Finding Fair and Efficient Allocations for Matroid Rank Valuations

Nawal Benabbou, Mithun Chakraborty, Ayumi Igarashi, Yair Zick

► To cite this version:

Nawal Benabbou, Mithun Chakraborty, Ayumi Igarashi, Yair Zick. Finding Fair and Efficient Allocations for Matroid Rank Valuations. ACM Transactions on Economics and Computation, 2021. hal-03271772

HAL Id: hal-03271772

<https://hal.sorbonne-universite.fr/hal-03271772>

Submitted on 20 Aug 2021

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

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

Finding Fair and Efficient Allocations for Matroid Rank Valuations

NAWAL BENABBOU, Sorbonne Université, CNRS, Laboratoire d'Informatique de Paris 6, France

MITHUN CHAKRABORTY, University of Michigan, Ann Arbor

AYUMI IGARASHI, National Institute of Informatics, Tokyo, Japan

YAIR ZICK, University of Massachusetts, Amherst, USA

In this paper, we present new results on the fair and efficient allocation of indivisible goods to agents whose preferences correspond to *matroid rank functions*. This is a versatile valuation class with several desirable properties (such as monotonicity and submodularity), which naturally lends itself to a number of real-world domains. We use these properties to our advantage; first, we show that when agent valuations are matroid rank functions, a socially optimal (i.e. utilitarian social welfare-maximizing) allocation that achieves envy-freeness up to one item (EF1) exists and is computationally tractable. We also prove that the Nash welfare-maximizing and the leximin allocations both exhibit this fairness/efficiency combination, by showing that they can be achieved by minimizing any symmetric strictly convex function over utilitarian optimal outcomes. To the best of our knowledge, this is the first valuation function class not subsumed by additive valuations for which it has been established that an allocation maximizing Nash welfare is EF1. Moreover, for a subclass of these valuation functions based on maximum (unweighted) bipartite matching, we show that a leximin allocation can be computed in polynomial time. Additionally, we explore possible extensions of our results to fairness criteria other than EF1 as well as to generalizations of the above valuation classes.

ACM Reference Format:

Nawal Benabbou, Mithun Chakraborty, Ayumi Igarashi, and Yair Zick. 2021. Finding Fair and Efficient Allocations for Matroid Rank Valuations. *ACM Transactions on Economics and Computation* 37, 4, Article 111 (August 2021), 43 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

Suppose that we are interested in allocating seats in courses to prospective students. How should this be done? On the one hand, courses offer limited seats and have scheduling conflicts; on the other, students have preferences over the classes that they take, which must be accounted for. In addition, students might have exogenous constraints, such as a hard limit on the number of classes they may take. Course allocation can be thought of as a problem of allocating a set of *indivisible goods* (course slots) to *agents* (students). One thing that immediately stands out is that the problem of assigning courses to students is very well-structured: students are either willing or unwilling to sign up for a class; this can be thought of as having either a value of 1 or of 0 for being assigned a class — it makes no sense to assign a student to a class they did not sign up for. In addition, if we assign a set of classes S to a student, and they are able to take it, then they would be able to

Authors' addresses: Nawal Benabbou, Sorbonne Université, CNRS, Laboratoire d'Informatique de Paris 6, France, nawal.benabbou@lip6.fr; Mithun Chakraborty, University of Michigan, Ann Arbor, dcsmc@umich.edu; Ayumi Igarashi, National Institute of Informatics, Tokyo, Japan, ayumi_igarashi@nii.ac.jp; Yair Zick, University of Massachusetts, Amherst, USA, yzick@umass.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2021 Association for Computing Machinery.

XXXX-XXXX/2021/8-ART111 \$15.00

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

take any subset of S as well. Finally (this is not trivial to prove, but is indeed true), given two sets of feasible course assignments S, T such that $|S| < |T|$, we can find some class $o \in T$ such that $S \cup \{o\}$ is also a feasible course assignment. Such “well-behaved” structures are also known as *matroids*. How should we divide goods among agents with subjective valuations? Can we find a “good” allocation in polynomial time? Can we exploit the structure of certain problems to efficiently find good allocations?

These questions have been the focus of intense study in the CS/Econ community in recent years; several justice criteria, as well as methods for computing allocations that satisfy them have been investigated. Generally speaking, justice criteria fall into two categories: *efficiency* and *fairness*. Efficiency criteria are chiefly concerned with lowering some form of *waste*, maximizing some notion of item utilization, or agent utilities. For instance, *Pareto optimality* (PO) is a popular efficiency concept which ensures that the value realized by no agent can be improved without diminishing that of another agent. Fairness criteria require that agents do not perceive the resulting allocation as mistreating them compared to others; for example, one might want to ensure that no agent prefers another agent’s assigned bundle (i.e. subset of goods) to her own bundle – this criterion is known as *envy-freeness* (EF) [Foley 1967]. However, envy-freeness is not always achievable when items are indivisible: consider a stylized setting, where there is just one course with one seat for which two students are competing; any student receiving this slot would be envied by the other. A simple solution ensuring envy-freeness would be to withhold the seat altogether, not assigning it to either student. Withholding items, however, violates most efficiency criteria.¹

As illustrated above and also observed by Budish [2011], envy-freeness is not always achievable, even under *completeness*, a very weak efficiency criterion requiring that each item is allocated to some agent. However, a less stringent fairness notion – *envy-freeness up to one good* (EF1) – can be attained. An allocation is EF1 if for any two agents i and j , there is some item in j ’s bundle whose removal results in i not envying j . Complete, EF1 allocations always exist for monotone valuations, and in fact, can be found in polynomial time, thanks to the now-classic *envy graph algorithm* due to Lipton et al. [2004].

It is already challenging to individually achieve strong allocative justice criteria; hence, computationally efficient methods that produce allocations satisfying multiple such criteria simultaneously are of particular interest. Caragiannis et al. [2019b] show that when agent valuations are *additive* – i.e. every agent i values its allocated bundle as the sum of values of individual items – there exist allocations that are both PO and EF1. Specifically, these are allocations that maximize the product of agents’ utilities – also known as the *Nash welfare* (MNW). Further work [Barman et al. 2018a] shows that such allocations can be found in pseudo-polynomial time. While encouraging, these results are limited to agents with additive valuations. In particular, they do not apply to settings such as the course allocation problem described above (e.g. being assigned two courses with conflicting schedules will not result in additive gain), or other settings we describe later on. In fact, Caragiannis et al. [2019b] left it open whether their result extends to other natural classes of valuation functions, such as the class of submodular valuations.² At present, little is known 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 (or dichotomous) marginal gains, which are also known as *matroid rank valuations* [Oxley 2011]. In this setting, the added benefit

¹One of the coauthors applied this solution when his children were fighting over a single toy; the method was ultimately deemed unsuccessful.

²Caragiannis et al. [2019b] do provide an instance of two agents with monotone supermodular/subadditive valuations where no allocation is PO and EF1.

of receiving another item is binary and obeys the law of diminishing marginal returns. This is equivalent to the class of valuations that can be captured by *matroid* constraints. Matroids are mathematical structures that generalize the concept of linear independence beyond vector spaces; we refer the interested reader to [Oxley 2011] for further details. Matroids have proven to be a versatile framework for describing a variety of problem domains. In our fair allocation domain, each agent has a different matroid constraint over the collection of items, and her value for a bundle is determined by the size of a maximum independent set included in the bundle.

Matroid rank valuations naturally arise in many practical applications, beyond the course allocation problem described above (where students are limited to either approving/disapproving a class). For example, suppose that a government body wishes to fairly allocate public goods to individuals of different minority groups (say, in accordance with a diversity-promoting policy). This could apply to the assignment of kindergarten slots to children from different neighborhoods/socioeconomic classes³ or of flats in public housing estates to applicants of different ethnicities [Benabbou et al. 2019, 2020]. A possible way of achieving group fairness in this setting is to model each minority group as an agent consisting of many individuals: each agent's valuation function is based on *optimally matching* items to its constituent individuals; envy naturally captures the notion that no group should believe that other groups were offered better bundles (this is the fairness notion studied by Benabbou et al. [2019]). Such assignment/matching-based valuations (known as OXS valuations [Leme 2017]) are non-additive in general, and constitute an important subclass of submodular valuations.

Matroid rank functions correspond to submodular functions with binary (i.e. $\{0, 1\}$) marginal gains. The binary marginal gains assumption is best understood in context of matching-based valuations described above — in this scenario, it simply means that individuals either approve or disapprove of items, and do not distinguish between items they approve (we call OXS functions with binary individual preferences $(0, 1)$ -OXS valuations). This is a reasonable assumption in kindergarten slot allocation (all approved/available slots are identical), and is implicitly made in some public housing mechanisms; for instance, Singapore housing applicants are required to effectively approve a subset of flats by selecting a block, and are precluded from expressing a more refined preference model. A similar assumption is made in student course selection, where students de-facto approve certain classes by signing up for them (and are thus precluded from expressing more refined preferences).

What if we further assume that there are exogenous capacity constraints? This is the case in course selection (students may only approve at most a fixed number of classes), and in housing allocation (ethnic minorities in Singapore may only receive a fixed number of flats [Benabbou et al. 2020]). 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 [Benabbou et al. 2020; Chua 1991; Deng et al. 2013; Parliament of Singapore. Parliament Debates: Official Report 1989; Phang and Kim 2013; Sim et al. 2003; Wong 2014]; socioeconomic status-based quotas in certain U.S. public school admission systems such as Chicago Public Schools [Benabbou et al. 2020; Chicago Public Schools 2017; Quick 2016; U.S. Department of Education, Office of Elementary and Secondary Education 2017] on the number of items each group is able or allowed to receive, then agents' valuations are *truncated* matching-based valuations. Such valuation functions are not OXS, but are still matroid rank functions (i.e. submodularity is preserved). Since agents still have binary/dichotomous preferences over items even with the quotas in place, our results apply to this broader class as well.

³see, e.g. <https://www.ed.gov/diversity-opportunity>.

Valuation class	MNW	Leximin	max-USW+EF1
(0, 1)-OXS	poly (Th. 4.1)	poly (Th. 4.1)	poly (Th. 3.4)
matroid rank	poly (BEF20)	poly (BEF20)	poly (Th. 3.4; BEF20)

Table 1. Summary of our computational complexity results: “poly” denotes polynomial; BEF20 refers to Babaioff et al. [2021]

Fig. 1. (a) Subsumption relations and intersections among the various valuation function classes defined in this paper: BM denotes the class of valuation functions with binary marginal gains, MR the matroid rank valuation function class (shaded) — the class of main interest to us, and BA denotes the class of binary additive valuations. (b) Equivalences and intersections among properties of clean allocations under the matroid rank valuation class. For both (a) and (b), sizes of blocks have no significance.

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

- For matroid rank valuations, we show that an EF1 allocation that also maximizes the *utilitarian social welfare* or USW (hence is Pareto optimal) always exists and can be computed in polynomial time by a simple greedy algorithm.
- For matroid rank valuations, we show that leximin⁴ and MNW allocations both possess the EF1 property.
- For matroid rank valuations, we provide a characterization of the leximin allocations; we show that they are identical to the minimizers of *any* symmetric strictly convex function over utilitarian optimal allocations (equivalently, the maximizers of any symmetric strictly concave function over utilitarian optimal allocations). We obtain the same characterization for MNW allocations.
- For (0, 1)-OXS valuations, we show that both leximin and MNW allocations can be computed efficiently.

Result (a) is remarkably positive: the EF1 and utilitarian welfare objectives are incompatible in general, even for additive valuations, as shown by Example A.1 in Appendix A. In fact, maximizing the utilitarian social welfare among all EF1 allocations is NP-hard for general valuations [Barman et al. 2019].

Result (b) is reminiscent of the Theorem 3.2 in Caragiannis et al. [2019b], showing that any MNW allocation is PO and EF1 under *additive* valuations; they also showed that a PO and EF1

⁴Roughly 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.

allocation may not exist under subadditive/supermodular valuations (Theorem 3.3) and MNW does not imply EF1 for arbitrary, real-valued submodular functions (Appendix C Example C.3) 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. [Plaut and Roughgarden 2018]) but the PO+EF1 existence issue beyond additive valuations is yet to be settled. To our knowledge, the class of matroid rank valuations is the first valuation class not subsumed by additive valuations for which the EF1 property of the MNW allocation 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+EF even for the seemingly simple class of binary additive valuations which is subsumed by our matroid rank class (Bouveret and Lang [2008] Proposition 21); leximin [Bezáková and Dani 2005]; and MNW [Nguyen et al. 2014]. Moreover, although previous work on binary additive valuations established 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 [Barman et al. 2018b], we extend this result to the strictly larger $(0, 1)$ -OXS class by uncovering connections to the rich literature on combinatorial optimization. For example, the result (d) will exploit the network flow technique, drawing the equivalence between leximin allocations and balanced network flows [Frank and Murota 2018].

Our analysis makes extensive use of tools and concepts from matroid theory [Oxley 2011]. While some papers have explored 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 provide with fairness and efficiency guarantees for a matroid-based valuation function class.

In addition to the above main results (a)–(d), we present in the appendices a discussion on non-envy-based fairness criteria for matroid rank valuations as well as our attempts at generalizing our results beyond this valuation class:

- In Appendix B, we show that for the more general submodular valuations with *subjective* binary marginal gains (for which adding an item to an agent’s bundle increases her valuation either by zero or by an item-invariant but agent-specific constant), MNW and leximin allocations no longer coincide – in particular, an MNW allocation retains the EF1 property whereas leximin allocations do not.
- In Appendix C, we formulate a heuristic extension of the fair allocation algorithm from Result (a) above that applies to general assignment (OXS) valuations and evaluate its efficiency in experiments based on a real-world data set *MovieLens-ml-1m* [Harper and Konstan 2015].
- In Appendix D, we explore (approximate) proportionality, equitability, and the maximin share guarantee for matroid rank valuations in particular and submodular valuations in general.

1.2 Related Work

1.2.1 Fair Division with General Valuations. There is a vast and growing literature on fairness and efficiency issues in resource allocation. Early work on divisible resource allocation provides an elegant result: an allocation that satisfies envy-freeness and Pareto optimality always exists under mild assumptions on valuations [Varian 1974], and can be computed via the convex programming of Eisenberg and Gale [1959] for additive valuations. In the domain of the allocation of indivisible goods (see Bouveret et al. [2016]; Markakis [2017] for an overview), Budish [2011] was the first to formalize the notion of EF1 as an approximation to envy-freeness; but, it implicitly appears in Lipton

et al. [2004]. More recently, Caragiannis et al. [2019b] prove the discrete analogue of Eisenberg and Gale [1959]: MNW allocation satisfies EF1 and Pareto optimality for additive valuations. Barman et al. [2018a] provide a pseudo-polynomial-time algorithm for computing allocations satisfying EF1 and PO.

