

Fairness Towards Groups of Agents in the Allocation of Indivisible Items

Nawal Benabbou¹, Mithun Chakraborty², Edith Elkind³, Yair Zick²

¹Sorbonne Université, CNRS, Laboratoire d’Informatique de Paris 6, LIP6 F-75005 Paris, France

²Department of Computer Science, National University of Singapore, Singapore

³Department of Computer Science, University of Oxford, United Kingdom

nawal.benabbou@lip6.fr, {mithun, zick}@comp.nus.edu.sg, elkind@cs.ox.ac.uk

Abstract

In this paper, we study the problem of matching a set of items to a set of agents partitioned into types so as to balance fairness towards the types against overall utility/efficiency. We extend multiple desirable properties of indivisible goods allocation to our model and investigate the possibility and hardness of achieving combinations of these properties, e.g. we prove that maximizing utilitarian social welfare under constraints of typewise envy-freeness up to one item (TEF1) is computationally intractable. We also define a new concept of waste for this setting, show experimentally that augmenting an existing algorithm with a marginal utility maximization heuristic can produce a TEF1 solution with reduced waste, and also provide a polynomial-time algorithm for computing a non-wasteful TEF1 allocation for binary agent-item utilities.

1 Introduction

Consider an academic department looking to assign incoming graduate students to advisers. It has computed a score for each potential advisor-advisee pair and cares not just about the overall score of the matching but also about a fair distribution of scores across the research clusters into which the faculty is divided. A similar trade-off between efficiency/welfare and fairness might be desirable in other planning/allocation scenarios such as public housing allocation; e.g. the Singapore Housing and Development Board imposes ethnicity quotas on the matching of tenant households to flats in its housing blocks to promote diversity [Benabbou *et al.*, 2018] but an alternative approach could be to aim for a fair distribution of overall utility across ethnic groups.

We can model these problems as a variant of fair allocation of indivisible goods (see e.g. [Bouveret *et al.*, 2016]): We have an underlying weighted bipartite matching problem, the nodes on the two sides corresponding to *items* and *agents*; however, the parties we are trying to be fair towards are not individual agents but subsets forming a partition over agents – we call these subsets *types*. An important aspect of this problem is that all agents in a type do not derive utility from

all items in the *bundle* allocated to that type (unlike the public goods scenario [Fain *et al.*, 2016]), each agent being assigned at most one item. Hence, the parties (types) under consideration end up violating the additive bundle-valuation assumption present in much of recent work. There are some approaches (e.g. [Lipton *et al.*, 2004]) that achieve good fairness guarantees under mild assumptions on the (non-negative, non-decreasing) valuation function but, for that very reason, can result in allocations that are wasteful/inefficient in some way given the structure of our problem, if used naively. These considerations necessitate novel solution concepts and techniques for our setting.

1.1 Our Contributions

We describe a new model of typewise fair allocation and define our desirable properties in Section 2 – in particular, *non-wastefulness* and *typewise envy-freeness up to one item* (TEF1) as well as a marginal envy-based variant of the latter. In Sections 3 and 4, we explore the problems of determining Pareto optimal, TMEF1 allocations and TEF1 allocations that maximize overall sum of weights/utilities respectively. In Section 5, we show experimentally that the classic algorithm of Lipton *et al.* [2004] equipped with a simple heuristic can produce TEF1 allocations with significantly reduced waste. Section 6 details a polynomial-time algorithm for computing a non-wasteful TEF1 allocation for binary agent-item utilities. We conclude with directions for future research in Section 7.

1.2 Related Work

There is a rich body of work on approaches towards the fair allocation of indivisible goods [Bouveret *et al.*, 2016; Markakis, 2017]. A popular fairness concept is *envy-freeness* [Foley, 1967]. A complete envy-free (EF) allocation may not exist but a relaxation that always does is one that is *envy-free up to one item* (EF1) [Budish, 2011] where any envy towards an agent can be eliminated by removing an item from its bundle. The bounded-envy, polynomial-time algorithm due to Lipton *et al.* [2004] also produces an EF1 allocation for general valuation functions [Plaut and Roughgarden, 2018].

Extensions of envy-freeness to groups include *strict envy-freeness* [Zhou, 1992]; *coalition fairness* or *group envy-freeness* [Lahaie and Parkes, 2009] as well as *envy-freeness of an individual/a group towards a group* [Todo *et al.*, 2011], both under monetary transfers; *group fairness* as defined

by Conitzer *et al.* [2019]. Barman *et al.* [2018a] recently defined a *groupwise* extension to the *maximin share*-based fairness concept. A major difference of these contributions with our model is that they deal with a σ -algebra of subsets of agents rather than an exogenously defined partition over agents. Notable papers that define fairness with respect to pre-defined groups of multiple players include Manurangsi and Suksompong [2017], Suksompong [2018], Segal-Halevi and Suksompong [2018], and Elzayn *et al.* [2019] whose concepts of utility and/or fairness are significantly different from ours. Recent work [Fain *et al.*, 2016; Fain *et al.*, 2018] has also explored non-envy-based fairness criteria in *public goods* allocation under additive valuations.

We must also mention the literature on *statistical fairness* (also called *group fairness*) in the fundamentally different problem domain of classification in machine learning: the equalization of some statistical property of the classifier across groups of data instances based on sensitive/protected attributes ([Dwork *et al.*, 2012; Hardt *et al.*, 2016; Kearns *et al.*, 2018] and references therein); we, on the other hand, are interested in fairness notions in terms of subjective valuations of items from the economics/social choice literature.

2 Model And Definitions

Throughout the paper, $[r]$ will denote the set $\{1, 2, \dots, r\}$ for any positive integer r . Our model, an extension of the classic framework of matching on a weighted bipartite graph [Lovász and Plummer, 2009], has the following ingredients:

- (i) a set N of n vertices called *agents* partitioned into k types N_1, \dots, N_k ,
- (ii) a set M of m vertices called *items*,
- (iii) a weight/utility $u(i, j) \in \mathbb{R}_+$ for each agent-item edge $(i, j) \in N \times M$, such that for at least one $i \in N$ (resp. at least one $j \in M$), there is at least one $j \in M$ (resp. each $i \in N$) with $u(i, j) > 0$.

