



HAL
open science

Finding Fair and Efficient Allocations When Valuations Don't Add Up

Nawal Benabbou, Mithun Chakraborty, Ayumi Igarashi, Yair Zick

► **To cite this version:**

Nawal Benabbou, Mithun Chakraborty, Ayumi Igarashi, Yair Zick. Finding Fair and Efficient Allocations When Valuations Don't Add Up. The 13th Symposium on Algorithmic Game Theory (SAGT 2020), Sep 2020, Augsburg (Virtual Conference), Germany. pp.32-46, 10.1007/978-3-030-57980-7_3 . hal-02955635

HAL Id: hal-02955635

<https://hal.sorbonne-universite.fr/hal-02955635>

Submitted on 2 Oct 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Finding Fair and Efficient Allocations When Valuations Don't Add Up^{*}

Nawal Benabbou¹, Mithun Chakraborty², Ayumi Igarashi³, and Yair Zick⁴

¹ Sorbonne Université, CNRS, Laboratoire d'Informatique de Paris 6, LIP6 F-75005, Paris, France nawal.benabbou@lip6.fr

² University of Michigan, Ann Arbor, USA dcsmc@umich.edu

³ National Institute of Informatics, Tokyo, Japan ayumi_igarashi@nii.ac.jp

⁴ University of Massachusetts, Amherst, USA yzick@umass.edu

Abstract. In this paper, we present new results on the fair and efficient allocation of indivisible goods to agents whose preferences correspond to *matroid rank functions*. This is a versatile valuation class with several desirable properties (monotonicity, submodularity) which naturally models a number of real-world domains. We use these properties to our advantage; first, we show that when agent valuations are matroid rank functions, a socially optimal (i.e. utilitarian social welfare-maximizing) allocation that achieves envy-freeness up to one item (EF1) exists and is computationally tractable. We also prove that the Nash welfare-maximizing and the leximin allocations both exhibit this fairness/efficiency combination, by showing that they can be achieved by minimizing any symmetric strictly convex function over utilitarian optimal outcomes. Moreover, for a subclass of these valuation functions based on maximum (unweighted) bipartite matching, we show that a leximin allocation can be computed in polynomial time.

Keywords: Fair Division · Envy-Freeness · Submodularity · Dichotomous preferences · Matroid rank functions · Optimal welfare

1 Introduction

Suppose that we are interested in allocating seats in courses to prospective students. How should this be done? On the one hand, courses offer limited seats and have scheduling conflicts; on the other, students have preferences over the classes that they take, which must be accounted for. Course allocation can be thought of as a problem of allocating a set of *indivisible goods* (course slots) to *agents* (students). How should we divide goods among agents with subjective valuations? Can we find a “good” allocation in polynomial time?

^{*} This research was funded by an MOE Grant (no. R-252-000-625-133) and a Singapore NRF Research Fellowship (no. R-252-000-750-733). Most of this work was done when Chakraborty and Zick were employed at National University of Singapore (NUS), and Igarashi was a research visitor at NUS.

These questions have been the focus of intense study in the CS/Econ community in recent years; several justice criteria as well as methods for computing allocations that satisfy them have been investigated. Generally speaking, there are two types of justice criteria: *efficiency* and *fairness*. *Efficiency* criteria are chiefly concerned with maximizing some welfare criterion, e.g. *Pareto optimality* (PO). *Fairness* criteria require that agents do not perceive the resulting allocation as mistreating them; for example, one might want to ensure that no agent wants another agent’s assigned bundle [18]. This criterion is known as *envy-freeness* (EF); however, envy-freeness is not always achievable with indivisibilities: consider, for example, two students competing for a single course slot. Any student receiving this slot would envy the other (in our stylized example, there is just the one course with the one seat).

A simple solution ensuring envy-freeness would be to withhold the seat altogether, not assigning it to either student. This solution, however, violates most efficiency criteria. Indeed, as observed by Budish [12], envy-freeness is not always achievable, even with the weakest efficiency criterion of completeness requiring that each item is allocated to some agent. However, a less stringent fairness notion — *envy-freeness up to one good* (EF1) — can be attained. An allocation is EF1 if for any two agents i and j , there is some item in j ’s bundle whose removal results in i not envying j . EF1 complete allocations always exist, and in fact, can be found in polynomial time [26].

While trying to efficiently achieve individual criteria is challenging in itself, things get really interesting when trying to simultaneously achieve multiple justice criteria. Caragiannis et al. [13] show that when agent valuations are *additive* — i.e. every agent i values its allocated bundle as the sum of values of individual items — there exist allocations that are both PO and EF1. Specifically, these are allocations that maximize the product of agents’ utilities — also known as the *max Nash welfare* (MNW). Further work [6] shows that such allocations can be found in pseudo-polynomial time. While encouraging, these results are limited to agents with additive valuations. In particular, they do not apply to settings such as the course allocation problem described above (e.g. being assigned two courses with conflicting schedules will not result in additive gain), or other settings we describe later on. In fact, Caragiannis et al. [13] left it open whether their result extends to other natural classes of valuation functions, such as the class of submodular valuations.⁵ At present, little is known about other classes valuation functions; this is where our work comes in.

