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

Nawal Benabbou
Sorbonne Université - LIP6
nawal.benabbou@lip6.fr

Mithun Chakraborty*
University of Michigan
dcsmc@umich.edu

Ayumi Igarashi
University of Tokyo
igarashi@mist.i.u-tokyo.ac.jp

Yair Zick
National University of Singapore
dcsyaz@nus.edu.sg

ABSTRACT

In this paper, we present new results on the fair and efficient allocation of indivisible goods to agents with monotone, submodular, *non-additive* valuation functions over bundles — a natural valuation model for several real-world domains despite its simple structure. We show that, if such a valuation function has binary marginal gains, a utilitarian social welfare-maximizing allocation that also achieves envy-freeness up to one item exists and is computationally tractable; also, Nash welfare-maximizing and leximin allocations (where no allocated bundle includes a 'wasted' item) are envy-free up to one item as well. For a subclass of these valuation functions based on maximum (unweighted) bipartite matching, we show that leximin and MNW allocations coincide with minimizers of any symmetric strictly convex function of agents' valuations among utilitarian optimal outcomes, and can also be computed in polynomial time.

1 INTRODUCTION

How should a collection of goods be divided amongst a population of agents with subjective valuations? Are there *computationally efficient* methods for finding *good* allocations? These questions have been the focus of intense study in the CS/Econ community in recent years. Several criteria of justice have been proposed in the literature. Some criteria focus on agent welfare: e.g. *Pareto optimality*, a popular criterion of allocative efficiency, stipulates that there is no other allocation that improves one agent's valuation without hurting another. Other criteria consider how agents perceive their bundles as compared to others' allocation; a key concept here is one of *envy*: an agent envies another if she believes that her bundle is worth less than that of another's [16]. Envy-free (EF) allocations that are also Pareto optimal (PO) or even complete (i.e. each item is allocated to at least one agent) are not guaranteed to exist when the items under consideration are indivisible: consider the case of two agents and one valuable item — assigning the item to any one of them results in envy by the other. This naturally leads to the notion of *envy-freeness up to one good* (EF1) [11]: for every pair of agents i and j, j's bundle contains some item whose removal results in i not envying j; Lipton et al. [28] provide a polynomial-time algorithm for computing a complete EF1 allocation. Of particular interest are methods that simultaneously achieve several desiderata. When agent valuations are *additive*, i.e. the value for a bundle is the sum of the values for its individual items, Caragiannis et al. [12] show that allocations that satisfy both envy-freeness up to one item and Pareto optimality (PO) exist, specifically the ones that maximize the product of agents' utilities — also known as *max Nash welfare* (MNW). Barman et al. [4, 5] show that an allocation with these properties can be computed in (pseudo-)polynomial time. Indeed, most work on the fair and efficient allocation of indivisible goods has focused on the additive setting; at present, little is known in this respect about other classes of valuation functions. This is where our work comes in.

1.1 Our contributions

We focus on monotone submodular valuations with binary marginal gains that we refer to as $(0, 1)$-SUB valuations. This class of valuations naturally arises in many practical applications. 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 classesor of flats in public housing estates to applicants of different ethnicities [7]. 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 – then, 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. [6]). Another possible domain is the assignment of courses to students [11]: each student has a set of courses she would like to take, which fit her schedule in a certain manner. Thus, a student's valuation function is induced by a *maximal matching* of courses to schedule slots. Such assignment/matching-based valuations are non-additive in general and, in fact, constitute a significant subclass of submodular valuations called OXS valuations [27].

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 among 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

*This work was done while the author was at the National University of Singapore.

Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), B. An, N. Yorke-Smith, A. El Fallah Seghrouchni, G. Sukthankar (eds.), May 2020, Auckland, New Zealand

https://doi.org/doi

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).

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 still submodular. Our results apply to this broader class, as agents still have binary/dichotomous preferences over items. For $(0, 1)$-SUB valuations, we establish the following existential and computational results on the compatibility of (approximate) envy-freeness with welfare-based allocation concepts.

(a) For $(0, 1)$-SUB 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 $(0, 1)$-SUB valuations, we show that leximin[1] and MNW allocations both possess the EF1 property.
(c) For $(0, 1)$-SUB 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.

Our computational results are summarized in Table 1.

Result (a) is remarkably positive: the EF1 and utilitarian welfare objectives are incompatible in general, even for additive valuations.[2] In fact, maximizing the utilitarian social welfare among all EF1 allocations is NP-hard for general valuations [3]. Such strong welfare/fairness guarantees are not known even for simple classes such as *binary* additive valuations (i.e. the value of a subset of items is the sum of the values of individual items which are, in turn, valued at either 0 or 1 each), which are subsumed by the $(0, 1)$-SUB class. Result (b) is reminiscent of the main theorem in Caragiannis et al. [12], showing that any MNW allocation is PO and EF1 under *additive* valuations; they also showed that a PO and EF1 allocation may not exist under subadditive/supermodular valuations (Theorem 3.3) and MNW does not imply EF1 for arbitrary, real-valued submodular functions but left the PO+EF1 existence question open for the submodular class. The open questions in this paper have received substantial attention in recent literature (for instance, progress has been made on EFX or envy-freeness up to the least valued item, see e.g. [32]) but the PO+EF1 existence issue beyond additive valuations is yet to be settled. To our knowledge, $(0, 1)$-SUB is the first valuation class not subsumed by additive valuations for which the EF1 property of the MNW allocation and multiple alternative ways of achieving the PO+EF1 combination have been established. The other properties of the MNW principle that we have uncovered for this valuation class (results (b) and (c)) may be of independent interest (see the discussion in Section 5).

