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

Nawal Benabbou
Sorbonne Université
nawal.benabbou@lip6.fr

Mithun Chakraborty
National University of Singapore
mithun@comp.nus.edu.sg

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

Yair Zick
National University of Singapore
zick@comp.nus.edu.sg

Abstract

In this paper, we present new results on the fair and efficient allocation of indivisible goods to agents that have monotone, submodular, *non-additive* valuation functions over bundles. In particular, we show that, if such a valuation function additionally has binary marginal gains, a socially optimal (i.e. utilitarian social welfare-maximizing) allocation that achieves envy-freeness up to one item exists and can be computed efficiently. We also prove that the Nash welfare-maximizing and the leximin allocations both exhibit this fairness-efficiency combination, by showing that they can be achieved by minimizing any symmetric strictly convex function over utilitarian optimal outcomes. Moreover, for a subclass of these valuation functions based on maximum (unweighted) bipartite matching, we show that a leximin allocation can be computed in polynomial time.

1 Introduction

How should a bundle of goods be divided amongst a group of agents with subjective valuations? Are there *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; for example, *Pareto-optimality* 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 (Foley 1967). Efficient envy-free allocations are not guaranteed to exist (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) (Budish 2011): 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. (2004) provide an efficient algorithm for computing an 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. (2016) show that allocations that satisfy both approximate envy-freeness and the popular efficiency criterion of Pareto-optimality (PO) exist, specifically the ones that maximize the product of agents' utilities — also known as *max Nash welfare* (MNW). Barman, Krishnamurthy, and Vaish; Barman, Krishnamurthy, and Vaish (2018a; 2018b) show that an allocation with these properties can be computed in (pseudo-)polynomial time.

Indeed, most works on the fair allocation of indivisible goods have focused on the additive setting; at present, little is known about other classes of valuation functions. This is where our work comes in.

Our contributions We focus on monotone submodular valuations with binary marginal gains (referred to as $(0,1)$-SUB valuations). This class of valuations naturally arises in many practical applications. In particular they include assignment valuations with binary marginal gains (called $(0,1)$-OXS valuations), which capture scenarios such as the allocation of public housing or kindergarten slots, where each agent is a group of individuals to whom we wish to match items in a groupwise fair way. For $(0,1)$-SUB valuations, we establish the following existential and computational results on the compatibility of envy-freeness with welfare-based allocation concepts.

- For $(0,1)$-SUB valuations, we show that an EF1 allocation that also maximizes the utilitarian social welfare (hence is Pareto-optimal) always exists and can be computed in polynomial time.
- For $(0,1)$-SUB valuations, we show that leximin and MNW allocations both exhibit the EF1 property.
- 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.
- For $(0,1)$-OXS valuations, we show that both leximin and MNW allocations can be computed efficiently.

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

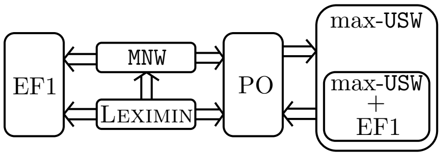

Figure 1: Equivalences and implications among properties for the $(0, 1)$-SUB valuation class.

non-wasteful (also called clean) bundles forms the set of independent sets of a matroid. While some papers explore the application of matroid theory to the fair division problem (Biswas and Barman 2018; Gourvès and Monnot 2017), we believe that ours is the first to demonstrate its strong connection with fairness and efficiency guarantees. Figure 1 describes the relation between properties for $(0, 1)$-SUB valuations. Our computational results are summarized in Table 1.

Related work Our paper is related to the active area of research on the fair allocation of indivisible goods. Budish (2011) was the first to formalize the notion of EF1; but, it implicitly appeared in Lipton et al. (2004), who designed a polynomial time algorithm that returns an EF1 allocation for arbitrary monotone valuations (called the *envy graph algorithm*). Caragiannis et al. (2016) proved the seminal existence result of EF1 and Pareto-optimal allocations for non-negative additive valuations; Barman, Krishnamurthy, and Vaish (2018a) subsequently provided a pseudo-polynomial time algorithm for computing such allocations. Closely related to ours is the work of Barman, Krishnamurthy, and Vaish (2018b), who developed 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 (since bundles of the same size may have different values under the latter class); hence their polynomial-time complexity result does not imply our Theorem 4.1.

One motivation for our paper is recent work by Benabbou et al. (2019) on promoting diversity in assignment problems through efficient, EF1 allocations of bundles to attribute-based groups in the population. Other work in this vein includes fairness/diversity through quotas (Aziz et al. 2019; Suzuki, Tamura, and Yokoo 2018; Benabbou et al. 2018, and references therein), or by the optimization of carefully constructed functions (Dickerson et al. 2019; Ahmed, Dickerson, and Fuge 2017; Lang and Skowron 2016, 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 $v_i(S) \subseteq v_i(T)$ 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.

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 (2011): 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 (Sen 1970). 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, 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

	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.

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 (1958) and is identical to the OXS valuation class (Lehmann, Lehmann, and Nisan 2006). Fair allocation in this setting was explored by Benabbou et al. (2019). 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 (Munkres 1957); 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 (Barman, Krishnamurthy, and Vaish 2018b) and of assignment valuations where the weight is binary (Benabbou et al. 2019). 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 $v(S) \leq |S|$).*

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. (2016) 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.[1] We show that the "unreasonable fairness" of maximum Nash welfare (Caragiannis et al. 2016) 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

[1] Indeed, Caragiannis et al. (2016) Example C.3 is an instance with submodular valuations where MNW does not imply EF1.

$o \in A_j$. Let $A_j = \{o_1, o_2, \ldots, o_r\}$. Define $S_0 = \emptyset$ and $S_t = \{o_1, o_2, \ldots, o_t\}$ 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, as Example A.1 in the supplementary material illustrates.

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 will now 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 $X \subseteq Y \in \mathcal{I}$, 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 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 $\{o, i\} \in X$. 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 $\mathcal{I}_i = \{\, X \subseteq E \mid r_i(X) = |X| \,\}$, is the set of independent sets of a matroid. Taking the union $\mathcal{I} = \mathcal{I}_1 \cup \cdots \cup \mathcal{I}_n$, the pair $(E, \mathcal{I})$ is known to form a matroid (Korte and Vygen 2006), 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* (Korte and Vygen 2006).