Closely related to ours is the work of Biswas and Barman [2018] who consider fair division under matroid constraints. In Section 6 of Biswas and Barman [2018], the authors consider a specific setting where agents have identical additive valuations, and the aim is to find a fair allocation under the same matroid constraint. They show that if a *feasible* allocation satisfying the matroid constraint exists, it can be transformed into an EF1 complete feasible allocation. Although our setting is different from theirs, the proof of Theorem 3.4 in our work uses a similar technique to the one used in Theorem 3 of Biswas and Barman [2018] in order to transfer an item from the envious bundle to the envied one.

In this paper, we admit allocations that may be incomplete (i.e. not all items are allocated to the agents under consideration) but satisfy strong fairness and efficiency guarantees. This brings us close to recent work on fairness with “charity” by Caragiannis et al. [2019a] and Chaudhury et al. [2020].

1.2.2 Binary Additive Valuations. Under binary preferences, agents either approve or disapprove of each good. Due to the advantage of simple elicitation, the use of binary preferences is widespread in social choice literature [Lackner and Skowron 2020]. In the context of fair division, there has been recent progress on the study of fair and efficient allocations in this domain.

In particular, several positive computational results have been obtained for binary additive valuations. Darmann and Schauer [2015] and Barman et al. [2018b] show that the maximum Nash welfare can be computed efficiently for binary additive valuations while computing MNW allocations of indivisible items is hard in general. The work of Barman et al. [2018b] further develops 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 number of approved items may have different values under the latter class;⁵ hence the polynomial-time complexity result of Barman et al. [2018a] does not imply our Theorem 4.1.

Independently of our work, Aziz and Rey [2020] show the equivalence between leximin and MNW in the context of binary additive valuations (Lemma 4 of Aziz and Rey [2020]), which is a special case of our result (c).

Halpern et al. [2020] show that for binary additive valuations, there is a group strategy-proof mechanism that returns an allocation satisfying utilitarian optimality and EF1 (Theorem 1 of Halpern et al. [2020]); dropping strategy-proofness, we generalize this result to the class of matroid rank valuations.

1.2.3 Matroid Rank Valuations. Babaioff et al. [2021] independently present a set of results similar to our own; moreover, they explore strategy-proof deterministic and randomized mechanisms for matroid rank valuations, showing that such mechanisms exist. Below, we compare our results with theirs in greater detail.

- Babaioff et al. [2021] show that a mechanism returning a special MNW allocation (called Prioritized Egalitarian mechanism, or PE mechanism for short) achieves strategy-proofness,

⁵Consider 3 items, o_1, o_2, o_3 , and a group of members $S = \{1, 2, 3\}$ with member 1 assigning weight 1 to items o_1 and o_3 , and members 2 and 3 assigning weight 1 to item o_2 only. The value of a maximum matching between $\{o_1, o_2\}$ and S is 2 while the value of a maximum matching between $\{o_1, o_3\}$ and S is 1.

EFX_0 ⁶, and maximizes the social welfare for matroid rank valuations. In addition, their mechanism runs in polynomial time (Theorem 1 of Babaioff et al. [2021]). In this work, we present a simple algorithm that finds a social welfare-maximizing allocation satisfying EF1; however, the returned allocation may not be MNW or EFX_0 (see Examples 3.9 and A.5). In addition, we provide a polynomial-time algorithm returning a MNW allocation for $(0, 1)$ -OXS valuations.

- Babaioff et al. [2021] show that for binary additive valuations, the PE mechanism achieves MMS (Proposition 5 of Babaioff et al. [2021]). They further prove that for general matroid rank functions, although PE mechanism may not satisfy full MMS, it achieves $\frac{1}{2}$ MMS (Proposition 6 of Babaioff et al. [2021]). We also show that for binary additive valuations, PO and EF1 allocation is MMS (Proposition D.7). In addition, we show that there is an instance with $(0, 1)$ -OXS valuations such that no PO and EF1 allocation is MMS.
- Babaioff et al. [2021] show that for matroid rank functions, PE mechanism minimizes the sum of squares. In this work, we prove a more general statement: for general matroid rank functions, any MNW/leximin minimizes a symmetric convex function of agents' valuations.
- While Babaioff et al. [2021] do not show the equivalence between MNW and leximin (as we do), they prove a lemma (Lemma 17 of Babaioff et al. [2021]) that corresponds to Lemma 3.12 of this work, which is crucial in establishing the equivalence.
- Babaioff et al. [2021] study the setting where the desirable items are allowed to take an arbitrary value in the range $[1, 1 + \epsilon]$ (called ϵ -leveled valuations), aiming to generalize the positive existence result of a fair and strategy-proof mechanism. Note that the marginal can vary over items so the class of ϵ -leveled valuations strictly generalizes the class of submodular functions with subjective binary marginal gains we study in Appendix B. However, there is little overlap with respect to this domain. We show MNW still satisfies EF1 but the equivalence between MNW and leximin is lost for the class of submodular functions with subjective binary marginal gains, which has not been addressed in Babaioff et al. [2021].

To conclude, our work was developed independently, and is conceptually different from Babaioff et al. [2021] in that the main focus of ours is on the fairness and efficiency compatibility and the properties of such allocations.

1.2.4 Other related work. One motivation for this 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. Similar works study quota-based fairness/diversity [Aziz et al. 2019; Benabbou et al. 2020; Suzuki et al. 2018, and references therein], or by the optimization of carefully constructed functions [Ahmed et al. 2017; Dickerson et al. 2019; 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 $\{1, 2, \dots, r\}$. We are given a set $N = [n]$ of *agents*, and a set $O = \{o_1, \dots, o_m\}$ of *items* or *goods*. Subsets of O are referred to as *bundles*, and each agent $i \in N$ has a *valuation function* $v_i : 2^O \rightarrow \mathbb{R}_+$ over bundles where $v_i(\emptyset) = 0$. We further assume polynomial-time oracle access to the valuation v_i of all agents. Given a valuation function $v_i : 2^O \rightarrow \mathbb{R}$, we define the *marginal gain* of an item $o \in O$ w.r.t. a bundle $S \subseteq O$, as $\Delta_i(S; o) \triangleq v_i(S \cup \{o\}) - v_i(S)$. A valuation function v_i is *monotone* if $v_i(S) \leq v_i(T)$ whenever $S \subseteq T$.

⁶We provide the definition of EFX_0 in Remark 2 in Section 3.

An *allocation* A of items to agents is a collection of n disjoint bundles $A_1, \dots, A_n$, such that $\bigcup_{i \in N} A_i \subseteq O$; the bundle A_i is allocated to agent i . Given an allocation A , we denote by A_0 the set of unallocated items, also referred to as *withheld items*. We may refer to agent i 's valuation of its bundle $v_i(A_i)$ under the allocation A as its *realized valuation* under A . An allocation is *complete* if every item is allocated to *some* agent, i.e. $A_0 = \emptyset$. We admit incomplete, but *clean* allocations: a bundle $S \subseteq O$ is *clean* for $i \in N$ if it contains no item $o \in S$ for which agent i has zero marginal gain (i.e., $\Delta_i(S \setminus \{o\}; o) = 0$); allocation A is *clean* if each allocated bundle A_i is clean for the agent i that receives it. It is easy to 'clean' any allocation without changing any realized valuation by iteratively revoking items of zero marginal gain from respective agents and placing them in A_0 . For example, if for agent i , $v_i(\{1\}) = v_i(\{2\}) = v_i(\{1, 2\}) = 1$, then the bundle $A_i = \{1, 2\}$ is not clean for agent i (and neither is any allocation where i receives items 1 and 2) but it can be cleaned by moving item 1 (or item 2 but not both) to A_0 .

2.1 Fairness and Efficiency Criteria

Our fairness criteria are based on the concept of *envy*. Agent i *envies* agent j under an allocation A if $v_i(A_i) < v_i(A_j)$. An allocation A is *envy-free* (EF) if no agent envies another. We will use the following relaxation of the EF property due to Budish [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.

Closely related to the concept of efficiency is the welfare of an allocation which can be measured in several ways [Sen 1970]. Specifically, given an allocation A ,

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

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

Since it is possible that the maximum attainable Nash welfare is 0 (e.g. if there are fewer items than agents, then one agent must have an empty bundle), we use the following refinement of the maximum Nash social welfare (MNW) criterion used in [Caragiannis et al. 2019b]: we find a largest subset of agents, say $N_{\max} \subseteq N$, to which we can allocate bundles of positive values, and compute an allocation to agents in $N_{\max}$ that maximizes the product of their realized valuations. If $N_{\max}$ is not unique, we choose the one that results in the highest product of realized valuations.

The *leximin* welfare is a lexicographic refinement of the maximin welfare concept, i.e. egalitarian optimality. Formally, for real n -dimensional vectors $\mathbf{x}$ and $\mathbf{y}$, $\mathbf{x}$ is *lexicographically greater than or equal to* $\mathbf{y}$ (denoted by $\mathbf{x} \geq_L \mathbf{y}$) if and only if $\mathbf{x} = \mathbf{y}$, or $\mathbf{x} \neq \mathbf{y}$ 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 $\mathbf{s}(A)$ the vector of the components $v_i(A_i)$ ($i \in N$) arranged in non-decreasing order. A *leximin* allocation A is an allocation that maximizes the egalitarian welfare in a lexicographic sense, i.e., $\mathbf{s}(A) \geq_L \mathbf{s}(A')$ for any other allocation A' .

2.2 Submodular Valuations

In this paper, agents' valuation functions are not necessarily additive but *submodular*. A valuation function v_i is *submodular* if each single item contributes more to a smaller set than to a larger one, namely, for all $S \subseteq T \subseteq O$ and all $o \in O \setminus T$, $\Delta_i(S; o) \geq \Delta_i(T; o)$.

One important sub-class of submodular valuations is the class of *assignment valuations*. This class of valuations was introduced by Shapley [1962] and is synonymous with the OXS valuation class [Balcan et al. 2012; Lehmann et al. 2006; Leme 2017]. 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,

$$v_h(S) = \max\left\{ \sum_{i \in N_h} u_{i,\pi(i)} \mid \pi \in \Pi(N_h, S) \right\},$$

where $\Pi(N_h, S)$ is the set of matchings $\pi : N_h \rightarrow S$ in the complete bipartite graph with bipartition (N_h, S) .

Our particular focus is on submodular functions with *binary marginal gains*. We say that v_i has *binary marginal gains* if $\Delta_i(S; o) \in \{0, 1\}$ for all $S \subseteq O$ and $o \in O \setminus S$. The class of submodular valuations with binary marginal gains includes the classes of binary additive valuations [Barman et al. 2018b] and of assignment valuations where the weight is binary [Benabbou et al. 2019]. We say that v_i is a *matroid rank valuation* if it is a submodular function with binary marginal gains (these are equivalent definitions [Oxley 2011]), and (0, 1)-OXS if it is an assignment valuation with binary marginal gains.⁷ The constrained assignment valuations discussed in the fourth paragraph of Section 1.1 are examples of matroid rank valuations that are not (0, 1)-OXS.

3 MATROID RANK VALUATIONS

The main theme of all results in this section is that, when all agents have matroid rank valuations, fairness (EF1) and efficiency (PO) properties are compatible with each other and also with all three optimal welfare criteria we consider. Lemma 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 (the proof is available in Appendix A).

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

We start the analysis of matroid rank valuations by introducing the basics of matroid theory. Formally, a *matroid* is an ordered pair $(E, \mathcal{I})$, where E is some finite set and $\mathcal{I}$ is a family of its subsets (referred to as the *independent sets* of the matroid), which satisfies the following three axioms:

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

The rank function $r : 2^E \rightarrow \mathbb{Z}$ of a matroid returns the *rank* of each set X , i.e. the maximum size of an independent subset of X . Another equivalent way to define a matroid is to use the axiom systems for a rank function. We require that (R1) $r(X) \leq |X|$, (R2) r is monotone, and (R3) r is submodular. Then, the pair $(E, \mathcal{I})$ where $\mathcal{I} = \{X \subseteq E \mid r(X) = |X|\}$ is a matroid [Oxley 2011]. In other words, if r satisfies properties (R1)–(R3) then it induces a matroid.

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

⁷ (0, 1)-OXS valuations coincide with rank functions of *transversal* matroids [Balcan et al. 2012].

PROPOSITION 3.2. *A valuation function v_i with binary marginal gains is monotone and takes values in $[|S|]$ for any bundle S (hence $v_i(S) \leq |S|$).*

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

PROPOSITION 3.3. *For matroid rank valuations, A is a clean allocation if and only if $v_i(A_i) = |A_i|$ for each $i \in N$.*

Example A.3 in Appendix A.5 shows that Lipton et al. [2004]’s classic envy graph algorithm does not guarantee a Pareto optimal allocation under matroid rank valuations (although the output allocation is complete and EF1), and thus underscores the difficulty of finding the PO+EF1 combination under this valuation class. Moreover, note that in the simple example of one good and two agents each valuing the good at 1, both agents’ valuation functions belong to the class under consideration — this shows that an envy-free and Pareto optimal allocation may not exist even under this class, and further justifies our quest for EF1 and Pareto-optimal allocations.

3.1 Finding a Utilitarian Optimal and EF1 Allocation

We will now establish that the existence of a PO+EF1 allocation, proved for additive valuations by Caragiannis et al. [2019b], extends to the class of matroid rank valuations. In fact, we provide a stronger — and surprisingly strong — relation between efficiency and fairness: utilitarian optimality (stronger than Pareto optimality) and EF1 turn out to be mutually compatible under this valuation class. Moreover, such an allocation can be computed in polynomial time!

THEOREM 3.4. *For matroid rank valuations, a utilitarian optimal allocation that is also EF1 exists and can be computed in polynomial time.*

Our result is constructive: we provide a way of computing the above allocation in Algorithm 1. The proof of Theorem 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 matroid rank 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 with respect to A_i .*

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 $A_j = \{o_1, o_2, \dots, o_r\}$. As in the proof of Proposition 3.2, if we define $S_0 = \emptyset$ and $S_t = \{o_1, o_2, \dots, o_t\}$ for each $t \in [r]$, we can write 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. $\square$

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 — see Example A.4 in Appendix A.6.

Below, we show that if i 's envy towards j under a clean allocation cannot be eliminated by removing one item from the latter's bundle, then the two agents' valuations for their respective bundles differ by at least two (in fact, we establish a stronger version of the result that does not require the envious agent i 's bundle to be clean). Formally, we say that agent i envies j up to more than 1 item if $A_j \neq \emptyset$ and $v_i(A_i) < v_i(A_j \setminus \{o\})$ for every $o \in A_j$.

LEMMA 3.6. *For submodular functions with binary marginal gains, if agent i envies agent j up to more than 1 item under an allocation A and j 's bundle A_j is clean, then $v_j(A_j) \geq v_i(A_i) + 2$.*

PROOF. From 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_j is a clean bundle for j , Proposition 3.3 implies that $v_j(A_j) = |A_j|$. Combining these, we get

$$v_i(A_i) < v_i(A_j \setminus \{o\}) \leq |A_j| - 1 = v_j(A_j) - 1 \quad \Rightarrow \quad v_j(A_j) > v_i(A_i) + 1,$$

which proves the theorem statement since all valuations are integers. $\square$

Next, we show that under matroid rank valuations, utilitarian social welfare maximization is polynomial-time solvable (3.7).