For any $T \subseteq N$ and any $S \subseteq M$, a (T, S) -*matching* is defined as a subset of the edges $T \times S$ such that every vertex in $T \cup S$ is incident on at most one of the edges or, equivalently, as a binary matrix $X = (x_{ij})_{i \in T, j \in S}$ such that for each agent i (resp. item j), there is at most one item j (resp. agent i) with $x_{ij} = 1$, i.e. each item is assigned to at most one agent and each agent is assigned at most one item. The *utilitarian social welfare* $\text{USW}(X)$ (or *total weight*) of a matching X is defined as the sum of the realized utilities of all agents under that matching: $\text{USW}(X) \triangleq \sum_{i \in T} \sum_{j \in S} x_{ij} u(i, j)$. An *optimal matching* is one that maximizes the corresponding USW .

We are interested in a (N, M) -matching that trades off some welfare/efficiency concept against some fairness criterion defined with respect to the agent types. More precisely, for every type $p \in [k]$, we are given a *type-value function* $v_p : 2^M \rightarrow \mathbb{R}_+$ which quantifies some concept of overall welfare derived by N_p from some *bundle* or subset of items $S \subseteq M$ in terms of the weights $u(i, j)$, $(i, j) \in N_p \times S$. In this paper, we will use the following specific type-value function for every type:

Definition 1 (Utilitarian type-value function). *The utilitarian type-value of any type $p \in [k]$ for any bundle $S \subseteq M$ is*

defined as the total weight of an optimal (N_p, S) -matching:

$$v_p(S) \triangleq \begin{cases} \max_{X \in \mathcal{X}(N_p, S)} \text{USW}(X), & \text{if } S \neq \emptyset; \\ 0, & \text{otherwise.} \end{cases}$$

where $\mathcal{X}(N_p, S)$ is the collection of all (N_p, S) -matchings.

We will define the *marginal utility* $\Delta_p(S; j)$ of an item $j \in M$ for a type $p \in [k]$ and a bundle $S \subseteq M$ as:

$$\Delta_p(S; j) \triangleq \begin{cases} v_p(S \cup \{j\}) - v_p(S), & \text{if } j \notin S; \\ v_p(S) - v_p(S \setminus \{j\}), & \text{otherwise.} \end{cases}$$

Given a type p and a bundle S , there can be multiple optimal (N_p, S) -matchings with the same type-value but possibly differing in other efficiency and/or fairness properties (see Example 1 at the end of this section) – with this in mind, we define an *allocation* in our setting as follows:

Definition 2 (Allocation). *An allocation \mathcal{A} is a collection of bundles $M_1^{\mathcal{A}}, \dots, M_k^{\mathcal{A}}$, such that $M_1^{\mathcal{A}} \cup \dots \cup M_k^{\mathcal{A}} \subseteq M$ and $M_p^{\mathcal{A}} \cap M_q^{\mathcal{A}} = \emptyset$ for all $p, q \in [k]$ with $p \neq q$, along with an optimal matching between each type N_p and the corresponding bundle $M_p^{\mathcal{A}}$ for all $p \in [k]$, thereby inducing a unique (N, M) -matching $X^{\mathcal{A}} = (x_{ij}^{\mathcal{A}})_{i \in N, j \in M}$.*

We call $M_p^{\mathcal{A}}$ the *allocated bundle* of type p under \mathcal{A} and $M_0^{\mathcal{A}} = M \setminus \cup_{p \in [k]} M_p^{\mathcal{A}}$ the set of *withheld items* (we will sometimes drop the superscript \mathcal{A} when there is no ambiguity). Thus, item assignment within a type is socially optimal (by design) and EF1 for individual agents (since each is assigned at most 1 item). Type p envies type q if $v_p(M_p^{\mathcal{A}}) < v_p(M_q^{\mathcal{A}})$; p envies q up to ν items, $\nu \in [|M_q^{\mathcal{A}}|]$, if there is a subset $C \subseteq M_q^{\mathcal{A}}$ such that $|C| = \nu$ and $v_p(M_p^{\mathcal{A}}) \geq v_p(M_q^{\mathcal{A}} \setminus C)$ and, for every subset $C' \subseteq M_q^{\mathcal{A}}$ with $|C'| < \nu$, $v_p(M_p^{\mathcal{A}}) < v_p(M_q^{\mathcal{A}} \setminus C')$. We can analogously define the envy of a type for a bundle (up to any number of items).

With these fundamentals in place, we now define the desiderata of an allocation \mathcal{A} that we investigate in this paper. The first three are concerned with efficiency; the rest are extensions of efficiency-agnostic fairness concepts introduced by Budish [2011] and Caragiannis *et al.* [2016] respectively.

Definition 3 (Type-completeness). *\mathcal{A} is type-complete if $\cup_{p \in [k]} M_p^{\mathcal{A}} = M$; otherwise it is type-incomplete.¹*

Definition 4 (Waste and non-wastefulness). *An item $j \in M$ is said to be wasted by an allocation \mathcal{A} if it has a positive marginal utility for some type $p \in [k]$ (i.e. $\Delta_p(M_p^{\mathcal{A}}; j) > 0$) but is either withheld (i.e. $j \in M_0^{\mathcal{A}}$) or belongs to the allocated bundle of some type $q \neq p$ for which it has zero marginal utility (i.e. $j \in M_q^{\mathcal{A}}$ and $\Delta_q(M_q^{\mathcal{A}}; j) = 0$). \mathcal{A} is called non-wasteful if it has no wasted item, and wasteful otherwise.*

Definition 5 (Typewise Pareto optimality). *Allocation \mathcal{A}_1 is said to typewise Pareto dominate another allocation \mathcal{A}_2 if $v_p(M_p^{\mathcal{A}_1}) \geq v_p(M_p^{\mathcal{A}_2})$ for all types $p \in [k]$ and $v_p(M_p^{\mathcal{A}_1}) > v_p(M_p^{\mathcal{A}_2})$ for some type $p \in [k]$. An allocation that is not typewise Pareto dominated by any other allocation is typewise Pareto optimal.*

¹Type-completeness does not preclude an item $j \in M_p^{\mathcal{A}}$ remaining unassigned in the $(N_p, M_p^{\mathcal{A}})$ -matching.