Now, observe that a common independent set of the two matroids $X \in \mathcal{O} \cap \mathcal{I}$ corresponds to a clean allocation A of our original instance where each agent i receives the items o with $\{o, i\} \in X$; indeed, each item o is allocated at most once because $|E_o \cap X| \leq 1$, and each A_i is clean because the realized valuation of agent i under A is exactly the size of the allocated bundle. Conversely, any clean allocation A of our instance corresponds to an independent set $X = \bigcup_{i \in N} X_i \in \mathcal{I} \cap \mathcal{O}$, where $X_i = \{\, \{o, i\} \mid o \in A_i \,\}$: for each $i \in N$, $r_i(X_i) = |X_i|$ by Proposition 3.3, 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 (Edmonds 1979). 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 $O(|E|^3 mnF)$. □

We are now ready to prove Theorem 3.4.

Proof. 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.

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

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 $\phi(A) := \sum_{i \in N} v_i(A_i)^2$.

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

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

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 runs in polynomial time. □

An interesting implication Algorithm 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 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. (2019)'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 A.2 in the supplementary material, 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 note, 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 outcomes coincide with utilitarian optimal outcomes 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 does not make any other agent worse off.

In the subsequent proof, we will use the following notions and results in 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 2006))**.** *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 $X \in \mathcal{I}$, 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, $\mathcal{I}_i$ 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 and 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., $C^{\mathcal{I}}(X, y) = C^{\mathcal{I}_i}(X, y)$.

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, which 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 (Edmonds 1979).

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 edges incident to $i(y_{t-1})$ by the observation made before, 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,

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

Theorem 3.11. *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. The proof is similar to that of Proposition 6.1 in Frank and Murota (2019), which shows the analogous equivalence over the integral base-polyhedron.

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 $\theta(A) \neq \theta(A')$ and hence $\theta(A) >_L \theta(A')$. By Theorem 3.10, we have $\texttt{USW}(A) = \texttt{USW}(A')$. Consider the set X of vectors $\boldsymbol{x} \in \mathbb{Z}^n$ with $x_1 \leq x_2 \leq \ldots \leq x_n$ where $\sum_{i=1}^n x_i = \alpha$ and

$$\theta(A') \leq_L \boldsymbol{x} \leq_L \theta(A).$$

Order the distinct elements $\boldsymbol{x}^{(1)}, \boldsymbol{x}^{(2)}, \ldots, \boldsymbol{x}^{(s)}$ in X in lexicographically increasing order such that

$$\boldsymbol{x}^{(1)} <_L \boldsymbol{x}^{(2)} <_L \cdots <_L \boldsymbol{x}^{(s)}.$$

So there is no $\boldsymbol{y} \in X$ such that $\boldsymbol{x}^{(t)} <_L \boldsymbol{y} <_L \boldsymbol{x}^{(t+1)}$ for each $t \in [s]$. Clearly, $\boldsymbol{x}^{(1)} = \theta(A')$ and $\boldsymbol{x}^{(s)} = \theta(A)$.

We will prove that $\Phi(\boldsymbol{x}^{(t)}) > \Phi(\boldsymbol{x}^{(t+1)})$ for each $t \in [s-1]$, which implies that $\Phi(\theta(A')) > \Phi(\theta(A))$ and gives us the desired contradiction. Now fix any $t \in [s-1]$. Consider the vectors $\boldsymbol{x}^{(t)}$ and $\boldsymbol{x}^{(t+1)}$. It is not difficult to see that $\boldsymbol{x}^{(t+1)}$ can be obtained from $\boldsymbol{x}^{(t)}$ by decreasing the value of the j-th element by 1 and increasing the value of the i-th element by 1, where $x_j^{(t)} \geq x_i^{(t)} + 2$. Namely,

$$\boldsymbol{x}^{(t+1)} = \boldsymbol{x}^{(t)} + \chi_i - \chi_j.$$

Let $\beta = x_j^{(t)} - x_i^{(t)} \geq 2$ and $\boldsymbol{y} = \boldsymbol{x}^{(t)} + \beta(\chi_i - \chi_j)$; the vector $\boldsymbol{y}$ only exchanges the i-th and j-th elements of $\boldsymbol{x}^{(t)}$, and thus $\Phi(\boldsymbol{x}^{(t)}) = \Phi(\boldsymbol{y})$ by symmetry of Φ. Define $\lambda = 1 - \frac{1}{\beta}$. We have $0 < \lambda < 1$ since $\beta \geq 2$. Observe that

$$\begin{aligned}\lambda \boldsymbol{x}^{(t)} + (1-\lambda)\boldsymbol{y} &= (1 - \frac{1}{\beta})\boldsymbol{x}^{(t)} + \frac{1}{\beta}(\boldsymbol{x}^{(t)} + \beta(\chi_i - \chi_j)) \\ &= \boldsymbol{x}^{(t)} + \chi_i - \chi_j = \boldsymbol{x}^{(t+1)},\end{aligned}$$

which gives us the following inequality (from the strict convexity of Φ):

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

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; see Example A.3 in the supplementary material.

Corollary 3.12. *For submodular valuations with binary marginal gains, any leximin or MNW allocation is EF1.*

Proof. Since both leximin and MNW allocations are Pareto-optimal, they maximize the utilitarian social welfare by Theorem 3.10. By Theorem 3.11 and the fact that the function

$\phi(A) = \sum_{i \in N} v_i(A_i)^2$ 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, and hence is EF1 by Corollary 3.8. □

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 3.10, 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. The problem of finding a leximin allocation can be reduced to the problem of finding an integral balanced flow in a network, which has been recently shown to be polynomial-time solvable (Frank and Murota 2019). 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 (2019) show that one can find a balanced flow in strongly polynomial time (see Section 7 in Frank and Murota (2019)).

Now, given an instance of assignment valuations with binary marginal gains, we build the following instance (V, A) of a network flow problem. We 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.

