Finding Fair and Efficient Allocations for Matroid Rank Valuations

Nawal Benabbou, Mithun Chakraborty, Ayumi Igarashi, Yair Zick

Earlier version published in proceedings of SAGT 2020

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 valuation class has several properties such as monotonicity and submodularity that make it naturally applicable to many real-world domains. We show that when agent valuations are matroid rank functions, an allocation that that maximizes utilitarian social welfare and also 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. To the best of our knowledge, this is the first valuation function class not subsumed by additive valuations for which it has been established that an allocation maximizing Nash welfare is EF1. 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.

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. In addition, students might have exogenous constraints, such as a hard limit on the number of classes they may take. Course allocation can be thought of as a problem of allocating a set of *indivisible goods* (course slots) to *agents* (students). The problem has an interesting structure to it. Students are either willing or unwilling to sign up for a class — this can be thought of as having a value of either 1 or 0 for a seat in the class. In addition, if a student is able to take a set of classes S, then she would be able to take any subset of S as well. Finally, given two sets of feasible course assignments S, T such that $|S| < |T|$, we can find some class $o \in T$ such that $S \cup \{o\}$ is also a feasible course assignment (this statement is non-trivial but follows a standard proof). Such "well-behaved" structures are also known as *matroids*. Can we exploit such structure of subjective valuations to find a way of distributing indivisible items among agents that satisfies multiple desiderata and has reasonable demands for computational resources?

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, justice criteria fall into two categories: *efficiency* and *fairness*. Efficiency criteria are chiefly concerned with lowering some form of *waste*, maximizing some notion of item utilization, or agent utilities. For instance, *Pareto optimality* (PO) is a popular efficiency concept which ensures that the value realized by no agent can be improved without diminishing that of another agent. Fairness criteria require that agents do not perceive the resulting allocation as mistreating them compared to others; for example, one might want to ensure that no agent prefers another agent's assigned bundle (i.e. subset of goods) to her own bundle – this criterion is known as *envy-freeness* (EF) [26]. However, envy-freeness is not always achievable when items are indivisible: consider a stylized setting, where there is just one course with one seat for which two students are competing; any student receiving this slot would be envied by the other. A simple solution ensuring envy-freeness would be to withhold the seat altogether, not assigning it to either student. Withholding items, however, violates most efficiency criteria.

As illustrated above and also observed by Budish [16], envy-freeness is not always achievable, even under *completeness*, a very weak efficiency criterion 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. Complete, EF1 allocations always exist for monotone valuations, and in fact, can be found in polynomial time, thanks to the now-classic *envy graph algorithm* due to Lipton et al. [41].

It is already challenging to individually achieve strong allocative justice criteria; hence, computationally efficient methods that produce allocations satisfying multiple such criteria simultaneously are of particular interest. Caragiannis et al. [18] 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 *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. This is where our work comes in.

1.1 Our Contributions

We focus on monotone submodular valuations with binary (or dichotomous) marginal gains, which are also known as *matroid rank valuations* [47]. 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 be captured by *matroid* constraints. Matroids are mathematical structures that generalize the concept of linear independence beyond vector spaces [47]. In our fair allocation domain, each agent has a different matroid constraint over the collection of items, and her value for a bundle is determined by the size of a maximum independent set included in the bundle.

Matroid rank valuations naturally arise 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 classes[1] or of flats in public housing estates to applicants of different ethnicities [9, 10]. 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. [9]). Such assignment/matching-based valuations, known as OXS valuations [40], are non-additive in general, and constitute an important subclass of submodular valuations.

The binary marginal gains assumption is best understood in context of matching-based valuations described above — in this scenario, it simply means that individuals either ap-

[1] see, e.g. `https://www.ed.gov/diversity-opportunity`.

Valuation class	MNW	Leximin	max-USW+EF1
$(0, 1)$-OXS	poly (Theorem 5)	poly (Theorem 5)	poly (Theorem 1)
matroid rank	poly ([4])	poly ([4])	poly (Theorem 1; [4])

Table 1: Summary of our computational complexity results: "poly" denotes polynomial.

prove 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), student course selection (one is usually only allowed to indicate interest in a course by signing up and not asked to provide a ranking) and is implicitly made in some public housing mechanisms (Singapore housing applicants effectively approve a subset of flats by selecting a block, and are precluded from expressing a more refined preference model).

In addition, imposing certain restrictions on the underlying matching problem retains the submodularity of the agents' induced valuation functions: if agents are subject to hard exogenous *capacity* or *budget* constraints (students may only approve at most a fixed number of classes) or the number of items each group is allowed to receive must respect pre-determined *quotas* (e.g. ethnicity-based quotas in Singapore public housing [48, 20, 56, 22, 49, 60, 10]; socioeconomic status-based quotas in certain U.S. public school admission systems such as Chicago Public Schools [51, 19, 58, 10] 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, hence 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:

(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 by a simple greedy algorithm.

(b) For matroid rank valuations, we show that leximin[2] 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 (equivalently, the maximizers of any symmetric strictly concave 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.

