

Finding Fair and Efficient Allocations When Valuations Don’t Add Up*

Nawal Benabbou¹, Mithun Chakraborty², Ayumi Igarashi³, Yair Zick⁴

¹Sorbonne Université, CNRS, Laboratoire d’Informatique de Paris 6, LIP6 F-75005, Paris, France ,

²University of Michigan, Ann Arbor, USA , ³National Institute of Informatics, Tokyo, Japan ,

⁴University of Massachusetts, Amherst, USA

¹nawal.benabbou@lip6.fr , ²dcsmc@umich.edu , ³ayumi_igarashi@nii.ac.jp , ⁴yzick@umass.edu

Abstract

In this paper, we present new results on the fair and efficient allocation of indivisible goods to agents whose preferences correspond to *matroid rank functions*. This is a versatile valuation class, with several desirable properties (monotonicity, submodularity) which naturally models several 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 of agents’ valuations over utilitarian optimal outcomes. Moreover, for a subclass of these valuation functions based on maximum (unweighted) bipartite matching, we show that a leximin allocation can be computed in polynomial time.

1 Introduction

What is a *good* way of distributing a collection of indivisible goods amongst a population of agents who value them subjectively? This question brings to mind two broad issues — *welfare* and *efficiency* (overall utility of recipients/utilization of goods) on the one hand, and *fairness* (how each agent perceives her own share relative to those of others) on the other. These notions can be formalized in many ways and it is natural to ask whether an allocation satisfying several of such desiderata exists at all and, if yes, whether it can be computed efficiently. The answer depends on the agents’ *valuation functions* for bundles (subsets) of goods. Most existing literature focusing on the existence and computational

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

tractability of simultaneous fairness-efficiency guarantees in the allocation of indivisible goods [Caragiannis *et al.*, 2016; Barman *et al.*, 2018a; Barman *et al.*, 2018b, etc.] assumes that valuations are *additive*, i.e. the marginal benefit to an agent for adding an item to her bundle is a fixed constant for that agent-item pair, independent of the current bundle. At present, little is known in this respect beyond the additive setting. This is where our work comes in.

1.1 Our contributions and related work

In this paper, we consider valuations are given by *matroid rank functions*, i.e. each agent has an idiosyncratic matroid constraint [Oxley, 2011] over the items, and her value for a bundle is the size of a maximum independent set of the matroid included in the bundle. The alternative defining characteristics of this valuation class are: monotonicity, submodularity, and binary marginal gains, i.e. adding an item to an agent’s bundle either keeps her value unchanged or increases the value by exactly 1; and, if the marginal gain of adding an item to a bundle is zero, it must be zero if the item is added to a superset of that bundle. This class can naturally arise in many practical problems. Suppose that a government body wishes to assign public goods to individuals of different minority groups (e.g. kindergarten slots to children from different neighborhoods/socioeconomic classes, as in certain U.S. public school admission systems; flats in public housing estates to applicants of different ethnicities, as in Singapore [Benabbou *et al.*, 2018]), where each individual either approves or disapproves each good (also called *dichotomous preferences*). If diversity policies require fairness towards such pre-defined groups, we can model each group as an agent whose valuation function is based on *optimally matching* approved goods to its constituent individuals [Benabbou *et al.*, 2019]. We call such valuation functions $(0, 1)$ -OXS valuations (see Section 3). Certain additional constraints, when imposed on such matching problems, induce valuations that are no longer $(0, 1)$ -OXS but retain submodularity, e.g. hard limits on items due to *budgets* or exogenous *quotas* (socioeconomic status-based, ethnicity-based, etc.). The problem of assigning of courses to students [Budish, 2011] can be handled in this manner when students have underlying dichotomous preferences: each student desires a set of (limited) seats in courses allowed by her personal schedule, and her valuation is the maximum number of courses she can be as-

signed subject to logistical/academic constraints such as time clashes between courses. All such valuation functions belong to the broader class of matroid rank valuations.