1.1 Our contributions

We focus on monotone submodular valuations with binary (or dichotomous) marginal gains, which we refer to as *matroid rank valuations*. In this setting, the added benefit of receiving another item is binary and obeys the law of diminishing marginal returns. This is equivalent to the class of valuations that can

⁵ There is an instance of two agents with monotone supermodular/subadditive valuations where no allocation is PO and EF1 [13].

be captured by matroid constraints; namely, each agent has a different matroid constraint over the items, and the value of a bundle is determined by the size of a maximum independent set included in the bundle.

Matroids offer a highly versatile framework for describing a variety of domains [29]. This class of valuations naturally arises in many practical applications, beyond the course allocation problem described above (where students are limited to either approving/disapproving a class). For example, suppose that a government body wishes to fairly allocate public goods to individuals of different minority groups (say, in accordance with a diversity-promoting policy). This could apply to the assignment of kindergarten slots to children from different neighborhoods/socioeconomic classes⁶ or of flats in public housing estates to applicants of different ethnicities [9, 8]. A possible way of achieving group fairness in this setting is to model each minority group as an agent consisting of many individuals: each agent's valuation function is based on *optimally matching* items to its constituent individuals; envy naturally captures the notion that no group should believe that other groups were offered better bundles (this is the fairness notion studied by Benabbou et al. [8]). Such assignment/matching-based valuations (known as OXS valuations [25]) are non-additive in general, and constitute an important subclass of submodular valuations. Matroid rank functions correspond to submodular valuations with binary (i.e. $\{0, 1\}$) marginal gains. The binary marginal gains assumption is best understood in context of matching-based valuations — in this scenario, it simply means that individuals either approve or disapprove of items, and do not distinguish between items they approve (we call OXS functions with binary individual preferences (0, 1)-OXS valuations). This is a reasonable assumption in kindergarten slot allocation (all approved/available slots are identical), and is implicitly made in some public housing mechanisms (e.g. Singapore housing applicants are required to effectively approve a subset of flats by selecting a block, and are precluded from expressing a more refined preference model).

In addition, imposing certain constraints on the underlying matching problem retains the submodularity of the agents' induced valuation functions: if there is a hard limit due to a *budget* or an exogenous *quota* (e.g. ethnicity-based quotas in Singapore public housing; socioeconomic status-based quotas in certain U.S. public school admission systems) on the number of items each group is able or allowed to receive, then agents' valuations are *truncated* matching-based valuations. Such valuation functions are not OXS, but are still matroid rank functions. Since agents still have binary/dichotomous preferences over items even with the quotas in place, our results apply to this broader class as well.

Using the matroid framework, we obtain a variety of positive existential and algorithmic results on the compatibility of (approximate) envy-freeness with welfare-based allocation concepts. The following is a summary of our main results (see also Table 1):

⁶ see, e.g. <https://www.ed.gov/diversity-opportunity>.

- (a) For matroid rank valuations, we show that an EF1 allocation that also maximizes the utilitarian social welfare or **USW** (hence is Pareto optimal) always exists and can be computed in polynomial time.
- (b) For matroid rank valuations, we show that leximin⁷ and MNW allocations both possess the EF1 property.
- (c) For matroid rank valuations, we provide a characterization of the leximin allocations; we show that they are identical to the minimizers of *any* symmetric strictly convex function over utilitarian optimal allocations. We obtain the same characterization for MNW allocations.
- (d) For (0, 1)-OXS valuations, we show that both leximin and MNW allocations can be computed efficiently.

	MNW	Leximin	max-USW+EF1
(0, 1)-OXS	poly-time (Th. 5)	poly-time (Th. 5)	poly-time (Th. 1)
matroid rank	?	?	poly-time (Th. 1)

Table 1. Summary of our computational complexity results.

All proofs omitted from the body of the paper due to space constraints as well as clarifying examples remarks, extensions, and additional references are available in the online full version with appendices at <https://git.io/JJYdW>.

Result (a) is remarkably positive: the EF1 and **USW** objectives are incompatible in general, even for additive valuations. Result (b) is reminiscent of Thm. 3.2 by Caragiannis et al. [13], showing that any MNW allocation is PO and EF1 under *additive* valuations. The PO+EF1 existence question beyond additive valuations, which they left open, has seen little progress. To our knowledge, the class of matroid rank valuations is the first valuation class not subsumed by additive valuations for which the EF1 property of the MNW allocation have been established.

1.2 Related work

Our paper is related to the vast literature on the fairness and efficiency issue in resource allocation. Early work on divisible resource allocation provides an elegant answer: an allocation that satisfies envy-freeness and Pareto optimality always exists under mild assumptions on valuations [34], and can be computed via convex programming of Eisenberg and Gale [17] for additive valuations. Four decades later, Caragiannis et al. [13] prove the discrete analogue of Eisenberg and Gale [17]: MNW allocation satisfies EF1 and Pareto optimality for additive

⁷ Roughly speaking, a leximin allocation is one that maximizes the realized valuation of the worst-off agent and, subject to that, maximizes that of the second worst-off agent, and so on.

valuations. Subsequently, Barman et al. [6] provide a pseudo-polynomial-time algorithm for computing allocations satisfying EF1 and PO.