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 and EF [10][3]; leximin [8]; and MNW [30]. Moreover, previous work on *binary* additive valuations establishes the polynomial time-solvability of MNW (and thus finding a PO and EF1 allocation) via a clever algorithm based on a subtle running time analysis [5] — we extend this result to the strictly larger $(0, 1)$-OXS class by uncovering deeper connections to the rich literature on combinatorial optimization.

From a technical perspective, our work makes extensive use of tools and concepts from matroid theory. For instance, most of our results are based on an observation that the set of *clean* bundles (i.e. bundles containing no item with zero marginal value for the agent under consideration) forms the set of independent sets of a matroid for every agent. While some papers have explored the application of matroid theory to the fair division problem [9, 20], we believe that ours is the first to demonstrate its strong connection with fairness and efficiency guarantees.

1.2 Related work

Our paper is related to the active research on the fair allocation of indivisible goods. Budish [11] was first to formalize the EF1 concept, but it implicitly appears in Lipton et al. [28], who design a poly-time algorithm that returns an EF1 allocation for monotone valuations. Caragiannis et al. [12] prove that EF1 and Pareto optimal allocations exist for non-negative additive valuations; Barman et al. [4] provide a pseudo-polynomial-time algorithm for computing such allocations. Barman et al. [5] establish the polynomial time-solvability of MNW allocations for binary additive valuations via a clever greedy algorithm based on a subtle running time analysis — we extend this result to the strictly larger $(0, 1)$-OXS class by uncovering deeper connections to the rich literature on combinatorial optimization; Barman et al. [5]'s algorithm extends to the setting when each agent's valuation is a concave function of the number of items she approves. We note that this valuation class does not subsume the $(0, 1)$-OXS class (since bundles with the same number of approved items may have different values under the latter class), hence their result does not imply our Theorem 4.1.

One motivation for our paper is recent work by Benabbou et al. [6] on promoting diversity in assignment problems through efficient, EF1 allocations of bundles to groups in a population. Other work in this vein includes fairness/diversity through quotas [2, 7, 36, and references therein], or by the optimization of carefully constructed functions [1, 14, 25, and references therein] in allocation/subset selection.

2 MODEL AND DEFINITIONS

Throughout the paper, given a positive integer r, let $[r]$ denote the set $\{1, 2, \dots, r\}$. We are given a set $N = [n]$ of *agents*, and a set

[1]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.

[2]Consider 3 items and 2 agents Alice and Bob with (additive) valuations 1/4, 3/8, 3/8 and 0, 1/2, 1/2 for items 1, 2, 3 respectively: the unique allocation maximizing USW gives item 1 to Alice and the rest to Bob but then Alice is not envy-free of Bob up to 1 item.

[3]The NP-hardness of PO and EF indeed holds even for binary additive valuations. See Proposition 21 in [10].

	MNW	Leximin	max-USW+EF1
(0, 1)-OXS	poly-time (Theorem 4.1)	poly-time (Theorem 4.1)	poly-time (Theorem 3.4)
(0, 1)-SUB	?	?	poly-time (Theorem 3.4)

Table 1: Summary of our computational complexity results.

$O = \{o_1, \ldots, o_m\}$ 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 $v_i(S) \subseteq v_i(T)$ 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., $\Delta_i(S \setminus \{o\}; o) = 0$) and an allocation is *clean* if each agent $i \in N$ receives a clean bundle. 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 [11]: 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.

There are several ways of measuring the welfare of an allocation [34]. Specifically, given an allocation A,

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

An allocation A is said to be *utilitarian optimal* (respectively, *egalitarian optimal*) if it maximizes $\mathtt{USW}(A)$ (respectively, $\mathtt{ESW}(A)$) among all allocations. Since it is possible that the maximum attainable Nash welfare is 0, we define the maximum Nash social welfare (MNW) allocation as follows: given a problem instance, 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. 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 $\theta(A)$ 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., $\theta(A) \geq_L \theta(A')$ for any other allocation A'.

2.2 Submodular valuations

The main focus of this paper is on fair allocation when agent valuations are 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 $o \in O \setminus T$, $\Delta_i(S; o) \geq \Delta_i(T; o)$.

Submodularity is known to arise in many real-life applications. One important class of submodular valuations is the class of *assignment valuations*. This class of valuations was introduced by Shapley [35] and is identical to the OXS valuation class [26]. Fair allocation in this setting was explored by Benabbou et al. [6]. 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 [29]; 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 [5] and of assignment valuations where the weight is binary [6]. We call a valuation function v_i $(0, 1)$-SUB if it is a submodular function with binary marginal gains, and $(0, 1)$-OXS if it is an assignment valuation with binary marginal gains.

3 SUBMODULARITY AND BINARY MARGINAL GAINS

The main theme of all results in this section is that, when all agents have $(0, 1)$-SUB valuations, fairness and efficiency properties are compatible with each other and also with the optimal values of all three welfare functions we consider. Lemma 3.1 below shows that Pareto optimality of optimal welfare is unsurprising; but, it is non-trivial to prove the EF1 property in each case.