Definition 6 (Typewise envy-freeness up to one item). *Allocation \mathcal{A} is typewise envy-free up to one item (TEF1) if for any two types $p, q \in [k]$, p either does not envy q or envies q up to one item, i.e. there exists an item $j \in M_q^{\mathcal{A}}$ such that $v_p(M_p^{\mathcal{A}}) \geq v_p(M_q^{\mathcal{A}} \setminus \{j\})$.*

Definition 7 (Typewise marginal envy-freeness up to one item). *Allocation \mathcal{A} is typewise marginally envy-free up to one item (TMEF1) if for any $p, q \in [k]$, there is an item $j \in M_q^{\mathcal{A}}$ such that $v_p(M_p^{\mathcal{A}}) \geq v_p(M_p^{\mathcal{A}} \cup M_q^{\mathcal{A}} \setminus \{j\}) - v_p(M_p^{\mathcal{A}})$.*

We are now ready to formulate and analyze specific problems that approach ‘good’ allocations in different ways. But first, we provide two problem instances that we will use as running examples in the rest of the paper.

Example 1. *The types N_1, N_2 and items M are as shown in Figure 1 (left): $u(i, j)$ is 1 if there is an edge between agent i and item j , and 0 otherwise. If bundles $M_1 = \{1, 2, 6\}$ and $M_2 = M \setminus M_1$ are allocated to N_1 and N_2 , there is a unique optimal (N_2, M_2) -matching with 3, 4, 5 assigned to b_1, b_2, b_3 respectively; but there are two optimal (N_1, M_1) -matchings in both of which 2 is assigned to a_2 : If 1 is assigned to a_1 and 6 remains unassigned, then we have a wasteful allocation since 6 could be assigned to b_4 so that $\Delta_2(M_2; 6) = 1$; but if 6 is assigned to a_1 instead, the allocation is non-wasteful since no agent in N_2 has a positive utility for item 1.*

Example 2. *The types and items are as shown in Figure 1 (right): All agents of the same type have the same utility for an item, equal to the integer with which the edge between that type and the item is labeled; no edge indicates zero utility.*

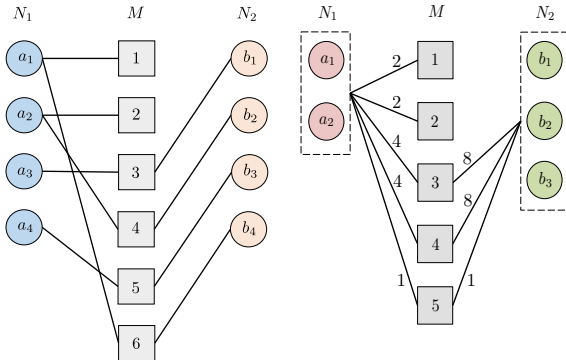


Figure 1: Instances of Examples 1 (left) and 2 (right).

3 Typewise Pareto Optimal TMEF1 Allocation

In this section, we will adapt a result of Caragiannis *et al.* [2016] to our setting. To that end, we first state without proof a property of our type-value function (Definition 1) that is part of the folklore of weighted bipartite matching.

Theorem 1. *The utilitarian type-value function $v_p(S)$ is a non-additive, monotone submodular function of $S \in 2^M$.*

Moreover, we define the *Nash type-welfare* of an allocation \mathcal{A} as the product of the type-values, i.e. $\prod_{p=1}^k v_p(M_p^{\mathcal{A}})$.

Theorem 2. *Every allocation that maximizes the Nash type-welfare is typewise Pareto optimal and TMEF1.*

Proof. Since each valuation function $v_p(\cdot)$ is monotone submodular by Theorem 1, a natural extension of Theorem 3.5 of [Caragiannis *et al.*, 2016] suffices to prove the result. \square

The next result establishes that non-wastefulness is a weakening of typewise Pareto optimality in our setting.

Lemma 1. *Any typewise Pareto optimal allocation is non-wasteful but the converse is not true.*

Proof. If an allocation \mathcal{A} had a wasted item j with positive marginal utility for a type q , then, by the definition of waste, we could augment $M_q^{\mathcal{A}}$ with j for an improved type-value without reducing any other type-value, resulting in an allocation that Pareto dominates \mathcal{A} . For the converse, Example 1 provides a counterexample: Take again the allocation \mathcal{A}_1 with bundles $M_1^{\mathcal{A}_1} = \{1, 2, 6\}$, $M_2^{\mathcal{A}_1} = M \setminus M_1^{\mathcal{A}_1}$, item 1 remaining unassigned and 2, \dots , 6 assigned to a_2, b_1, b_2, b_3, a_1 respectively. This is non-wasteful (and incidentally also TEF1) but is Pareto dominated by allocation \mathcal{A}_2 with bundles $M_1^{\mathcal{A}_2} = \{1, 2, 3\}$, $M_2^{\mathcal{A}_2} = M \setminus M_1^{\mathcal{A}_2}$ since $v_1(M_1^{\mathcal{A}_2}) = 3 > 2 = v_1(M_1^{\mathcal{A}_1})$ and $v_2(M_2^{\mathcal{A}_2}) = v_2(M_2^{\mathcal{A}_1}) = 3$. \square