See Figure 2 in the supplementary material for an illustration of the network. We will show that an integral balanced flow $f : A \to \mathbb{Z}$ of the constructed network corresponds to a leximin allocation. Consider an allocation A^f where each group receives the items o for which some member i of the group has positive flow $f(i, o) > 0$. It is easy to see that the allocation A^f maximizes the utilitarian social welfare since the flow f is a maximum integral feasible flow. Thus, by Theorem 3.10, A^f has the same utilitarian social welfare as any leximin allocation. To see balancedness, observe that the amount of flow from the source s to each group h is the valuation of h for bundle A^f_h, i.e., $f(s, h) = \sum_{i \in N_h} f(h, i) = v_h(A^f_h)$. Indeed if $v_h(A^f_h) > f(s, h)$, then it would contradict the optimality of the flow f; and if $v_h(A^f_h) < f(s, h)$, it would contradict the fact that $v_h(A^f_h)$ is the value of a maximum-size matching between A^f_h and N_h. 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 3.11, 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.*

5 Discussion

We have studied allocations of indivisible goods under submodular valuations with binary marginal gains in terms of envy, efficiency, and various welfare concepts. Our work presents an interesting relation between fairness and efficiency: 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.

Future work includes investigating which of our findings extend to more general submodular valuations. Indeed, while we have not obtained any theoretical guarantees for the non-binary case, our experiments on real-world data (see Appendix B) reveal that approximate envy-freeness can often be achieved simultaneously with good efficiency guarantees.

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 (Gul and Stacchetti 1999; Kelso and Crawford 1982). A promising research direction is to investigate the co-existence of Pareto-optimality and approximate fairness for GS valuations.

Finally, the analysis of submodular valuations ties in with existing works on diversity in allocation (see our related work section). $(0, 1)$-OXS valuations with *budgets* (a capacity on each agent's bundle size) correspond nicely to the public housing allocation problem studied by Benabbou et al. (2018). Adding these caps results in $(0, 1)$-SUB valuations. That several fairness/justice criteria still hold under item quotas bodes well for ensuring fair allocation of goods under hard diversity constraints (e.g. Singapore's ethnic integration policy in public housing). Moreover, Theorem 3.11 shows that, for $(0, 1)$-SUB valuations, the MNW or leximin principle maximizes commonly used *diversity indices* such as the Shannon entropy and the Gini-Simpson index (see e.g. Jost (2006)) applied to shares of the agents in the optimal utilitarian social welfare — it will be interesting to explore connections of this concept to recent work on *soft* diversity framed as convex function optimization (Ahmed, Dickerson, and Fuge 2017).

References

Ahmed, F.; Dickerson, J. P.; and Fuge, M. 2017. Diverse weighted bipartite b-matching. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, 35–41.

Aziz, H.; Gaspers, S.; Sun, Z.; and Walsh, T. 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.

Barman, S.; Krishnamurthy, S. K.; and Vaish, R. 2018a. Finding fair and efficient allocations. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*, 557–574. ACM.

Barman, S.; Krishnamurthy, S. K.; and Vaish, R. 2018b. Greedy algorithms for maximizing Nash social welfare. In *Proceedings of the 17th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 7–13.

Benabbou, N.; Chakraborty, M.; Xuan, V. H.; Sliwinski, J.; and Zick, Y. 2018. Diversity constraints in public housing allocation. *Proceedings of the 17th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)* 973–981.

Benabbou, N.; Chakraborty, M.; Elkind, E.; and Zick, Y. 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.

Biswas, A., and Barman, S. 2018. Fair division under cardinality constraints. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*, 91–97.

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

Caragiannis, I.; Kurokawa, D.; Moulin, H.; Procaccia, A. D.; Shah, N.; and Wang, J. 2016. The unreasonable fairness of maximum Nash welfare. In *Proceedings of the 17th ACM Conference on Economics and Computation (EC)*, 305–322. ACM.

Dickerson, J. P.; Sankararaman, K. A.; Srinivasan, A.; and Xu, P. 2019. Balancing relevance and diversity in online bipartite matching via submodularity. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, 1877–1884.

Edmonds, J. 1979. Matroid intersection. *Annals of Discrete Mathematics* 4:39–49. Discrete Optimization I.

Foley, D. 1967. Resource allocation and the public sector. *Yale Economics Essays 7* 45–98.

Frank, A., and Murota, K. 2019. Discrete Decreasing Minimization, Part III: Network Flows. *arXiv e-prints* arXiv:1907.02673.

Freeman, R.; Sikdar, S.; Vaish, R.; and Xia, L. 2019. Equitable allocations of indivisible goods. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, 280–286.

Gourvès, L., and Monnot, J. 2017. Approximate maximin share allocations in matroids. In Fotakis, D.; Pagourtzis, A.; and Paschos, V. T., eds., *Algorithms and Complexity*, 310–321. Springer International Publishing.

Gul, F., and Stacchetti, E. 1999. Walrasian equilibrium with gross substitutes. *Journal of Economic Theory* 87:95–124.

Harper, F. M., and Konstan., J. A. 2015. The movielens datasets: History and context. *ACM Transactions on Interactive Intelligent Systems (TiiS) 5, 4, Article 19 (December 2015), 19 pages.*

Jost, L. 2006. Entropy and diversity. *Oikos* 113(2):363–375.

Kelso, A. S., and Crawford, V. P. 1982. Job matching, coalition formation, and gross substitutes. *Econometrica* 50(6):1483–1504.

Korte, B., and Vygen, J. 2006. *Combinatorial Optimization: Polyhedra and Efficiency*. Algorithms and Combinatorics. Springer.

Lang, J., and Skowron, P. K. 2016. Multi-attribute proportional representation. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)*, 530–536.

Lehmann, B.; Lehmann, D.; and Nisan, N. 2006. Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behavior* 55(2):270–296.

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

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

Nguyen, T. T.; Roos, M.; and Rothe, J. 2013. 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.

Ramezani, S., and Endriss, U. 2010. Nash social welfare in multiagent resource allocation. In David, E.; Gerding, E.; Sarne, D.; and Shehory, O., eds., *Agent-Mediated Electronic Commerce. Designing Trading Strategies and Mechanisms for Electronic Markets*, 117–131. Berlin, Heidelberg: Springer Berlin Heidelberg.

Sen, A. 1970. *Collective choice and social welfare*. San Francisco: Holden Day.

Shapley, L. S. 1958. Complements and substitutes in the optimal assignment problem. *ASTIA Document*.

Suzuki, T.; Tamura, A.; and Yokoo, M. 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.