Lemma 3.1. *For monotone valuations, every utilitarian optimal, MNW, and leximin allocation is Pareto optimal.*

Before proceeding further, we state some useful properties of the (0, 1)-SUB valuation class.

Proposition 3.2. *A valuation function v_i with binary marginal gains is always monotone and takes values in $[|S|]$ for any input 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 valuation subclass under consideration — a crucial component in 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 3.3. *For submodular valuations with binary marginal gains, A is a clean allocation if and only if $v_i(A_i) = |A_i|$ for each $i \in N$.*

A simple example of one good and two agents shows that an envy-free and Pareto optimal allocation may not exist even under (0, 1)-SUB valuations. This justifies our quest for EF1 and Pareto-optimal allocations.

3.1 Utilitarian optimal and EF1 allocation

For non-negative additive valuations, Caragiannis et al. [12] prove that every MNW allocation is Pareto optimal and EF1. However, the existence question of an allocation satisfying PO and EF1 remains open for submodular valuations. We show that the existence of a PO+EF1 allocation [12] extends to a class of submodular valuations with binary marginal gains. In fact, we provide a surprising relation between efficiency and fairness: both utilitarian optimality and EF1 turn out to be compatible under (0, 1)-SUB valuations.

Theorem 3.4. *For submodular valuations with binary marginal gains, 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 3.4 and those of the latter theorems utilize Lemmas 3.5 and 3.6 which shed light on the interesting interaction between envy and (0, 1)-SUB valuations.

Lemma 3.5 (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.*

Proof. Assume that agent i envies agent j under an allocation A, i.e. $v_i(A_i) < v_i(A_j)$, but no item $o \in A_j$ has a positive marginal gain, i.e., $\Delta_i(A_i; o) = 0$ for each $o \in A_j$. Let $A_j = \{o_1, o_2, \dots, o_r\}$. Define $S_0 = \emptyset$ and $S_t = \{o_1, o_2, \dots, o_t\}$ for each $t \in [r]$. This gives us the following telescoping series:

$$v_i(A_i \cup A_j) - v_i(A_i) = \sum_{t=1}^{r} \Delta_i(A_i \cup S_{t-1}; o_t).$$

However, submodularity implies that for each $t \in [r]$, $\Delta_i(A_i \cup S_{t-1}; o_t) \leq \Delta_i(A_i; o_t) = 0$, meaning that

$$v_i(A_i \cup A_j) - v_i(A_i) = \sum_{t=1}^{r} \Delta_i(A_i \cup S_{t-1}; o_t) = 0.$$

Together with monotonicity, this yields $v_i(A_j) \leq v_i(A_i \cup A_j) = v_i(A_i) < v_i(A_j)$, a contradiction. □

Note that Lemma 3.5 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 cannot be eliminated by removing one item, then the sizes of their "clean" bundles differ at least by 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 3.6. *For submodular functions with binary marginal gains, if agent i envies agent j up to more than 1 item under a clean allocation A, then $|A_j| \geq |A_i| + 2$.*

Proof. From the definition: $A_j \neq \emptyset$ and $v_i(A_i) < v_i(A_j \setminus \{o\})$ for every $o \in A_j$. Consider one such o. From Proposition 3.2, $v_i(A_j \setminus \{o\}) \leq |A_j \setminus \{o\}| = |A_j| - 1$. Since A is clean, $v_i(A_i) = |A_i|$. Combining these, we get

$$|A_i| = v_i(A_i) < v_i(A_j \setminus \{o\}) \leq |A_j| - 1,$$

which proves the theorem statement. □

We are now ready to show that under (0, 1)-SUB valuations, utilitarian social welfare maximization is polynomial-time solvable (3.7). To this end, we will exploit the fact that the set of *clean* bundles forms the set of independent sets of a matroid. 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 $r : 2^E \to \mathbb{Z}$ 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.

Theorem 3.7. *For submodular functions with binary marginal gains, 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 3.2 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 $\mathcal{I}_i = \{ X \subseteq E_i \mid r_i(X) = |X| \}$, is the set of independent sets of a matroid. Taking the union

$$\mathcal{I} = \{ X_1 \cup \cdots \cup X_n \mid X_i \in \mathcal{I}_i, \ \forall i \in N \},$$

the pair $(E, \mathcal{I})$ is known to form a matroid [24], often referred to as a *union matroid*. By definition, 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* [24]. 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 a common independent set $X \in \mathcal{I} \cap \mathcal{O}$. To see this, given a clean allocation A, let $X = \{ \{o, i\} \mid o \in A_i \wedge i \in N \}$. Then, $X_i = \{ \{o, i\} \mid o \in A_i \}$ for each $i \in N$. By Proposition 3.3, $r_i(X_i) = |X_i|$ for each $i \in N$, 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 [15]. 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 $O(|E|^3 mnF)$. This completes the proof. □

We are now ready to prove Theorem 3.4.