THEOREM 3.7. *If all agents have 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. $\mathcal{I}_i = \{X \subseteq E \mid r_i(X) = |X|\}$, is the set of independent sets of a matroid. Taking the union $\mathcal{I} = \mathcal{I}_1 \cup \dots \cup \mathcal{I}_n$, the pair $(E, \mathcal{I})$ is known to form a matroid [Korte and Vygen 2006], often referred to as a *union matroid*. By definition, $\mathcal{I} = \{\bigcup_{i \in N} X_i \mid X_i \in \mathcal{I}_i \wedge i \in N\}$, so any independent set in $\mathcal{I}$ corresponds to a union of clean bundles for each $i \in N$ and vice versa. To ensure that each item is assigned at most once (i.e. bundles are disjoint), we will define another matroid $(E, \mathcal{O})$ where the set of independent sets is given by

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

Here, $E_o = \{e = \{o, i\} \mid i \in N\}$ for $o \in O$. The pair $(E, \mathcal{O})$ is known as a *partition matroid* [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 $O(|E|^3 \gamma)$ where γ 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 $O(|E|^3 mnF)$. $\square$

Finally, we are ready to prove Theorem 3.4.

PROOF OF THEOREM 3.4. Consider Algorithm 1. This algorithm maintains optimal USW as an invariant and terminates on an EF1 allocation. Specifically, we first compute a clean allocation that maximizes the utilitarian social welfare. The EIT subroutine in the algorithm iteratively diminishes envy by transferring an item from the envied bundle to the envious agent; Lemma 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

```

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

Correctness: Each EIT step maintains the optimal utilitarian social welfare as well as cleanliness: an envied agent's valuation diminishes exactly by 1 while that of the envious agent increases by exactly 1. Specifically, recall that for matroid rank valuations, an allocation A is clean if and only if $v_i(A_i) = |A_i|$ for all $i \in N$ by Proposition 3.3. This means that if the previous allocation A is clean, then we have $v_i(A_i \cup \{o\}) = |A_i \cup \{o\}|$, and $v_j(A_j \setminus \{o\}) = |A_j \setminus \{o\}|$. Hence the new allocation after each EIT step remains clean. Thus, if the algorithm terminates, the EIT subroutine retains the initial (optimal) USW and, by the stopping criterion, induces the EF1 property.

To show that the algorithm terminates (in polynomial time), we define the potential function $\Phi(A) \triangleq \sum_{i \in N} v_i(A_i)^2$. At each step of the algorithm, $\Phi(A)$ strictly decreases by 2 or a larger integer. To see this, let A' denote the resulting allocation after reallocation of item o from agent j to i . Since A is clean, we have $v_i(A'_i) = v_i(A_i) + 1$ and $v_j(A'_j) = v_j(A_j) - 1$; since all other bundles are untouched, $v_k(A'_k) = v_k(A_k)$ for every $k \in N \setminus \{i, j\}$. Also, since i envies j up to more than one item under allocation A , $v_i(A_i) + 2 \leq v_j(A_j)$ by Lemma 3.6. Combining these, we get

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

Complexity: By Theorem 3.7, a clean utilitarian optimal allocation can be computed in polynomial time. The value of the non-negative potential function has a polynomial upper bound: $\sum_{i \in N} v_i(A_i)^2 \leq (\sum_{i \in N} v_i(A_i))^2 \leq m^2$. Thus, Algorithm 1 terminates in polynomial time. $\square$

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

COROLLARY 3.8. *For matroid rank valuations, any clean, utilitarian optimal allocation A that minimizes*

$$\Phi(A) \triangleq \sum_{i \in N} v_i(A_i)^2$$

among all utilitarian optimal allocations is EF1.

REMARK 1 (CHOICE OF THE POTENTIAL FUNCTION). In the proof of Theorem 3.4, we used the sum of squared valuations as the potential function to prove termination in polynomial time mainly for ease of exposition. Additionally, it shows that the EIT subroutine terminates after $O(m^2)$ iterations. However, any *symmetric, strictly convex,⁸ polynomial* function Φ of the realized valuations strictly decreases with each EIT step and, as such, it is sufficient to use any such function as our potential function if we just wish to establish termination in a polynomial number of iterations. Moreover, Corollary 3.8 holds for any such function Φ as well — we elaborate on this theme in Section 3.2. ♦

Despite its simplicity, Algorithm 1 significantly generalizes that of Benabbou et al. [2019]’s Theorem 4 (which ensures the existence of a non-wasteful EF1 allocation for $(0, 1)$ -OXS valuations) to matroid rank valuations. We note, however, that the resulting allocation may be neither MNW nor leximin even when agents have $(0, 1)$ -OXS valuations: Example 3.9 below illustrates this and also shows that the converse of Corollary 3.8 does not hold.

Example 3.9. The instance we use is identical to Example 1 in Benabbou et al. [2019]. There are two groups (i.e. agents with $(0, 1)$ -OXS valuations) and six items $o_1, o_2, o_3, o_4, o_5, o_6$. The first group N_1 (identical to agent 1) contains four members a_1, a_2, a_3, a_4 and the second group N_2 (identical to agent 2) contains four members b_1, b_2, b_3, b_4 ; each individual has utility (weight) 1 for an item o if and only if she is adjacent to o in the graph depicted in Figure 2:

Fig. 2. An instance where Algorithm 1 produces an allocation that is not MNW or leximin.

The valuation function of each group for any bundle X is defined as the value (equivalently, the size) of a maximum-size matching of X to the group’s members. The algorithm may initially compute a utilitarian optimal allocation A that assigns items o_1, o_2, o_3, o_5 to the group N_1 (with these items assigned to a_1, a_2, a_3, a_4 respectively), and the remaining items to group N_2 (with o_4, o_6

⁸See Section 3.2 for the definition of a symmetric, strictly convex function. For the proof of Theorem 3.4, it suffices for the function to be strictly convex only over the non-negative orthant since valuations are always non-negative.

assigned to b_2, b_4 respectively). Then, $v_1(A_1) = 4 > 2 = v_1(A_2)$ and $v_2(A_2) = 2 = v_2(A_1)$, hence the allocation A is EF1 — in fact, envy-free! So, the EIT subroutine will not be invoked and the output of Algorithm 1 will be 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. ■

REMARK 2 (EFX ALLOCATION). It is worthwhile at this point to comment on the implications of our results for a stronger version of the EF1 property that has received considerable attention in recent literature: envy-freeness up to *any* item, often called the EFX condition. There are two definitions in the literature:

- (1) Caragiannis et al. [2019b] who introduced this concept (for additive valuations) called it *envy-freeness up to the least (positively) valued good*; we can naturally extend their definition to general valuations as follows: an allocation A is EFX if, for every pair of agents $i, j \in N$ such that i envies j , $v_i(A_i) \geq v_i(A_j \setminus \{o\})$ for every item $o \in A_j$ satisfying $\Delta_i(A_j \setminus \{o\}) > 0$. We will call this property EFX_+ for clarity (see Remark 4 in Appendix A.9 for issues with defining EFX_+ beyond additive valuations).
- (2) Plaut and Roughgarden [2018] defined an allocation A to be EFX if, for every pair of agents i, j , $v_i(A_i) \geq v_i(A_j \setminus \{o\}) \forall o \in A_j$ or equivalently $v_i(A_i) \geq \max_{o \in A_j} v_i(A_j \setminus \{o\})$ — this stronger definition favors allocations where more agents are envy-free of others since $v_i(A_j \setminus \{o\}) = v_i(A_j)$ whenever o is of zero marginal value to agent i with respect to the bundle A_j ; the authors show that under this definition, no EFX allocation can be Pareto optimal even for two agents with additive valuations or general but identical valuations.⁹ Caragiannis et al. [2019a] and Chaudhury et al. [2020] use this definition as well. Following Kyropoulou et al. [2020], who studied both the (above) weaker and (this) stronger variants of approximate envy-freeness under a different valuation model, we call this stronger property EFX_0 .

For matroid rank valuations, all items with non-zero marginal values for an agent are also valued identically at 1, hence EF1 trivially implies EFX_+ ; Theorem 3.4 and Corollary 3.8 further guarantee the existence of an EFX_+ and PO allocation for any instance under this valuation class. However, we demonstrate with Example A.5 with $(0, 1)$ -OXS valuations in Appendix A.8 that even an EF1 and utilitarian optimal (hence PO) allocation may not satisfy the EFX_0 condition. ♦

Also, see Remark 5 in Appendix A.10 for connections to fair allocation “with charity” [Caragiannis et al. 2019a; Chaudhury et al. 2020].

3.2 MNW and Leximin Allocations

We saw in Section 3.1 that under matroid rank valuations, a simple iterative procedure allows us to reach an EF1 allocation while preserving utilitarian optimality. However, as we previously noted, such allocations are not necessarily leximin or MNW. In this subsection, we characterize the set of leximin and MNW allocations under matroid rank valuations. We start by showing that Pareto optimal allocations coincide with utilitarian optimal allocations when agents have matroid rank valuations. Intuitively, if an allocation is not utilitarian optimal, one can always find an ‘augmenting’ path that makes at least one agent happier but no other agent worse off.

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.

⁹However, both their examples establishing negative results for these sets of conditions on the valuation functions involve eliminating items with zero marginal value; their second example (for identical valuations) uses a non-submodular valuation function.

LEMMA 3.10 (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 any $X \in \mathcal{I}$ and $y \in E \setminus X$ such that $X \cup \{y\} \notin \mathcal{I}$.

THEOREM 3.11. *For matroid rank valuations, 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 3.10 implies that the circuit $C(\mathcal{I}, X, y)$ is contained in E_i , i.e.,

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

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 [Edmonds 1979].

To formally define an augmenting path, we define an auxiliary graph $G_{X^*} = (E, B_{X^*}^{(1)} \cup B_{X^*}^{(2)})$ where the set of arcs is given by

$$\begin{aligned} B_{X^*}^{(1)} &= \{(x, y) \mid y \in E \setminus X^* \wedge x \in C(\mathcal{O}, X^*, y) \setminus \{y\}\}, \\ B_{X^*}^{(2)} &= \{(y, x) \mid y \in E \setminus X^* \wedge x \in C(\mathcal{I}, X^*, y) \setminus \{y\}\}. \end{aligned}$$

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 $P = (y_0, x_1, y_1, \dots, x_s, y_s)$ in G_{X^*} with $y_0, y_1, \dots, y_s \notin X^*$ and $x_1, x_2, \dots, x_s \in X^*$, where X^* can be augmented by one element along the path, i.e.,

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

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

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 $t = 1, 2, \dots, s$, 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, $i(x_1), i(x_2), \dots, i(x_s)$, 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. $\square$

Theorem 3.11 above, along with Lemma 3.1, implies that both leximin and MNW allocations are utilitarian optimal. Next, we show that for the class of matroid rank valuations, leximin and MNW allocations are identical to each other; further, they can be characterized as the minimizers of any symmetric strictly convex function among all utilitarian optimal allocations.

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

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

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

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

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

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

It is not difficult to see that $\Phi : \mathbb{Z}^n \rightarrow \mathbb{R}$ is strictly convex if and only if $-\Phi$ is strictly concave. Examples of symmetric, strictly convex functions are the following:

$$\begin{aligned} \Phi(z_1, z_2, \dots, z_n) &\triangleq \sum_{i=1}^n z_i^2 && \text{for } z_i \in \mathbb{Z} \quad \forall i \in [n]; \\ \Phi(z_1, z_2, \dots, z_n) &\triangleq \sum_{i=1}^n z_i \ln z_i && \text{for } z_i \in \mathbb{Z}_{\geq 0} \quad \forall i \in [n]. \end{aligned}$$

For an allocation A , we define $\Phi(A) \triangleq \Phi(v_1(A_1), v_2(A_2), \dots, v_n(A_n))$.

We start by showing that given a non-leximin socially optimal allocation A , there exists an adjacent socially optimal allocation A' which is the result of transferring one item from a ‘happy’ agent j to a less ‘happy’ agent i . The underlying submodularity guarantees the existence of such allocation. We denote by χ_i the n -dimensional incidence vector where the j -th component of χ_i is 1 if $j = i$, and it is 0 otherwise.

LEMMA 3.12. *Suppose that agents have matroid rank valuations. Let A be a utilitarian optimal allocation. If A is not a leximin allocation, then there is another utilitarian optimal allocation A' such that*

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

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

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 3.11. 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 $\sum_{i \in N} |A'_i \Delta A_i^*|$ over all clean allocations with $\mathbf{s}(A') = \mathbf{s}(A)$. Assume also w.l.o.g. that $v_1(A'_1) \leq v_2(A'_2) \leq \dots \leq v_n(A'_n)$. We let $v_{j_1}(A_{j_1}^*) \leq v_{j_2}(A_{j_2}^*) \leq \dots \leq v_{j_n}(A_{j_n}^*)$. Since A^* lexicographically dominates A' , for the minimum index k with $v_k(A'_k) \neq v_{j_k}(A_{j_k}^*)$,

$$v_k(A'_k) < v_{j_k}(A_{j_k}^*). \quad (2)$$

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

$$v_i(A'_i) < v_i(A_i^*). \quad (3)$$

Indeed, if for all $i \in [k]$, $v_i(A'_i) \geq v_i(A_i^*)$, the k -th smallest value of realized valuations under A' is at least $v_{j_k}(A_{j_k}^*)$, contradicting with (2). Take the minimum index i satisfying (3). Since both A' and A^* are clean allocations, we have

$$|A'_i| = v_i(A'_i) < v_i(A_i^*) = |A_i^*|. \quad (4)$$

By minimality, for all $h \in [i - 1]$, $v_h(A'_h) \geq v_h(A_h^*)$. In fact, the equality

$$v_h(A'_h) = v_h(A_h^*) \quad (5)$$

holds for all $h \in [i - 1]$. Indeed if $v_h(A'_h) > v_h(A_h^*)$ for some $h \in [i - 1]$, then h -th smallest value of the realized valuations under A' would be strictly greater than that under A^* , yielding $s(A') >_L s(A^*)$, a contradiction.

Now, recall that the family of clean bundles $\mathcal{I}_h = \{ S \subseteq O \mid v_h(S) = |S| \}$ for $h \in N$ forms a family of independent sets of a matroid. By (I3) of the independent-set matroid axioms and by the inequality (4), 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 is allocated to some agent, i.e., $o_1 \in A'_{i_1}$ for some $i_1 \neq i$. Consider the following three cases:

- Suppose $v_{i_1}(A'_{i_1}) \geq v_i(A'_i) + 2$. Then, we obtain a desired allocation by transferring o_1 from i_1 to i .
- Suppose $v_{i_1}(A'_{i_1}) = v_i(A'_i) + 1$. Then by transferring o_1 from i_1 to i , we get another utilitarian optimal allocation with the same vector as $s(A')$, which has a smaller symmetric difference than $\sum_{i \in N} |A'_i \Delta A_i^*|$, a contradiction.
- Suppose $v_{i_1}(A'_{i_1}) \leq v_i(A'_i)$. We will first show that $v_{i_1}(A'_{i_1}) \leq v_{i_1}(A_{i_1}^*)$. By (5), this clearly holds if $i_1 \leq i$. Also, when $i_1 > i$, this means that $v_{i_1}(A'_{i_1}) = v_i(A'_i)$; thus $v_{i_1}(A'_{i_1}) \leq v_{i_1}(A_{i_1}^*)$, as otherwise the i -th smallest value of realized valuations under A' would be greater than that under A^* , contradicting that A^* is leximin. Further by the facts that $|A'_{i_1} \setminus \{o_1\}| < |A_{i_1}^*|$ and that both $A'_{i_1} \setminus \{o_1\}$ and $A_{i_1}^*$ are clean (i.e., independent sets of a matroid), there exists an item $o_2 \in A_{i_1}^* \setminus A'_{i_1}$ such that $v_{i_1}(A'_{i_1} \cup \{o_2\} \setminus \{o_1\}) = v_{i_1}(A'_{i_1})$. Again by utilitarian optimality of A' , o_2 is allocated to some agent, i.e., $o_2 \in A'_{i_2}$ for some $i_2 \neq i_1$.

Repeating the same argument and letting $i_0 = i$, we obtain a sequence of items and agents $(i_0, o_1, i_1, o_2, i_2, \dots, o_t, i_t)$ such that

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

See Figure 3 for an illustration of the sequence. If the same agent appears again, i.e., $i_h = i_{h'}$ for some $h < h' \leq t$, 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 i_t with $v_{i_t}(A'_{i_t}) \geq v_i(A'_i) + 2$. Exchanging items along the path, we get a desired allocation. $\square$

