECAI)*, pages 141–146, 2014.

13. S. Bouveret and M. Lemaître. Characterizing conflicts in fair division of indivisible goods using a scale of criteria. In *Proceedings of the 13th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 1321–1328, 2014.

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

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

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

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

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

19. S. Demko and T. P. Hill. Equitable distribution of indivisible obejcts. *Mathematical Social Sciences*, 16:145–158, 1988.

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

21. A. Erdil and H. Ergin. Two-sided matching with indifferences. June 2015.

22. E. Fujita, J. Lesca, A. Sonoda, T. Todo, and M. Yokoo. 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)*, pages 907–913, 2015.

23. D. Golovin. Max-min fair allocation of indivisible goods. Technical Report CMU-CS-05-144, Carnegie Mellon University, 2005.

24. A.-K. Katta and J. Sethuraman. A solution to the random assignment problem on the full preference domain. *Journal of Economic Theory*, 131(1):231–250, 2006.

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

26. J. Lesca and P. Perny. LP solvable models for multiagent fair allocation problems. In *Proceedings of the 19th European Conference on Artificial Intelligence (ECAI)*, pages 393–398, 2010.

27. R. J. Lipton, E. Markakis, E. Mossel, and A. Saberi. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce (ACM-EC)*, pages 125–131, 2004.

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

29. A. McLennan. Ordinal efficiency and the polyhedral separating hyperplane theorem. *Journal of Economic Theory*, 105(2):435–449, 2002.

30. H. Moulin. *Fair Division and Collective Welfare.* The MIT Press, 2003.

31. T. T. Nguyen, M. Roos, and J. Rothe. A survey of approximability and inapproximability results for social welfare optimization in multiagent resource allocation. *Annals of Mathematics and Artificial Intelligence*, pages 1–26, 2013.

32. A. D. Procaccia and J. Wang. Fair enough: Guaranteeing approximate maximin shares. In *Proceedings of the 15th ACM Conference on Economics and Computation (ACM-EC)*, pages 675–692, 2014.

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

34. T. Todo, H. Sun, and M. Yokoo. Strategyproof exchange with multiple private endowments. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*, pages 805–811, 2014.

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

Haris Aziz
Data61 and UNSW
Sydney, Australia
Email: haris.aziz@nicta.com.au

Péter Biró
Hungarian Academy of Sciences
Budapest, Hungary
Email: biro.peter@krtk.mta.hu

Jérôme Lang
LAMSADE, Université Paris-Dauphine
Paris, France
Email: lang@lamsade.dauphine.fr

Julien Lesca
LAMSADE, Université Paris-Dauphine
Paris, France
Email: julien.lesca@dauphine.fr

Jérôme Monnot
LAMSADE, Université Paris-Dauphine
Paris, France
Email: jerome.monnot@dauphine.fr

Appendix

Below we provide the proof of Theorem 2.

Proof. The reduction is done from 2-NUMERICAL MATCHING WITH TARGET SUMS (2NMTS in short). The input of 2NMTS is a sequence $a_1, \dots, a_k$ of k positive integers such that $\sum_{i=1}^{k} a_i = k(k+1)$ and $1 \le a_i \le 2k-1$ for $i = 1, \dots, k$, and $a_1 \le a_2 \dots \le a_k$. We want to decide if there are two permutations π and θ of the integers $\{1, \dots, k\}$ such that $\pi(i) + \theta(i) = a_i$ for $i = 1, \dots, k$. 2NMTS is known to be strongly NP-complete [35].

The reduction from an instance of 2NMTS is as follows. There are $3k+1$ agents $N = L \cup C \cup R \cup \{d\}$ where $L = \{\ell_1, \dots, \ell_k\}$, $R = \{r_1, \dots, r_k\}$ and $C = \{c_1, \dots, c_k\}$ and $6k+2$ objects $O = F \cup G \cup H \cup \{o\}$ where $F = \{f_i^L, f_i^R : i = 1, \dots, k\}$, $G = \{g_i^L, g_i^R : i = 1, \dots, k\} \cup \{g^C\}$, $H = \{h_i^{CL}, h_i^{CR} : i = 1, \dots, k\}$. Let ε be a positive value strictly lower than $1/2$. The following table summarizes the non-zero utilities provided by the different objects, 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 with non-zero utility 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 $a_j \ge i+1$	i
f_i^R	r_i	1	c_j with $a_j \ge i+1$	$3k+i$
g_i^R	r_i	3	r_{i+1} if $i<k$ d if $i=k$	$3+\varepsilon$ $3+\varepsilon$
g_i^L	ℓ_i	3	ℓ_{i-1} if $i>1$ r_1 if $i=1$	$3-\varepsilon$ $3+\varepsilon$
g^C	d	3	ℓ_k	$3-\varepsilon$
o	d	1		

The initial assignment provides the following utilities to the agents: $u_{c_i}(\{h_i^{CL}, h_i^{CR}\}) = 3k + a_i$, $u_{\ell_i}(\{f_i^L, g_i^L\}) = 4$ and $u_{r_i}(\{f_i^R, g_i^R\}) = 4$ for $i = 1 \dots k$, and $u_d(\{g^C, o\}) = 4$.

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

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

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

because $\pi(i) \ge 1$ and $\theta(i) \ge 1$. 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. Using (1), the utility of agent c_i is $3k + \pi(i) + \theta(i) = 3k + a_i$.
- $\{o, g_k^R\}$ is assigned to d with utility $4 + \varepsilon$.