Result (a) is remarkably positive: the EF1 and utilitarian welfare objectives are incompatible in general, even for additive valuations. In fact, maximizing the utilitarian social welfare among all EF1 allocations is NP-hard for general valuations [8]. Result (b) is reminiscent of the Theorem 3.2 in Caragiannis et al. [18], showing that any MNW allocation is PO and EF1 under *additive* valuations; they left the PO+EF1 existence question open for the submodular class. To our knowledge, the class of matroid rank valuations is the

[2]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.

first valuation class not subsumed by additive valuations for which the EF1 property of the MNW allocation has been established. Our computational tractability results (d) are significant since we know that for arbitrary real valuations, it is NP-hard to compute the following types of allocations: PO+EF even for the seemingly simple class of binary additive valuations which is subsumed by our matroid rank class (Bouveret and Lang [14] Proposition 21); leximin [12]; and MNW [45].

1.2 Related Work

There is a vast and growing literature on fairness and efficiency issues in resource allocation. Early work on divisible resource allocation provides an elegant result: an allocation that satisfies envy-freeness and Pareto optimality always exists under mild assumptions on valuations [59], and can be computed via the convex programming of Eisenberg and Gale [25] for additive valuations. In the domain of the allocation of indivisible goods (see Bouveret et al. [15], Markakis [42] for an overview), Budish [16] was the first to formalize the notion of EF1 as an approximation to envy-freeness but it implicitly appears in Lipton et al. [41]. More recently, Caragiannis et al. [18] prove the discrete analogue of Eisenberg and Gale [25]: MNW allocation satisfies EF1 and Pareto optimality for additive valuations. Barman et al. [6] provide a pseudo-polynomial-time algorithm for computing allocations satisfying EF1 and PO. Closely related to ours is the work of Biswas and Barman [13] who consider fair division under matroid constraints; our setting is fundamentally different from theirs (are valuation functions are themselves matroid rank functions and we care about efficiency as well), some of our proof techniques (e.g. item transfer for our Theorem 1) are similar to theirs. Moreover, we admit fair and efficient allocations that may be incomplete (i.e. not all items are allocated to the agents under consideration), bringing us close to recent work on fairness with "charity" [17**?**]. Many of our results generalize existing results on allocative fairness and efficiency under *binary additive preferences*, widespread in the social choice literature [37], that is a sub-class of our $(0, 1)$-OXS class: For this sub-class, Darmann and Schauer [21] and Barman et al. [7] prove that the maximum Nash welfare can be computed efficiently — generlized by our result (d); Aziz and Rey [2] (Lemma 4) establish the equivalence between leximin and MNW — a special case of our result (c); Halpern et al. [32] (Theorem 1) design a group strategy-proof mechanism that returns an allocation satisfying utilitarian optimality and EF1 — dropping strategy-proofness, we generalize this result to the matroid rank valuation class. Barman et al. [7] develop an efficient greedy algorithm to find an MNW allocation when the valuation of each agent is a concave function that depends on the number of items approved by her — we note that this class of valuations does not subsume the class of $(0, 1)$-OXS valuations,[3] hence the polynomial-time complexity result of Barman et al. [6] does not imply our Theorem 5.

Recently, Babaioff et al. [4] presented a set of results similar to ours, and established the existence of strategy-proof deterministic and randomized mechanisms for fair allocation allocation under matroid rank valuations. Our work was developed independently, and has many conceptual differences from Babaioff et al. [4]: our algorithms are based on fundamentally different principles, and our main focus is on the fairness and efficiency compatibility as well as other properties of such allocations and possible extensions beyond matroid rank valuations. We defer a more detailed comparison to the full version of the paper [11].

One motivation for this paper is recent work by Benabbou et al. [9] on promoting diversity in assignment problems through efficient, EF1 allocations of bundles to attribute-based groups in the population. Similar works study quota-based fairness/diversity [3, 10, 57, and

[3]Consider 3 items, o_1, o_2, o_3, and a group of members $S = \{1, 2, 3\}$ with member 1 assigning weight 1 to items o_1 and o_3, and members 2 and 3 assigning weight 1 to item o_2 only. The value of a maximum matching between $\{o_1, o_2\}$ and S is 2 while the value of a maximum matching between $\{o_1, o_3\}$ and S is 1.

references therein], or by the optimization of carefully constructed functions [1, 23, 38, and references therein] in allocation/subset selection.

2 Model and definitions

Throughout the paper, given a positive integer r, let $[r]$ denote the set $\{1, 2, \ldots, r\}$. We are given a set $N = [n]$ of *agents*, and a set $O = \{o_1, \ldots, o_m\}$ of *items* or *goods*. Subsets of O are referred to as *bundles*, and each agent $i \in N$ has a *valuation function* $v_i : 2^O \to \mathbb{R}_+$ over bundles where $v_i(\emptyset) = 0$.We further assume polynomial-time oracle access to the valuation v_i of all agents. Given a valuation function $v_i : 2^O \to \mathbb{R}$, 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 $A_1, \ldots, A_n$, 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., $\Delta_i(S \setminus \{o\}; o) = 0$); 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. For example, if for agent i, $v_i(\{1\}) = v_i(\{2\}) = v_i(\{1, 2\}) = 1$, then the bundle $A_i = \{1, 2\}$ is not clean for agent i (and neither is any allocation where i receives items 1 and 2) but it can be cleaned by moving item 1 (or item 2 but not both) to A_0.