This implies, in conjunction with Theorem 2, that a maximum Nash type-welfare allocation is also non-wasteful. Despite the above existence results, maximizing Nash welfare with indivisible items is known to be hard in general, and the above guarantees may break down for constant-factor approximations [Caragiannis *et al.*, 2016]. Moreover, marginal envy-freeness up to one item (MEF1) is a relatively new fairness concept that is less intuitive than envy-freeness up to one item and not extensively used as yet. In fact, it is easy to show that EF1 implies MEF1 for monotone submodular bundle-valuation functions. Hence, we will henceforth focus on efficient TEF1 allocations. Note that an allocation maximizing Nash type-welfare is not necessarily TEF1: In Example 2, any allocation with $M_1 = \{1, 2\}$ and $M_2 = \{3, 4, 5\}$ maximizes the Nash type-welfare but is not TEF1 since $v_1(M_1) = 4$ whereas $v_1(M_2 \setminus \{3\}) = v_1(M_2 \setminus \{4\}) = 5$, $v_1(M_2 \setminus \{5\}) = 8$.

4 Assignment under TEF1 Constraints

We first study the problem of finding a TEF1 allocation \mathcal{A} that maximizes the sum of weights of the induced matching:

$$\text{USW}(X^{\mathcal{A}}) \triangleq \sum_{i \in N} \sum_{j \in M} x_{ij}^{\mathcal{A}} u(i, j) = \sum_{p \in [k]} v_p(M_p^{\mathcal{A}}).$$

This is equivalent to the *assignment* problem [Munkres, 1957] under TEF1 constraints specified in Definition 6. We define the decision version of the problem as follows:

Definition 8 (ASSIGNTEF1). *An instance of the Assignment under TEF1 constraints (ASSIGNTEF1) problem is given by parameters (i) to (iii) of Section 2 as well as a value $U \in \mathbb{R}_+$; it is a ‘yes’-instance iff it admits a TEF1 allocation \mathcal{A} with a utilitarian social welfare at least U .*

A TEF1 allocation always exists under our Definition 2 (see Section 5 for further details) but we prove next that it is hard to compute one with the maximum USW for more than 2 types. The hardness question for 2 types is open.

Theorem 3. *The AssignTEF1 problem is NP-complete, even with only 3 types.*

Proof. The problem is in NP: given an allocation \mathcal{A} , we need to evaluate k type-value functions, solving the polynomial-time unconstrained assignment problem (see e.g. Kuhn [1955]) each time, and can hence verify that \mathcal{A} satisfies all requirements in polynomial time.

We will now describe a polynomial-time reduction to ASSIGNTTEF1 from the NP-complete *partition* problem [Garey and Johnson, 1979]. An instance of the latter is given by a set $S = \{s_j\}_{j \in [l]}$ of l positive integers that sum to σ ; it is a ‘yes’-instance iff S can be partitioned into two subsets S_1 and S_2 such that both sum to $\sigma/2$. Given an instance of the partition problem, we construct an ASSIGNTTEF1 instance as follows. We have a set of $l+2$ items $M = [l+2]$ and a set of $2l+4$ agents N partitioned into $k=3$ types $N_1 = \{a_i\}_{i \in [l+1]}$, $N_2 = \{b_i\}_{i \in [l+1]}$, and $N_3 = \{c_1, c_2\}$. The utilities are given by $u(a_j, j) = u(b_j, j) = s_j, \forall j \in [l]$; $u(a_{l+1}, j) = u(b_{l+1}, j) = \sigma/2, \forall j \in \{l+1, l+2\}$; $u(c_1, l+1) = u(c_2, l+2) = \kappa$ for an arbitrarily large constant $\kappa > \sigma$; $u(i, j) = 0$ for every other (i, j) -pair. Finally, let $U = 2\kappa + \sigma$.

First, we prove that, for any ‘yes’-instance of the partition problem, so is the corresponding ASSIGNTTEF1 instance we constructed. Given the two parts S_1 and S_2 of S as above, consider the bundles $M_1 = \{j \in M : s_j \in S_1\}$, $M_2 = \{j \in M : s_j \in S_2\}$, and $M_3 = \{l+1, l+2\}$ allocated to N_1 , N_2 , and N_3 respectively, with no withheld items; evidently, assigning item j to agent a_j (resp. b_j) for every s_j in S_1 (resp. S_2), $l+1$ to c_1 , and $l+2$ to c_2 constitutes the unique optimal matching between each type and its allocated bundle, inducing an allocation \mathcal{A} (Definition 2). We want to prove that \mathcal{A} is TEF1. Note that for the above utilities, $v_3(M_3) = 2\kappa > 0 = v_3(M_q)$ for all $q \in \{1, 2\}$; $v_1(M_1) = v_2(M_1) = \sum_{s_j \in S_1} s_j = \sigma/2$ and also $v_2(M_2) = v_1(M_2) = \sum_{s_j \in S_2} s_j = \sigma/2$ from the definition of a ‘yes’-instance of the partition problem; $v_1(M_3 \setminus \{j\}) = v_2(M_3 \setminus \{j\}) = \sigma/2$ for all $j \in M_3$. Hence, each type envies any other type up to at most one item under \mathcal{A} . Finally, $\text{USW}(X^{\mathcal{A}}) = \sigma/2 + \sigma/2 + 2\kappa = 2\kappa + \sigma = U$.