A Supplementary Material

A.1 Example of non-submodular valuation functions violate the transferability property

Example A.1. Agent 1 wants to have a pair of matching shoes; her current allocated bundle is a single red shoe, whereas agent 2 has a matching pair of blue shoes. Agent 1 clearly envies agent 2, but cannot increase the value of her bundle by taking any one of agent 2's items. More formally, suppose $N = [2]$ and $O = \{r_L, b_L, b_R\}$; agent 1's valuation function is: $v_1(S) = 1$ only if $\{b_L, b_R\} \subseteq S$, $v_1(S) = 0$ otherwise. Under the allocation $A_1 = \{r_L\}$ and $A_2 = \{b_L, b_R\}$, $v_1(A_1) < v_1(A_2)$ but $\Delta_1(A_1; o) = 0$ for all $o \in A_2$.

A.2 Example where the output of Algorithm 1 may not be either MNW or leximin

Example A.2. We introduce the same instance as Example 2 in Benabbou et al. (2019). There are two groups and six items $o_1, o_2, o_3, o_4, o_5, o_6$. The first group N_1 contains four agents a_1, a_2, a_3, a_4 and the second group N_2 contains four agents b_1, b_2, b_3, b_4; each individual has utility 1 for an item o if and only if she is adjacent to o in the following graph:

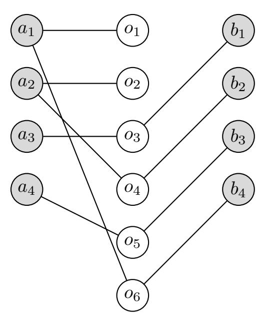

The valuation functions of the groups for each bundle X are defined as the value of a maximum-size matching of X to its members. The algorithm may initially compute an utilitarian optimal allocation A that assigns items o_1, o_2 to the first group, and the remaining items to the second group; the allocation A also satisfies EF1, so the output of Algorithm 1 is A. However, the (unique) leximin and MNW allocation assigns items o_1, o_2, o_3 to the first group, and the remaining items to the second group – this is also the (unique) utilitarian optimal allocation with the minimum sum of squares of the agents' valuations.

A.3 Example where neither leximin nor MNW allocation is utilitarian optimal

Example A.3. Consider an instance with assignment valuations given as follows. Suppose there are three groups, each of which contains a single agent, Alice, Bob, and Charlie, respectively, and three items with weights given in Table 2. The unique leximin and MNW allocation is the allocation that assigns Alice to the first item, Bob to the second item, and Charlie to the third item; each agent has positive utility at the allocation and the total utilitarian social welfare is 3.1. However, the utilitarian optimal allocation assigns Alice to nothing, Bob to the first item, and Charlie to the second item, which yields the total utilitarian social welfare 4.9.

	①	②	③
Alice:	2	1	0
Bob:	2	1	0
Charlie:	0	2.9	0.1

Table 2: An instance where neither leximin nor MNW allocation is utilitarian optimal.

A.4 The network flow instance constructed in the proof of Theorem 4.1.

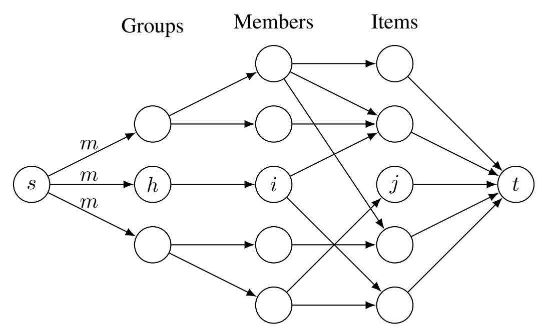

Figure 2: An illustrative network flow instance constructed in the proof of Theorem 4.1: each edge is either labeled with its capacity or has unit capacity.

A.5 Proof of Theorem 4.2

Proof. The reduction is similar to the hardness reduction for two agents with identical additive valuations (Nguyen, Roos, and Rothe 2013; Ramezani and Endriss 2010). We give a Turing reduction from PARTITION. Recall that an instance of PARTITION is given by a set of positive integers $W = \{w_1, w_2, \ldots, w_m\}$; it is a 'yes'-instance if and only if it can be partitioned into two subsets S_1 and S_2 of W such that the sum of the numbers in S_1 equals the sum of the numbers in S_2.

Consider an instance of PARTITION $W = \{w_1, w_2, \ldots, w_m\}$. We create m items $1, 2, \ldots, m$, two groups 1 and 2, and m individuals for each group where every individual has a weight w_j for item j. Observe that fore each group, the value of each bundle X is the sum $\sum_{w_j \in X} w_j$: the number of members in the group exceeds the number of items in X, and thus one can fully assign each item to each member of the group.

Suppose we had an algorithm which finds a leximin allocation. Run the algorithm on the allocation problem constructed above to obtain a leximin allocation A. It can be easily verified that the instance of PARTITION has a solution if and only if $v_1(A_1) = v_2(A_2)$. Similarly, suppose we had an algorithm which finds an MNW allocation, and run the algorithm to find an MNW allocation A'. Since the valuations are identical, the utilitarian social welfare of the MNW allocation is the sum $\sum_{w_j \in W} w_j$, which means that the product of the valuations is maximized when both groups have the same realized valuation. Thus, the instance of PARTITION has a solution if and only if $v_1(A'_1) = v_2(A'_2)$. □

B General assignment valuations

In this section, we introduce an extension of the envy-induced transfers (Algorithm 1) to general valuations, as described in Algorithm 2. Recall that the general principle is to iteratively eliminate envy by transfering an item from the envied bundle to the envied bundle until no agent is envied up to more than one item. Here at each iteration step, we transfer the item that induces the maximal increase in overall utility (see lines 3-6). Note that, since agent j loses one of her items, she may develop a positive marginal utility for a currently withheld item; in that case, the item in A_0 for she has maximal marginal utility is given to her (see lines 6-8). Note also that, for arbitrary utilities, if an agent i acquires a new item o due to an envy-induced transfer, at most one of her previous items, say o^*, may become unused, e.g. if i's positive marginal utility for o with her previous bundle A_i was due to the fact that the individual who was assigned item o^* has a higher utility for o than o^* and no other indivual prefers o^* to its assigned item. In that case, item o^* is revoked from agent i and allocated to the agent with maximal and strictly positive marginal utility for it (see lines 10-13). If this creates another unused item, we repeat the process until there are no unused items or the unused item has zero marginal utility for all agents – in the latter case, the unused item is added to the withheld set (see lines 14-18).