Fig. 3. The path $(i_0, o_1, i_1, o_2, i_2, \dots, o_t, i_t)$

We further observe that such adjacent allocation decreases the value of any symmetric strictly convex function (equivalently, increases the value of any symmetric strictly concave function). The proof is similar to that of Proposition 6.1 in Frank and Murota [2018], which shows the analogous equivalence over the integral base-polyhedron.

LEMMA 3.13. *Let $\Phi : \mathbb{Z}^n \rightarrow \mathbb{Z}$ be a symmetric strictly convex function and $\Psi : \mathbb{Z}^n \rightarrow \mathbb{Z}$ be a symmetric strictly concave function. Let A be a utilitarian optimal allocation. Let A' be another utilitarian optimal allocation such that $s(A') = s(A) + \chi_i - \chi_j$ for some $i, j \in [n]$ with $s(A)_j \geq s(A)_i + 2$. Then, $\Phi(A) > \Phi(A')$ and $\Psi(A) < \Psi(A')$.*

PROOF. Let $\beta = s(A)_j - s(A)_i \geq 2$, and $\mathbf{y} = \mathbf{s}(A) + \beta(\chi_i - \chi_j)$. Thus $\Phi(\mathbf{s}(A)) = \Phi(\mathbf{y})$ by symmetry of Φ . Define $\lambda = 1 - \frac{1}{\beta}$. We have $0 < \lambda < 1$ since $\beta \geq 2$. Observe that

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

which gives us the following inequality (from the strict convexity of Φ): $\Phi(\mathbf{s}(A)) = \lambda \Phi(\mathbf{s}(A)) + (1 - \lambda) \Phi(\mathbf{s}(A')) > \Phi(\mathbf{s}(A'))$. Since $-\Psi$ is symmetric strictly convex, the analogous proof shows $\Psi(A) < \Psi(A')$. $\square$

Now we are ready to prove the following.

THEOREM 3.14. *Let $\Phi : \mathbb{Z}^n \rightarrow \mathbb{R}$ be a symmetric strictly convex function, and $\Psi : \mathbb{Z}^n \rightarrow \mathbb{R}$ be a symmetric strictly concave function. Let A be some allocation. For matroid rank valuations, the following statements are equivalent:*

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

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

Assume towards a contradiction that $\mathbf{s}(A) \neq \mathbf{s}(A')$. By Theorem 3.11, we have $\text{USW}(A) = \text{USW}(A')$. By Lemma 3.12, 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 $\mathbf{s}(A')$ by 1 and increasing the value of the i -th element of $\mathbf{s}(A')$ by 1, where $\mathbf{s}(A')_j \geq \mathbf{s}(A')_i + 2$. Applying Lemma 3.13, we get $\Phi(\mathbf{s}(A')) > \Phi(\mathbf{s}(A''))$, which gives us the desired contradiction.

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

To prove $3 \Leftrightarrow 4$, let A be a leximin allocation, and let A' be an MNW allocation. Again, we will show that $\mathbf{s}(A')$ is the same as $\mathbf{s}(A)$, which by the uniqueness of the leximin valuation vector and symmetry of 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 $\bar{\mathbf{s}}(A)$ (respectively, $\bar{\mathbf{s}}(A')$) the vector of the non-zero components $v_i(A_i)$ (respectively, $v_i(A'_i)$) arranged in non-decreasing order. Assume towards a contradiction that $\bar{\mathbf{s}}(A) >_L \bar{\mathbf{s}}(A')$. Since A' maximizes the product $\text{NW}(A')$ when focusing on $N_{>0}(A')$ only, the value $\sum_{i \in N_{>0}(A')} \log v_i(A'_i)$ is maximized. However, $\Psi(\mathbf{x}) = \sum_{i=1}^{n'} \log x_i$ is a symmetric concave function for $\mathbf{x} \in \mathbb{Z}^n$ with each $x_i > 0$. Thus, by a similar argument as before, one can show that $\Psi(\bar{\mathbf{s}}(A')) > \Psi(\bar{\mathbf{s}}(A))$, a contradiction. This completes the proof. $\square$

REMARK 3 (THE PIGOU-DALTON PRINCIPLE). Maximizing any symmetric strictly concave function (equivalently, minimizing any symmetric strictly convex function) of valuations among all utilitarian optimal allocations is consistent with the well-known *Pigou-Dalton transfer principle* (see, e.g., Moulin [2004]) of welfare economics in a strong sense, and the proof is similar to that of Lemma 3.12. We formalize this notion in Appendix A.11.¹⁰ $\blacklozenge$

¹⁰We thank an anonymous reviewer for pointing out this connection.

The above theorem does not generalize to the non-binary case: Example 3.15 represents an instance where neither leximin nor MNW allocation is utilitarian optimal.

Example 3.15. 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

	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.

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. ■

Combining the above characterization with the results of Section 3.1, we get the following fairness-efficiency guarantee for matroid rank valuations.

COROLLARY 3.16. *For matroid rank valuations, any clean leximin or MNW allocation is EF1.*

PROOF. Since both leximin and MNW allocations are Pareto-optimal, they maximize the utilitarian social welfare, by Theorem 3.11. By Theorem 3.14 and the fact that the function $\Phi(A) \triangleq \sum_{i \in N} v_i(A_i)^2$ is a symmetric strictly convex function, any leximin or MNW allocation is a utilitarian optimal allocation that minimizes $\Phi(A)$ among all utilitarian optimal allocations; hence, if such an allocation is clean, it must be EF1 by Corollary 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 3.11, one can find a leximin or MNW allocation in polynomial time, by a reduction to the network flow problem.

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 under the $(0, 1)$ -OXS valuation class can be reduced to that of finding an integral balanced flow (or increasingly-maximal integer-valued 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 $c : A \rightarrow \mathbb{Z}$, 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. 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.

Fig. 4. 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.

See Figure 4 for an illustration of the network. We will show that an integral balanced flow $f : A \rightarrow \mathbb{Z}$ 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 3.11, 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_h^f , i.e., $f(s, h) = \sum_{i \in N_h} f(h, i) = v_h(A_h^f)$. Indeed if $v_h(A_h^f) > f(s, h)$, then it would contradict the optimality of the flow f ; and if $v_h(A_h^f) < f(s, h)$, it would contradict the fact that $v_h(A_h^f)$ is the value of a maximum-size matching between A_h^f 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 3.14, the leximin allocation A_f is also MNW. $\square$

In contrast with assignment valuations with binary marginal gains, the problem of computing a leximin or MNW allocation becomes intractable for weighted assignment valuations even when there are only two agents, the following theorem shows (the proof is in Appendix A.12).

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

5 DISCUSSION

We studied allocations of indivisible goods under submodular valuations with binary marginal gains in terms of the interplay among envy, efficiency, and various welfare concepts. We showed that three seemingly disjoint outcomes — minimizers of arbitrary symmetric strictly convex functions among utilitarian optimal allocations, the leximin allocation, and the MNW allocation — coincide in this class of valuations. Since the class of matroid rank functions is rather broad, our results can be applied to settings where agents' valuations are induced by a matroid structure. Beyond the domains described in this work, these include several others. For example, *partition matroids* model instances where agents' have access to different item types, but can only hold a limited number of each type (their utility is the total number of items they hold); a variety of other domains, such as spanning trees, independent sets of vectors, coverage problems and more admit a matroid structure (see Oxley [2011] for an overview). Indeed, a well-known result in combinatorial optimization states that *any* agent valuation structure where the greedy algorithm can be used to find the (weighted) optimal bundle, is induced by some matroid [Oxley 2011, Theorem 1.8.5].

We will conclude with additional implications of this work, some work in progress, and directions for further research.

Complete vs clean EF1, PO allocations. In Section 3.1, we showed that focusing on *clean* allocations helps us design an elegantly simple and polynomial-time algorithm (Algorithm 1) that computes an EF1 and utilitarian optimal (but potentially *incomplete*) allocation under matroid rank valuations — in other words, cleaning followed by further processing of a utilitarian optimal allocation is *sufficient* for achieving the EF1+PO combination. It is still an interesting open problem whether cleaning (and hence withholding some items) is *necessary* for achieving the desired fairness-efficiency combination. Intuitively, a complete, utilitarian optimal allocation may induce *avoidable* envy levels among agents¹¹ — agent j 's bundle A_j might include a subset S that contributes nothing to the overall welfare, i.e. $v_j(A_j \setminus S) = v_j(A_j)$, but makes another agent i envious of j up to more than 1 items, i.e. $v_i(A_i) < v_i(A_j \setminus \{o\})$ for every $o \in A_j$ but $v_i(A_i) \geq v_i(A_j \setminus (S \cup \{o'\}))$ for some $o' \in A_j \setminus S$. However, an instance that admits no complete, EF1, utilitarian optimal allocation (thus proving the necessity of cleaning in general) remains elusive. Please refer to Appendix A.4 for further comments on the interplay between cleanness and completeness.

More general valuation functions. An imperative line of future work is investigating which of our findings extend to more general valuation functions. There are several known extensions to matroid structures, with deep connections to submodular optimization [Oxley 2011, Chapter 11]. Matroid rank functions are submodular functions with binary marginal gains; however, general submodular functions (i.e. those with non-negative real marginal gains) admit some matroid structure which may potentially be used to extend our results to more general settings.

An obvious generalization of the matroid rank 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 matroid rank valuations) but the leximin and MNW allocations no longer coincide and leximin no longer implies EF1 (the details are in Appendix B).

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 data set, comparing

¹¹The observation that retaining items with zero marginal gain in an agent's bundle can make other agents envious is also a barrier to extending our results to a setting without *free disposal*, i.e. a setting where an agent incurs a cost for having an item removed from her bundle.

the performance of a heuristic extension of Algorithm 1 (Section 3.1) to real-valued individual-item utilities (weights) with Lipton et al. [2004]’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, described in detail in Appendix C, 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 matroids is a subclass of the well-known *gross substitutes* (GS) valuations [Gul and Stacchetti 1999; Kelso and Crawford 1982]. A promising research direction is to investigate PO+EF1 existence for GS valuations.

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 that she would receive her least favorite part), *equitability* (all agents have equal realized valuations), etc. (see, e.g. Caragiannis et al. [2019b]; Freeman et al. [2019] and references therein for further details) for matroid rank valuations. We present our results from a preliminary exploration of these questions in Appendix D. It is worthwhile to summarize here one of these results that extends a recent paper by Freeman et al. [2019]. This paper shows 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 generalize this result to $(0, 1)$ -OXS valuations: we first show that any EQ1 and PO allocation under the matroid rank valuation class, if it exists, is leximin, then invoke Corollary 3.16 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 in Appendix D.3).

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. [2016, 2018]; Jost [2006]). A popular measurement for how diverse a solution 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 3.14 also shows that, for matroid rank 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 interpretation to recent work on *soft* diversity framed as convex function optimization [Ahmed et al. 2017].

ACKNOWLEDGMENTS

Benabbou was supported by the ANR project 14-CE24-0007-01-Cocorico-CoDec, Chakraborty and Zick by a Singapore MOE grant (no. R-252-000-625-133) and a Singapore NRF Research Fellowship (no. R-252-000-750-733), and Ayumi Igarashi by the KAKENHI Grant-in-Aid for JSPS Fellows no. 18J00997 and JST, ACT-X. Most of this work was done when Chakraborty and Zick were employed at the National University of Singapore (NUS), and Igarashi was a research visitor at NUS, supported by Singapore MOE Grant R-252-000-625-133. The authors would like to thank Edith Elkind, Warut

Suksompong, Dominik Peters, and Tushant Jha for valuable insights and feedback on earlier versions of the paper. Thanks are also due to the anonymous reviewers of GAIW 2020, the Harvard CRCS Workshop on AI for Social Good 2020, and SAGT 2020 for their feedback on earlier versions of this paper.

REFERENCES

- 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.
- Haris Aziz, Serge Gaspers, Zhaozhong 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.
- Haris Aziz and Simon Rey. 2020. Almost Group Envy-free Allocation of Indivisible Goods and Chores. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI)*. 39–45.
- Moshe Babaioff, Tomer Ezra, and Uriel Feige. 2021. Fair and Truthful Mechanisms for Dichotomous Valuations. *Proceedings of the AAAI Conference on Artificial Intelligence* 35, 6 (May 2021), 5119–5126. <https://ojs.aaai.org/index.php/AAAI/article/view/16647>
- Maria Florina Balcan, Florin Constantin, Satoru Iwata, and Lei Wang. 2012. Learning valuation functions. In *Conference on Learning Theory*. 4–1.
- 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.
- Siddharth Barman, Sanath Kumar Krishnamurthy, and Rohit Vaish. 2018a. Finding fair and efficient allocations. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*. ACM, 557–574.
- Siddharth Barman, Sanath Kumar Krishnamurthy, and Rohit Vaish. 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.
- 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.
- Nawal Benabbou, Mithun Chakraborty, Xuan-Vinh Ho, Jakub Sliwinski, and Yair Zick. 2020. The Price of Quota-based Diversity in Assignment Problems. *ACM Transactions on Economics and Computation (TEAC)* 8, 3 (2020), 1–32.
- Ivona Bezáková and Varsha Dani. 2005. Allocating indivisible goods. *ACM SIGecom Exchanges* 5, 3 (2005), 11–18.
- 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.
- Sylvain Bouveret, Yann Chevaleyre, and Nicolas Maudet. 2016. Fair Allocation of Indivisible Goods. In *Handbook of Computational Social Choice*, Brandt, Felix and Conitzer, Vincent and Endriss, Ulle and Lang, Jérôme and Procaccia, Ariel D. (Ed.). Cambridge University Press, Chapter 12, 284–310.
- 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.
- Eric Budish. 2011. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy* 119, 6 (2011), 1061–1103.
- Ioannis Caragiannis, Nick Gravin, and Xin Huang. 2019a. Envy-Freeness Up to Any Item with High Nash Welfare: The Virtue of Donating Items. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*. ACM, 527–545.
- Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. 2019b. The Unreasonable Fairness of Maximum Nash Welfare. *ACM Transactions on Economics and Computation* 7, 3 (2019), 12:1–12:32.
- Elisa Celis, Amit Deshpande, Tarun Kathuria, and Nisheeth K. Vishnoi. 2016. How to be Fair and Diverse? (2016). arXiv:cs.LG/1610.07183 <https://arxiv.org/abs/1610.07183>
- Elisa Celis, Vijay Keswani, Damian Straszak, Amit Deshpande, Tarun Kathuria, and Nisheeth K. Vishnoi. 2018. Fair and diverse DPP-based data summarization. In *International Conference on Machine Learning*. PMLR, 716–725.
- Bhaskar Ray Chaudhury, Telikepalli Kavitha, Kurt Mehlhorn, and Alkmini Sgouritsa. 2020. A little charity guarantees almost envy-freeness. In *Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 2658–2672.
- Chicago Public Schools. 2017. Chicago Public Schools Policy Manual: Admissions Policy for Magnet, Selective Enrollment and Other Options For Knowledge Schools and Programs. (26 April 2017). Section 202.2. Board Report 17-0426-PO2. Available at <http://policy.cps.edu/Policies.aspx>.

- Beng-Huat Chua. 1991. Race relations and public housing policy in Singapore. *J. Arch. Plan. Res.* 8, 4 (1991), 343–354.
- Andreas Darmann and Joachim Schauer. 2015. Maximizing Nash product social welfare in allocating indivisible goods. *European Journal of Operational Research* 247, 2 (2015), 548–559.