Now let us prove that, assuming our constructed ASSIGNTTEF1 instance to be a ‘yes’-instance, so is the partition instance. Let \mathcal{A} be an allocation verifying all desiderata. Since $\kappa > \sigma$ and the maximum sum of realized utilities that can be achieved from items $j \in [l]$ is σ , the only way for $\text{USW}(X^{\mathcal{A}})$ to be at least $U = 2\kappa + \sigma$ is to have $\{l+1, l+2\} \subseteq M_3^{\mathcal{A}}$ with c_1 (resp. c_2) assigned to $l+1$ (resp. $l+2$). Moreover, since the other items must contribute a sum of realized utilities at least σ , it is also clear from the utility structure that each $j \in [l]$ must be assigned to either a_j or b_j ; hence, $M_3 = \{l+1, l+2\}$, there are no withheld items, and $v_1(M_1^{\mathcal{A}}) + v_2(M_2^{\mathcal{A}}) = \sigma$. Now consider the sets $S_1 = \{s_j \in S : j \in M_1^{\mathcal{A}}\}$ and $S_2 = \{s_j \in S : j \in M_2^{\mathcal{A}}\}$: It is evident that they form a partition of S such that the sum of the values in S_1 (resp. S_2) equals $v_1(M_1^{\mathcal{A}})$ (resp. $v_2(M_2^{\mathcal{A}})$). Since \mathcal{A} is TEF1, we must have $v_1(M_1^{\mathcal{A}}) \geq v_1(M_3^{\mathcal{A}} \setminus \{j\})$ for some $j \in M_3^{\mathcal{A}}$; but, from our utility structure, $v_1(M_3^{\mathcal{A}} \setminus \{j\}) = \sigma/2$ for all $j \in M_3^{\mathcal{A}}$ and so the inequality $v_1(M_1^{\mathcal{A}}) \geq \sigma/2$ holds. Arguing similarly, $v_2(M_2^{\mathcal{A}}) \geq \sigma/2$. But since $v_1(M_1^{\mathcal{A}}) + v_2(M_2^{\mathcal{A}}) = \sigma$, then

$$\sum_{s_j \in S_1} s_j = v_1(M_1^{\mathcal{A}}) = \sum_{s_j \in S_2} s_j = v_2(M_2^{\mathcal{A}}) = \sigma/2. \quad \square$$

Since $v_p(\cdot)$ is a particular submodular function, Theorem 3 implies the following result for the traditional indivisible item allocation setting where each agent receives a bundle.

Corollary 1. *For monotone submodular agent valuation functions over bundles, it is NP-hard to compute the EF1 allocation that maximizes the sum of valuations.*

One might conjecture that the maximum-USW TEF1 allocation is non-wasteful. But the following surprising result belies this intuition, and raises the question: Does a non-wasteful, TEF1 allocation always exist?

Proposition 1. *The TEF1 allocation that maximizes the utilitarian social welfare may waste items, even in a problem instance that admits a non-wasteful TEF1 allocation.*

Proof. In Example 2, any allocation with bundles $M_1 = \{1, 2\}$ and $M_2 = \{3, 4\}$, and no other allocation, maximizes USW under TEF1 constraints with $\text{USW} = 20$. But, such an allocation is wasteful since item 5 is withheld although $\Delta_2(M_2; 5) = 1$. However, any allocation with bundles $M'_1 = \{1, 3\}$ and $M'_2 = \{2, 4, 5\}$ is non-wasteful and TEF1 but has $\text{USW} = 15$. \square

5 TEF1 Allocation with Reduced Waste

In our quest for a TEF1 allocation with no (or, at least, low) waste, we note that it is possible to obtain a type-complete TEF1 allocation in polynomial time by a natural extension (called L hereafter) of the algorithm due to Lipton *et al.* [2004]: Iterate over the items $j \in M$, giving item j to a type, say p , that is currently not envied by any other type for its bundle M_p ; compute an optimal matching with the augmented bundle $M_p \cup \{j\}$; construct the *envy graph* where there is a directed edge from a type q to a type r whenever q envies r ; eliminate any cycle in this graph by transferring the bundle of every envied type to its envying type on this cycle (to ensure that there is an unenvied type in each iteration), followed by re-matching within each such type. Although no item is withheld, it is possible for the final allocation to be wasteful: An item may be allocated to a type which has zero marginal utility for it or may *become* unassigned after a bundle is transferred between types.

One heuristic that could reduce waste is to allocate the item to the unenvied type that has the maximum marginal utility for it, breaking further ties uniformly at random, rather than to an arbitrary unenvied type – we call L , augmented with this heuristic, H . Unfortunately, Example 1 shows that, H can be wasteful in general. Consider the order $1, 2, \dots, 6$ over items: 1 and 2 are obviously allocated to N_1 while, depending on how ties are broken, 3, 4, 5 can all be given to N_2 . With this allocation of item 5, envy appears for the first time and N_1 is the only unenvied type. Hence, 6 must go to N_1 , although $\Delta_1(\{1, 2\}; 6) = 0$, and is wasted. Notice further that if 6 were allocated to N_2 , it would increase N_2 ’s own type-value but make N_1 envy N_2 up to 2 items although N_1 does not want 6 in conjunction with its current bundle! This is especially disappointing since Example 1 admits three non-wasteful TEF1 allocations, which are also typewise Pareto

optimal and maximize USW, with bundles $M_1^* = \{1, 2, 3\}$, $\{1, 2, 5\}$, or $\{1, 2, 3, 5\}$, and $M_2^* = M \setminus M_1^*$ (each resulting in a unique optimal matching for each type).

Nevertheless, to see how the marginal utility maximization heuristic performs in practice, we experimentally compared procedures L and H using the percentage of items wasted as our performance metric. We simulated two sets of problem instances with $n = 100$ agents partitioned into $k = 3$ types:

UNEQUAL: $|N_1| = 74, |N_2| = 13, |N_3| = 13$.²

EQUAL: $|N_p| \approx n/k$ for all types $p \in [k]$.

For each, we used $m \in \{50, 100\}$ items; for each agent, we sampled m numbers uniformly at random from $[0, 1]$ and normalized them to generate utilities for all m items. We report results averaged over 100 runs each.