We now provide numerical results to compare the loss incurred by Algorithm 2 and that of P_L the extension of (Lipton et al. 2004) algorithm to groups of agents, as presented in (Benabbou et al. 2019). The performances of these procedure are estimated both in terms of percentage of items wasted (Waste) and price of fairness (PoF) which is formally defined as follows:

$$PoF(P) = \frac{\max\{\texttt{USW}(A) \mid A \text{ is an allocation}\}}{\texttt{USW}(A(P))}$$

where $A(P)$ is the allocation returned by a given procedure P. In our experiments, we use the dataset called "MovieLens-ml-1m" which contains approximately 1,000,000 ratings (from 0 to 5) of 4,000 movies made by 6,000 users (Harper and Konstan. 2015). To generate an instance of our allocation problem, we select 200 movies at random ($|O| = 200$) and then we only consider the users that rated at least one of these movies. We use the user ratings as individual utilities. We consider the following user partitions:

- Gender: 2 agents (male or female),
- Age: 7 agents representing the 7 age-groups,

and the following agents' valuations:

- Ratings:

$$v_i(S) = \max\{ \sum_{u \in N_i} r_{u,\pi(u)} \mid \pi \in \Pi(N_i, S) \}$$

where r_{uo} is the rating of movie o made by user u,

- Norm: $v'_i(S) = v_i(S)/v_i(O)$,

for all agents $i \in N$ representing a group of users and for all bundles of movies $S \subseteq O$. The results given in Table 3 are averaged over 50 runs.

In Table 3, we observe that P_L the extension of (Lipton et al. 2004) algorithm to groups may be wasteful but in practice it has almost zero-waste and also good PoF (the lower the better). In comparison, Algorithm 2 the extension of the envy-induced transfers to general valuations is guaranteed to be non-wasteful but in practice it has larger PoF than P_L.

Algorithm 2: Envy-Induced Transfers for general valuations

1 Compute a clean, socially optimal allocation.
2 **/**Envy-Induced Transfers and Reallocations**/**
3 **while** $\exists i, j \in N$ *such that* i *envies* j *up to more than* 1 *item* **do**
4   Pick i, j, o maximizing $\Delta_i(A_i; o) + \Delta(A_j; o)$ over all $i, j \in N$ and all $o \in O$ such that i envies j more than 1 item and $\Delta_i(A_i; o) > 0$.
5   $A_j \leftarrow A_j \backslash \{o\}$; $A_i \leftarrow A_i \cup \{o\}$.
6   **if** $\exists o \in A_0$ *such that* $\Delta_j(A_j; o) > 0$ **then**
7   Pick $o \in A_0$ that maximizes $\Delta_j(A_j; o)$.
8     $A_j \leftarrow A_j \cup \{o\}$.
9   **end**
10   **if** $\exists o^* \in A_i$ *that is unused* **then**
11     $A_i \leftarrow A_i \backslash \{o^*\}$; revoked = **true**.
12   **while** *revoked* = ***true and*** $\exists k$ *s.t.* $\Delta(A_k; o^*) > 0$ **do**
13   Allocate o^* to agent k maximizing $\Delta(A_k; o^*)$.
14   **if** $\exists o \in A_k$ *that is unused* **then**
15         $A_k \leftarrow A_k \backslash \{o\}$; $o^* \leftarrow o$.
16   **else** revoked = **false**.
17   **end**
18   **if** *revoked* = ***true*** **then** $A_0 \leftarrow A_0 \cup \{o^*\}$.
19   **end**
20 **end**

		P_L		Algorithm 2	
	Types	Ratings	Norm	Ratings	Norm
PoF	Age	1.01	1.15	1.05	1.19
Waste	Age	1.25	0.20	0.00	0.00
PoF	Gender	1.00	1.02	1.00	1.03
Waste	Gender	0.00	0.00	0.00	0.00

Table 3: Performances of allocation procedures.

C Submodularity with subjective binary gains

An obvious generalization of the $(0, 1)$-SUB valuation function class is the class of submodular valuation functions $v(\cdot)$ with *subjective* binary marginal gains: agent i's bundle-valuation function $v_i(\cdot)$ is said to have subjective binary marginal gains if $\Delta_i(S; o) \in \{0, \lambda_i\}$ for some agent-specific constant $\lambda_i > 0$, for every $i \in N$. We define clean bundles and clean allocations for this function class exactly as we did for $(0, 1)$-SUB valuations in Section 2.

Understandably, most of the properties of allocations under $(0, 1)$-SUB valuations do not extend to this more general setting. It is obvious that Pareto-optimality does not imply utilitarian optimality (e.g. consider an instance with two agents and one item which the agents value at 1 and 2 respectively: assigning the item to agent 1 is PO but not utilitarian optimal). Moreover, the leximin allocation may not be EF1, as shown by the following example where both agents have additive valuations.

Example C.1. Suppose $N = [2]$; $O = \{o_1, o_2, o_3, o_4\}$; the valuations are additive with $v_1(\{o_1\}) = 0$, $v_1(\{o\}) = 1$ $\forall o \in O \setminus \{o_1\}$, and $v_2(\{o\}) = 3$ $\forall o \in O$. It is straightforward to check that the unique leximin allocation is $A_1 = \{o_1, o_2, o_3\}$, $A_2 = \{o_4\}$. Under this allocation, $v_1(A_2) = 0 < 3 = v_1(A_1)$, but $v_2(A_1 \setminus \{o\}) = 6 > 3 = v_2(A_2)$ for every $o \in A_1$ — in fact, at least two (any two) items must be removed from A_1 for agent 2 to stop envying agent 1.

Note another difference of this valuation class from $(0, 1)$-SUB that is also evidenced by Example C.1: the leximin and MNW allocations may not coincide. In this example, any allocation A that gives two of the items $\{o_2, o_3, o_4\}$ to agent 1 and the rest to agent 2 is MNW, with $v_1(A_1) = 2$ and $v_2(A_2) = 6$, so that $\texttt{NW}(A) = 12$; such an allocation is also EF1 (in fact, envy-free) since $v_1(A_2) = 1 < 2 = v_1(A_1)$ and $v_2(A_1) = 6 = v_2(A_2)$. This is not an accident, as the following theorem shows.