- Yongheng Deng, Tien Foo Sing, and Chaohun Ren. 2013. The story of Singapore’s public housing: From a nation of home-seekers to a nation of homeowners. In *The Future of Public Housing: Ongoing Trends in the East and the West*, Jie Chen, Mark Stephens, and Yanyun Man (Eds.). Springer, Chapter 7, 103–121.
- 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.
- Jack Edmonds. 1979. Matroid Intersection. *Annals of Discrete Mathematics* 4 (1979), 39–49. Discrete Optimization I.
- Edmund Eisenberg and David Gale. 1959. Consensus of Subjective Probabilities: The Pari-Mutuel Method. *The Annals of Mathematical Statistics* 30, 1 (1959), 165–168.
- Duncan Foley. 1967. Resource allocation and the public sector. *Yale Economics Essays* 7 (1967), 45–98.
- András Frank and Kazuo Murota. 2018. Discrete Decreasing Minimization, Part I: Base-polyhedra with Applications in Network Optimization. (2018). arXiv:math.CO/1808.07600v3 <https://arxiv.org/abs/1808.07600v3>
- András Frank and Kazuo Murota. 2019. Discrete Decreasing Minimization, Part III: Network Flows. (2019). arXiv:math.CO/1907.02673v2 <https://arxiv.org/abs/1907.02673v2>
- 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.
- Mohammad Ghodsi, MohammadTaghi HajiAghayi, Masoud Seddighin, Saeed Seddighin, and Hadi Yami. 2018. Fair allocation of indivisible goods: Improvements and generalizations. In *Proceedings of the 2018 ACM Conference on Economics and Computation*. ACM, 539–556.
- 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.
- Frank Gul and Ennio Stacchetti. 1999. Walrasian equilibrium with Gross Substitutes. *Journal of Economic Theory* 87 (1999), 95–124. Issue 1.
- Daniel Halpern, Ariel D Procaccia, Alexandros Psomas, and Nisarg Shah. 2020. Fair division with binary valuations: One rule to rule them all. In *Proceedings of the 16th Conference on Web and Internet Economics (WINE)*. Springer, 370–383.
- F. Maxwell Harper and Joseph A. Konstan. 2015. The MovieLens Datasets: History and Context. *ACM Transactions on Interactive Intelligent Systems (TiiS)* 5, 4, Article 19 (2015), 19 pages.
- Lou Jost. 2006. Entropy and diversity. *Oikos* 113, 2 (2006), 363–375.
- Alexander S. Kelso and Vincent P. Crawford. 1982. Job Matching, Coalition Formation, and Gross Substitutes. *Econometrica* 50, 6 (1982), 1483–1504.
- Bernhard Korte and Jens Vygen. 2006. *Combinatorial Optimization: Theory and Algorithms*. Number 3 in Algorithms and Combinatorics. Springer-Verlag Berlin Heidelberg.
- Maria Kyropoulou, Warut Suksompong, and Alexandros A Voudouris. 2020. Almost envy-freeness in group resource allocation. *Theoretical Computer Science* 841 (2020), 110–123.
- Martin Lackner and Piotr Skowron. 2020. Approval-Based Committee Voting: Axioms, Algorithms, and Applications. (2020). arXiv:cs.GT/2007.01795 <https://arxiv.org/abs/2007.01795>
- Jérôme Lang and Piotr Krzysztof Skowron. 2016. Multi-Attribute Proportional Representation. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)*. 530–536.
- Benny Lehmann, Daniel Lehmann, and Noam Nisan. 2006. Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behavior* 55, 2 (2006), 270–296.
- Renato Paes Leme. 2017. Gross substitutability: An algorithmic survey. *Games and Economic Behavior* 106 (2017), 294 – 316.
- 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.
- Evangelos Markakis. 2017. Approximation Algorithms and Hardness Results for Fair Division with Indivisible Goods. In *Trends in Computational Social Choice*, Ulle Endriss (Ed.). AI Access, Chapter 12, 231–247.
- Hervé Moulin. 2004. *Fair division and collective welfare*. MIT press.
- James Munkres. 1957. Algorithms for the assignment and transportation problems. *J. Soc. Indust. Appl. Math.* 5, 1 (1957), 32–38.
- 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.
- 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.

- James Oxley. 2011. *Matroid Theory* (2nd ed.). Oxford University Press.
- Parliament of Singapore. Parliament Debates: Official Report. 1989. Better racial mix in HDB housing estates. (16 February 1989). Vol. 52, cols. 650-668.
- Sock Yong Phang and Kyunghwan Kim. 2013. Singapore's Housing Policies: 1960-2013. *Frontiers in Development Policy: Innovative Development Case Studies* (2013), 123-153.
- 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.
- Kimberly Quick. 2016. Chicago Public Schools: Ensuring Diversity in Selective Enrollment and Magnet Schools. (14 October 2016). The Century Foundation. <https://tcf.org/content/report/chicago-public-schools>.
- Sara Ramezani and Ulle Endriss. 2010. Nash Social Welfare in Multiagent Resource Allocation. In *Proceedings of the 11th Agent-Mediated Electronic Commerce. Designing Trading Strategies and Mechanisms for Electronic Markets (AMEC/TADA)*. 117-131.
- Erel Segal-Halevi and Warut Suksompong. 2019. Democratic fair allocation of indivisible goods. *Artificial Intelligence* 277 (2019), 103167.
- Amartya Sen. 1970. *Collective choice and social welfare*. Holden Day, San Francisco.
- Lloyd S. Shapley. 1962. Complements and substitutes in the optimal assignment problem. *Naval Research Logistics Quarterly* 9, 1 (1962), 45-48.
- Loo Lee Sim, Shi Ming Yu, and Sun Sheng Han. 2003. Public housing and ethnic integration in Singapore. *Habitat International* 27, 2 (2003), 293-307.
- 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.
- U.S. Department of Education, Office of Elementary and Secondary Education. 2017. Improving Outcomes for All Students: Strategies and Considerations to Increase Student Diversity. (19 January 2017). Washington, D.C. <https://www.ed.gov/diversity-opportunity>.
- Hal R Varian. 1974. Equity, envy, and efficiency. *Journal of Economic Theory* 9, 1 (1974), 63 - 91.
- Maisy Wong. 2014. Estimating the distortionary effects of ethnic quotas in Singapore using housing transactions. *Journal of Public Economics* 115 (2014), 131-145.

APPENDICES

A OMITTED PROOFS, EXAMPLES, AND REMARKS

A.1 Example where no EF1 allocation is utilitarian optimal

Example A.1. Consider an instance with three items and two agents Alice and Bob having additive valuations as described in Table 3. The unique allocation maximizing USW is the one which

	①	②	③
Alice:	1/4	3/8	3/8
Bob:	0	1/2	1/2

Table 3. An instance where no EF1 allocation is utilitarian optimal.

gives item 1 to Alice and items 2 and 3 to Bob for respective valuations $1/4$ and 1. However, even if either 2 or 3 were dropped from Bob's bundle, Alice would value it at $3/8 > 1/4$; hence, the unique utilitarian optimal allocation is not EF1 and so no EF1 allocation can be utilitarian optimal. ■

A.2 Proof of Lemma 3.1

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

PROOF. If an allocation A were not Pareto optimal, then there would be an allocation A' with $v_i(A'_i) > v_i(A_i)$ for at least one agent $i \in N$ and $v_j(A'_j) \geq v_j(A_j)$ for every other agent $j \in N \setminus \{i\}$. This implies that $\text{USW}(A') > \text{USW}(A)$ so that A cannot be utilitarian optimal. Moreover, the vector of realized valuations under A' arranged in non-decreasing order must also be lexicographically strictly greater than that of A (since the vectors must differ at the coordinate corresponding to agent i , if not earlier), hence A cannot be leximin either. Finally, suppose that A is an MNW allocation if possible, $N_{\max}$ being the subset of agents with strictly positive realized valuations: if $i \in N_{\max}$, then $\prod_{i \in N_{\max}} v_i(A'_i) > \prod_{i \in N_{\max}} v_i(A_i)$, contradicting the optimality of $NW(A)$, and if $i \in N \setminus N_{\max}$, that would contradict the maximality of $N_{\max}$. □

A.3 Proof of Proposition 3.2

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

PROOF. Consider subsets of items $T \subset S \subseteq O$ such that $S \setminus T = \{o_1, o_2, \dots, o_r\}$ where $r = |S \setminus T|$. 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:

$$\begin{aligned}
 v_i(S) - v_i(T) &= \sum_{t=1}^r (v_i(T \cup S_t) - v_i(T \cup S_{t-1})) \\
 &= \sum_{t=1}^r (v_i(T \cup S_{t-1} \cup \{o_t\}) - v_i(T \cup S_{t-1})) \\
 &= \sum_{t=1}^r \Delta_i(T \cup S_{t-1}; o_t).
 \end{aligned}$$

Since all marginal gains are binary, $\Delta_i(T \cup S_{t-1}; o_t) \geq 0$ for every $t \in [r]$, hence the above identity implies $v_i(S) - v_i(T) \geq 0$ for $S \supset T$, i.e. v_i is monotone.

Moreover, by setting $T = \emptyset$ and noting that $\Delta_i(T \cup S_{t-1}; o_t) \leq 1$ for every $t \in [r]$, we get $v_i(S) \leq v_i(\emptyset) + r = 0 + |S \setminus \emptyset| = |S|$. $\square$

A.4 Proof of Proposition 3.3

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

PROOF. The “if” part: Suppose, there is an item $o \in S$ such that $\Delta_i(S \setminus \{o\}; o) = 0$. Now, by Proposition 3.2, $v_i(S \setminus \{o\}) \leq |S \setminus \{o\}| = |S| - 1$ since $o \in S$. This implies that $v_i(S) = v_i(S \setminus \{o\}) + \Delta_i(S \setminus \{o\}; o) < |S|$. Thus, by contraposition, if $v_i(S) = |S|$, then $\Delta_i(S \setminus \{o\}; o) = 1 \forall o \in S$ since the marginal gain can be either 0 or 1, i.e. S is a clean bundle for i .

The “only if” part: As in the proof of Proposition 3.2, let $S = \{o_1, o_2, \dots, o_r\}$; define $S_0 = \emptyset$ and $S_t = \{o_1, o_2, \dots, o_t\}$ for each $t \in [r]$. By the definition of cleanness, $\Delta_i(S \setminus \{o_t\}; o_t) = 1 \forall t \in [r]$. Since $S_{t-1} \subseteq S \setminus \{o_t\}$ for every $t \in [r]$, $\Delta_i(S_{t-1}; o_t) \geq \Delta_i(S \setminus \{o_t\}; o_t) = 1$; moreover due to marginal gains in $\{0, 1\}$, we must have $\Delta_i(S_{t-1}; o_t) = 1$ for every $t \in [r]$. Hence, $v_i(S) = \sum_{t=1}^r \Delta_i(S_{t-1}; o_t) = r = |S|$. $\square$

At this point, it is worthwhile to discuss the interplay between cleanness and completeness. It is obvious that every instance admits a utilitarian optimal allocation that is complete (since, given any incomplete utilitarian optimal allocation, we can add withheld items to arbitrary bundles until completeness is reached, keeping the utilitarian social welfare unchanged) and also one that is clean (for analogous reasons). But, it may not be possible to achieve these two properties simultaneously under optimal utilitarian social welfare, as the next result shows.

COROLLARY A.2. *For an instance where all agents have matroid rank valuations, every utilitarian optimal allocation is clean as well as complete if and only if the maximum utilitarian social welfare for the instance under consideration is equal to the number of items.*

PROOF. Since $\sum_{i \in N} |A_i| \leq m$ for any allocation A , it follows readily from Proposition 3.2 that the maximum utilitarian social welfare of any instance cannot exceed the total number of items m under binary marginal gains.

If an instance admits a utilitarian optimal allocation A^* that is incomplete, i.e. $|A_0| > 0$, then $\text{USW}(A^*) = \sum_{i \in N} v_i(A_i) \leq \sum_{i \in N} |A_i| < m$. Likewise, if there is an A^* that is complete but not clean, then there is at least one agent i and one item $o \in A_i^*$ such that $\Delta_i(A_i^* \setminus \{o\}; o) = 0$; hence, $v_i(A_i^*) = v_i(A_i^* \setminus \{o\}) \leq |A_i^* \setminus \{o\}| = |A_i^*| - 1$. Thus, $\text{USW}(A^*) \leq |A_i^*| - 1 + \sum_{j \in N \setminus \{i\}} |A_j^*| \leq m - 1 < m$. Taking the contraposition proves the “if” part.

For the “only if” part, note that if an allocation A is complete, then $\sum_{i \in N} |A_i| = m$, and if it is clean under matroid rank valuations, then $v_i(A_i) = |A_i|$ for every $i \in N$ by Proposition 3.3. So, if a clean, complete allocation exists under matroid rank valuations, then $\text{USW}(A) = m$ which is the highest feasible utilitarian social welfare regardless of the specific valuation functions. Hence, A must be utilitarian optimal with a USW of m . $\square$

A simple example where each utilitarian optimal allocation is either not complete or not clean: $N = \{1, 2\}$; $O = \{o_1, o_2, o_3, o_4\}$; $v_1(S) = 1$ and $v_2(S) = \max\{2, |S|\}$ for every $S \in 2^O \setminus \emptyset$. It is easy to see that both v_1 and v_2 are both matroid rank valuations, and the USW cannot exceed $1 + 2 = 3 < 4$. In any utilitarian optimal allocation, any one item goes to agent 1 for a valuation of 1, any two of the remaining items go to agent 2 for a valuation of 2, and the final item may be arbitrarily allocated to either agent (not clean) or withheld (incomplete).

A.5 Example showing that Lipton et al. [2004]’s algorithm may not produce a Pareto optimal allocation under matroid rank valuations

The algorithm under consideration works as follows: in each iteration, a new item is allocated to an arbitrary agent not currently envied by any other agent; the envy graph is constructed by drawing a directed edge from an agent to every agent it envies; if a cycle forms in the graph, it is eliminated by transferring bundles from envied to envious agent on the (reverse) cycle, starting with the smallest cycle in case of overlapping cycles. We can augment the first step of the above algorithm with a natural heuristic: allocate the item under consideration to an agent that has the maximum marginal gain from it, breaking ties arbitrarily — for valuations with binary marginal gains, this is equivalent to giving the item to an agent whose marginal gain for it (given its current bundle) is 1, and to an arbitrary agent if none has non-zero marginal gain for it.

Example A.3. Consider 2 agents and 2 items such that $v_1(o_1) = v_1(o_2) = v_1(o_1, o_2) = 1$, $v_2(o_1) = v_2(o_1, o_2) = 1$ and $v_2(o_2) = 0$. Lipton et al. [2004]’s algorithm may assign o_1 to agent 1; then o_2 will be arbitrarily allocated, resulting in an allocation A with $v_1(A_1) = 1$ and $v_1(A_2) = 0$. This is Pareto dominated by $A'_1 = \{o_2\}$, $A'_2 = \{o_1\}$ where each agent realizes a valuation of 1. The myopic¹², sequential nature of the algorithm results in this undesirable outcome. ■

A.6 Example of a non-submodular valuation function that violates the transferability property

Example A.4. 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.7 Proof of Corollary 3.8

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

PROOF. Let A be a clean utilitarian optimal allocation that minimizes the sum of squares of the realized valuations among all utilitarian optimal allocations. We will show that A is EF1. Assume towards a contradiction that A is not EF1. Then, there is a pair of agents i, j such that i envies j up to more than 1 item. By Lemma 3.5, there is an item $o \in A_j$ such that $\Delta_i(A_j; o) = 1$. Let A^* be the allocation achieved by transferring o from j to i , everything else remaining the same. By Lemma 3.6 and the fact that A_j is clean, we have

$$v_i(A_i) + 2 \leq v_j(A_j),$$

which implies $\sum_{i \in N} v_i(A_i^*)^2 < \sum_{i \in N} v_i(A_i)^2$ proceeding exactly as in the proof of Theorem 3.4 — another contradiction. Hence, A must be EF1. □

A.8 Example of an EF1 and utilitarian optimal allocation violating EFX₀

Example A.5. There are two groups and four items o_1, o_2, o_3, o_4 . The first group N_1 has two members a_1, a_2 and the second group N_2 has three b_1, b_2, b_3 ; each individual has utility 1 for an item if and only if she is adjacent to it in the graph in Figure 5:

¹²The algorithm is *myopic* in the sense that the allocation of a new item in each iteration does not take into account its downstream impact on efficiency and is only geared towards maintaining the EF1 invariant.

Fig. 5. An instance that admits an EF1, utilitarian optimal allocation that is not EFX₀

The $(0, 1)$ -OXS valuation functions of groups N_1 and N_2 are denoted by $v_1(\cdot)$ and $v_2(\cdot)$ respectively. The allocation A where $A_1 = \{o_1\}$ and $A_2 = \{o_2, o_3, o_4\}$ is utilitarian optimal; it is also EF1 since $v_1(A_1) = 1 = v_1(A_2 \setminus \{o\})$ for $o \in \{o_2, o_3\}$, with $\Delta_1(A_2 \setminus \{o_4\}; o_4) = 0$, and $v_2(A_2) = 3 > 0 = v_2(A_1)$. A could be the output of Algorithm 1 and is clean and complete. However, A is not EFX₀ since $v_1(A_2 \setminus \{o_4\}) = 2 > 1 = v_1(A_1)$. ■

A.9 Remark on the definition of EFX₊ for general valuations