For the UNEQUAL set, 39% (resp. 13%) of the items are wasted on average for $m = 100$ (resp. $m = 50$) by L whereas *no* item is wasted by H ; for the EQUAL set, L wastes only 0.005% of the items with $m = 100$ and *no* item with $m = 50$, and again *no* item is wasted by H for either value of m . Thus, we can conclude that the performance of the natural extension of [Lipton *et al.*, 2004] strongly depends on parameters such as type proportions and the number of items, whereas augmenting it with the heuristic under consideration gives surprisingly good results over a variety of input instances.

6 Binary Utilities: Non-Wasteful TEF1 Allocation

In this section, we will focus on the *binary utility model*: $u(i, j) \in \{0, 1\}, \forall i \in N, \forall j \in M$. This captures scenarios where each agent either approves or disapproves of an item but does not distinguish among its approved items. There exists prior work on fair allocation algorithms producing allocations with binary item utilities [Barman *et al.*, 2018b] but most assume additive bundle valuations.

Theorem 4. *For any problem instance with a binary utility model, there exists a non-wasteful TEF1 allocation that can be computed in polynomial time.*

Our proof is constructive: We provide and analyze an allocation algorithm for the problem (Algorithm 1). Like Lipton *et al.* [2004], we iterate over the items; however, we augment a type's bundle with an item not based on (the absence of) envy towards it but on the marginal utility of the type for the item. Let us call any algorithm that starts with empty bundles and follows the principle of augmenting the current bundle M_p of a type p with an item j *only if* p has positive marginal utility (i.e. $\Delta_p(M_p; j) > 0$) a PMU algorithm.

Proposition 2. *Under the binary utility model, we have $v_p(S) \leq \min\{|N_p|, |S|\}, \forall p \in [k], \forall S \subseteq M$. In particular, any type's value for its allocated bundle at any stage of a PMU algorithm is equal to the cardinality of the bundle, i.e. $v_p(M_p) = |M_p|, \forall p \in [k]$.*

²These numbers roughly follow the proportions of Chinese, Malay, and Indian/Other residents of Singapore according to the 2010 census report [Department of Statistics, Singapore, 2010].

Algorithm 1: PMURR($\{N_p\}_{p \in [k]}, M, (u(i, j))_{i \in N, j \in M}$)

```

Withheld set  $M_0 \leftarrow M$ ;
Temporary set  $M_T \leftarrow \emptyset$ ;
Allocated bundles  $M_p \leftarrow \emptyset \forall p \in [k]$ .
repeat
  for  $j \in M_0$  do
    if  $\Delta_p(M_p; j) = 0 \forall p \in [k]$  then
       $M_0 \leftarrow M_0 \setminus \{j\}; M_T \leftarrow M_T \cup \{j\}$ .
    else
      Find a type  $p$  such that  $\Delta_p(M_p; j) > 0$ .
       $M_0 \leftarrow M_0 \setminus \{j\}; M_p \leftarrow M_p \cup \{j\}$ .
      if some type envies  $p$  up to more than 1 item
        then
          /*Revocation and reallocation*/
          repeat
            Find types  $q, r$  such that  $r$  envies  $q$ 
              up to more than 1 item.
            Find item  $j' \in M_q$  such that
               $\Delta_r(M_r; j') = 1$ .
             $M_q \leftarrow M_q \setminus \{j'\}; M_r \leftarrow M_r \cup \{j'\}$ .
          until no type envies another up to more
            than 1 item.
      end
       $M_0 \leftarrow M_T; M_T \leftarrow \emptyset$ .
  until  $M_0 = \emptyset$  or  $\Delta_p(M_p; j) = 0 \forall j \in M_0, \forall p \in [k]$ .

```

Proof. The first part follows directly from the fact that, for binary utilities, any item can contribute either 1 or 0 to the type-value of a bundle, i.e. $\Delta_p(S; j) \in \{0, 1\}$ for any $p \in [k], S \subseteq M$, and $j \notin S$. For the second part, note that each type starts with an empty bundle (hence zero type-value) and increases its type-value by 1 every time it acquires an item under a PMU algorithm; moreover, since all positive utilities are equal, the only way for an item j to improve the type-value of p is to ensure that every item is assigned to an agent in every $(N_p, M_p \cup \{j\})$ -matching – hence, if an item is revoked, both the bundle-size and type-value diminish by 1. \square

Corollary 2. *A type $p \in [k]$ with $|M_p| = |N_p|$ under a PMU algorithm cannot envy any bundle $S \subseteq M$ since $v_p(S) \leq |N_p|$.*

If after receiving a new item, a type is still envied by all other types up to at most 1 item, no further action is necessary in that iteration. But a PMU approach by itself cannot ensure that no type will start envying the recipient up to more than 1 item if the latter was already envied up to 1 item. If envy does exceed the acceptable limit, we execute a special revocation and reallocation (RR) subroutine repeatedly until we restore the TEF1 property – hence, we call our algorithm PMURR. The functioning of the RR subroutine depends on the following result:

Proposition 3. *At any stage of a PMU algorithm, if type p envies type q up to ν or more items, then $|M_q| \geq |M_p| + \nu$.*

Proof. It follows from the definition of envy up to ν (or more) items: For any subset $C \subseteq M_q$ with $|C| = \nu - 1, v_p(M_p) < v_p(M_q \setminus C)$. But, from Proposition 2 (and since $C \subseteq M_q$), we have $v_p(M_q \setminus C) \leq |M_q \setminus C| = |M_q| - |C| = |M_q| - \nu + 1$.

Combining these with Proposition 2 ($v_p(M_p) = |M_p|$), we get $|M_p| < |M_q| - \nu + 1$, i.e. $|M_p| \leq |M_q| - \nu$. \square

Corollary 3. *Type p can envy another type q only if q has an allocated bundle of a larger size than p which, in turn, implies that there can be no cycles in the envy graph among types for binary utilities under a PMU algorithm.*

Next, we prove that the existence of envy is sufficient for having an item that can be revoked and reallocated.