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 [16]: 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* (or PO for short) if it is not Pareto dominated by any other allocation.

Closely related to the concept of efficiency is the welfare of an allocation which can be measured in several ways [54]. Specifically, given an allocation A,

- its *utilitarian social welfare* is $\texttt{USW}(A) \triangleq \sum_{i=1}^{n} v_i(A_i)$;
- its *egalitarian social welfare* is $\texttt{ESW}(A) \triangleq \min_{i \in N} v_i(A_i)$;
- its *Nash welfare* is $\texttt{NW}(A) \triangleq \prod_{i \in N} v_i(A_i)$.

An allocation A is said to be *utilitarian optimal* (respectively, *egalitarian optimal*) if it maximizes $\texttt{USW}(A)$ (respectively, $\texttt{ESW}(A)$) among all allocations.

Since it is possible that the maximum attainable Nash welfare is 0 (e.g. if there are fewer items than agents, then one agent must have an empty bundle), we use the following refinement of the maximum Nash social welfare (MNW) criterion used in [18]: we find a largest 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 the maximin welfare concept, i.e. egalitarian optimality. Formally, for real n-dimensional vectors $\boldsymbol{x}$ and $\boldsymbol{y}$, $\boldsymbol{x}$ is *lexicographically greater than or equal to* $\boldsymbol{y}$ (denoted by $\boldsymbol{x} \geq_L \boldsymbol{y}$) if and only if $\boldsymbol{x} = \boldsymbol{y}$, or $\boldsymbol{x} \neq \boldsymbol{y}$ 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 $\boldsymbol{s}(A)$ the vector of the components $v_i(A_i)$ $(i \in N)$ arranged in non-decreasing order. A *leximin* allocation A is an allocation that maximizes the egalitarian welfare in a lexicographic sense, i.e., $\boldsymbol{s}(A) \geq_L \boldsymbol{s}(A')$ for any other allocation A'.

2.2 Submodular Valuations

In this paper, agents' valuation functions are not necessarily additive but *submodular*. A valuation function v_i is *submodular* if each single item contributes more to a smaller set than to a larger one, 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 sub-class of submodular valuations is the class of *assignment valuations*. This class of valuations was introduced by Shapley [55] and is synonymous with the OXS valuation class [39, 40, 5]. Fair allocation in this setting was explored by Benabbou et al. [9]. 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*) having 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 [44]; 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 $\pi : N_h \to S$ 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 [9]. We say that v_i is a *matroid rank* valuation if it is a submodular function with binary marginal gains (these are equivalent definitions [47]), and $(0,1)$-OXS if it is an assignment valuation with binary marginal gains.[4] The constrained assignment valuations discussed in the fourth paragraph of Section 1.1 are examples of matroid rank valuations that are not $(0,1)$-OXS.

3 Matroid rank valuations

The main theme of all results in this section is that, when all agents have matroid rank valuations, fairness (EF1) and efficiency (PO) properties are compatible with each other and also with all three optimal welfare criteria we consider. Lemma 1 below shows that Pareto optimality of optimal welfare is unsurprising; but, it is non-trivial to prove the EF1 property in each case.

Lemma 1. *For monotone valuations, every utilitarian optimal, leximin, and MNW allocation is Pareto optimal.*

[4] $(0,1)$-OXS valuations coincide with rank functions of *transversal* matroids [5].

We start the analysis of matroid rank valuations by introducing the basics of 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 $r : 2^E \to \mathbb{Z}$ 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 [47]. In other words, if r satisfies properties (R1)–(R3) then it induces a matroid.

Within the fair allocation context, if an agent has a matroid rank valuation, then the set of *clean* bundles forms the set of independent sets of a matroid. The following are useful properties of matroid rank valuations.

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

This property leads us to the following equivalence between the size and realized valuation of every *clean* allocated bundle for the matroid rank valuation class — a crucial component of all our proofs. Note that cleaning any optimal-welfare allocation leaves the welfare unaltered and ensures that each resulting withheld item is of zero marginal gain to each agent; hence it preserves the PO condition.

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

Lipton et al. [41]'s classic envy graph algorithm does not guarantee a Pareto optimal allocation under matroid rank valautions (although the output allocation is complete and EF1), and thus underscores the difficulty of finding the PO+EF1 combination under this valuation class. Moreover, note that in the simple example of one good and two agents each valuing the good at 1, both agents' valuation functions belong to the class under consideration — this shows that an envy-free and Pareto optimal allocation may not exist even under this class, and further justifies our quest for EF1 and Pareto-optimal allocations.