Our fairness criteria are based on *envy*: an agent envies another if she believes that her bundle is worth less than that of the latter [Foley, 1967]. Envy-free (EF) allocations that are also Pareto optimal or PO (i.e. there is no reallocation that improves the valuation of one agent without worsening that of at least one other agent) or even complete (i.e. each item is allocated to at least one agent) are not guaranteed to exist for indivisible items. This leads to a relaxation called *envy-freeness up to one good* (EF1): for every pair of agents i and j , j 's bundle contains some item whose removal results in i not envying j . Budish [2011] was the first to formalize the EF1 concept, but it implicitly appears in Lipton *et al.* [2004], who design a polynomial-time algorithm that returns a complete, EF1 allocation for monotone valuations.

When all agents have additive valuations, Caragiannis *et al.* [2016] show that allocations that are both EF1 and PO exist, specifically the ones that maximize the product of agents' utilities — also known as *max Nash welfare* (MNW). Barman *et al.* [2018a] and Barman *et al.* [2018b] show that an allocation with these properties can be computed in (pseudo-)polynomial time. We establish similar guarantees for our non-additive valuation class of interest. Here is a summary of our main results:

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

Our results imply some known results for *binary additive valuations* (subsumed by matroid rank functions): Aziz and Rey [2019] show that the algorithm proposed by Darmann and Schauer [2015] with an MNW allocation in mind outputs a leximin allocation — in particular, this implies that the leximin and MNW solutions coincide for binary additive valuations; similar results are established by Halpern *et al.* [2020], who also show that the leximin/MNW allocation is group-strategyproof for agents with binary additive valuations. Finally, Babaioff *et al.* [2020] present a set of results similar to our own in addition to showing the existence of a strategyproof mechanism for matroid rank valuations. Our work was developed independently, and is different from a technical perspective.

Our omitted proofs, clarifying examples, and further discussions are available in the online full version of our paper [Benabbou *et al.*, 2020].

2 Model and definitions

For a positive integer r , let $[r]$ denote the set $\{1, 2, \dots, r\}$. We have 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 that is *normalized* (i.e. $v_i(\emptyset) = 0$) and *monotone* (i.e. $v_i(S) \leq v_i(T)$ whenever $S \subseteq T$); we assume polynomial-time oracle access to v_i for every $i \in N$.

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 and we call $v_i(A_i)$ agent i 's *realized valuation* under the allocation A . Given an allocation A , we denote by A_0 the set of unallocated or *withheld* items, i.e. $A_0 \triangleq O \setminus \bigcup_{i \in N} A_i$. 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 agent $i \in N$ if it contains no item $o \in S$ for which agent i has zero marginal gain (i.e., $\Delta_i(S \setminus \{o\}; o) = 0$, or equivalently $v_i(S \setminus \{o\}) = v_i(S)$); an allocation A is *clean* if each allocated bundle A_i is clean for the agent i that receives it. It is easy to ‘clean’ any allocation without changing any realized valuation by iteratively revoking items of zero marginal gain from respective agents and placing them in A_0 .

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

An allocation A' is said to *Pareto dominate* the allocation A if $v_i(A'_i) \geq v_i(A_i)$ for all agents $i \in N$ and $v_j(A'_j) > v_j(A_j)$ for some agent $j \in N$. An allocation is *Pareto optimal* (PO) if it is not Pareto dominated by any other allocation. There are several ways of measuring the welfare of an allocation [Sen, 1970]. Specifically, given an allocation A , (i) its *utilitarian social welfare* is $\text{USW}(A) \triangleq \sum_{i=1}^n v_i(A_i)$; (ii) its *egalitarian social welfare* is $\text{ESW}(A) \triangleq \min_{i \in N} v_i(A_i)$; (iii) its *Nash welfare* is $\text{NW}(A) \triangleq \prod_{i \in N} v_i(A_i)$. An allocation A is said to be *utilitarian optimal* if it maximizes $\text{USW}(A)$ among all allocations. Since the maximum attainable $\text{NW}(A)$ may be 0, we use the following refinement of the maximum Nash social welfare (MNW) criterion used in [Caragiannis *et al.*, 2016]: we find a maximal subset of agents, say $N_{\max} \subseteq N$, to which we can allocate bundles of positive values, and compute an allocation to agents in N_{\max} that maximizes the product of their realized valuations. If N_{\max} is not unique, we choose the one that results in the highest product of realized valuations. The *leximin* criterion is a lexicographic refinement of the maximum egalitarian welfare concept. 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 $\theta(A)$ the vector of the components $v_i(A_i)$ ($i \in N$) arranged in non-decreasing order. A *leximin* allocation A is one that maximizes the egalitarian welfare in a lexicographic sense, i.e., $\theta(A) \geq_L \theta(A')$ for any other allocation A' .