While computing leximin/MNW allocations of indivisible items is hard in general, several positive results are known when agents have binary additive valuations. Darmann and Schauer [14] and Barman et al. [7] show that the maximum Nash welfare can be computed efficiently for binary additive valuations. Further, the equivalence between leximin and MNW for binary additive valuations has been obtained in several recent papers. Aziz and Rey [3] show that the algorithm proposed by Darmann and Schauer outputs a leximin optimal allocation; in particular this implies that the leximin and MNW solutions coincide for binary additive valuations. This is implied by our results. Similar results are shown by Halpern et al. [21], who also show that the leximin/MNW optimal allocation is group-strategyproof for agents with binary additive valuations. In the context of *divisible* goods, Aziz and Ye [4] show the leximin and MNW solutions also coincide for dichotomous preferences.

From a technical perspective, our work makes extensive use of matroid theory; while some papers have explored the application of matroid theory to the fair division problem [10, 20], we believe that ours is the first to demonstrate its strong connection with fairness and efficiency guarantees.

One motivation for our paper is recent work by Benabbou et al. [8] on promoting diversity in assignment problems through efficient, EF1 allocations of items to groups in a population. Similar works study quota-based fairness/diversity [2, 9, 33, and references therein], or by the optimization of carefully constructed functions [1, 15, 23, and references therein] in allocation/subset selection.

Finally, Babaioff et al. [5] present a set of results similar to our own; they further explore strategyproof mechanisms for matroid rank valuations, showing that such mechanisms exist. Our work was developed independently, and is very different from a technical perspective.

2 Model and Definitions

Throughout the paper, given a positive integer r , let $[r]$ denote the set $\{1, 2, \dots, r\}$. We are given a set $N = [n]$ of *agents*, and a set $O = \{o_1, \dots, o_m\}$ of *items* or *goods*. Subsets of O are referred to as *bundles*, and each agent $i \in N$ has a *valuation function* $v_i : 2^O \rightarrow \mathbb{R}_+$ over bundles where $v_i(\emptyset) = 0$, i.e. all valuations are *normalized*. We further assume polynomial-time oracle access to the valuation v_i of all agents. Given a valuation function $v_i : 2^O \rightarrow \mathbb{R}$, we define the *marginal gain* of an item $o \in O$ w.r.t. a bundle $S \subseteq O$, as $\Delta_i(S; o) \triangleq v_i(S \cup \{o\}) - v_i(S)$. A valuation function v_i is *monotone* if $v_i(S) \leq v_i(T)$ whenever $S \subseteq T$.

An *allocation* A of items to agents is a collection of n disjoint bundles A_1, \dots, A_n , such that $\bigcup_{i \in N} A_i \subseteq O$; the bundle A_i is allocated to agent i . Given an allocation A , we denote by A_0 the set of unallocated items, also referred to as *withheld items*. We may refer to agent i 's valuation of its bundle $v_i(A_i)$ under the allocation A as its *realized valuation* under A . An allocation is *complete* if every item is allocated to *some* agent, i.e. $A_0 = \emptyset$. We admit incomplete, but *clean*

allocations: a bundle $S \subseteq O$ is *clean* for $i \in N$ if it contains no item $o \in S$ for which agent i has zero marginal gain (i.e., $\Delta_i(S \setminus \{o\}; o) = 0$, or equivalently $v_i(S \setminus \{o\}) = v_i(S)$); an allocation A is *clean* if each allocated bundle A_i is clean for the agent i that receives it. It is easy to ‘clean’ any allocation without changing any realized valuation by iteratively revoking items of zero marginal gain from respective agents and placing them in A_0 (see Example 1 in Appendix A).

2.1 Fairness and efficiency criteria

Our fairness criteria are based on the concept of *envy*. Agent i *envies* agent j under an allocation A if $v_i(A_i) < v_i(A_j)$. An allocation A is *envy-free* (EF) if no agent envies another. We will use the following relaxation of the EF property due to Budish [12]: we say that A is *envy-free up to one good* (EF1) if, for every $i, j \in N$, i does not envy j or there exists o in A_j such that $v_i(A_i) \geq v_i(A_j \setminus \{o\})$. The efficiency concept that we are primarily interested in is *Pareto optimality*. An allocation A' is said to *Pareto dominate* the allocation A if $v_i(A'_i) \geq v_i(A_i)$ for all agents $i \in N$ and $v_j(A'_j) > v_j(A_j)$ for some agent $j \in N$. An allocation is *Pareto optimal* (PO) if it is not Pareto dominated by any other allocation.

There are several ways of measuring the welfare of an allocation [31]. Specifically, given an allocation A , (i) its *utilitarian social welfare* is $\text{USW}(A) \triangleq \sum_{i=1}^n v_i(A_i)$; (ii) its *egalitarian social welfare* is $\text{ESW}(A) \triangleq \min_{i \in N} v_i(A_i)$; and (iii) its *Nash welfare* is $\text{NW}(A) \triangleq \prod_{i \in N} v_i(A_i)$. An allocation A is said to be *utilitarian optimal* (respectively, *egalitarian optimal*) if it maximizes $\text{USW}(A)$ (respectively, $\text{ESW}(A)$) among all allocations. Since it is possible that the maximum attainable Nash welfare is 0 (say, if there are less items than agents then one agent must have an empty bundle), we use the following refinement of the maximum Nash social welfare (MNW) used in [13]: we find a maximal subset of agents, say $N_{\max} \subseteq N$, to which we can allocate bundles of positive values, and compute an allocation to agents in N_{\max} that maximizes the product of their realized valuations. If N_{\max} is not unique, we choose the one that results in the highest product of realized valuations.