3.1 Finding a Utilitarian Optimal and EF1 Allocation

We will now establish that the existence of a PO+EF1 allocation, proved for additive valuations by Caragiannis et al. [18], extends to the class of matroid rank valuations. In fact, we provide a stronger — and surprisingly strong — relation between efficiency and fairness: utilitarian optimality (stronger than Pareto optimality) and EF1 turn out to be mutually compatible under this valuation class. 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 2 and 3 which shed light on the interesting interaction between envy and matroid rank valuations.

Lemma 2 (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.*

Note that Lemma 2 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.

Below, we show that if i's envy towards j under a clean allocation cannot be eliminated by removing one item from the latter's bundle, then the two agents' valuations for their respective bundles differ by at least two (in fact, we establish a stronger version of the result that does not require the envious agent i's bundle to be clean). 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 3. *For submodular functions with binary marginal gains, 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 $v_j(A_j) \geq v_i(A_i) + 2$.*

Next, we show that under matroid rank valuations, utilitarian social welfare maximization is polynomial-time solvable (2).

Theorem 2. *If all agents have submodular functions with binary marginal gains, one can compute a clean utilitarian optimal allocation in polynomial time.*

Finally, we are ready to prove Theorem 1.

Proof of Theorem 1. Consider Algorithm 1. This algorithm 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 2 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

```
Compute a clean, utilitarian optimal allocation A.
/*Envy-Induced Transfers (EIT)*/
while there are two agents i, j such that i envies j more than 1 item do
    Find item o ∈ A_j with Δ_i(A_i; o) = 1.
    A_j ← A_j \ {o}; A_i ← A_i ∪ {o}.
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. Specifically, recall that for matroid rank valuations, an allocation A is clean if and only if $v_i(A_i) = |A_i|$ for all $i \in N$ by Proposition 2. This means that if the previous allocation A is clean, then we have $v_i(A_i \cup \{o\}) = |A_i \cup \{o\}|$, and $v_j(A_j \setminus \{o\}) = |A_j \setminus \{o\}|$. Hence the new allocation after each `EIT` step remains clean. Thus, if the algorithm 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 3. Combining these, we get

$$\begin{aligned}\Phi(A') - \Phi(A) &= (v_i(A_i) + 1)^2 + (v_j(A_j) - 1)^2 - v_i(A_i)^2 - v_j(A_j)^2 \\ &= 2(1 + v_i(A_i) - v_j(A_j)) \leq 2(1 - 2) = -2.\end{aligned}$$

Complexity: By Theorem 2, a clean utilitarian optimal allocation can be computed 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. □

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.

Remark 1 (Choice of the potential function)**.** In the proof of Theorem 1, the use of the sum of squared valuations as the potential function shows that the `EIT` subroutine terminates after $O(m^2)$ iterations. However, to establish polynomial time complexity it suffices to use any *symmetric, strictly convex, polynomial* function Φ of the realized valuations (See Section 3.2) as our potential function. Moreover, Corollary 1 holds for any such function Φ as well — we elaborate on this theme in Section 3.2. ♦

Despite its simplicity, Algorithm 1 significantly generalizes that of Benabbou et al. [9]'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: In the full version, we illustrate this and also show that the converse of Corollary 1 does not hold. Also, we discuss the implications of our results for a stronger version of the EF1 property called EFX (Remark 2) and fair allocation "with charity" [17?] (Remark 3).

3.2 MNW and Leximin Allocations

We saw in Section 3.1 that under matroid rank valuations, a simple iterative procedure allows us to reach an EF1 allocation while preserving utilitarian optimality. However, as we previously noted, such allocations are not necessarily leximin or MNW. In this subsection, 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 always find an 'augmenting' path that makes at least one agent happier but no other agent worse off.

Theorem 3. *For matroid rank valuations, any Pareto optimal allocation is utilitarian optimal.*

Theorem 3 above, along with Lemma 1, implies that both leximin and MNW allocations are utilitarian optimal. 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 $\Phi : \mathbb{Z}^n \to \mathbb{R}$ is *symmetric* if for any permutation $\pi : [n] \to [n]$,

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

and is *strictly convex* if for any $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{Z}^n$ with $\boldsymbol{x} \neq \boldsymbol{y}$ and $\lambda \in (0, 1)$ where $\lambda \boldsymbol{x} + (1 - \lambda)\boldsymbol{y}$ is an integral vector,

$$\lambda \Phi(\boldsymbol{x}) + (1 - \lambda)\Phi(\boldsymbol{y}) > \Phi(\lambda \boldsymbol{x} + (1 - \lambda)\boldsymbol{y}).$$