Proof of Theorem 3.4. 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 eliminates envy by transferring an item from the envied bundle to the envious agent; Lemma 3.5 ensures that there is always an item in the envied bundle for which the envious agent has a positive marginal gain.
Correctness: Each `EIT` step maintains the optimal utilitarian social welfare as well as cleanliness: 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 $\phi(A) := \sum_{i \in N} v_i(A_i)^2$.

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 $o \in A_j$ with $\Delta_i(A_i; o) = 1$.
    $A_j \leftarrow A_j \setminus \{o\}$; $A_i \leftarrow A_i \cup \{o\}$.
end

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 $v_j(A'_j) = v_j(A_j) - 1$. 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 3.6. Combining these, we get

$$(v_i(A_i) + 1)^2 + (v_j(A_j) - 1)^2 - v_i(A_i)^2 - v_j(A_j)^2$$
$$= 2(1 + v_i(A_i) - v_j(A_j)) \leq -2.$$

Complexity: It remains to analyze the running time of the algorithm. By Theorem 3.7, 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. □

An interesting implication ofAlgorithm 1, specifically the above potential function argument, is that a utilitarian optimal allocation that minimizes $\sum_{i \in N} v_i(A_i)^2$ is always EF1.

Corollary 3.8. *For submodular valuations with binary marginal gains, any clean utilitarian optimal allocation A that minimizes $\phi(A) := \sum_{i \in N} v_i(A_i)^2$ among all utilitarian optimal allocations is EF1.*

Despite its simplicity, Algorithm 1 significantly generalizes that of Benabbou et al. [6]'s Theorem 4 (which ensures the existence of a non-wasteful EF1 allocation for $(0, 1)$-OXS valuations) to $(0, 1)$-SUB valuations. We note, however, that the resulting allocation may be neither MNW nor leximin even when agents have $(0, 1)$-OXS valuations: see Example 1 in Benabbou et al. [6] (2 groups with $(0, 1)$-OXS valuations), which also shows that the converse of Corollary 3.8 does not hold.

3.2 MNW and Leximin allocation

We saw that under $(0, 1)$-SUB 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 $(0, 1)$-SUB valuations. We start by showing that Pareto optimal allocations coincide with utilitarian optimal allocations when agents have $(0, 1)$-SUB 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.

In the subsequent proof, we will use the following notions and results from matroid theory: Given a matroid $(E, \mathcal{I})$, the sets in $2^E \setminus \mathcal{I}$ are called *dependent*, and a minimal dependent set of a matroid is called a *circuit*. The following is a crucial property of circuits.

Lemma 3.9 (Korte and Vygen [24]). *Let $(E, \mathcal{I})$ be a matroid, $X \in \mathcal{I}$, and $y \in E \setminus X$ such that $X \cup \{y\} \notin \mathcal{I}$. Then the set $X \cup \{y\}$ contains a unique circuit.*

Given a matroid $(E, \mathcal{I})$, we denote by $C(\mathcal{I}, X, y)$ the unique circuit contained in $X \cup \{y\}$ for any $X \in \mathcal{I}$ and $y \in E \setminus X$ such that $X \cup \{y\} \notin \mathcal{I}$.

Theorem 3.10. *For submodular valuations with binary marginal gains, any Pareto optimal allocation is utilitarian optimal.*

Proof. Define E, X_i, E_i, $\mathcal{I}_i$ for $i \in N$, $\mathcal{I}$, and $\mathcal{O}$ as in the proof of Theorem 3.7. We first observe that for each $X \in \mathcal{I}$ and each $y \in E \setminus X$, if $X \cup \{y\} \notin \mathcal{I}$, then there is agent $i \in N$ whose corresponding items in X_i together with y is not clean, i.e., $X_i \cup \{y\} \notin \mathcal{I}_i$, which by Lemma 3.9 implies that the circuit $C(\mathcal{I}, X, y)$ is contained in E_i, i.e.,

$$C(\mathcal{I}, X, y) = C(\mathcal{I}_i, X, y). \tag{1}$$

Now to prove the claim, let A be a Pareto optimal allocation. Without loss of generality, we assume that A is clean. Then, as we have seen before, A corresponds to a common independent set X^* in $\mathcal{I} \cap \mathcal{O}$ given by

$$X^* = \bigcup_{i \in N} \{\, e = \{o, i\} \in E \mid o \in A_i \,\}.$$

Suppose towards a contradiction that A does not maximize the utilitarian social welfare. This means that X^* is not a largest common independent set of $\mathcal{I}$ and $\mathcal{O}$. It is known that given two matroids and their common independent set, if it is not a maximum-size common independent set, then there is an 'augmenting' path [15]. To formally define an augmenting path, we define an auxiliary graph $G_{X^*} = (E, B^{(1)}_{X^*} \cup B^{(2)}_{X^*})$ where the set of arcs is given by

$$B^{(1)}_{X^*} = \{\, (x, y) \mid y \in E \setminus X^* \wedge x \in C(\mathcal{O}, X^*, y) \setminus \{y\} \,\},$$
$$B^{(2)}_{X^*} = \{\, (y, x) \mid y \in E \setminus X^* \wedge x \in C(\mathcal{I}, X^*, y) \setminus \{y\} \,\}.$$