3 Matroid rank valuations

For technical details on matroids and matroid rank functions, the interested reader is referred to Oxley [2011]; here, we provide an alternative definition in terms of axioms that are easy to understand and more relevant to our allocation problem.

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 function v_i is *sub-modular* if single items contribute more to smaller sets than to larger ones, i.e. for all $S \subseteq T \subseteq O$ and all $o \in O \setminus T$, $\Delta_i(S; o) \geq \Delta_i(T; o)$. We say that v_i has *binary marginal gains* if $\Delta_i(S; o) \in \{0, 1\}$ for all $S \subseteq O$ and $o \in O \setminus S$. A *matroid rank valuation* can be defined as a monotone, sub-modular valuation function with binary marginal gains.

One important subclass of submodular valuations is the class of *assignment valuations*, also called OXS valuations [Shapley, 1958; Lehmann *et al.*, 2006; Leme, 2017]. Here, each agent $h \in N$ represents a group of individuals N_h (such as ethnic groups and genders); each individual $i \in N_h$ (also called a *member*) has a fixed non-negative weight $u_{i,o}$ for each item o . An agent h values a bundle S via a *matching* of the items to its members (i.e. each item is assigned to at most one member and vice versa) that maximizes the sum of weights [Munkres, 1957]; namely, $v_h(S) = \max\{\sum_{i \in N_h} u_{i,\pi(i)} \mid \pi \in \Pi(N_h, S)\}$, where $\Pi(N_h, S)$ is the set of matchings $\pi : N_h \rightarrow S$ in the complete bipartite graph with bipartition (N_h, S) . If $u_{i,o} \in \{0, 1\}$ for every $i \in N$ and $o \in O$, we call the corresponding subclass of OXS valuations $(0, 1)$ -OXS or assignment valuations with binary marginal gains. Fair allocation under such valuations was explored by Benabbou *et al.* [2019]; these functions form a subclass of matroid rank valuations and, in turn, subsume binary additive valuations [Barman *et al.*, 2018b].

3.1 Utilitarian optimal and EF1 allocation

Our first main contribution is to show that the existence of a PO+EF1 allocation [Caragiannis *et al.*, 2016] extends to the class of matroid rank valuations. In fact, we provide a surprisingly strong relation between efficiency and fairness: utilitarian optimality (stronger than Pareto optimality) and EF1 turn out to be compatible under matroid rank valuations. Moreover, such an allocation can be computed in polynomial time!

Theorem 3.1. *For matroid rank valuations, a utilitarian optimal allocation that is also EF1 exists and can be computed in polynomial time.*

Our result is constructive: we provide a way of computing the above allocation in Algorithm 1. The proof of Theorem 3.1 utilizes Lemmas 3.2 and 3.3 which shed light on the interesting interaction between envy and matroid rank valuations.

Lemma 3.2 (Transferability property). *For monotone sub-modular 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.*

Note that Lemma 3.2 holds for submodular functions with arbitrary real-valued marginal gains, and is trivially true for additive valuations. However, there exist non-submodular valuation functions that violate the transferability property, even when they have binary marginal gains.

It is easy to see that, for matroid rank valuations, $v_i(S)$ takes values in $\{0\} \cup [|S|]$ for any bundle S (hence $v_i(S) \leq |S|$) in general; in particular, S is a clean bundle for agent $i \in N$ if and only if $v_i(S) = |S|$. Using this insight, we can show the following result which holds when i 's envy towards j cannot be eliminated by removing one item.