REMARK 4 (EFX₊ ALLOCATION). In this remark, we will point out an issue with extending the definition of envy-freeness up to the least positively valued good to settings with general (possibly non-additive) valuation functions. Recall that in Remark 2 in Section 3, we defined an allocation A of indivisible items to be EFX₊ if for every pair of agents i, j such that i envies j , $v_i(A_i) \geq v_j(A_j \setminus o)$ for every $o \in A_j$ such that i 's marginal valuation of o given A_j is *strictly positive*, i.e. $\Delta_i(A_j \setminus \{o\}) > 0$. This definition can make certain non-EF1 allocations trivially EFX₊, even under matroid rank valuations, if there is no single item in the envied agent's bundle for which the envious agent has positive marginal valuation. Consider two agents with matroid rank valuations, $v_1(S) = \min\{|S|, 2\}$ and $v_2(S) = |S|$, and any allocation A where agent 1 gets one item and agent 2 three items. Agent 2 does not envy agent 1 but $v_1(A_2 \setminus \{o\}) = 2 > 1 = v_1(A_1)$ for every $o \in A_2$, hence the allocation is not EF1. However, the condition for EFX₊ is vacuously satisfied since $\Delta_i(A_j \setminus \{o\}) = 0$ for every $o \in A_2$. This is inconsistent with the property that EFX is stronger than (i.e. implies) EF1, which holds for additive valuations. Note that, as we argued in Remark 2, EF1 implies EFX₊ under matroid rank valuations; we just showed with the above example that the converse is not true, implying that EF1 is a stronger property than EFX₊ for this valuation class. This is problematic since the EFX property was originally introduced as a strengthening of EF1 and it is reasonable to want to retain this relation beyond additive valuations!

A possible way of fixing this issue is by requiring that, for an allocation A to be EFX₊ under general valuations, it must first be EF1, i.e. if i is envious of j under A , then i must be EF1 of j and additionally the envy of i towards j can be eliminated by removing i 's least positively marginally valued item from j 's bundle. However, more fundamentally, this issue calls into question the applicability of EFX₊ as a fairness concept for non-additive valuations.

On the other hand, EFX₀ is stronger than EF1 for any monotone valuation $v(\cdot)$ with $v(\emptyset) = 0$. Example A.5 in Appendix A.8 already demonstrates that an EF1 allocation may not be EFX₀ even under $(0, 1)$ -OXS valuations. Conversely, if an allocation A is EFX₀, then for any pair of agents i and j , there is always an item $o \in A_j$ regardless of $\Delta_i(A_j \setminus \{o\})$ such that $v_i(A_i) \geq v_i(A_j \setminus \{o\})$, i.e. EFX₀ always implies EF1. ♦

A.10 Remark on cleanness and fair allocation “with charity”

REMARK 5 (THE WITHHELD SET). A remark about the withheld set is in order. The output A of Algorithm 1 is clean by construction and hence may be incomplete (Corollary A.2), leaving us with a non-empty withheld set A_0 . If completeness is not a requirement (i.e. there is free disposal and/or no stipulation of the form “All items must go!”), we can view the withheld set as “surplus” that can be set aside for future use or “charity.” This is similar in spirit to recent work on *EFX-with-charity* [Caragiannis et al. 2019a] and *EFX-with-bounded-charity* [Chaudhury et al. 2020].

For additive valuations, Caragiannis et al. [2019a] design an algorithm that computes an EFX_0 , *partial* allocation with at least half the optimal Nash welfare, running in polynomial time if it has oracle access to an MNW allocation; Chaudhury et al. [2020] consider general valuations and provide a way to set aside a charity bundle of size strictly less than the number of agents such that the allocation of the remaining items to the agents is EFX_0 and every agent weakly prefers her own bundle to the charity bundle — this allocation retains the high Nash welfare guarantee of Caragiannis et al. [2019a] for additive valuations.

It is easy to identify similar desirable properties in the output A of Algorithm 1: the (potentially partial) allocation to agents has optimal utilitarian social welfare and is $EF1$ (although not necessarily EFX_0); utilitarian optimality further dictates that every agent $i \in N$ weakly prefers its allocated bundle to the withheld set, i.e. $v_i(A_i) \geq v_i(A_0)$, and also has zero marginal value for any subset of withheld items given its bundle, i.e. $v_i(A_i \cup S) = v_i(A_i) \forall S \subseteq A_0; |A_0| = m - U^*$, where U^* is the optimal social welfare which is equal to the number of items allocated under any clean, utilitarian optimal allocation (by Proposition 3.3).

Although we did not start with the problem of withholding items to satisfy desiderata, this outcome emerged from our search for $EF1+PO$ allocations beyond additive valuations. ♦

A.11 Connections between Theorem 3.14 and the Pigou-Dalton principle

Here, we will elaborate on Remark 3 from Section 3.2. We begin with an introduction to the Pigou-Dalton principle (PDP). Consider a fair allocation instance — where agents have arbitrary valuation functions — which admits two allocations A and A' with the following properties: $v_1(A_1) + v_2(A_2) = v_1(A'_1) + v_2(A'_2)$; $v_i(A_i) = v_i(A'_i) \forall i \in N \setminus \{1, 2\}$; $v_1(A_1) < v_1(A'_1) < v_2(A_2)$, the first and last conditions implying $v_1(A_1) < v_2(A'_2) < v_2(A_2)$ as well. In other words, A' can be obtained from A by transferring some utility from the “richer” agent 2 to the “poorer” agent 1, everything else (including the utilitarian social welfare) remaining the same and without making agent 2 poorer under A' than agent 1 was under A . Then, PDP stipulates an ordering that (strictly) prefers A' , which reduces the inequality between agents 1 and 2 to A (the choice of agents is w.l.o.g.).

We know that the utilitarian social welfare function itself satisfies PDP weakly. We will now show that, if there are multiple utilitarian optimal allocations, one that is preferred to another by PDP also has a *strictly* higher value of an arbitrary symmetric, strictly concave function $\Psi : \mathbb{Z}^n \rightarrow \mathbb{R}$ applied to agents’ utilities.

PROPOSITION A.6. *Let $\Psi : \mathbb{Z}^n \rightarrow \mathbb{R}$ be an arbitrary symmetric, strictly concave function applied to agents’ valuations. If an instance where all agents have arbitrary monotone valuation functions admits two utilitarian optimal allocations A and A' with respective valuation vectors $\mathbf{r}(A)$ and $\mathbf{r}(A')$ such that $v_1(A_1) < v_1(A'_1) < v_2(A_2)$ and $v_i(A_i) = v_i(A'_i) \forall i \in N \setminus \{1, 2\}$, then $\Psi(\mathbf{r}(A')) > \Psi(\mathbf{r}(A))$.*

PROOF. Since both allocations have the same (optimal) utilitarian social welfare, it is obvious that

$$\begin{aligned} v_1(A_1) + v_2(A_2) &= v_1(A'_1) + v_2(A'_2) \\ \Rightarrow v_1(A'_1) - v_1(A_1) &= v_2(A_2) - v_2(A'_2). \end{aligned} \quad (6)$$

By definition, PDP prefers A' to A . Moreover,

$$\begin{aligned} v_1(A_1) &< v_1(A'_1) < v_2(A_2) \\ \Rightarrow 0 &< v_1(A'_1) - v_1(A_1) < v_2(A_2) - v_1(A_1). \end{aligned} \quad (7)$$

Let $\alpha \triangleq v_1(A'_1) - v_1(A_1) = v_2(A_2) - v_2(A'_2)$ by (6); $\beta \triangleq v_2(A_2) - v_1(A_1)$. By (7), $0 < \frac{\alpha}{\beta} < 1$.

Since the vector $\mathbf{r}(A')$ can be obtained from the vector $\mathbf{r}(A)$ by increasing the first entry by $v_1(A'_1) - v_1(A_1)$ and reducing the second by $v_2(A_2) - v_2(A'_2)$, we can write $\mathbf{r}(A') = \mathbf{r}(A) + \alpha(\chi_1 - \chi_2)$ where χ_i the n -dimensional incidence vector whose j -th component 1 if $j = i$, and 0 otherwise, as in Section 3.2.

Let us further define a vector $\mathbf{y} \triangleq \mathbf{r}(A) + \beta(\chi_1 - \chi_2)$. It is easy to see that $\mathbf{y}$ is a permutation of $\mathbf{r}(A)$ with entries 1 and 2 swapped, hence $\Psi(\mathbf{r}(A)) = \Psi(\mathbf{y})$ due to the symmetry of Ψ .

Finally, let $\lambda \triangleq 1 - \frac{\alpha}{\beta} \in (0, 1)$. Simple algebra shows that $\lambda \mathbf{r}(A) + (1 - \lambda)\mathbf{y} = \mathbf{r}(A')$. Hence, by the strict concavity of Ψ , we get $\Psi(\mathbf{r}(A')) > \lambda \Psi(\mathbf{r}(A)) + (1 - \lambda)\Psi(\mathbf{y}) = \lambda \Psi(\mathbf{r}(A)) + (1 - \lambda)\Psi(\mathbf{r}(A)) = \Psi(\mathbf{r}(A))$. $\square$

A.12 Proof of Theorem 4.2

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

PROOF. The reduction is similar to the hardness reduction for two agents with identical additive valuations [Nguyen et al. 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 $W = \{w_1, w_2, \dots, w_m\}$; 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 $W = \{w_1, w_2, \dots, w_m\}$. We create m items $1, 2, \dots, m$, two groups 1 and 2, and m individuals for each group where every individual has a weight w_j for item j . Observe that for 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 $v_1(A'_1) = v_2(A'_2)$. $\square$

B SUBMODULARITY WITH SUBJECTIVE BINARY GAINS

An obvious generalization of the matroid rank valuation function class is the class of submodular valuation functions 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 matroid rank valuations in Section 2.

Understandably, most of the properties of allocations under matroid rank 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 B.1. Suppose $N = [2]$; $O = \{o_1, o_2, o_3, o_4\}$; the valuations are additive with $v_1(\{o_4\}) = 0$, $v_1(\{o\}) = 1 \forall o \in O \setminus \{o_1\}$, 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, $v_2(A_1 \setminus \{o\}) = 6 > 3 = v_2(A_2)$ 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 the class of matroid rank valuations that is also evidenced by Example B.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 $NW(A) = 12$; 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 B.2. *For agents having submodular valuation functions 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 B.2 are natural extensions of Propositions 3.3 and Lemma 3.6 — Proposition B.3 and Lemma B.4 below, respectively:

PROPOSITION B.3. *For submodular valuations with subjective binary marginal gains defined by agent-specific positive constants $\lambda_i \forall i \in N$, 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 $S = \{o_1, o_2, \dots, o_r\}$ 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

$$v_i(S) = \sum_{t=1}^r \Delta_i(S_{t-1}; o_t) \leq \sum_{t=1}^r \lambda_i = \lambda_i r = \lambda_i |S|. \quad (8)$$

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 (8) 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. $\square$

LEMMA B.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 (8) in the proof of Proposition B.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. $\square$

We are now ready to prove Theorem B.2.

PROOF OF THEOREM B.2. Our proof non-trivially extends that of Theorem 3.2 of Caragiannis et al. [2019b]. 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, $\Delta_2(A_2 \setminus \{o\}; o) > 0$, i.e. $\Delta_2(A_2 \setminus \{o\}; o) = \lambda_2$.

Let us convert A to a new allocation A' by only transferring this item o from agent 2 to agent 1. Hence, $v_1(A'_1) = v_1(A_1) + \Delta_1(A_1; o) = v_1(A_1) + \lambda_1$, $v_2(A'_2) = v_2(A_2) - \Delta_2(A_2 \setminus \{o\}; o) = v_2(A_2) - \lambda_2$,

$v_i(A'_i) = v_i(A_i)$ for each $i \in N \setminus \{1, 2\}$. $NW(A)$ is positive since A is MNW and $N_{\max} = N$. Hence,

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

Here, the third equality comes from Proposition B.3 since A is clean, and the first inequality from Lemma B.4 due to our assumption. But $NW(A') > NW(A)$ 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. [2019b]'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 $i \in N \setminus N_{\max}$ 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 B.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 $i \in N \setminus N_{\max}$ must be envy-free up to one item towards any $j \in N_{\max}$. $\square$

C GENERAL ASSIGNMENT VALUATIONS

In this section, we address the fair and efficient allocation of items to agents who have general assignment or OXS valuations, as defined in Section 2.2. Recall that an agent with such a valuation function is equivalent to a group with multiple members each having an arbitrary non-negative weight for each item. As such, we will henceforth use the terms "group" and "agent" interchangeably.

We know that, for arbitrary non-negative monotone valuations, the classic *envy graph algorithm* due to Lipton et al. [2004] produces a complete, EF1 allocation that does not, however, come with any efficiency guarantee (except completeness, of course). The trick is to iterate over the items and allocate each to an agent that is currently not envied by any other agent (the existence of such an unenvied agent can be guaranteed by *de-cycling*, if necessary, the graph induced by a directed edge from every envious agent to every agent that it envies: see Lipton et al. [2004] for details).

Benabbou et al. [2019] focus on fair allocation to *types* that are, in fact, agents/groups with OXS valuations; they use a natural extension of this procedure that they denote by Algorithm **H**. In an iteration of Algorithm **H**, we do not give an arbitrary unallocated item to an arbitrary unenvied agent; instead, we find an item-agent pair having the maximum marginal utility among all currently unenvied agents and all unallocated items (breaking further ties uniformly at random, say), and allocate that item to that agent. Although this modification should, intuitively, improve efficiency, Benabbou et al. [2019] provide no formal guarantee in this regard; they evaluate the performance of Algorithm **H** in experiments where all agents have OXS valuations in terms of *waste* which they define as follows: under a complete allocation A , an item o is said to be *wasted* if it has positive marginal utility for some group h under A (i.e. $v_h(A_h \cup \{o\}) > v_h(A_h)$) but is allocated to another group h' (i.e. $o \in A_{h'}$) where it is either unassigned or assigned to a member $i \in N_{h'}$ with zero weight for it (i.e. $u_{i,o} = 0$), under the particular optimal matching of $A_{h'}$ to $N_{h'}$. The waste of a run of Algorithm **H** is defined as the percentage of the total number of items that are wasted under the complete allocation produced by Algorithm **H**.

Here, we ask whether the concept of *envy-induced transfers* (EIT) presented in Algorithm 1 for matroid rank valuations (Section 3.1) can be used to compute fair and efficient allocations (perhaps in some approximate sense) under more general monotone submodular valuation functions. This is motivated in part by the fact that the transferability property (Lemma 3.5), on which the EIT concept relies, characterizes any monotone submodular function and not just matroid rank valuations. In Algorithm 2, we delineate our work in progress in this vein: a heuristic scheme that extends Algorithm 1 to general OXS valuations.

Algorithm 2 retains the general principle of starting with a(n arbitrary) clean, utilitarian optimal allocation¹³ and iteratively eliminating envy by transferring an item from an envied bundle to an envious agent. For matroid rank valuations, the “donor” and the recipient of the transferred item have their valuations decreased and increased respectively by exactly 1 for any envy-induced transfer; this is no longer the case when we remove the binary marginal utilities restriction. Hence, such a transfer does not, in general, keep the utilitarian social welfare unchanged; the welfare is only constrained to never exceed its starting (optimal) value computed in line 1. As an approach to minimizing the loss in welfare/efficiency due to such transfers, we employ various heuristics in Algorithm 2:

- First, in each EIT step, we transfer the item that induces the minimal decrease in — or, equivalently, the maximal increase — in the welfare (see lines 3-6).
- Next, as the donor agent loses one of its items, it may develop a positive marginal utility for a currently withheld item; in that case, the item in A_0 for which it has maximal marginal utility is given to it (see lines 6-8).

¹³Maximizing the utilitarian social welfare is NP-hard when agents have general monotone submodular valuations but can be accomplished in polynomial time under the subclass gross substitutes valuations, assuming oracle access to each valuation function [Lehmann et al. 2006]. In particular, under OXS valuations (assuming that such a valuation function is specified in terms of the weights of each member of the group for all items), computing a utilitarian optimal allocation reduces to the polynomial-time solvable *assignment problem* or maximum sum-of-weights matching on a bipartite graph [Munkres 1957]; the result is automatically clean if we make sure that no item is assigned to an individual with zero weight for it.

Algorithm 2: Envy-Induced Transfers for general OXS valuations