A function $\Psi : \mathbb{Z}^n \to \mathbb{R}$ is *strictly concave* if for any $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{Z}^n$ with $\boldsymbol{x} \neq \boldsymbol{y}$ and $\lambda \in (0, 1)$ where $\lambda \boldsymbol{x} + (1 - \lambda)\boldsymbol{y}$ is an integral vector,

$$\lambda \Psi(\boldsymbol{x}) + (1 - \lambda)\Psi(\boldsymbol{y}) < \Psi(\lambda \boldsymbol{x} + (1 - \lambda)\boldsymbol{y}).$$

It is not difficult to see that $\Phi : \mathbb{Z}^n \to \mathbb{R}$ is strictly convex if and only if $-\Phi$ is strictly concave. Examples of symmetric, strictly convex functions include the following: $\Phi(z_1, z_2, \ldots, z_n) \triangleq \sum_{i=1}^n z_i^2$ for $z_i \in \mathbb{Z} \quad \forall i \in [n]$; $\Phi(z_1, z_2, \ldots, z_n) \triangleq \sum_{i=1}^n z_i \ln z_i$ for $z_i \in \mathbb{Z}_{\geq 0}$ $\forall i \in [n]$. For an allocation A, we define $\Phi(A) \triangleq \Phi(v_1(A_1), v_2(A_2), \ldots, v_n(A_n))$.

We start by showing that given a non-leximin socially optimal allocation A, there exists an adjacent socially 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 allocation. We denote by χ_i the n-dimensional incidence vector where the j-th component of χ_i is 1 if $j = i$, and it is 0 otherwise.

Lemma 4. *Suppose that agents have matroid rank valuations. Let A be a utilitarian optimal allocation. If A is not a leximin allocation, then there is another utilitarian optimal allocation A' such that*

$$\boldsymbol{s}(A') = \boldsymbol{s}(A) + \chi_i - \chi_j,$$

for $i, j \in [n]$ with $\boldsymbol{s}(A)_j \geq \boldsymbol{s}(A)_i + 2$.

We further observe that such adjacent allocation decreases the value of any symmetric strictly convex function (equivalently, increases the value of any symmetric strictly concave function). The proof is similar to that of Proposition 6.1 in Frank and Murota [27], which shows the analogous equivalence over the integral base-polyhedron.

Lemma 5. *Let $\Phi : \mathbb{Z}^n \to \mathbb{Z}$ be a symmetric strictly convex function and $\Psi : \mathbb{Z}^n \to \mathbb{Z}$ be a symmetric strictly concave function. Let A be a utilitarian optimal allocation. Let A' be another utilitarian optimal allocation such that $\boldsymbol{s}(A') = \boldsymbol{s}(A) + \chi_i - \chi_j$ for some $i, j \in [n]$ with $\boldsymbol{s}(A)_j \geq \boldsymbol{s}(A)_i + 2$. Then, $\Phi(A) > \Phi(A')$ and $\Psi(A) < \Psi(A')$.*

Now we are ready to prove the following.

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

1. *A is a minimizer of Φ over all the utilitarian optimal allocations; and*

2. *A is a maximizer of Ψ over all the utilitarian optimal allocations; and*

3. *A is a leximin allocation; and*

4. *A maximizes Nash welfare.*

Proof. To prove 1 ⇔ 2, let A be a leximin allocation, and let A' be a minimizer of Φ over all the utilitarian optimal allocations. We will show that $\boldsymbol{s}(A')$ is the same as $\boldsymbol{s}(A)$, which, by the uniqueness of the leximin valuation vector and symmetry of Φ, proves the theorem statement.

Assume towards a contradiction that $\boldsymbol{s}(A) \neq \boldsymbol{s}(A')$. By Theorem 3, we have $\texttt{USW}(A) = \texttt{USW}(A')$. By Lemma 4, we can obtain another utilitarian optimal allocation A'' that is a lexicographic improvement of A' by decreasing the value of the j-th element of $\boldsymbol{s}(A')$ by 1 and increasing the value of the i-th element of $\boldsymbol{s}(A')$ by 1, where $\boldsymbol{s}(A')_j \geq \boldsymbol{s}(A')_i + 2$. Applying Lemma 5, we get $\Phi(\boldsymbol{s}(A')) > \Phi(\boldsymbol{s}(A''))$, which gives us the desired contradiction.

The equivalence 2 ⇔ 3 immediately holds by the fact that $-\Psi$ is a symmetric strictly convex function.

To prove 3 ⇔ 4, let A be a leximin allocation, and let A' be an MNW allocation. Again, we will show that $\boldsymbol{s}(A')$ is the same as $\boldsymbol{s}(A)$, which by the uniqueness of the leximin valuation vector and symmetry of $\texttt{NW}$, proves the theorem statement. Let $N_{>0}(A)$ (respectively, $N_{>0}(A')$) be the agent subset to which we allocate bundles of positive values under leximin allocation A (respectively, MNW allocation A'). By definition, the number n' of agents who get positive values under leximin allocation A is the same as that of MNW allocation A'. Now we denote by $\bar{\boldsymbol{s}}(A)$ (respectively, $\bar{\boldsymbol{s}}(A')$) the vector of the non-zero components $v_i(A_i)$ (respectively, $v_i(A'_i)$) arranged in non-decreasing order. Assume towards a contradiction that $\bar{\boldsymbol{s}}(A) >_L \bar{\boldsymbol{s}}(A')$. Since A' maximizes the product $\texttt{NW}(A')$ when focusing on $N_{>0}(A')$ only, the value $\sum_{i \in N_{>0}(A')} \log v_i(A'_i)$ is maximized. However, $\Psi(\boldsymbol{x}) = \sum_{i=1}^{n'} \log x_i$ is a symmetric concave function for $\boldsymbol{x} \in \mathbb{Z}^n$ with each $x_i > 0$. Thus, by a similar argument as before, one can show that $\Psi(\bar{\boldsymbol{s}}(A')) > \Psi(\bar{\boldsymbol{s}}(A))$, a contradiction. This completes the proof. □