Since X^* is not a maximum common independent set of $\mathcal{O}$ and $\mathcal{I}$, the set X^* admits an *augmenting* path, which is an alternating path $P = (y_0, x_1, y_1, \ldots, x_s, y_s)$ in G_{X^*} with $y_0, y_1, \ldots, y_s \notin X^*$ and $x_1, x_2, \ldots, x_s \in X^*$, where X^* can be augmented by one element along the path, i.e.,

$$X' = (X^* \setminus \{x_1, x_2, \ldots, x_s\}) \cup \{y_0, y_1, \ldots, y_s\} \in \mathcal{I} \cap \mathcal{O}.$$

Now let's write the pairs of agents and items that correspond to y_t and x_t as follows:

- $y_t = \{i(y_t), o(y_t)\}$ where $i(y_t) \in N$ and $o(y_t) \in O$ for $t = 0, 1, \ldots, s$; and
- $x_t = \{i(x_t), o(x_t)\}$ where $i(x_t) \in N$ and $o(x_t) \in O$ for $t = 1, 2, \ldots, s$.

Since each x_t ($t \in [s]$) belongs to the unique circuit $C(\mathcal{I}, X^*, y_{t-1})$, which is contained in the set of edges incident to $i(y_{t-1})$ by the observation made in (1), we have $i(x_t) = i(y_{t-1})$ for each $t \in [s]$. This means that along the augmenting path P, each agent $i(x_t)$ receives a new item $o(y_{t-1})$ and discards the old item $o(x_t)$.

Now consider the reallocation corresponding to X' where agent $i(x_t)$ receives a new item $o(y_{t-1})$ but loses the item $o(x_t)$ for each $t = 1, 2, \ldots, s$, and agent $i(y_s)$ receives the item $o(y_s)$. Such a reallocation increases the valuation of agent $i(y_s)$ by 1, while it does not decrease the valuations of all the intermediate agents, $i(x_1), i(x_2), \ldots, i(x_s)$, as well as the other agents whose agents do not appear on P. We thus conclude that A is Pareto dominated by the new allocation, a contradiction. □

Theorem 3.10 above, along with Lemma 3.1, implies that both leximin and MNW allocations are utilitarian optimal. Next, we show that for the class of (0, 1)-SUB 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,

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

Examples of symmetric, strictly convex functions are the following: $\Phi(z) \triangleq \sum_{i=1}^n z_i^2$ for $z_i \in \mathbb{Z}\ \forall i \in [n]$ and $\Phi(z) \triangleq \sum_{i=1}^n z_i \ln z_i$ for $z_i \in \mathbb{Z}_{\geq 0}\ \forall i \in [n]$.

We start by showing that given a non-leximin socially optimal allocation A, there exists an adjacent socially optimal allocation Afi 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 3.11. *Suppose that agents have (0,1)–SUB 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*

$$\theta(A') = \theta(A) + \chi_i - \chi_j,$$

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

Proof. Let A be an arbitrary utilitarian optimal allocation which is not leximin, and let A^* be a leximin allocation. Recall that A^* is utilitarian optimal by Theorem 3.10. Without loss of generality, we assume that both A and A^* are clean allocations. Now take a clean allocation A' that minimizes the symmetric difference $\sum_{i \in N} |A'_i \triangle A^*_i|$ over all clean allocation with $\theta(A') = \theta(A)$. Assume also w.l.o.g. that $v_1(A'_1) \leq v_2(A'_2) \leq \cdots \leq v_n(A'_n)$. We let $v_{i_1}(A^*_{i_1}) \leq v_{i_2}(A^*_{i_2}) \leq \cdots \leq v_{i_n}(A^*_{i_n})$. Since $\theta(A')$ is leximin, for the minimum index k with $v_j(A'_k) \neq v_{i_k}(A^*_{i_k})$,

$$v_k(A'_k) < v_{i_k}(A^*_{i_k}). \tag{2}$$

We note that $v_h(A'_h) = v_{i_h}(A^*_{i_h})$ for all $1 \leq h \leq k - 1$. By (2), there exists $i \in [k]$ with

$$v_k(A'_i) < v_k(A^*_i). \tag{3}$$

Indeed, if for all $i \in [k]$, $v_i(A'_i) \geq v_k(A^*_i)$, the k-th smallest value of realized valuations under A' is at least $v_{i_k}(A^*_{i_k})$, contradicting with (2). Take the minimum index i satisfying (3). By minimality, for all

$h \in [i-1]$, $v_h(A'_h) \geq v_h(A^*_h)$; indeed, the equality holds for all $h \in [i-1]$. Otherwise, it would contradict the fact that A^* is leximin. Now, by (I3) of the independent-set matroid axioms, there exists an item $o_1 \in A^*_i \setminus A'_i$ with positive contribution to A'_i, i.e., $v_i(A'_i \cup \{o_1\}) = v_i(A'_i)+1$. By utilitarian optimality of A', $o_1 \in A'_{j_1}$ for some $j_1 \neq i$. If $v_{j_1}(A'_{j_1}) \geq v_i(A'_i)+2$, we obtain a desired allocation by transferring o_1 from j_1 to i. Thus, consider the case when $v_{j_1}(A'_{j_1}) \leq v_i(A'_i)+1$. If $v_{j_1}(A'_{j_1}) = v_i(A'_i)+1$, then by transferring o_1 from j_1 to i, we get another utilitarian optimal allocation with the same vector as $\theta(A')$, which has a smaller symmetric difference than $\sum_{i \in N} |A'_i \triangle A^*_i|$, a contradiction. Thus, $v_{j_1}(A'_{j_1}) \leq v_i(A'_i)$. Since A^* is leximin, $v_{j_1}(A'_{j_1}) \leq v_{j_1}(A^*_{j_1})$. Further by $A'_{j_1} \neq A^*_{j_1}$, there exists an item $o_2 \in A^*_{j_1} \setminus A'_{j_1}$ such that $v_{j_1}(A'_{j_1} \cup \{o_2\} \setminus \{o_1\}) = v_i(A'_{j_1})$. Again by utilitarian optimality of A', $o_2 \in A'_{j_2}$ for some $j_2 \neq j_1$. Repeating the same argument, we obtain a sequence of items and agents $(j_0, o_1, j_1, o_2, j_2, \ldots, o_t, j_t)$ such that