Theorem C.2. *For agents having submodular valuation functions $v(\cdot)$ with subjective binary marginal gains, any clean, MNW allocation is EF1.*

Since our valuation functions are still submodular, the transferability property (Lemma 3.5) still holds. Two other components of the proof of Theorem C.2 are natural extensions of Propositions 3.3 and Lemma 3.6 — Proposition C.3 and Lemma C.4 below, respectively:

Proposition C.3. *For submodular valuations with subjective binary marginal gains defined by agent-specific positive constants λ_i $\forall i \in N$, A is a clean allocation if and only if $v_i(A_i) = \lambda_i |A_i|$ for each $i \in N$.*

Proof. Consider an arbitrary bundle $S \subseteq O$ such that $S = \{o_1, o_2, \ldots, o_r\}$ for some $r \in [m]$ w.l.o.g. Let $S_0 = \emptyset$ and $S_t = S_{t-1} \cup \{o_t\}$ for every $t \in [r]$. Then, an arbitrary agent i's valuation of bundle S under marginal gains in $\{0, \lambda_i\}$ is

$$v_i(S) = \sum_{t=1}^{r} \Delta_i(S_{t-1}; o_t) \leq \sum_{t=1}^{r} \lambda_i = \lambda_i r = \lambda_i |S|. \quad (1)$$

Now, if agent i's allocated bundle under an allocation A has a valuation $v_i(A_i) = \lambda_i |A_i|$, then her marginal gain for any item in $o \in A_i$ is given by

$$\begin{aligned} v_i(A_i) - v_i(A_i \setminus \{o\}) &= \lambda_i |A_i| - v_i(A_i \setminus \{o\}) \\ &\geq \lambda_i |A_i| - \lambda_i (|A_i| - 1) \\ &= \lambda_i > 0, \end{aligned}$$

where the first inequality follows from Inequality (1) and the fact that $|A_i \setminus \{o\}| = |A_i| - 1$. This means that the bundle A_i is clean and, since this holds for every i, the allocation is clean. This completes the proof of the "if" part.

If allocation A is clean, then we must have $\Delta_i(A_i \setminus \{o\}; o) > 0$ for every $o \in A_i$ for every $i \in N$. Let us define an arbitrary agent i's bundle A_i as S above, so that $|A_i| = r$. Then, since $S_{t-1} \subseteq A_i \setminus \{o_t\}$ for every $t \in [r]$, submodularity dictates that

$$\Delta_i(S_{t-1}; o_t) \geq \Delta_i(A_i \setminus \{o_t\}; o_t) > 0 \quad \forall t \in [r].$$

Since $\Delta_i(S_{t-1}; o_t) \in \{0, \lambda_i\}$ with $\lambda_i > 0$, the above inequality implies that $\Delta_i(S_{t-1}; o_t) = \lambda_i$ $\forall t \in [r]$. Hence,

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

This completes the proof of the "only if" part. □

Lemma C.4. *For submodular functions with subjective binary marginal gains, if agent i envies agent j up to more than 1 item under clean allocation A, then $|A_j| \geq |A_i| + 2$.*

Proof. Since i envies j under A up to more than 1 item, we must have $A_j \neq \emptyset$ and $v_i(A_i) < v_i(A_j \setminus \{o\})$ for every $o \in A_j$. Consider one such o. From Inequality (1) in the proof of Proposition C.3, $v_i(A_j \setminus \{o\}) \leq \lambda_i |A_j \setminus \{o\}| = \lambda_i (|A_j| - 1)$. Since A is clean, $v_i(A_i) = \lambda_i |A_i|$. Combining these, we get

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

Since $\lambda_i > 0$, we have $|A_i| < |A_j| - 1$, i.e. $|A_i| \leq |A_j| - 2$ because $|A_i|$ and $|A_j|$ are integers. □

We are now ready to prove Theorem C.2.

Proof of Theorem C.2. Our proof non-trivially extends that of Theorem 3.2 of Caragiannis et al. (2016). We will first address the case when it is possible to allocate items in such a way that each agent has a positive realized valuation for its bundle, i.e. $N_{\max} = N$ in the definition of an MNW allocation, and then tackle the scenario $N_{\max} \subsetneq N$.

Consider a pair of agents $1, 2 \in N$ w.l.o.g. such that 1 envies 2 up to two or more items, if possible, under an MNW allocation A. Since every agent has a positive realized valuation under A, we have $v_i(A_i) = \lambda_i |A_i| > 0$, i.e. $|A_i| > 0$ for each $i \in \{1, 2\}$. From Lemma 3.5, we know that there is an item in A_2 for which agent 1 has positive marginal utility – consider any one such item $o \in A_2$. Thus, $\Delta_1(A_1; o) > 0$, i.e. $\Delta_1(A_1; o) = \lambda_1$; also, since A_2 is a clean bundle, $\Delta_1(A_2 \setminus \{o\}; o) > 0$, i.e. $\Delta_1(A_2 \setminus \{o\}; o) = \lambda_2$.

Let us convert A to a new allocation A' by only transferring this item o from agent 2 to agent 1. Hence, $v_1(A'_1) = v_1(A_1) + \Delta_1(A_1; o) = v_1(A_1) + \lambda_1$, $v_2(A'_2) = v_2(A_2) - \Delta_1(A_2 \setminus \{o\}; o) = v_2(A_2) - \lambda_2$, $v_i(A'_i) = v_i(A_i)$ for each $i \in N \setminus \{1, 2\}$. $\texttt{NW}(A)$ is positive since A is MNW and

$N_{\max} = N$. Hence,

$$\begin{aligned}
\frac{\texttt{NW}(A')}{\texttt{NW}(A)} &= \left[\frac{v_1(A_1)+\lambda_1}{v_1(A_1)}\right]\left[\frac{v_2(A_2)-\lambda_2}{v_2(A_2)}\right] \\
&= \left[1+\frac{\lambda_1}{v_1(A_1)}\right]\left[1-\frac{\lambda_2}{v_2(A_2))}\right] \\
&= \left[1+\frac{\lambda_1}{\lambda_1|A_1|}\right]\left[1-\frac{\lambda_2}{\lambda_2|A_2|}\right] \\
&= \left[1+\frac{1}{|A_1|}\right]\left[1-\frac{1}{|A_2|}\right] \\
&= 1+\frac{|A_2|-|A_1|-1}{|A_1||A_2|}, \\
&\geq 1+\frac{(|A_1|+2)-|A_1|-1}{|A_1||A_2|}, \\
&\geq 1+\frac{1}{|A_1||A_2|}, \\
&> 1.
\end{aligned}$$

Here, the third equality comes from Proposition C.3 since A is clean, and the first inequality from Lemma C.4 due to our assumption. But $\texttt{NW}(A') > \texttt{NW}(A)$ contradicts the optimality of A, implying that any agent can envy another up to at most 1 item under A.