The *leximin* welfare is a lexicographic refinement of egalitarian optimality. Formally, for real n -dimensional vectors x and y , x is *lexicographically greater than or equal to* y (denoted by $x \geq_L y$) if and only if $x = y$, or $x \neq y$ and for the minimum index j such that $x_j \neq y_j$ we have $x_j > y_j$. For each allocation A , we denote by $\theta(A)$ the vector of the components $v_i(A_i)$ ($i \in N$) arranged in non-decreasing order. A *leximin* allocation A is one that maximizes the egalitarian welfare in a lexicographic sense, i.e., $\theta(A) \geq_L \theta(A')$ for any other allocation A' .

2.2 Submodular Valuations

The main focus of this paper is on fair allocation when agent valuations are *submodular*. A valuation function v_i is *submodular* if single items contribute more to smaller sets than to larger ones, namely: for all $S \subseteq T \subseteq O$ and all $o \in O \setminus T$, $\Delta_i(S; o) \geq \Delta_i(T; o)$.

One important subclass of submodular valuations is *assignment valuations*, introduced by Shapley [32] and also called OXS valuations [24]. Fair allocation in this setting was explored by Benabbou et al. [8]. Here, each agent $h \in N$ represents a group of individuals N_h (such as ethnic groups and genders); each individual $i \in N_h$ (also called a *member*) has a fixed non-negative weight $u_{i,o}$ for each item o . An agent h values a bundle S via a *matching* of the items to its individuals (i.e. each item is assigned to at most one member and vice versa) that maximizes the sum of weights [27]; namely, $v_h(S) = \max\{\sum_{i \in N_h} u_{i,\pi(i)} \mid \pi \in \Pi(N_h, S)\}$, where $\Pi(N_h, S)$ is the set of matchings $\pi : N_h \rightarrow S$ in the complete bipartite graph with bipartition (N_h, S) .

Our particular focus is on submodular functions with *binary marginal gains*. We say that v_i has *binary marginal gains* if $\Delta_i(S; o) \in \{0, 1\}$ for all $S \subseteq O$ and $o \in O \setminus S$. The class of submodular valuations with binary marginal gains includes the classes of binary additive valuations [7] and of assignment valuations where the weight is binary [8]. We say that v_i is a *matroid rank valuation* if it is a submodular function with binary marginal gains (these are equivalent definitions [29]), and $(0, 1)$ -OXS if it is an assignment valuation with binary marginal gains.

3 Matroid Rank Valuations

The main theme of all results in this section is that, when all agents have matroid rank valuations, fairness and efficiency properties are compatible with one another, and there exist allocations that satisfy all three welfare criteria we consider. We start by introducing some notions from matroid theory. Formally, a *matroid* is an ordered pair (E, \mathcal{I}) , where E is some finite set and \mathcal{I} is a family of its subsets (referred to as the *independent sets* of the matroid), which satisfies the following three axioms:

- (I1) $\emptyset \in \mathcal{I}$,
- (I2) if $Y \in \mathcal{I}$ and $X \subseteq Y$, then $X \in \mathcal{I}$, and
- (I3) if $X, Y \in \mathcal{I}$ and $|X| > |Y|$, then there exists $x \in X \setminus Y$ such that $Y \cup \{x\} \in \mathcal{I}$.

The rank function $r : 2^E \rightarrow \mathbb{Z}$ of a matroid returns the *rank* of each set X , i.e. the maximum size of an independent subset of X . Another equivalent way to define a matroid is to use the axiom systems for a rank function. We require that (R1) $r(X) \leq |X|$, (R2) r is monotone, and (R3) r is submodular. Then, the pair (E, \mathcal{I}) where $\mathcal{I} = \{X \subseteq E \mid r(X) = |X|\}$ is a matroid [29]. In other words, if r satisfies properties (R1)–(R3) then it induces a matroid. In the fair allocation terminology, if an agent has a matroid rank valuation, then the set of *clean* bundles forms the set of independent sets of a matroid. Before proceeding further, we state some useful properties of the matroid rank valuation class.

Proposition 1. *A valuation function v_i with binary marginal gains is monotone and takes values in $[|S|]$ for any bundle S (hence $v_i(S) \leq |S|$).*

Proposition 2. *For matroid rank valuations, A is a clean allocation if and only if $v_i(A_i) = |A_i|$ for each $i \in N$.*

Even for binary additive valuations, EF and PO allocations may not exist (as a simple example of two agents and a single good valued at 1 by each of them demonstrates); thus, we turn our attention to EF1 and PO allocations.

3.1 Utilitarian optimal and EF1 allocation

For non-negative additive valuations, Caragiannis et al. [13] prove that every MNW allocation is Pareto optimal and EF1. However, the existence question of an allocation satisfying both the PO and EF1 properties remains open for submodular valuations. We show that the existence of a PO+EF1 allocation [13] extends to the class of matroid rank valuations. In fact, we provide a surprisingly strong relation between efficiency and fairness: utilitarian optimality (stronger than Pareto optimality) and EF1 turn out to be compatible under matroid rank valuations. Moreover, such an allocation can be computed in polynomial time!