- $v_{j_h}(A'_{j_h}) = v_{j_h}(A'_{j_h} \cup \{o_{h+1}\} \setminus \{o_h\})$ for all $1 \leq h \leq t-1$; and
- $o_h \in A^*_{j_{h-1}} \setminus A'_{j_h}$ for all $1 \leq h \leq t$.

Here $j_0 = i$. If the same agent appears again, i.e., $j_t = j_h$ for some $h < t$, then by transferring items along the cycle, we can decrease the symmetric difference with A^*, a contradiction. Thus, the sequence must terminate when we reach the agent j_t with $v_{j_t}(A'_{j_t}) \geq v_i(A'_i)+2$. Exchanging items along the path, we get a desired allocation. □

We further observe that such adjacent allocation improves the value of any symmetric strictly convex function.

Lemma 3.12. *Let $\Phi : \mathbb{Z}^n \to \mathbb{Z}$ be a symmetric strictly convex function. Let A be a utilitarian optimal allocation. Let A' be another utilitarian optimal allocation such that $\theta(A') = \theta(A) + \chi_i - \chi_j$ for some $i, j \in [n]$ with $\theta(A)_j \geq \theta(A)_i + 2$. Then $\Phi(A) > \Phi(A')$.*

Proof. The proof is similar to that of Proposition 6.1 in Frank and Murota [17], which shows the analogous equivalence over the integral base-polyhedron. Let $\beta = \theta(A)_j - \theta(A)_i \geq 2$, and $\boldsymbol{y} = \theta(A) + \beta(\chi_i - \chi_j)$. Thus $\Phi(\theta(A)) = \Phi(\boldsymbol{y})$ by symmetry of Φ. Define $\lambda = 1 - \frac{1}{\beta}$. We have $0 < \lambda < 1$ since $\beta \geq 2$. Observe that

$$\begin{aligned}\lambda\theta(A) + (1-\lambda)\boldsymbol{y} &= (1 - \frac{1}{\beta})\theta(A) + \frac{1}{\beta}(\theta(A) + \beta(\chi_i - \chi_j)) \\ &= \theta(A) + \chi_i - \chi_j = \theta(A'),\end{aligned}$$

which gives us the following inequality (from the strict convexity of Φ): $\Phi(\theta(A)) = \lambda\Phi(\theta(A)) + (1-\lambda)\Phi(\theta(A)) > \Phi(\theta(A'))$. □

Now we are ready to prove the following.

Theorem 3.13. *Let $\Phi : \mathbb{Z}^n \to \mathbb{R}$ be a symmetric strictly convex function; let A be some allocation. For submodular valuations with binary marginal gains, the following statements 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.

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

Assume towards a contradiction that $\theta(A) \neq \theta(A')$. By Theorem 3.10, we have $\texttt{USW}(A) = \texttt{USW}(A')$. By Lemma 3.11, 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 $\theta(A')$ by 1 and increasing the value of the i-th element of $\theta(A')$ by 1, where $\theta(A')_j \geq \theta(A')_i + 2$. Applying Lemma 3.12, we get $\Phi(\theta(A')) > \Phi(\theta(A''))$, which gives us the desired contradiction.

To prove $2 \Leftrightarrow 3$, let A be a leximin allocation, and let A' be an MNW allocation. Again, we will show that $\theta(A')$ is the same as $\theta(A)$, 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 $\bar{\theta}(A)$ (respectively, $\bar{\theta}(A')$) 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 $\bar{\theta}(A) >_L \bar{\theta}(A')$. 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, $\phi(\boldsymbol{x}) = -\sum_{i=1}^{n'} \log x_i$ is a symmetric convex function for $\boldsymbol{x} \in \mathbb{Z}^n$ with each $x_i > 0$. Thus, by a similar argument as before, one can show that $\phi(\bar{\theta}(A')) < \phi(\bar{\theta}(A))$, a contradiction. This completes the proof. □

The above statement does not generalize to the non-binary case: there is an instance where neither leximin nor MNW allocation is utilitarian optimal. Theorem 3.13 and Corollary 3.8 imply the following result.

Corollary 3.14. *For submodular valuations with binary marginal gains, any clean leximin or MNW allocation is EF1.*

4 ASSIGNMENT VALUATIONS WITH BINARY GAINS

We now consider the special but practically important case when valuations come from maximum matchings. For this class, we show that invoking Theorem 3.10, 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 submodular valuations with binary marginal gains.