This completes the proof for the $N_{\max} = N$ case. The rest of the proof mirrors the corresponding part of the proof of Caragiannis et al. (2016)'s Theorem 3.2. If $N_{\max} \subsetneq N$, it is easy to see that there can be no envy towards any $i \notin N_{\max}$: this is because we must have $v_i(A_i) = 0$ for any such i from the definition of $N_{\max}$, which in turn implies that $A_i = \emptyset$ since A is clean; hence, $v_j(A_i) = 0$ for every $j \in N$. Also, for any $i, j \in N_{\max}$, we can show exactly as in the proof for the $N_{\max} = N$ case above that there cannot be envy up to more than one item between them, since A maximizes the Nash welfare over this subset of agents $N_{\max}$. Suppose for contradiction that an agent $i \in N \backslash N_{\max}$ envies some $j \in N_{\max}$ up to more than one item under A. Then, from Lemma 3.5, there is one item $o_1 \in A_j$ w.l.o.g. such that $v_i(\{o_1\}) = \Delta_i(\emptyset; o_1) = \Delta_i(A_i; o_1) > 0$. Moreover, since A is clean,

$$\begin{aligned}
v_j(A_j \setminus \{o_1\}) &= v_j(A_j) - \Delta_j(A_j \setminus \{o_1\}; o_1) \\
&= \lambda_j|A_j| - \lambda_j \\
&= \lambda_j(|A_j|-1) \\
&\geq \lambda_j(|A_i|+1) \\
&= \lambda_j > 0,
\end{aligned}$$

where the first inequality comes from Lemma C.4. Thus, if we transfer o_1 from j to i and leave all other bundles unchanged, then every agent in $N_{\max} \cup \{i\}$ will have a positive valuation under the new allocation. This contradicts the maximality of $N_{\max}$. Hence, any $i \in N \backslash N_{\max}$ must be envy-free up to one item towards any $j \in N_{\max}$. □

D Implications for approximate equitability

An allocation A is said to be *equitable* or EQ if the realized valuations of all agents are equal under it, i.e. for every pair of agents $i, j \in N$, $v_i(A_i) = v_j(A_j)$; an allocation A is *equitable up to one item* or EQ1 if, for every pair of agents $i, j \in N$ such that $A_j \neq \emptyset$, there exists some item $o \in A_j$ such that $v_i(A_i) \geq v_j(A_j \setminus \{o\})$ (Freeman et al. 2019).[2] We can further relax the equitability criterion up to an arbitrary number of items: an allocation A is said to be *equitable up to c items* or EQc if, for every pair of agents $i, j \in N$ such that $|A_j| > c$, there exists some subset $S \in A_j$ of size $|S| = c$ such that $v_i(A_i) \geq v_j(A_j \setminus S)$.[3]

Freeman et al. (2019) show[4] that, even for binary additive valuations (which is a subclass of the $(0,1)$-OXS valuation class), an allocation that is both EQ1 and PO may not exist; however, in Theorem 4, they establish that 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 under binary additive valuations. We will show that the above positive result about computational tractability extends to the $(0,1)$-OXS valuation class. We will begin by proving that an EQ1 and PO allocation, if it exists, is also EF1 under $(0,1)$-SUB valuations — we achieve this by combining Theorem D.1 below with Corollary 3.8. This simplifies the problem of finding an EQ1, EF1 and PO allocation to that of finding an EQ1 and PO allocation.

Theorem D.1. *For submodular valuations with binary marginal gains, any EQ1 and PO allocation, if it exists, coincides with an allocation satisfying the properties in Theorem 3.11, i.e. is leximin, MNW, and a minimizer of any symmetric strictly convex function of agents' realized valuations among all utilitarian optimal allocations.*

Proof. Let the optimal USW for a problem instance under this valuation class be $\widehat{U}$; also, suppose this instance admits an EQ1 and PO allocation A. The EQ1 property implies that for every pair of agents $i, j \in \widehat{N}$,

$$\begin{aligned}
v_i(A_i) &\geq v_j(A_j \setminus \{o\}) \quad \text{for some } o \in A_j \\
&= v_j(A_j) - \Delta_j(A_j \setminus \{o\}; o) \\
&\geq v_j(A_j) - 1,
\end{aligned}$$

since $\Delta_j(A_j \setminus \{o\}; o) \in \{0,1\}$. Let i (resp., j) be an agent with the minimum (resp., maximum) realized valuation $v^A_{\min}$ (resp., $v^A_{\max}$) under A, i.e. $i \in \arg\min_{k \in N} v_i(A_i)$

[2] Note that if $A_j = \emptyset$, $v_i(A_i) \geq v_j(A_j)$ trivially hence the ordered pair (i, j) for any $i \in N$ could never prevent the allocation from being EQ1.

[3] Again, if $|A_j| \leq c$, no ordered pair (i, j) for any $i \in N$ could get in the way of the allocation being EQ1.