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.*

Proof. Since both leximin and MNW allocations are Pareto-optimal, they maximize the utilitarian social welfare, by Theorem 3. By Theorem 4 and the fact that the function $\Phi(A) \triangleq \sum_{i \in N} v_i(A_i)^2$ is a symmetric strictly convex function, any leximin or MNW allocation is a utilitarian optimal allocation that minimizes $\Phi(A)$ among all utilitarian optimal allocations; hence, if such an allocation is clean, it must be EF1 by Corollary 1. □

Theorem 4 does not generalize to the non-binary case: There is an instance (with assignment valuations based on non-binary, real-valued weights) where neither leximin nor MNW allocation is utilitarian optimal. Moreover, for connections between Theorem 4 and the *Pigou-Dalton principle* Moulin [43].

4 Assignment valuations with binary gains

We now consider the special but practically important case when valuations come from maximum matchings. For this class of valuations, 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. The problem of finding a leximin allocation under the (0, 1)-OXS valuation class can be reduced to that of computing an integral balanced flow (or increasingly-maximal integer-valued flow) in a network, which has been recently shown to be polynomial-time solvable [28].

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

In contrast with assignment valuations with binary marginal gains, the problem of computing a leximin or MNW allocation becomes intractable for weighted assignment valuations even when there are only two agents.

Theorem 6. *For two agents with general assignment valuations, it is NP-hard to compute a leximin or MNW allocation.*

5 Discussion

We studied allocations of indivisible goods under submodular valuations with binary marginal gains in terms of the interplay among envy, efficiency, and various welfare concepts. We showed that three seemingly disjoint outcomes — minimizers of arbitrary symmetric strictly convex functions among utilitarian optimal allocations, the leximin allocation, and the MNW allocation — coincide in this class of valuations. We will conclude with additional implications of this work and directions for further research.

In Section 3.1, we showed that cleaning followed by further processing of a utilitarian optimal allocation (Algorithm 1) is *sufficient* for achieving the EF1+PO combination. It is still an interesting open problem whether cleaning (and hence withholding some items) is *necessary* for this purpose, i.e. can we achieve a *complete* allocation with the desired fairness-efficiency combination for matroid rank valuations.

Another imperative line of future work is investigating which of our findings extend to *more general valuation functions.* There are several known extensions to matroid structures, with deep connections to submodular optimization [47, Chapter 11]. We have already made some progress to that end. Consider, for instance, the class of submodular valuation functions with *subjective* binary marginal gains, i.e. $\Delta_i(S; o) \in \{0, \lambda_i\}$ for some agent-specific constant $\lambda_i > 0$, for every $i \in N$. For this valuations class that we call $(0, \lambda_i)$-SUB, we show that any clean, MNW allocation is still EF1 (clean bundles being defined the same way as for matroid rank valuations) but the leximin and MNW allocations no longer coincide and leximin no longer implies EF1. We have also empirically delved into *general assignment valuations* (i.e. when group members have positive real weights for items) — we report in the full version [11], experiments on a real-world data set, comparing the performance of a heuristic extension of Algorithm 1 (Section 3.1) to this valuation class with Lipton et al. [41]'s envy graph algorithm in terms of a natural measurement of *waste*, demonstrating that approximate envy-freeness can often be achieved in practice simultaneously with good efficiency guarantees even for this larger valuation class. A promising direction is to investigate PO+EF1 existence for the class well-known *gross substitutes* (GS) valuations [31, 34] which subsumes matroid rank valuations.

The fairness concept we consider here is (approximate) envy-freeness. An obvious next step is to explore other popular fairness criteria such as *proportionality* he *maximin share guarantee* or MMS *equitability*, etc. (see, e.g. Caragiannis et al. [18], Freeman et al. [29] and references therein for further details) for matroid rank valuations. We present our results from a preliminary exploration of these questions in the full version [11]. In particular, Freeman et al. [29] show that, for binary additive valuations, it can be verified in polynomial time whether an EQ1 (a relaxation of equitability in the same spirit as EF1), EF1 and PO allocation exists and, whenever it does exist, it can also be computed in polynomial time (for the time complexity result, they show that such an allocation is MNW).

References

[1] Faez Ahmed, John P. Dickerson, and Mark Fuge. Diverse weighted bipartite b-matching. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 35–41, 2017.