```

1  Compute a clean, utilitarian optimal allocation.
2  /*Envy-Induced Transfers (EIT)*/
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  $\Delta_i(A_i; o) + \Delta(A_j \setminus \{o\}; o)$  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       $A_j \leftarrow A_j \setminus \{o\}; A_i \leftarrow A_i \cup \{o\}$ .
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          $A_i \leftarrow A_i \setminus \{o^*\}$ ; 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                  $A_k \leftarrow A_k \setminus \{o\}; o^* \leftarrow o$ .
16             else revoked = false.
17         end
18         if revoked = true then  $A_0 \leftarrow A_0 \cup \{o^*\}$ .
19     end
20 end

```

- Finally, if an agent (group) i acquires a new item o due to an envy-induced transfer, at most one of its previous items, say o^* , may become *unused*, i.e. it is no longer assigned to a member of the group under the new matching. This happens, for example, if i 's positive marginal utility for o with its previous bundle A_i was due to the fact that the member who was assigned item o^* has a higher weight for o than for o^* and no other member prefers o^* to its assigned item. In such a 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 do not yet have theoretical guarantees for Algorithm 2; but, if the EIT subroutine terminates, then the final allocation is EF1 and has zero waste (as defined above) by construction. To estimate the efficiency properties of our scheme, we ran numerical tests with it on a set of fair allocation instances based on a real-world data set.

In our experiments, we measure and compare the performances of Algorithm 2 and the procedure **H** as described above in terms of waste (as defined above) as well as the *price of fairness* (PoF) which we formally define as follows

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

where $A(P)$ is the allocation returned by a given procedure P (Algorithm 2 or Algorithm **H**) on a problem instance. Obviously, the PoF is bounded below by 1 for any instance and lower values are better.

The data set we use is *MovieLens-ml-1m* [Harper and Konstan 2015] which contains approximately 1,000,000 ratings (from 0 to 5) of 4,000 movies made by 6,000 users. To generate an instance of our allocation problem, we select 200 movies uniformly at random ($|O| = 200$) and then we only consider the users that rated at least one of these movies. Each such sample of 200 movies defines one run of our experiments. The users are our group-members and the movies our items. We generate agents/groups by partitioning users based on a demographic attribute; in fact, we use two attributes recorded in the data set, giving us two sets of allocation problem instances for each run:

- Gender: 2 agents (male or female, as recorded in the data set);
- Age: 7 agents representing the 7 age-groups recognized in the data set.

Moreover, for each such set (with 2 and 7 agents respectively), we adopt two models for the member-item weights or, equivalently, agents' valuation functions (raw and normalized ratings), giving us a 2×2 experimental design:

- RATINGS:

$$v_h(S) = \max \left\{ \sum_{u \in N_h} r_{u, \pi(u)} \mid \pi \in \Pi(N_h, S) \right\}$$

where r_{uo} is the user u 's rating of movie o ;

- NORM: $v'_h(S) = v_h(S)/v_h(O)$,

for every agent h (group N_h) and for any bundle of movies $S \subseteq O$. We provide the results, averaged over 50 runs, in Table 4.

	Attribute (#groups)	H [Benabbou et al. 2019; Lipton et al. 2004]		Algorithm 2	
		RATINGS	NORM	RATINGS	NORM
PoF	Age (7)	1.01	1.15	1.05	1.19
Waste		1.25%	0.20%	0%	0%
PoF	Gender (2)	1.00	1.02	1.00	1.03
Waste		0.00%	0.00%	0%	0%

Table 4. Experimental assessment of allocation procedures under OXS valuations.

We observe that Algorithm **H** has no guarantees on waste but, in practice, has negligible waste; and the waste appears to be lower for a lower number of agents in our experiments. In comparison, Algorithm 2 (which always terminated on its own for all our instances) is waste-free by design but has at least as much average PoF as Algorithm **H** in all our experiments.

D OTHER FAIRNESS CRITERIA UNDER MATROID RANK VALUATIONS

In the main paper, we have focused on Pareto optimal and EF1 allocations for the matroid rank valuation class. However, many other concepts have been defined and studied in the literature that formalize different intuitive ideas for what it means for an allocation of indivisible goods to be fair. In this appendix, we will investigate the implications of our results from the main paper for some alternative fairness notions.

D.1 Marginal envy-freeness up to one item

Caragiannis et al. [2019b] define an allocation A to be *marginally envy-free up to one item* or MEF1 if for every pair of agents $i, j \in N$ such that i envies j , there is an item $o \in A_j$ such that

$$v_i(A_i) \geq v_i(A_i \cup A_j \setminus \{o\}) - v_i(A_i).$$

Clearly, MEF1 is in general a relaxation of EF1: it coincides with EF1 for additive valuations and is implied by EF1 for submodular valuations, as Proposition D.1 below shows. Thus, our main results trivially prove the existence (and computational tractability) of MEF1 allocations with optimal welfare guarantees (hence Pareto optimality) for matroid rank valuations. However, Caragiannis et al. [2019b] already established the more general existence result that, for submodular valuations, every MNW allocation is MEF1.

PROPOSITION D.1. *For submodular valuation functions, an EF1 allocation is always MEF1 but the converse is not true, even if the functions have binary marginal gains.*

PROOF. Consider an EF1 allocation A , and pick an arbitrary pair of envious and envied agents $i, j \in N$. Let $o \in A_j$ be an item for which $v_i(A_i) \geq v_i(A_j \setminus \{o\})$. Then, since A_i and A_j are disjoint, submodularity implies

$$\begin{aligned} v_i(A_i \cup A_j \setminus \{o\}) - v_i(A_i) &\leq [v_i(A_i) + v_i(A_j \setminus \{o\})] - v_i(A_i) \\ &= v_i(A_j \setminus \{o\}) \\ &\leq v_i(A_i). \end{aligned}$$

Thus, A satisfies the definition of an MEF1 allocation.

The following is a counterexample to the converse. Suppose, $N = [2]$, $O = \{o_1, o_2, o_3, o_4\}$, and the valuation functions of the agents are

$$\begin{aligned} v_1(S) &= \begin{cases} |S| - 1 & \text{whenever } o_1 \in S \text{ and } |S| \geq 2; \\ |S| & \text{otherwise.} \end{cases} \\ v_2(S) &= |S| \quad \forall S \subseteq O. \end{aligned}$$

It is easy to see that both v_1 and v_2 are matroid rank functions. Under an allocation with bundles $A_1 = \{o_1\}$ and $A_2 = \{o_2, o_3, o_4\}$, $v_1(A_1) = 1$ and $v_2(A_2) = 3$. Moreover, $v_2(A_2 \cup A_1) - v_2(A_2) = 4 - 3 = 1 < 3 = v_2(A_2)$ and $v_1(A_1 \cup A_2 \setminus \{o_2\}) - v_1(A_1) = v_1(\{1, 3, 4\}) - 1 = 2 - 1 = 1 = v_1(A_1)$. So, the allocation is also MEF1. Agent 2 does not envy agent 1 since $v_2(A_1) = 1 < 3 = v_2(A_2)$; but, $v_1(A_2 \setminus \{o\}) = 2 > 1 = v_1(A_1)$ for every $o \in A_2$ so that agent 1 envies agent 2 up to more than 1 item. Thus, the allocation is not EF1. $\square$

D.2 Approximate proportionality

An allocation A of indivisible items O to agents N is said to be *proportional* if $v_i(A_i) \geq \frac{1}{n}v_i(O)$ for every agent $i \in N$. Since such an allocation may not be achievable in general,¹⁴ a relaxation is defined along the same lines as EF1: *proportionality up to one item*. Under additive valuations, i.e. when $\Delta_i(S; o) = v_i(\{o\})$ for every $i \in N$, every item o , and every bundle $S \subseteq O \setminus \{o\}$, an allocation A is defined to be proportional up to one item or PROP1 if, for every agent $i \in N$, there is at least one item not allocated to i , i.e. $o \in O \setminus A_i$, such that $v_i(A_i) \geq \frac{1}{n}v_i(O) - v_i(\{o\})$. For general non-additive valuations, it is debatable what the (subtractive) relaxation term on the right-hand side should be. For submodular valuations, we know that $\Delta_i(O \setminus \{o\}; o) \leq \Delta_i(S; o) \leq v_i(\{o\})$ for every agent $i \in N$, every item $o \in O$, and every bundle $S \subseteq O \setminus \{o\}$; as such, we could pick either $\Delta_i(O \setminus A_i \setminus \{o\}; o)$ or $v_i(\{o\})$ as the relaxation term for the strongest and weakest possible definition of PROP1 respectively. It is easy to see that, for general submodular valuations, any PO and EF1 allocation — if it exists — satisfies the weakest definition of proportionality up to 1

¹⁴This non-existence can be demonstrated even for matroid rank (in fact, binary additive) valuations by the following simple example: there are two agents and one item which each agent values at 1. Hence, $v_1(O) = v_2(O) = 1$, so that each agent must realize a value of at least $1/2$ for the allocation to be proportional. But in any allocation, at least one agent gets no item and hence realizes zero value.

item (Proposition D.3 below); the implications of the PO+EF1 property for stronger definitions of approximate proportionality are not clear yet, even for matroid rank valuations.

Definition D.2. For general valuation functions, an allocation A is called weakly proportional up to one item or WPROP1 if, for every agent $i \in N$,

$$v_i(A_i) \geq \frac{1}{n}v_i(O) - \max_{o \in O \setminus A_i} v_i(\{o\}).$$

PROPOSITION D.3. *If an instance of fair allocation with indivisible goods under general submodular valuations admits a Pareto optimal and EF1 allocation, then such an allocation is also WPROP1.*

PROOF. Consider a PO, EF1 allocation A admitted by an instance with submodular valuations. Hence, for any agent i ,

$$\begin{aligned} v_i(A_i) &\geq v_i(A_j \setminus \{o\}) && \text{for every } j \in N \setminus \{i\}, \text{ for some } o \in A_j \\ &= v_i(A_j) - \Delta_i(A_j \setminus \{o\}; o) \\ &\geq v_i(A_j) - v_i(\{o\}) && \text{since } \Delta_i(A_j \setminus \{o\}; o) \leq v_i(\{o\}) \text{ due to the submodularity of } v_i(\cdot). \end{aligned}$$

For every $j \in A_j$, consider an arbitrary such item o_j satisfying the above inequality. Since $o_j \in A_j \subseteq O \setminus A_i$, we have $v_i(\{o_j\}) \leq \max_{o \in O \setminus A_i} v_i(\{o\})$. Moreover, if $A_0 \neq \emptyset$, we must also have $v_i(A_i \cup A_0) = v_i(A_i)$, otherwise we could augment agent i 's bundle A_i with the withheld set A_0 and thereby increase her realized valuation without diminishing that of any other agent, contradicting the Pareto optimality of A . Thus, combining the above $n - 1$ inequalities together with this equality, we get