Theorem 1. *For matroid rank valuations, a utilitarian optimal allocation that is also EF1 exists and can be computed in polynomial time.*

Our result is constructive: we provide a way of computing the above allocation in Algorithm 1. The proof of Theorem 1 and those of the latter theorems utilize Lemmas 1 and 2 which shed light on the interesting interaction between envy and matroid rank valuations.

Lemma 1 (Transferability property). *For monotone submodular valuation functions, if agent i envies agent j under an allocation A , then there is an item $o \in A_j$ for which i has a positive marginal gain with respect to A_i .*

Lemma 1 holds for submodular functions with arbitrary real-valued marginal gains, and is trivially true for (non-negative) additive valuations. However, there exist non-submodular valuation functions that violate the transferability property, even when they have binary marginal gains (see Example 2 in Appendix A). Below, we show that if i 's envy towards j cannot be eliminated by removing one item, then the sizes of their *clean* bundles differ by at least two. Formally, we say that agent i envies j up to more than 1 item if $A_j \neq \emptyset$ and $v_i(A_i) < v_i(A_j \setminus \{o\})$ for every $o \in A_j$.

Lemma 2. *For matroid rank valuations, if agent i envies agent j up to more than 1 item under an allocation A and j 's bundle A_j is clean, then $v_j(A_j) \geq v_i(A_i) + 2$.*

We are now ready to show that under matroid rank valuations, utilitarian social welfare maximization is polynomial-time solvable (2).

Theorem 2. *For matroid rank valuations, one can compute a clean utilitarian optimal allocation in polynomial time.*

Proof. We prove the claim by a reduction to the matroid intersection problem. Let E be the set of pairs of items and agents, i.e., $E = \{ \{o, i\} \mid o \in O \wedge i \in N \}$.

For each $i \in N$ and $X \subseteq E$, we define X_i to be the set of edges incident to i , i.e., $X_i = \{\{o, i\} \in X \mid o \in O\}$. Note that taking $E = X$, E_i is the set of all edges in E incident to $i \in N$. For each $i \in N$ and for each $X \subseteq E$, we define $r_i(X)$ to be the valuation of i , under function $v_i(\cdot)$, for the items $o \in O$ such that $\{o, i\} \in X_i$; namely,

$$r_i(X) = v_i(\{o \in O \mid \{o, i\} \in X_i\}).$$

Clearly, r_i is also a submodular function with binary marginal gains; combining this with Proposition 1 and the fact that $r_i(\emptyset) = 0$, it is easy to see that each r_i is a rank function of a matroid. Thus, the set of clean bundles for i , i.e. $\mathcal{I}_i = \{X \subseteq E \mid r_i(X) = |X|\}$, is the set of independent sets of a matroid. Taking the union $\mathcal{I} = \mathcal{I}_1 \cup \dots \cup \mathcal{I}_n$, the pair (E, \mathcal{I}) is known to form a matroid [22], often referred to as a *union matroid*. By definition, $\mathcal{I} = \{\bigcup_{i \in N} X_i \mid X_i \in \mathcal{I}_i \wedge i \in N\}$, so any independent set in \mathcal{I} corresponds to a union of clean bundles for each $i \in N$ and vice versa. To ensure that each item is assigned at most once (i.e. bundles are disjoint), we will define another matroid (E, \mathcal{O}) where the set of independent sets is given by

$$\mathcal{O} = \{X \subseteq E \mid |X \cap E_o| \leq 1, \forall o \in O\}.$$

Here, $E_o = \{e = \{o, i\} \mid i \in N\}$ for $o \in O$. The pair (E, \mathcal{O}) is known as a *partition matroid* [22].

Now, observe that a common independent set of the two matroids $X \in \mathcal{O} \cap \mathcal{I}$ corresponds to a clean allocation A of our original instance where each agent i receives the items o with $\{o, i\} \in X$; indeed, each item o is allocated at most once because $|E_o \cap X| \leq 1$, and each A_i is clean because the realized valuation of agent i under A is exactly the size of the allocated bundle. Conversely, any clean allocation A of our instance corresponds to an independent set $X = \bigcup_{i \in N} X_i \in \mathcal{I} \cap \mathcal{O}$, where $X_i = \{\{o, i\} \mid o \in A_i\}$: for each $i \in N$, $r_i(X_i) = |X_i|$ by Proposition 2, and hence $X_i \in \mathcal{I}_i$, which implies that $X \in \mathcal{I}$; also, $|X \cap E_o| \leq 1$ as A is an allocation, and hence $X \in \mathcal{O}$.

Thus, the maximum utilitarian social welfare is the same as the size of a maximum common independent set in $\mathcal{I} \cap \mathcal{O}$. It is well known that one can find a largest common independent set in two matroids in time $O(|E|^3\theta)$ where θ is the maximum complexity of the two independence oracles [16]. Since the maximum complexity of checking independence in two matroids (E, \mathcal{O}) and (E, \mathcal{I}) is bounded by $O(mnF)$ where F is the maximum complexity of the value query oracle, we can find a set $X \in \mathcal{I} \cap \mathcal{O}$ with maximum $|X|$ in time $O(|E|^3mnF)$. \square

We are now ready to prove Theorem 1.