[4] Freeman et al. (2019) use an example with 3 agents having binary additive valuations (Example 1). But it is easy to construct a (ono-degenerate) fair allocation instance with only two agents having binary additive valuations that does not admit an EQ1 and PO allocation: $N = [2]$; $O = \{o_1, o_2, o_3, o_4\}$; $v_1(o) = 1$ for every $o \in O$; $v_2(o_1) = 1$ and $v_2(o) = o$ for every $o \in O \setminus \{o_1\}$. Obviously, any PO allocation must give $\{o_2, o_3, o_4\}$ to agent 1 so that this agent's realized valuation is at least 2 even after dropping one its items; even if agent 2 receives o_1, her realized valuation of 1 will always be less than the above.

and $j \in \arg\max_{k \in N} v_j(A_j)$. Then, the above unequality implies that

$$v^A_{\min} \le v^A_{\max} \le v^A_{\min} + 1.$$

Since each $v_i(A_i)$ must be a non-negative integer for our valuation function class, there are only two possibilities consistent with the above: either **(I)** $v_i(A_i) = \alpha$ for every $i \in N$ (which corresponds to an EQ allocation), or **(II)** there is a subset of agents $N_0 \subset N$, with $|N_0| = n_0 < n$, such that $v_i(A_i) = \alpha$ for every $i \in N_0$ and $v_i(A_i) = \alpha + 1$ for every $i \in N \setminus N_0$.

For Case **(I)**, notice that for any utilitarian optimal allocation A', $\frac{1}{n}\sum_{i=1}^n v_i(A'_i)^2 \ge \left(\frac{1}{n}\sum_{i=1}^n v_i(A'_i)\right)^2 = \frac{\widehat{U}^2}{n^2}$, i.e. $\sum_{i=1}^n v_i(A'_i)^2 \ge \frac{\widehat{U}^2}{n}$ due to the well-known generalized mean inequalities. Since allocation A is PO under $(0,1)$-SUB valuations, Theorem 3.10 implies that it is also utilitarian optimal so that $n\alpha = \widehat{U}$. Hence, $\sum_{i=1}^n v_i(A_i)^2 = n\left(\frac{\widehat{U}}{n}\right)^2 = \frac{\widehat{U}^2}{n}$, i.e. A minimizes the sum of squares of agents' realized valuations — a symmetric strictly convex function — among all utilitarian optimal allocations.

For Case **(II)**, it is sufficient to prove that A must be leximin and we will show this by contradiction. Note that $\widehat{U} = n_0\alpha + (n - n_0)(\alpha + 1)$ due to utilitarian optimality of A. Hence, we cannot have an allocation in which the worst-off agent realizes a valuation strictly higher than α — if that were possible, each agent would have a valuation of at least $\alpha + 1$ (each valuation being an integer), hence the `USW` would be at least $n(\alpha + 1) > \widehat{U}$, which is absurd. This implies that the minimum possible agent valuation in a leximin allocation is α.

Let A^* be a leximin allocation; let $N^*_0 \subset N$ be the subset of agents that have a valuation of α each under A^*, with $|N^*_0| = n^*_0 < n$. Note that we cannot that $N^*_0 = N$ — if all agents in $N \setminus N^*_0$ have valuations α each, A would dominate A^* lexicographically as well as in terms of the `USW`. Thus, each agent in the non-empty set $N \setminus N^*_0$ must have a valuation at least $\alpha + 1$. Moreover, if A is not a leximin allocation, then there must exist at least one agent in $N \setminus N^*_0$ that has a valuation strictly greater than $\alpha + 1$ under A^*. In other words, we must have $\sum_{i \in N \setminus N^*_0} v_i(A^*_i) > (n - n^*_0)(\alpha + 1)$. But, as we will now show, this implies that the sum of the squares of the agents' realized valuations under A^* is strictly higher than that under A.

From Theorem 3.11, we know that A^* is also utilitarian optimal, hence

$$n^*_0\alpha + \sum_{i \in N \setminus N^*_0} v_i(A^*_i) = \widehat{U} = n_0\alpha + (n - n_0)(\alpha + 1).$$

This gives us two useful results. First, plugging the above strict inequality involving $\sum_{i \in N \setminus N^*_0} v_i(A^*_i)$ into this equality, we get using simple algebra that $n^*_0 > n_0$.

Secondly, we obtain

$$\sum_{i \in N \setminus N^*_0} v_i(A^*_i) = (n - n^*_0)\alpha + (n - n_0).$$

Since it is well-known that, for non-negative numbers, the quadratic mean is at least as much as the arithmetic mean,

$$\begin{aligned}
&\frac{1}{n - n^*_0}\sum_{i \in N \setminus N^*_0} v_i(A^*_i)^2 \ge \left(\frac{1}{n - n^*_0}\sum_{i \in N \setminus N^*_0} v_i(A^*_i)\right)^2 \\
\Rightarrow \quad &\sum_{i \in N \setminus N^*_0} v_i(A^*_i)^2 \ge \frac{1}{n - n^*_0}\left(\sum_{i \in N \setminus N^*_0} v_i(A^*_i)\right)^2 \\
&= \frac{((n - n^*_0)\alpha + (n - n_0))^2}{n - n^*_0} \\
&= (n - n^*_0)\alpha^2 + 2(n - n_0)\alpha \\
&\quad + \frac{(n - n_0)^2}{n - n^*_0} \\
&> (n - n^*_0)\alpha^2 + 2(n - n_0)\alpha \\
&\quad + (n - n_0),
\end{aligned}$$

where the strict inequality follows from the inequalities $n > n^*_0 > n_0 > 0$ that we have already established. Hence,

$$\begin{aligned}
\sum_{i \in N} v_i(A^*_i)^2 &= n^*_0\alpha^2 + \sum_{i \in N \setminus N^*_0} v_i(A^*_i)^2 \\
&> (n^*_0 + n - n^*_0)\alpha^2 \\
&\quad + 2(n - n_0)\alpha + (n - n_0) \\
&= n\alpha^2 + (n - n_0)(2\alpha + 1) \\
&= n\alpha^2 + (n - n_0)((\alpha + 1)^2 - \alpha^2) \\
&= (n - n + n_0)\alpha^2 + (n - n_0)(\alpha + 1)^2 \\
&= n_0\alpha^2 + (n - n_0)(\alpha + 1)^2 \\
&= \sum_{i \in N} v_i(A_i)^2.
\end{aligned}$$

But this absurd because we know from Theorem 3.11 and Corollary 3.12 that the leximin allocation A^* must have the minimum sum of squares of realized valuations among all utilitarian optimal allocations. Hence, we cannot have an agent with a valuation strictly higher than $\alpha + 1$ in a leximin allocation, so A must be leximin. □

Theorem D.1, together with Corollary 3.12, implies that if an EQ1 and PO allocation exists under $(0,1)$-SUB valuations, it must be EF1.