$$\begin{aligned} nv_i(A_i) &\geq v_i(A_i \cup A_0) + \sum_{j \in N \setminus \{i\}} v_i(A_j) - (n - 1) \times \max_{o \in O \setminus A_i} v_i(\{o\}) \\ &\geq v_i(\cup_{i=0}^n A_i) - (n - 1) \times \max_{o \in O \setminus A_i} v_i(\{o\}) && \text{due to the submodularity, hence subadditivity, of } v_i(\cdot) \\ &= v_i(O) - (n - 1) \times \max_{o \in O \setminus A_i} v_i(\{o\}) \\ &\geq v_i(O) - n \times \max_{o \in O \setminus A_i} v_i(\{o\}) && \text{since } \max_{o \in O \setminus A_i} v_i(\{o\}) \geq 0. \end{aligned}$$

Dividing both sides by n , we conclude that A is WPROP1. $\square$

Since a PO+EF1 allocation always exists and can be efficiently computed under matroid rank valuations (Theorem 3.4), the existence and computational tractability of a WPROP1 allocation for this function class immediately follows from Proposition D.3.

D.3 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].¹⁵ 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 EQ c 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 $v_i(A_i) \geq v_j(A_j \setminus S)$.¹⁶

¹⁵Note that if $A_j = \emptyset$ for some j , $v_i(A_i) \geq v_j(A_j)$ trivially. Hence the ordered pair (i, j) for any $i \in N \setminus \{j\}$ could never prevent the allocation from being EQ1.

¹⁶Again, if $|A_j| \leq c$ for some j , no ordered pair (i, j) for any $i \in N \setminus \{j\}$ could get in the way of the allocation being EQ c .

Freeman et al. [2019] show¹⁷ 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 under matroid rank valuations, an EQ1 and PO allocation, if it exists, is also EF1 — we achieve this by combining Theorem D.4 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.4. *For submodular valuations with binary marginal gains, any EQ1 and PO allocation, if it exists, is a leximin allocation.*

Hence, from Theorem 3.14, we further obtain that if an EQ1 and PO allocation exists under matroid rank valuations, it is also MNW and a minimizer of any symmetric strictly convex function of agents' realized valuations among all utilitarian optimal allocations. Moreover, Theorem D.4, together with Corollary 3.16, implies that if an EQ1 and PO allocation exists under matroid rank valuations, it must be EF1.

PROOF. Let the optimal USW for a problem instance under this valuation class be U^* ; also, suppose this instance admits an EQ1 and PO allocation A . The EQ1 property implies that for every pair of agents $i, j \in N$ such that $A_j \neq \emptyset$,

$$\begin{aligned} v_i(A_i) &\geq v_j(A_j \setminus \{o\}) \quad \text{for some } o \in A_j \\ &= v_j(A_j) - \Delta_j(A_j \setminus \{o\}; o) \\ &\geq v_j(A_j) - 1, \quad \text{since } \Delta_j(A_j \setminus \{o\}; o) \in \{0, 1\}. \end{aligned}$$

This inequality holds trivially and strictly if $A_j = \emptyset$. Thus, $\max_{i \in N} v_i(A_i) \leq \min_{i \in N} v_i(A_i) + 1$. In other words, there exist a non-negative integer $\alpha \leq U^*$ and a positive integer $n_0 \in [n]$ such that n_0 agents have valuations α each and the remaining agents, if any, have valuations $\alpha + 1$ each under allocation A , with $U^* = n_0\alpha + (n - n_0)(\alpha + 1) = n\alpha + n - n_0$. We can write the agents' realized valuations under A (with arbitrary tie-breaking) as the n -dimensional vector

$$s(A) = \left(\underbrace{\alpha, \alpha, \dots, \alpha}_{n_0 \text{ entries}}, \underbrace{\alpha + 1, \alpha + 1, \dots, \alpha + 1}_{n - n_0 \text{ entries}} \right).$$

If A were not leximin, there would be another allocation A' for which the corresponding valuation vector $s(A')$ would have an entry strictly higher than that of A at the same position, say $n' \in [n]$. If $n' \leq n_0$, then every entry of $s(A')$ from position n' is at least $\alpha + 1$, so the USW under A' is

$$\begin{aligned} U' &\geq (n' - 1)\alpha + (n - n' + 1)(\alpha + 1) \\ &= n\alpha + n - n' + 1 \\ &\geq n\alpha + n - n_0 + 1 \\ &= U^* + 1. \end{aligned}$$

¹⁷Freeman et al. [2019] use an example with 3 agents having binary additive valuations (Example 1). But it is easy to construct a 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 $v_2(o) = 0$ 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 of its items; even if agent 2 receives o_1 , her realized valuation of 1 will always be less than the above.

If $n_0 < n' \leq n$, then similarly,

$$\begin{aligned}
 U' &\geq n_0\alpha + (n' - n_0 - 1)(\alpha + 1) + (n - n' + 1)(\alpha + 2) \\
 &= n\alpha + 2n - n' - n_0 + 1 \\
 &\geq n\alpha + 2n - n - n_0 + 1 \\
 &= n\alpha + n - n_0 + 1 \\
 &= U^* + 1.
 \end{aligned}$$

In either case, we have a contradiction since U^* is the optimal utilitarian social welfare for this instance. Hence, A must be leximin. $\square$

We conjecture that a stronger result holds: under matroid rank valuations, the leximin allocation is optimally EQc for $c \in \{0, 1, \dots, m\}$ among all PO allocations. A proof or a counterexample remains elusive. We present this more formally as follows:

CONJECTURE D.5. *For a problem instance where all agents have submodular valuations with binary marginal gains, if c^* is the smallest $c \in [m]$ for which a leximin allocation is EQc, then the instance admits no PO allocation that is EQc for any $c < c^*$.*

D.4 The MMS guarantee

Let $\Pi(O)$ denote the collection of all n -partitions of the set of items O . Then, the *maximin share* of agent i is defined as:

$$\text{MMS}_i \triangleq \max_{(A_1, A_2, \dots, A_n) \in \Pi(O)} \min_{j \in N} v_i(A_j).$$

An allocation A is called MMS if $v_i(A_i) \geq \text{MMS}_i$ for every $i \in N$; for any approximation ratio $\alpha \in (0, 1]$, A is called α -MMS if $v_i(A_i) \geq \alpha \cdot \text{MMS}_i$ for every $i \in N$.

Ghodsi et al. [2018] showed that for agents with (general) submodular valuations, a $1/3$ -MMS allocation always exists and can be computed in polynomial time; it is not yet clear what the connections of the MMS guarantee with PO+EF1 allocations are, even for the subclass of submodular valuations with binary marginal gains, and whether a better approximation ratio can be achieved for this subclass.

In general, a PO and EF1 allocation may not be MMS even for $(0, 1)$ -OXs valuations, which we show by an example below:

PROPOSITION D.6. *For assignment valuations (hence submodular valuations) with binary marginal gains, a Pareto optimal and EF1 allocation may not be MMS.*

PROOF. Consider an augmentation to Example 3.9 from Section 3.1: we now have eight items and two groups $N_1 = \{a_1, a_2, \dots, a_6\}$ and $N_2 = \{b_1, b_2, b_3, b_4\}$ with $(0, 1)$ -OXs valuations; an agent's utility for an item is 1 if there is an edge between the nodes representing them in Figure 6, otherwise it is 0. The maximin share of group/agent N_2 is $\text{MMS}_2 = 3$ corresponding to the partition $\{A_1, A_2\}$ where $A_1 \supseteq \{o_3, o_4, o_5\}$, $A_2 \supseteq \{o_6, o_7, o_8\}$, and each of o_1, o_2 is arbitrarily included in either A_1 or A_2 : with A_1 (resp. A_2) allocated to N_2 , the members b_1, b_2, b_3 (resp. b_1, b_3, b_4) are assigned items o_3, o_4, o_5 (resp. o_7, o_8, o_6) respectively.

However, take the allocation A' with bundles $A'_2 = \{o_4, o_6\}$ and $A'_1 = O \setminus A'_2$: $v_1(A'_1) = 6$ with $o_1, o_2, o_3, o_5, o_7, o_8$ assigned to $a_1, a_2, a_3, a_4, a_5, a_6$ respectively; $v_2(A'_2) = 2$ with o_4, o_6 assigned to b_2, b_4 respectively; $v_1(A'_2) = 2$ with o_4, o_6 assigned to a_2, a_1 respectively; $v_2(A'_1) = 2$ with o_1, o_2 unassigned, either o_3 or o_7 assigned to b_1 , and either o_5 or o_8 assigned to b_3 respectively. Since $v_1(A_1) + v_2(A_2) = 8 = |O|$, the allocation is utilitarian optimal (hence PO); it is also envy-free since $v_1(A'_1) > v_1(A'_2)$ and $v_2(A'_2) = v_2(A'_1)$. But $v_2(A'_2) < \text{MMS}_2$. $\square$

Fig. 6. Problem instance used in the proof of Proposition D.6.

There are a few follow-up questions that are yet to be settled:

- Does an MNW/leximin allocation under matroid rank valuations (Theorem 3.14 and Corollary 3.16) have the MMS guarantee?¹⁸
- Is a PO and EF1 allocation under matroid rank valuations guaranteed to provide a constant approximation to the MMS guarantee, and if so, what is a tight bound on this approximation ratio?

We conclude with a positive existential result on a subclass of $(0, 1)$ -OXS valuations: binary additive valuations, i.e. for every $i \in N$, we have $v_i(\{o\}) \in \{0, 1\}$ for every $o \in O$, and $v_i(S) = \sum_{o \in S} v_i(\{o\})$ for every $S \subseteq O$. For this valuation subclass, we can safely assume that there are no *redundant* items in O , i.e. there is no $o \in O$, such that $v_i(\{o\}) = 0$ for every $i \in N$.¹⁹

PROPOSITION D.7. *If all agents have binary additive valuations, any Pareto optimal, EF1 allocation is also MMS.*

PROOF. It is easy to see that a PO allocation must be complete for this valuation class: if not, we would have an item $o^* \in A_0$ and, by our non-redundancy assumption, an agent i^* with $v_{i^*}(\{o^*\}) = 1$, hence we could improve agent i^* 's valuation at no cost to the other agents by allocating the item o^* to agent i^* . We will prove the required result by contraposition. Let A be an allocation that is PO but not MMS for an instance where all agents have binary additive valuations. Then, there is an agent i such that $v_i(A_i) < \mu$ where $\mu = \text{MMS}_i$. Since valuations are integers, $v_i(A_i) \leq \mu - 1$.

If $v_i(A_j) \leq \mu$ for every other agent $j \in N \setminus \{i\}$, then due to additivity and the emptiness of A_0 ,

$$v_i(O) = v_i(A_i) + \sum_{j \in N \setminus \{i\}} v_i(A_j) = v_i(A_i) + \sum_{j \in N \setminus \{i\}} v_i(A_j) < \mu + (n-1)\mu = n\mu.$$

¹⁸Notice that, in Example 3.9, $A_1 = \{o_1, o_2, o_3\}$ and $A_2 = \{o_4, o_5, o_6\}$ constitute the unique MNW/leximin allocation that is also MMS.

¹⁹Proposition D.7 also follows from Segal-Halevi and Suksompong [2019] who extend many of the fairness notions we study in this paper to a setting where agents are partitioned in k groups (ours is a special case with $k = n$ since each agent can be viewed as constituting its own group): their Lemmas 2.3 and 2.7(c) jointly imply that under binary additive valuations, every complete, EF1 allocation is MMS. We thank an anonymous reviewer for pointing us to this connection.

But, from the definition of the maximin share, there must exist a complete allocation A' such that $v_i(A'_j) \geq v_i(A'_i) = \mu$ for every $j \in N \setminus \{i\}$. This, together with additivity, implies that

$$v_i(O) = \sum_{j \in N} v_i(A'_j) \geq n\mu,$$

a contradiction. Hence, under the allocation A , there must be at least one agent, say $k \in N \setminus \{i\}$, such that $v_i(A_k) > \mu$.

Thus, for any item $o \in A_k$, we have

$$v_i(A_k \setminus \{o\}) = v_i(A_k) - v_i(\{o\}) \geq v_i(A_k) - 1 > \mu - 1 > v_i(A_i).$$

In other words, agent i is not envy-free of agent k up to 1 item, i.e. A cannot be EF1.

Hence, any PO and EF1 allocation must also be MMS. □