Proof (Proof of Theorem 1). Algorithm 1 maintains optimal USW as an invariant and terminates on an EF1 allocation. Specifically, we first compute a clean allocation that maximizes the utilitarian social welfare. The EIT subroutine in the algorithm iteratively diminishes envy by transferring an item from the envied bundle to the envious agent; Lemma 1 ensures that there is always an item in the envied bundle for which the envious agent has a positive marginal gain.

Algorithm 1: Algorithm for finding utilitarian optimal EF1 allocation

```

1 Compute a clean, utilitarian optimal allocation  $A$ .
2 /*Envy-Induced Transfers (EIT)*/
3 while there are two agents  $i, j$  such that  $i$  envies  $j$  more than 1 item. do
4   | Find item  $o \in A_j$  with  $\Delta_i(A_i; o) = 1$ .
5   |  $A_j \leftarrow A_j \setminus \{o\}$ ;  $A_i \leftarrow A_i \cup \{o\}$ .
6 end

```

Correctness: Each EIT step maintains the optimal utilitarian social welfare as well as cleanness: an envied agent’s valuation diminishes exactly by 1 while that of the envious agent increases by exactly 1. Thus, if it terminates, the EIT subroutine retains the initial (optimal) USW and, by the stopping criterion, induces the EF1 property. To show that the algorithm terminates in polynomial time, we define the potential function $\phi(A) \triangleq \sum_{i \in N} v_i(A_i)^2$. At each step of the algorithm, $\phi(A)$ strictly decreases by 2 or a larger integer. To see this, let A' denote the resulting allocation after reallocation of item o from agent j to i . Since A is clean, we have $v_i(A'_i) = v_i(A_i) + 1$ and $v_j(A'_j) = v_j(A_j) - 1$; since all other bundles are untouched, $v_k(A'_k) = v_k(A_k)$ for every $k \in N \setminus \{i, j\}$. Also, since i envies j up to more than one item under allocation A , $v_i(A_i) + 2 \leq v_j(A_j)$ by Lemma 2. Combining these, simple algebra gives us $\phi(A') - \phi(A) \leq -2$.

Complexity: By Theorem 2, computing a clean utilitarian optimal allocation can be done in polynomial time. The value of the non-negative potential function has a polynomial upper bound: $\sum_{i \in N} v_i(A_i)^2 \leq (\sum_{i \in N} v_i(A_i))^2 \leq m^2$. Thus, Algorithm 1 terminates in polynomial time. \square

An interesting implication of the above analysis is that a utilitarian optimal allocation that minimizes $\sum_{i \in N} v_i(A_i)^2$ is always EF1.

Corollary 1. *For matroid rank valuations, any clean, utilitarian optimal allocation A that minimizes $\phi(A) \triangleq \sum_{i \in N} v_i(A_i)^2$ among all utilitarian optimal allocations is EF1.*

Despite its simplicity, Algorithm 1 significantly generalizes that of Benabbou et al. [8]’s Theorem 4 (which ensures the existence of a *non-wasteful* EF1 allocation for (0,1)-OXS valuations) to matroid rank valuations. We note, however, that the resulting allocation may be neither MNW nor leximin even when agents have (0,1)-OXS valuations: Example 3 in Appendix A illustrates this and also shows that the converse of Corollary 1 does not hold.

3.2 MNW and Leximin Allocations for Matroid Rank Functions

We characterize the set of leximin and MNW allocations under matroid rank valuations. We start by showing that Pareto optimal allocations coincide with utilitarian optimal allocations when agents have matroid rank valuations. Intuitively, if an allocation is not utilitarian optimal, one can find an ‘augmenting’

path that makes at least one agent happier but no other agent worse off. The full proof, which is more involved and relies on the concept of *circuits* of matrices, is available online in Appendix A.

Theorem 3. *For matroid rank valuations, PO allocations are utilitarian optimal.*

Since leximin and MNW allocations are Pareto optimal [13, 11], Theorem 3 implies that such allocations are utilitarian optimal as well. Next, we show that for the class of matroid rank valuations, leximin and MNW allocations are identical to each other; further, they can be characterized as the minimizers of any symmetric strictly convex function among all utilitarian optimal allocations.

A function $\Phi : \mathbb{Z}^n \rightarrow \mathbb{R}$ is *symmetric* if for any permutation $\pi : [n] \rightarrow [n]$,

$$\Phi(z_1, z_2, \dots, z_n) = \Phi(z_{\pi(1)}, z_{\pi(2)}, \dots, z_{\pi(n)}),$$

and is *strictly convex* if for any $x, y \in \mathbb{Z}^n$ with $x \neq y$ and $\lambda \in (0, 1)$ where $\lambda x + (1 - \lambda)y$ is an integral vector, $\lambda\Phi(x) + (1 - \lambda)\Phi(y) > \Phi(\lambda x + (1 - \lambda)y)$. Examples of symmetric, strictly convex functions include: $\Phi(z_1, z_2, \dots, z_n) \triangleq \sum_{i=1}^n z_i^2$ for $z_i \in \mathbb{Z} \ \forall i$; $\Phi(z_1, z_2, \dots, z_n) \triangleq \sum_{i=1}^n z_i \ln z_i$ for $z_i \in \mathbb{Z}_{\geq 0} \ \forall i$. For an allocation A , we define $\phi(A) \triangleq \phi(v_1(A_1), v_2(A_2), \dots, v_n(A_n))$.