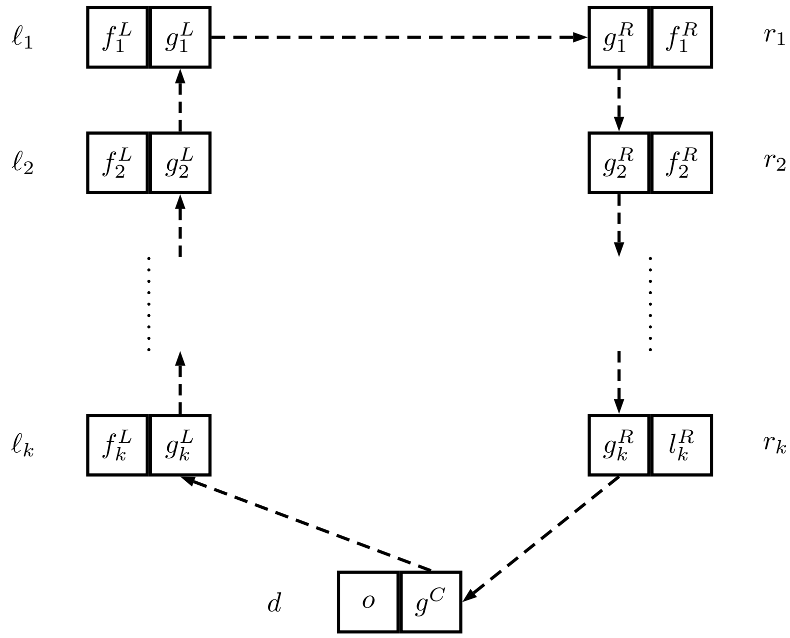

Figure 3: Initial assignment for agents of $L \cup R \cup \{d\}$.

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

Assume now that $\{a_i : i = 1 \dots k\}$ is a no-instance of 2NMTS. By contradiction, assume that there exists a Pareto improvement p of the initial assignment. Note first that any agent should receive in p at least two objects. Indeed there is no object which provides a utility greater than $3+\varepsilon$ to any agent of $L \cup R \cup \{d\}$, and any of those agents receives a utility of 4 in the initial assignment. Furthermore, any good f_i^R provides a utility of at most $3k + i$ to an agent c_j, which is strictly lower than her utility $3k + a_j$ 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. Those objects are the only ones which can provide a utility of at least $3 - \varepsilon$ to the agents of $L \cup R \cup \{d\}$. All other objects provide a utility of at most $1 + \varepsilon$ to the agents in $L \cup R \cup \{d\}$. So, to achieve a utility of at least 4 for all those agents in $L \cup R \cup \{d\}$, each of them should receive exactly one good from G (with non-zero utility for it) because $|L \cup R \cup \{d\}| = |G| = 2k+1$. Figure 3 illustrates the initial assignment for the agents of $L \cup R \cup \{d\}$. In this figure, a dashed arrow from an object of G means that this object can be reassigned to the agent pointed at with a non zero utility. Figure 3 illustrates the fact that the goods of G could be allocated in only two different manners in p to be a Pareto improvement of the initial endowment: either every good of G is assigned to the same agent as in the initial assignment, or every good of G is assigned to the agent pointed at by the corresponding arrow in Figure 3.

First, we consider the case where all goods 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 the object f_i^R to complete her bundle of two objects. This implies that those 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 every agent c_i should receive the object h_i^{CR} with utility $3k$. Furthermore no agent c_i can receive an object f_j^R to complete her bundle of two objects because this object would provide her a utility of at most $a_i - 1$. So, every agent c_i should receive the object h_i^{CL}. From this, we conclude that p should be exactly the same assignment as the initial assignment, which contradicts the assumption that p Pareto-dominates this initial assignment.

From the previous paragraphs, we know that any good of G should be assigned in p to the agent pointed at by the corresponding dotted arrow in Figure 3. To achieve a utility of at least 4, any agent ℓ_i should receive the good h_i^{CL} to complete her bundle of two objects. If an agent c_i did not receive at least one good 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. So, every agent c_i should receive exactly one good f_j^R such that $a_i \geq j+1$. Therefore no good f_i^R can be assigned to agent r_i. So, to achieve a utility of at least 4, any agent r_i should receive the good h_i^{CR} to complete her bundle of two objects. Then the good o should be assigned to agent d to complete her bundle of two goods. Finally it remains to assign to every agent c_i a good f_j^L such that $a_i \geq j+1$.

Now let us focus on the pair of goods assigned to agent c_i in p with $i = 1 \ldots k$. Note that those two objects belong to F. We know that the total amount of utility provided by the goods of F to the agents of C should be exactly equal to $3k^2 + k(k+1)$. Furthermore any agent c_i should receive a share of at least $3k + a_i$ of this total amount of utility. Since $\sum_{i=1}^{k}(3k + a_i) = 3k^2 + k(k+1)$, any 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 $\{1, \ldots, k\}$ such that for any $i = 1 \ldots k$, the objects $f_{\pi(i)}^L$ and $f_{\theta(i)}^R$ are assigned in p to agent c_i. Those two permutations are such that for any $i = 1 \ldots k$, $\pi(i) + \theta(i) = a_i$. This leads to a contradiction with the assumption that $\{a_i : i = 1 \ldots k\}$ is a no-instance. $\square$