Theorem 4.1. *For assignment valuations with binary marginal gains, one can find a leximin or MNW allocation in polynomial time.*

Proof Sketch. The problem of finding a leximin allocation can be reduced to that of finding an integral balanced flow in a network, which has been recently shown to be polynomial-time solvable [18]. Specifically, for a network $D = (V, A)$ with source s, sink t, and a capacity function $c : A \to \mathbb{Z}$, a *balanced flow* is a maximum integral feasible flow where the out-flow vector from the source s to the adjacent vertices h is lexicographically maximized among all maximum integral feasible flows; that is, the smallest flow-value on the edges (s, h) is as large as possible, the second smallest flow-value on the edges (s, h) is as large as possible, and so on. Frank and Murota [18] show that one can find a balanced flow in strongly polynomial time (see Section 7 in Frank and Murota [18]).

Now, given an instance of assignment valuations with binary marginal gains, we build the following instance (V, A) of a network flow problem. Let N_h denote the set of members in each group h. We first create a source s and a sink t. We create a vertex h for each group h, a vertex i for each member i of some group, and a vertex o for each item o. We construct the edges of the network as follows:

- for each group h, create an edge (s, h) with capacity m; and
- for each group h and member i in group h, create an edge (h, i) with unit capacity; and
- for each member i of some group and item o for which i has positive weight u_{io} (i.e. $u_{io} = 1$), create an edge (i, o) with unit capacity; and
- for each item o, create an edge (o, t) with unit capacity.

We can easily verify that an integral balanced flow $f : A \rightarrow \mathbb{Z}$ of this network corresponds to a leximin allocation. Thus, among all utilitarian optimal allocations, A^f lexicographically maximizes the valuation of each group, and hence A_f is a leximin allocation. By Theorem 3.13, the leximin allocation A_f is also MNW. □

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

Theorem 4.2. *For two agents with general assignment valuations, it is NP-hard to compute a leximin or MNW allocation.*

The proof is omitted owing to space constraints and deferred to a full version of the paper. To summarize, the reduction is similar to the hardness reduction for two agents with identical additive valuations [31, 33]: we give a Turing reduction from Partition.

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. In particular, Theorem 3.13 reduces the problem of finding a leximin/MNW allocation for $(0, 1)$-SUB valuations to the minimization of the sum of squared valuations subject to utilitarian optimality, which can be solved efficiently *in practice* using standard solvers for convex programming. We will conclude with additional implications of this work and further research directions.

Other fairness criteria. The fairness concept we consider here is (approximate) envy-freeness. An obvious next step is to explore other criteria such as *proportionality* (each agent gets at least $1/n$ of her valuation of the full collection of goods O), the *maximin share guarantee* or MMS (each agent gets at least as much value as she would realize if allowed to partition O completely among all agents knowing the she would receive her least favorite part), *equitability* (all agents have equal realized valuations), etc. (see, e.g. [12, 19] and references therein for further details) for $(0, 1)$-SUB valuations.

Freeman et al. [19] show that: an allocation that is equitable up to one item or EQ1 (a relaxation of equitability in the same spirit as EF1) and PO may not exist even for binary additive valuations; however, for this valuation class, it can be verified in polynomial time whether an EQ1, 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). We can extend this result to $(0, 1)$-OXS valuations: we first show that any EQ1 and PO allocation under the $(0, 1)$-SUB valuation class, if it exists, is leximin, then invoke Corollary 3.14 to conclude that it must be EF1 and finally Theorem 4.1 to establish its polynomial-time complexity for the $(0, 1)$-OXS class (the full proof is deferred to a full version of the paper).

More general valuation functions. Another imperative line of future work is investigating which of our findings extend to more general submodular valuations (i.e. those with positive real marginal gains). An obvious generalization of the $(0, 1)$-SUB valuation function class is 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 can show that any clean, MNW allocation is still EF1 (clean bundles being defined the same way as for $(0, 1)$-SUB valuations) but the leximin and MNW allocations no longer coincide and leximin no longer implies EF1 (see the full version for details).

For general assignment valuations (i.e. members have positive real weights for items), we have no theoretical guarantees yet. However, we ran experiments on a real-world dataset, comparing the performance of a heuristic extension of Algorithm 1 (Section 3.1) to real-valued individual-item utilities (weights) with Lipton et al. [28]'s envy graph algorithm in terms of the number of items *wasted* (left unassigned or assigned to individuals with zero utility for it although another agent has positive utility for the item). These experiments suggest that approximate envy-freeness can often be achieved in practice simultaneously with good efficiency guarantees even for this larger valuation class.

It is important to note that the class of rank functions of a matroid (equivalently, $(0, 1)$-SUB functions) is a subclass of the *gross substitutes* (GS) valuations [21, 23]. A promising research direction is to investigate PO+EF1 existence for GS valuations.