Lemma 2. *At any stage of a PMU algorithm, if type p envies type q , then there exists an item j in the bundle M_q which has a positive marginal utility for p , i.e. $\Delta_p(M_p; j) = 1$.*

Proof. By Proposition 2 and Corollary 2, if N'_p is the subset of agents who are assigned items in the current (N_p, M_p) -matching, then we must have $v_p(M_p) = |M_p| = |N'_p| < |N_p|$ for p to envy q . Now, if N''_p is the subset of agents who would be assigned items in a (N_p, M_q) -matching, we must have $|N''_p| = v_p(M_q) > v_p(M_p) = |N'_p|$, so that $|N''_p \setminus N'_p| \geq |N''_p| - |N'_p| > 0$. Since $N''_p \setminus N'_p \subseteq N_p \setminus N'_p$, each agent in $N''_p \setminus N'_p$ is assigned no item under the current matching but has a positive utility for a distinct item in M_q . \square

A revocation from q (resp. a reallocation to r) decrements (resp. increments) its type-value, thus reducing the envy of r towards q ; but it is not obvious how it affects all envy relations (types that already envied r up to 1 or more items, those that q might start envying up to 1 or more items, etc.) and whether we could trigger a never-ending chain reaction. Our final lemma dispels such doubts.

Lemma 3. *The revocation and reallocation subroutine produces a TEF1 allocation in a polynomial number of steps.*

Proof. First note that computing a type-value for any bundle is polynomial-time and so is checking whether r envies q up to more than 1 item (i.e. $v_r(M_r) < v_q(M_q \setminus \{j\})$ for some $j \in M_q$). Now, for an iteration of the RR subroutine to occur, we must have a type r envying a type q up to 2 or more items, so that $|M_q| \geq |M_r| + 2$ by Proposition 3. For any collection of bundles $\bar{M} = \{M_1, \dots, M_2\}$, define the potential function $\Phi(\bar{M}) \triangleq \sum_{p \in [k]} |M_p|^2$, and let $\bar{M}' = \{M'_1, \dots, M'_2\}$ be the collection of bundles after revocation from q and reallocation to r . Hence, $|M'_q| = |M_q| - 1$ and $|M'_r| = |M_r| + 1$ (by Proposition 2) but $|M'_p| = |M_p|$ for every other type p . Thus,

$$\begin{aligned} \Phi(\bar{M}') - \Phi(\bar{M}) &= |M'_q|^2 - |M_q|^2 + |M'_r|^2 - |M_r|^2 \\ &= 2(1 + |M_r| - |M_q|), \text{ on simplification} \\ &\leq 2(1 - 2) = -2, \end{aligned}$$

i.e. Φ strictly decreases with each RR iteration, and obviously lies between 0 and m^2 (since $\sum_{p \in [k]} |M_p| \leq m$). Hence, RR terminates after a polynomial number of iterations; by the stop criterion, the final allocation is TEF1. \square

We are now ready to prove the main result of this section.

Proof of Theorem 4. To prove that the final allocation is non-wasteful, note that by construction, no bundle-augmentation and revocation-and-reallocation allow an item in an allocated bundle to remain unassigned; by Corollary 3, no envy cycles are ever formed in the envy graph among types, hence bundles are never passed between types (unlike in the decycling procedure of [Lipton *et al.*, 2004]) and so no item once assigned can become unassigned (or transferred to the temporary set). However, within the “for” loop, the revocation(s) might result in an agent becoming unassigned who has a positive utility for a currently withheld item that was put in the temporary set because it previously had zero marginal utility for each type (hence this item now becomes a wasted item). This necessitates the outer “repeat” loop whose stop criterion ensures that there are no wasted items in the withheld set at the end of the algorithm.

Lemma 3 ensures that the allocation is TEF1 after every iteration of the “for” loop, hence upon termination.

For proving the polynomial time-complexity, recall that computing $v_p(\cdot)$ and $\Delta_p(\cdot)$, verifying envy up to more than 1 item, and finding an item to revoke and reallocate are all polynomial-time; Lemma 3 ensures that each iteration of the “for” loop takes polynomial time. With each iteration of the outer “repeat” loop, the size of the withheld set, which starts at m , strictly decreases (if it does not, it must mean that each withheld item has zero marginal utility for every type – a stop criterion), hence we have a linear number of iterations. \square

7 Discussion And Future Work

We have introduced and investigated typewise fairness and efficiency concepts for the allocation of indivisible goods. The major open question is whether a non-wasteful TEF1 allocation always exists for arbitrary agent-item utilities. Important properties (e.g. the equality of bundle size and type-value) do not carry over from binary to arbitrary utilities, hence extensions of our PMURR algorithm to general utility models remain elusive. An alternative would be to have heuristic approaches such as the one in Section 5 with *provable* upper bounds on the amount of waste.

We considered one type-value function $v_p(\cdot)$ (Definition 1) here; the *average utilitarian type-value* $\hat{v}_p(S) \triangleq v_p(S)/|N_p|$ is equivalent to $v_p(S)$ for all intents and purposes in this paper. More complex functions such as OWA operators [Yager, 1988] or those that address fairness within a type (e.g. [Bouweret and Lang, 2011]) merit further analysis.

Other possible directions for future research include non-envy-based fairness concepts (egalitarian type-welfare, proportionality, maximin share etc.) as well as strategic implications of typewise fair allocation algorithms.

Acknowledgments

Chakraborty and Zick were supported by the Singapore MOE grant R-252-000-625-133 and the Singapore NRF Research Fellowship R-252-000-750-733, Benabbou by the ANR project 14-CE24-0007-01-Cocorico-CoDec, and Elkind by the ERC grant 639945 (ACCORD). The authors would like to thank the anonymous reviewers of IJCAI 2019 and the reviewers and attendees of FAMAS 2019 for their feedback.