[2] Haris Aziz and Simon Rey. Almost group envy-free allocation of indivisible goods and chores. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 39–45, 7 2020.

[3] Haris Aziz, Serge Gaspers, Zhaohong Sun, and Toby Walsh. 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)*, pages 377–385, 2019.

[4] Moshe Babaioff, Tomer Ezra, and Uriel Feige. Fair and truthful mechanisms for dichotomous valuations, 2020. URL `https://arxiv.org/abs/2002.10704`.

[5] Maria Florina Balcan, Florin Constantin, Satoru Iwata, and Lei Wang. Learning valuation functions. In *Conference on Learning Theory*, pages 4–1, 2012.

[6] Siddharth Barman, Sanath Kumar Krishnamurthy, and Rohit Vaish. Finding fair and efficient allocations. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*, pages 557–574. ACM, 2018.

[7] Siddharth Barman, Sanath Kumar Krishnamurthy, and Rohit Vaish. Greedy algorithms for maximizing Nash social welfare. In *Proceedings of the 17th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 7–13, 2018.

[8] Siddharth Barman, Ganesh Ghalme, Shweta Jain, Pooja Kulkarni, and Shivika Narang. Fair division of indivisible goods among strategic agents. In *Proceedings of the 18th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 1811–1813, 2019.

[9] Nawal Benabbou, Mithun Chakraborty, Edith Elkind, and Yair Zick. Fairness towards groups of agents in the allocation of indivisible items. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 95–101, 2019.

[10] Nawal Benabbou, Mithun Chakraborty, Xuan-Vinh Ho, Jakub Sliwinski, and Yair Zick. The price of quota-based diversity in assignment problems. *ACM Transactions on Economics and Computation (TEAC)*, 8(3):1–32, 2020.

[11] Nawal Benabbou, Mithun Chakraborty, Ayumi Igarashi, and Yair Zick. Finding fair and efficient allocations when valuations don't add up, 2020. URL `https://arxiv.org/abs/2003.07060`.

[12] Ivona Bezáková and Varsha Dani. Allocating indivisible goods. *ACM SIGecom Exchanges*, 5(3):11–18, 2005.

[13] Arpita Biswas and Siddharth Barman. Fair division under cardinality constraints. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 91–97, 2018.

[14] Sylvain Bouveret and Jérôme Lang. Efficiency and envy-freeness in fair division of indivisible goods: Logical representation and complexity. *Journal of Artificial Intelligence Research*, 32:525–564, 2008.

[15] 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.

[16] Eric Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103, 2011.

[17] Ioannis Caragiannis, Nick Gravin, and Xin Huang. Envy-freeness up to any item with high Nash welfare: The virtue of donating items. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*, pages 527–545. ACM, 2019.

[18] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. The unreasonable fairness of maximum nash welfare. *ACM Transactions on Economics and Computation*, 7(3):12:1–12:32, 2019.

[19] Chicago Public Schools. Chicago public schools policy manual: Admissions policy for magnet, selective enrollment and other options for knowledge schools and programs, April 2017. Section 202.2. Board Report 17-0426-PO2. Available at `http://policy.cps.edu/Policies.aspx`.

[20] Beng-Huat Chua. Race relations and public housing policy in singapore. *J. Arch. Plan. Res.*, 8(4):343–354, 1991.

[21] Andreas Darmann and Joachim Schauer. Maximizing Nash product social welfare in allocating indivisible goods. *European Journal of Operational Research*, 247(2):548–559, 2015.

[22] Yongheng Deng, Tien Foo Sing, and Chaoqun Ren. The story of Singapore's public housing: From a nation of home-seekers to a nation of homeowners. In Jie Chen, Mark Stephens, and Yanyun Man, editors, *The Future of Public Housing: Ongoing Trends in the East and the West*, chapter 7, pages 103–121. Springer, 2013. ISBN 978-3-642-41621-7.

[23] John P Dickerson, Karthik Abinav Sankararaman, Aravind Srinivasan, and Pan Xu. Balancing relevance and diversity in online bipartite matching via submodularity. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, pages 1877–1884, 2019.

[24] Jack Edmonds. Matroid intersection. *Annals of Discrete Mathematics*, 4:39–49, 1979. ISSN 0167-5060. Discrete Optimization I.

[25] Edmund Eisenberg and David Gale. Consensus of subjective probabilities: The pari-mutuel method. *The Annals of Mathematical Statistics*, 30(1):165–168, 1959.

[26] Duncan Foley. Resource allocation and the public sector. *Yale Economics Essays 7*, pages 45–98, 1967.

[27] András Frank and Kazuo Murota. Discrete Decreasing Minimization, Part I: Base-polyhedra with Applications in Network Optimization, 2018. URL `https://arxiv.org/abs/1808.07600v3`.

[28] András Frank and Kazuo Murota. Discrete Decreasing Minimization, Part III: Network Flows, 2019. URL `https://arxiv.org/abs/1907.02673v2`.

[29] Rupert Freeman, Sujoy Sikdar, Rohit Vaish, and Lirong Xia. Equitable allocations of indivisible goods. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 280–286, 2019.