Implications for diversity. Finally, the analysis of submodular valuations ties in with existing works on diversity in various fields from biology to machine learning (see, e.g. Celis et al. [13], Jost [22]). A popular measurement for how diverse a solution is is to apply one of several concave functions called *diversity indices* to the proportions of the different entities/attributes (with respect to which we wish to be diverse) in the solution, e.g. the Shannon entropy and the Gini-Simpson index: if we denote the maximum USW of one of the problem instances studied in this paper by U^* and agent i's realized valuation in a utilitarian optimal allocation as u_i, then the above two indices can be expressed as $-\sum_{i \in N}(u_i/U^*)\ln(u_i/U^*)$ and $1 - \sum_{i \in N}(u_i/U^*)^2$ respectively such that $\sum_{i \in N} u_i = U^*$. Thus, Theorem 3.13 also shows that, for $(0, 1)$-SUB valuations, the MNW or leximin principle maximizes among all utilitarian optimal allocations commonly used *diversity indices* applied to shares of the agents in the optimal USW. It will be interesting to explore potential connections of this concept to recent work on *soft* diversity framed as convex function optimization [1].

REFERENCES

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

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

[3] Siddharth Barman, Ganesh Ghalme, Shweta Jain, Pooja Kulkarni, and Shivika Narang. 2019. Fair Division of Indivisible Goods Among Strategic Agents. In *Proceedings of the 18th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*. 1811–1813.

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

[5] Siddharth Barman, Sanath Kumar Krishnamurthy, and Rohit Vaish. 2018. Greedy Algorithms for Maximizing Nash Social Welfare. In *Proceedings of the 17th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*. 7–13.

[6] Nawal Benabbou, Mithun Chakraborty, Edith Elkind, and Yair Zick. 2019. Fairness Towards Groups of Agents in the Allocation of Indivisible Items. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*. 95–101.

[7] Nawal Benabbou, Mithun Chakraborty, Vinh Ho Xuan, Jakub Sliwinski, and Yair Zick. 2018. Diversity Constraints in Public Housing Allocation. *Proceedings of the 17th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)* (2018), 973–981.

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

[9] Arpita Biswas and Siddharth Barman. 2018. Fair Division Under Cardinality Constraints. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*. 91–97.

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

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

[12] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. 2016. The unreasonable fairness of maximum Nash welfare. In *Proceedings of the 17th ACM Conference on Economics and Computation (EC)*. ACM, 305–322.

[13] L. Elisa Celis, Amit Deshpande, Tarun Kathuria, and Nisheeth K. Vishnoi. 2016. How to be fair and diverse?. In *3rd, arXiv preprint arXiv:1610.07183*.

[14] John P Dickerson, Karthik Abinav Sankararaman, Aravind Srinivasan, and Pan Xu. 2019. Balancing Relevance and Diversity in Online Bipartite Matching via Submodularity. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*. 1877–1884.

[15] Jack Edmonds. 1979. Matroid Intersection. *Annals of Discrete Mathematics* 4 (1979), 39–49. Discrete Optimization I.

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

[17] András Frank and Kazuo Murota. 2019. Discrete Decreasing Minimization, Part I: Base-polyhedra with Applications in Network Optimization. *arXiv e-prints*, Article arXiv:1808.07600v3 (Jul 2019).

[18] András Frank and Kazuo Murota. 2019. Discrete Decreasing Minimization, Part III: Network Flows. *arXiv e-prints*, Article arXiv:1907.02673v2 (Sep 2019).

[19] Rupert Freeman, Sujoy Sikdar, Rohit Vaish, and Lirong Xia. 2019. Equitable Allocations of Indivisible Goods. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*. 280–286.

[20] Laurent Gourvès and Jérôme Monnot. 2017. Approximate Maximin Share Allocations in Matroids. In *Algorithms and Complexity*, Dimitris Fotakis, Aris Pagourtzis, and Vangelis Th. Paschos (Eds.). Springer International Publishing, 310–321.

[21] Frank Gul and Ennio Stacchetti. 1999. Walrasian equilibrium with Gross Substitutes. *Journal of Economic Theory* 87 (1999), 95–124. Issue 1.

[22] Lou Jost. 2006. Entropy and diversity. *Oikos* 113, 2 (2006), 363–375.

[23] Alexander S. Kelso and Vincent P. Crawford. 1982. Job Matching, Coalition Formation, and Gross Substitutes. *Econometrica* 50, 6 (1982), 1483–1504.

[24] B. Korte and J. Vygen. 2006. *Combinatorial Optimization: Polyhedra and Efficiency*. Springer.

[25] Jérôme Lang and Piotr Krzysztof Skowron. 2016. Multi-Attribute Proportional Representation. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)*. 530–536.

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

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

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

[29] James Munkres. 1957. Algorithms for the assignment and transportation problems. *J. Soc. Indust. Appl. Math.* 5, 1 (1957), 32–38.

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

[31] Trung Thanh Nguyen, Magnus Roos, and Jörg Rothe. 2013. A survey of approximability and inapproximability results for social welfare optimization in multiagent resource allocation. *Annals of Mathematics and Artificial Intelligence* 68, 1 (01 Jul 2013), 65–90.

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

[33] Sara Ramezani and Ulle Endriss. 2010. Nash Social Welfare in Multiagent Resource Allocation. In *Agent-Mediated Electronic Commerce. Designing Trading Strategies and Mechanisms for Electronic Markets*, Esther David, Enrico Gerding, David Sarne, and Onn Shehory (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 117–131.

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

[35] Lloyd S Shapley. 1958. Complements and substitutes in the optimal assignment problem. *ASTIA Document* (1958).

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