Lemma 3.3. *For matroid rank valuations, if agent i envies agent j up to more than 1 item under an allocation A (i.e. $A_j \neq \emptyset$ and $v_i(A_i) < v_i(A_j \setminus \{o\})$ for every $o \in A_j$) and j 's bundle A_j is clean, then $v_j(A_j) \geq v_i(A_i) + 2$.*

The next important result is that, under matroid rank valuations, utilitarian social welfare maximization can be achieved in polynomial time.

Theorem 3.4. *For matroid rank valuations, one can compute a clean utilitarian optimal allocation in polynomial time.*

Roughly speaking, the key proof idea is the following: computing a clean utilitarian optimal allocation reduces to the problem of finding the largest common independent set of two matroids – one that is a *union matroid* representing collections of potentially overlapping clean bundles, one for each agent, and another that is a *partition matroid* modeling the constraint that every item goes to at most one agent or, equivalently, that no two bundles in an allocation overlap (see Korte and Vygen [2006] for details); this latter problem (the *matroid intersection problem* [Edmonds, 1979]) is known to be polynomial-time solvable, assuming polynomial-time value query oracles. We are now ready to prove Theorem 3.1.

Proof of Theorem 3.1. Algorithm 1 maintains optimal USW as an invariant and terminates on an EF1 allocation. Specifically, we first compute a clean allocation that maximizes the utilitarian social welfare. The EIT subroutine in the algorithm iteratively diminishes envy by transferring an item from the envied bundle to the envious agent; Lemma 3.2 ensures that there is always an item in the envied bundle for which the envious agent has a positive marginal gain.

Algorithm 1: Algorithm for finding utilitarian optimal EF1 allocation

```

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

```

Correctness: Each EIT step maintains the optimal utilitarian social welfare as well as cleanness: an envied agent's valuation diminishes exactly by 1 while that of the envious agent increases by exactly 1. Thus, if it terminates, the EIT subroutine retains the initial (optimal) USW and, by the stopping criterion, induces the EF1 property. To show that the algorithm terminates in polynomial time, we define the potential function $\Phi(A) \triangleq \sum_{i \in N} v_i(A_i)^2$. At each step of the algorithm, $\Phi(A)$ strictly decreases by 2 or a larger integer. To see this, let A' denote the resulting allocation after

reallocation of item o from agent j to i . Since A is clean, we have $v_i(A'_i) = v_i(A_i) + 1$ and $v_j(A'_j) = v_j(A_j) - 1$; since all other bundles are untouched, $v_k(A'_k) = v_k(A_k)$ for every $k \in N \setminus \{i, j\}$. Also, since i envies j up to more than one item under allocation A , $v_i(A_i) + 2 \leq v_j(A_j)$ by Lemma 3.3. Combining these, simple algebra gives us $\Phi(A') - \Phi(A) \leq -2$.

Complexity: By Theorem 3.4, computing a clean utilitarian optimal allocation can be done in polynomial time. The value of the non-negative potential function has a polynomial upper bound: $\sum_{i \in N} v_i(A_i)^2 \leq (\sum_{i \in N} v_i(A_i))^2 \leq m^2$. Thus, the EIT subroutine terminates in polynomial time. \square

Despite its simplicity, Algorithm 1 significantly generalizes Benabbou *et al.* [2019]’s PMURR algorithm (which ensures the existence of a non-wasteful EF1 allocation for $(0, 1)$ -OXS valuations) to matroid rank valuations. We note that Algorithm 1 may not produce an allocation that is MNW or leximin, even when agents have $(0, 1)$ -OXS valuations, but its above analysis readily gives us the following result.

Corollary 3.5. *For matroid rank valuations, any clean allocation A that minimizes $\Phi(A) \triangleq \sum_{i \in N} v_i(A_i)^2$ among all utilitarian optimal allocations is EF1.*

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 (the actual proof involves more subtle arguments in terms of the properties of *circuits* of matroids [Korte and Vygen, 2006]).