Theorem 4. *Let $\Phi : \mathbb{Z}^n \rightarrow \mathbb{R}$ be a symmetric strictly convex function; let A be some allocation. For matroid rank valuations, the following are equivalent:*

1. *A is a minimizer of Φ over all the utilitarian optimal allocations; and*
2. *A is a leximin allocation; and*
3. *A maximizes Nash welfare.*

The proof is highly technical and is hence relegated to Appendix A online. To summarize, we first establish the equivalence of statements 1 and 2 by showing: (i) Lemma 4: given a non-leximin utilitarian optimal allocation A , there exists an “adjacent” utilitarian optimal allocation A' which is the result of transferring one item from a ‘happy’ agent j to a less ‘happy’ agent i (the underlying submodularity guarantees the existence of such an allocation); (ii) Lemma 5: such an adjacent allocation A' has a strictly higher value of any symmetric strictly convex function than A . We complete the three-way equivalence by noting that maximizing Nash welfare is identical to minimizing the symmetric, strictly convex function $\phi(x) = -\sum_{i=1}^n \log x_i$ (carefully accounting for the possibility that some agents may realize zero valuations).

Theorem 4 does not generalize to the non-binary case: Example 5 in Appendix A presents an instance where the leximin and MNW allocation are not USW optimal. Combining the above characterization with the results of Section 3.1, we get the following fairness-efficiency guarantee for matroid rank valuations.

Corollary 2. *For matroid rank valuations, any clean leximin or MNW allocation is EF1.*

4 Assignment Valuations With Binary Gains

We now consider the practically important special case where valuations come from maximum matchings. For this valuation class, we show that invoking Theorem 3, one can find a leximin or MNW allocation in polynomial time, by a reduction to the network flow problem. We note that the complexity of the problem remains open for general matroid rank valuations.

Theorem 5. *For assignment valuations with binary marginal gains, one can find a leximin or MNW allocation in polynomial time.*

The proof, available in Appendix A, is based on the following key idea: given any instance with $(0, 1)$ -OXS valuations, we construct a flow network such that the problem of finding a leximin allocation in the original instance reduces to that of finding a *increasingly-maximal integer-valued flow* on the induced network for which Frank and Murota [19] recently gave a polynomial-time algorithm. In contrast with $(0, 1)$ -OXS valuations, computing a leximin or MNW allocation becomes NP-hard for weighted assignment valuations, even for two agents.

Theorem 6. *Computing a leximin/MNW allocation for two agents with general assignment valuations is NP-hard.*

The proof is available in Appendix A. We give a Turing reduction from PARTITION. The reduction is similar to the hardness reduction for two agents with identical additive valuations [28, 30].

5 Discussion

We study allocations of indivisible goods under matroid rank valuations in terms of the interplay among envy, efficiency, and various welfare concepts. Since the class of matroid rank functions is rather broad, our results can be immediately applied to settings where agents' valuations are induced by a matroid structure. Beyond the domains described in this work, these include several others. For example, *partition matroids* model instances where agents' have access to different item types, but can only hold a limited number of each type (their utility is the total number of items they hold); a variety of other domains, such as spanning trees, independent sets of vectors, coverage problems and more admit a matroid structure (see Oxley [29] for an overview). Indeed, a well-known result in combinatorial optimization states that *any* agent valuation structure where the greedy algorithm can be used to find the (weighted) optimal bundle, is induced by some matroid [29, Theorem 1.8.5].

There are several known extensions to matroid structures, with deep connections to submodular optimization [29, Chapter 11]. Matroid rank functions are submodular functions with binary marginal gains; however, general submodular functions admit some matroid structure which may potentially be used to extend our results to more general settings. Finally, it would be interesting to explore other fairness criteria such as proportionality, the maximin share guarantee, equitability. etc. (see, e.g. [11] and references therein) for matroid rank valuations. We present some of our attempts along these lines in Appendices B through D.