References

- [Barman *et al.*, 2018a] Siddharth Barman, Arpita Biswas, Sanath Kumar Krishnamurthy, and Yadati Narahari. Groupwise maximin fair allocation of indivisible goods. In *Proc. AAAI*, pages 917–924, 2018.
- [Barman *et al.*, 2018b] Siddharth Barman, Sanath Kumar Krishnamurthy, and Rohit Vaish. Greedy algorithms for maximizing Nash social welfare. In *Proc. AAMAS*, pages 7–13, 2018.
- [Benabbou *et al.*, 2018] Nawal Benabbou, Mithun Chakraborty, Vinh Ho Xuan, Jakub Sliwinski, and Yair Zick. Diversity constraints in public housing allocation. *Proc. AAMAS*, pages 973–981, 2018.
- [Bouveret and Lang, 2011] Sylvain Bouveret and Jérôme Lang. A general elicitation-free protocol for allocating indivisible goods. In *Proc. IJCAI*, pages 73–78, 2011.
- [Bouveret *et al.*, 2016] Sylvain Bouveret, Yann Chevaleyre, and Nicolas Maudet. Fair allocation of indivisible goods. In Brandt, Felix and Conitzer, Vincent and Endriss, Ulle and Lang, Jérôme and Procaccia, Ariel D., editor, *Handbook of Computational Social Choice*, chapter 12, pages 284–310. Cambridge University Press, 2016.
- [Budish, 2011] Eric Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103, 2011.
- [Caragiannis *et al.*, 2016] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. The unreasonable fairness of maximum Nash welfare. In *Proc. EC*, pages 305–322. ACM, 2016.
- [Conitzer *et al.*, 2019] Vincent Conitzer, Rupert Freeman, Nisarg Shah, and Jennifer Wortman Vaughan. Group fairness for indivisible good allocation. In *Proc. AAAI*, 2019. To appear.
- [Department of Statistics, Singapore, 2010] Department of Statistics, Singapore. Singapore 2010 Census: Key Indicators of the Resident Population, 2010.
- [Dwork *et al.*, 2012] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. Fairness through awareness. In *Proc. ITCS*, pages 214–226. ACM, 2012.
- [Elzayn *et al.*, 2019] Hadi Elzayn, Shahin Jabbari, Christopher Jung, Michael Kearns, Seth Neel, Aaron Roth, and Zachary Schutzman. Fair algorithms for learning in allocation problems. In *Proc. ACM FAT**, 2019. To appear.
- [Fain *et al.*, 2016] Brandon Fain, Ashish Goel, and Kamesh Munagala. The core of the participatory budgeting problem. In *Proc. WINE*, pages 384–399, 2016.
- [Fain *et al.*, 2018] Brandon Fain, Kamesh Munagala, and Nisarg Shah. Fair allocation of indivisible public goods. In *Proc. EC*, pages 575–592, 2018.
- [Foley, 1967] Duncan K. Foley. Resource allocation in the public sector. *Yale Economic Essays*, 7:45–98, 1967.
- [Garey and Johnson, 1979] Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., 1979.
- [Hardt *et al.*, 2016] Moritz Hardt, Eric Price, Nati Srebro, et al. Equality of opportunity in supervised learning. In *Proc. NIPS*, pages 3315–3323, 2016.
- [Kearns *et al.*, 2018] Michael Kearns, Seth Neel, Aaron Roth, and Zhiwei Steven Wu. Preventing fairness gerrymandering: Auditing and learning for subgroup fairness. In *Proc. ICML*, pages 2564–2572, 2018.
- [Kuhn, 1955] Harold W. Kuhn. The Hungarian method for the assignment problem. *Naval Research Logistics (NRL)*, 2(1-2):83–97, 1955.
- [Lahaie and Parkes, 2009] Sébastien Lahaie and David C Parkes. Fair package assignment. In *Proc. AMMA*, page 92, 2009.
- [Lipton *et al.*, 2004] Richard J. Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. On approximately fair allocations of indivisible goods. In *Proc. EC*, pages 125–131. ACM, 2004.
- [Lovász and Plummer, 2009] László Lovász and Michael D. Plummer. *Matching theory*, volume 367. American Mathematical Society, 2009.
- [Manurangsi and Suksompong, 2017] Pasin Manurangsi and Warut Suksompong. Asymptotic existence of fair divisions for groups. *Mathematical Social Sciences*, 89:100–108, 2017.
- [Markakis, 2017] Evangelos Markakis. Approximation algorithms and hardness results for fair division with indivisible goods. In *Trends in Computational Social Choice*, chapter 12, pages 231–247. AI Access, 2017.
- [Munkres, 1957] James Munkres. Algorithms for the assignment and transportation problems. *Journal of SIAM*, 5(1):32–38, 1957.
- [Plaut and Roughgarden, 2018] Benjamin Plaut and Tim Roughgarden. Almost envy-freeness with general valuations. In *Proc. SODA*, pages 2584–2603. SIAM, 2018.
- [Segal-Halevi and Suksompong, 2018] Erel Segal-Halevi and Warut Suksompong. Democratic fair allocation of indivisible goods. In *Proc. IJCAI*, pages 482–488, 2018.
- [Suksompong, 2018] Warut Suksompong. Approximate maximin shares for groups of agents. *Mathematical Social Sciences*, 92:40–47, 2018.
- [Todo *et al.*, 2011] Taiki Todo, Runcong Li, Xuemei Hu, Takayuki Mouri, Atsushi Iwasaki, and Makoto Yokoo. Generalizing envy-freeness toward group of agents. In *Proc. IJCAI*, pages 386–392, 2011.
- [Yager, 1988] Ronald R. Yager. On Ordered Weighted Averaging aggregation operators in multicriteria decision making. *IEEE Transactions on Systems, Man, and Cybernetics*, 18(1):183–190, 1988.
- [Zhou, 1992] Lin Zhou. Strictly fair allocations in large exchange economies. *Journal of Economic Theory*, 57(1):158 – 175, 1992.