[30] Mohammad Ghodsi, MohammadTaghi HajiAghayi, Masoud Seddighin, Saeed Seddighin, and Hadi Yami. Fair allocation of indivisible goods: Improvements and generalizations. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 539–556. ACM, 2018.

[31] Frank Gul and Ennio Stacchetti. Walrasian equilibrium with gross substitutes. *Journal of Economic Theory*, 87:95–124, 1999.

[32] Daniel Halpern, Ariel D Procaccia, Alexandros Psomas, and Nisarg Shah. Fair division with binary valuations: One rule to rule them all. In *Proceedings of the 16th Conference on Web and Internet Economics (WINE)*, pages 370–383. Springer, 2020.

[33] F. Maxwell Harper and Joseph A. Konstan. The movielens datasets: History and context. *ACM Transactions on Interactive Intelligent Systems (TiiS)*, 5(4), 2015.

[34] Alexander S. Kelso and Vincent P. Crawford. Job matching, coalition formation, and gross substitutes. *Econometrica*, 50(6):1483–1504, 1982.

[35] Bernhard Korte and Jens Vygen. *Combinatorial Optimization: Theory and Algorithms*. Number 3 in Algorithms and Combinatorics. Springer-Verlag Berlin Heidelberg, 2006.

[36] Maria Kyropoulou, Warut Suksompong, and Alexandros A Voudouris. Almost envy-freeness in group resource allocation. *Theoretical Computer Science*, 841:110–123, 2020.

[37] Martin Lackner and Piotr Skowron. Approval-based committee voting: Axioms, algorithms, and applications, 2020. URL `https://arxiv.org/abs/2007.01795`.

[38] Jérôme Lang and Piotr Krzysztof Skowron. Multi-attribute proportional representation. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)*, pages 530–536, 2016.

[39] Benny Lehmann, Daniel Lehmann, and Noam Nisan. Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behavior*, 55(2):270–296, 2006.

[40] Renato Paes Leme. Gross substitutability: An algorithmic survey. *Games and Economic Behavior*, 106:294 – 316, 2017.

[41] Richard J. Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce (EC)*, pages 125–131. ACM, 2004.

[42] Evangelos Markakis. Approximation algorithms and hardness results for fair division with indivisible goods. In Ulle Endriss, editor, *Trends in Computational Social Choice*, chapter 12, pages 231–247. AI Access, 2017.

[43] Hervé Moulin. *Fair division and collective welfare*. MIT press, 2004.

[44] James Munkres. Algorithms for the assignment and transportation problems. *Journal of the Society for Industrial and Applied Mathematics*, 5(1):32–38, 1957.

[45] Nhan-Tam Nguyen, Trung Thanh Nguyen, Magnus Roos, and Jörg Rothe. Computational complexity and approximability of social welfare optimization in multiagent resource allocation. *Autonomous agents and multi-agent systems*, 28(2):256–289, 2014.

[46] Trung Thanh Nguyen, Magnus Roos, and Jörg Rothe. 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.

[47] James Oxley. *Matroid Theory*. Oxford Graduate Texts in Mathematics. Oxford Univerity Press, 2nd edition, 2011.

[48] Parliament of Singapore. Parliament Debates: Official Report. Better racial mix in HDB housing estates, February 1989. Vol. 52, cols. 650-668.

[49] Sock Yong Phang and Kyunghwan Kim. Singapore's Housing Policies: 1960-2013. *Frontiers in Development Policy: Innovative Development Case Studies*, pages 123–153, 2013.

[50] Benjamin Plaut and Tim Roughgarden. Almost envy-freeness with general valuations. In *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2584–2603. Society for Industrial and Applied Mathematics, 2018.

[51] Kimberly Quick. Chicago public schools: Ensuring diversity in selective enrollment and magnet schools, October 2016. The Century Foundation. `https://tcf.org/content/report/chicago-public-schools`.

[52] Sara Ramezani and Ulle Endriss. Nash social welfare in multiagent resource allocation. In *Proceedings of the 11thAgent-Mediated Electronic Commerce. Designing Trading Strategies and Mechanisms for Electronic Markets (AMEC/TADA)*, pages 117–131, 2010.

[53] Erel Segal-Halevi and Warut Suksompong. Democratic fair allocation of indivisible goods. *Artificial Intelligence*, 277:103167, 2019.

[54] Amartya Sen. *Collective choice and social welfare*. Holden Day, San Francisco, 1970.

[55] Lloyd S. Shapley. Complements and substitutes in the opttmal assignment problem. *Naval Research Logistics Quarterly*, 9(1):45–48, 1962.

[56] Loo Lee Sim, Shi Ming Yu, and Sun Sheng Han. Public housing and ethnic integration in Singapore. *Habitat International*, 27(2):293–307, 2003.

[57] Takamasa Suzuki, Akihisa Tamura, and Makoto Yokoo. Efficient allocation mechanism with endowments and distributional constraints. In *Proceedings of the 17th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 50–58, 2018.

[58] U.S. Department of Education, Office of Elementary and Secondary Education. Improving Outcomes for All Students: Strategies and Considerations to Increase Student Diversity, January 2017. Washington, D.C. `https://www.ed.gov/diversity-opportunity`.

[59] Hal R Varian. Equity, envy, and efficiency. *Journal of Economic Theory*, 9(1):63 – 91, 1974. ISSN 0022-0531.

[60] Maisy Wong. Estimating the distortionary effects of ethnic quotas in Singapore using housing transactions. *Journal of Public Economics*, 115:131–145, 2014.