Bibliography

- [1] Ahmed, F., Dickerson, J.P., Fuge, M.: Diverse weighted bipartite b-matching. In: Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI). pp. 35–41 (2017)
- [2] Aziz, H., Gaspers, S., Sun, Z., Walsh, T.: From matching with diversity constraints to matching with regional quotas. In: Proceedings of the 18th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS). pp. 377–385 (2019)
- [3] Aziz, H., Rey, S.: Almost group envy-free allocation of indivisible goods and chores. CoRR abs/1907.09279 (2019)
- [4] Aziz, H., Ye, C.: Cake cutting algorithms for piecewise constant and piecewise uniform valuations. In: Proceedings of the 10th Conference on Web and Internet Economics (WINE). pp. 1–14 (2014)
- [5] Babaioff, M., Ezra, T., Feige, U.: Fair and truthful mechanisms for dichotomous valuations. CoRR abs/2002.10704 (2020)
- [6] Barman, S., Krishnamurthy, S.K., Vaish, R.: Finding fair and efficient allocations. In: Proceedings of the 19th ACM Conference on Economics and Computation (EC). pp. 557–574. ACM (2018)
- [7] Barman, S., Krishnamurthy, S.K., Vaish, R.: Greedy algorithms for maximizing Nash social welfare. In: Proceedings of the 17th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS). pp. 7–13 (2018)
- [8] Benabbou, N., Chakraborty, M., Elkind, E., Zick, Y.: Fairness towards groups of agents in the allocation of indivisible items. In: Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI). pp. 95–101 (2019)
- [9] Benabbou, N., Chakraborty, M., Xuan, V.H., Sliwinski, J., Zick, Y.: Diversity constraints in public housing allocation. Proceedings of the 17th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS) pp. 973–981 (2018)
- [10] Biswas, A., Barman, S.: Fair division under cardinality constraints. In: Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI). pp. 91–97 (2018)
- [11] Bouveret, S., Chevaleyre, Y., Maudet, N.: Fair allocation of indivisible goods. In: Brandt, Felix and Conitzer, Vincent and Endriss, Ulle and Lang, Jérôme and Procaccia, Ariel D. (ed.) Handbook of Computational Social Choice, chap. 12, pp. 284–310. Cambridge University Press (2016)
- [12] Budish, E.: The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy* 119(6), 1061–1103 (2011)
- [13] Caragiannis, I., Kurokawa, D., Moulin, H., Procaccia, A.D., Shah, N., Wang, J.: The unreasonable fairness of maximum Nash welfare. In: Proceedings

- of the 17th ACM Conference on Economics and Computation (EC). pp. 305–322. ACM (2016)
- [14] Darmann, A., Schauer, J.: Maximizing Nash product social welfare in allocating indivisible goods. *European Journal of Operational Research* 247(2), 548–559 (2015)
 - [15] Dickerson, J.P., Sankararaman, K.A., Srinivasan, A., Xu, P.: Balancing relevance and diversity in online bipartite matching via submodularity. In: *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*. pp. 1877–1884 (2019)
 - [16] Edmonds, J.: Matroid intersection. *Annals of Discrete Mathematics* 4, 39–49 (1979), discrete Optimization I
 - [17] Eisenberg, E., Gale, D.: Consensus of subjective probabilities: The parimutuel method. *The Annals of Mathematical Statistics* 30(1), 165–168 (1959)
 - [18] Foley, D.: Resource allocation and the public sector. *Yale Economics Essays* 7 pp. 45–98 (1967)
 - [19] Frank, A., Murota, K.: Discrete Decreasing Minimization, Part III: Network Flows. arXiv e-prints arXiv:1907.02673v2 (Sep 2019)
 - [20] Gourvès, L., Monnot, J.: Approximate maximin share allocations in matroids. In: Fotakis, D., Pagourtzis, A., Paschos, V.T. (eds.) *Algorithms and Complexity*. pp. 310–321. Springer International Publishing (2017)
 - [21] Halpern, D., Procaccia, A.D., Psomas, A., Shah, N.: Fair division with binary valuations: One rule to rule them all (2020), unpublished Manuscript
 - [22] Korte, B., Vygen, J.: *Combinatorial Optimization: Polyhedra and Efficiency*. Algorithms and Combinatorics, Springer (2006)
 - [23] Lang, J., Skowron, P.K.: Multi-attribute proportional representation. In: *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)*. pp. 530–536 (2016)
 - [24] Lehmann, B., Lehmann, D., Nisan, N.: Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behavior* 55(2), 270–296 (2006)
 - [25] Leme, R.P.: Gross substitutability: An algorithmic survey. *Games and Economic Behavior* 106, 294 – 316 (2017)
 - [26] Lipton, R.J., Markakis, E., Mossel, E., Saberi, A.: On approximately fair allocations of indivisible goods. In: *Proceedings of the 5th ACM Conference on Electronic Commerce (EC)*. pp. 125–131. ACM (2004)
 - [27] Munkres, J.: Algorithms for the assignment and transportation problems. *Journal of the Society for Industrial and Applied Mathematics* 5(1), 32–38 (1957)
 - [28] Nguyen, T.T., Roos, M., Rothe, J.: A survey of approximability and inapproximability results for social welfare optimization in multiagent resource allocation. *Annals of Mathematics and Artificial Intelligence* 68(1), 65–90 (Jul 2013)
 - [29] Oxley, J.: *Matroid Theory*. Oxford Graduate Texts in Mathematics, Oxford University Press, 2nd edn. (2011)

- [30] Ramezani, S., Endriss, U.: Nash social welfare in multiagent resource allocation. In: Proceedings of the 11th International Workshop on Agent-Mediated Electronic Commerce and Trading Agents Design and Analysis (AMEC/TADA). pp. 117–131 (2010)
- [31] Sen, A.: Collective choice and social welfare. Holden Day, San Francisco (1970)
- [32] Shapley, L.S.: Complements and substitutes in the optimal assignment problem. ASTIA Document (1958)
- [33] Suzuki, T., Tamura, A., Yokoo, M.: Efficient allocation mechanism with endowments and distributional constraints. In: Proceedings of the 17th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS). pp. 50–58 (2018)
- [34] Varian, H.R.: Equity, envy, and efficiency. *Journal of Economic Theory* 9(1), 63 – 91 (1974)