Theorem 3.6. *For matroid rank valuations, Pareto Optimal allocations are utilitarian optimal.*

Theorem 3.6 shows that every PO allocation is also utilitarian optimal; since leximin and MNW allocations are Pareto optimal [Caragiannis *et al.*, 2016; Bouveret *et al.*, 2016], they are utilitarian optimal as well. Next, we show that for the class of matroid rank valuations, leximin and MNW allocations are identical to each other; further, they can be characterized as the minimizers of any symmetric strictly convex function among all utilitarian optimal allocations.

A function $\Phi : \mathbb{Z}^n \rightarrow \mathbb{R}$ is *symmetric* if for any permutation $\pi : [n] \rightarrow [n]$, $\Phi(z_1, z_2, \dots, z_n) = \Phi(z_{\pi(1)}, z_{\pi(2)}, \dots, z_{\pi(n)})$, and is *strictly convex* if for any $\mathbf{x}, \mathbf{y} \in \mathbb{Z}^n$ with $\mathbf{x} \neq \mathbf{y}$ and $\lambda \in (0, 1)$ where $\lambda \mathbf{x} + (1 - \lambda)\mathbf{y}$ is an integral vector, $\lambda \Phi(\mathbf{x}) + (1 - \lambda)\Phi(\mathbf{y}) > \Phi(\lambda \mathbf{x} + (1 - \lambda)\mathbf{y})$. Examples of symmetric, strictly convex functions include: $\Phi(z_1, z_2, \dots, z_n) \triangleq \sum_{i=1}^n z_i^2$ for $z_i \in \mathbb{Z} \ \forall i$; $\Phi(z_1, z_2, \dots, z_n) \triangleq \sum_{i=1}^n z_i \ln z_i$ for $z_i \in \mathbb{Z}_{\geq 0} \ \forall i$.

Theorem 3.7. *Let $\Phi : \mathbb{Z}^n \rightarrow \mathbb{R}$ be a symmetric strictly convex function; let A be some allocation. For matroid rank valuations, the following are equivalent:*

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

Combining the above characterization with the Corollary 3.5, we get the following fairness-efficiency guarantee for matroid rank valuations, reminiscent of the result of Caragiannis *et al.* [2016] on MNW allocations for additive valuations.

Corollary 3.8. *For matroid rank valuations, any clean leximin or MNW allocation is EF1.*

4 Assignment valuations with binary gains

We now consider the special but practically important $(0, 1)$ -OXS subclass. For this class, we show that invoking Theorem 3.6, one can find a leximin or MNW allocation in polynomial time. The trick is to construct a flow network such that computing a leximin allocation in the original instance reduces to finding a *increasingly-maximal integer-valued flow* on the induced network; Frank and Murota [2019] recently gave a polynomial-time algorithm for this flow problem. We note that the complexity of the leximin allocation problem is open for the more general class of matroid rank valuations.

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

5 Discussion

We studied allocations of indivisible goods under matroid rank valuations in terms of the interplay among envy, efficiency, and various welfare concepts. The class of matroid rank functions is rather broad, and our results can be immediately applied to settings where agents’ valuations are induced by a matroid structure, even beyond the domains described in this work. 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. 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]. There are several known extensions to matroid structures, with deep connections to submodular optimization [Oxley, 2011, Chapter 11]. We focused here on submodular functions with binary marginal gains; however, general submodular functions admit some matroid structure which may potentially be used to extend our results to more general settings. Finally, it would be interesting to explore other fairness criteria such as proportionality, the maximin share guarantee, equitability. etc. (see, e.g. [Bouveret *et al.*, 2016] and references therein) for matroid rank valuations.

References

- [Aziz and Rey, 2019] H. Aziz and S. Rey. Almost group envy-free allocation of indivisible goods and chores. *CoRR*, abs/1907.09279, 2019.
- [Babaioff *et al.*, 2020] M. Babaioff, T. Ezra, and U. Feige. Fair and truthful mechanisms for dichotomous valuations. *arXiv preprint arXiv:2002.10704*, 2020.
- [Barman *et al.*, 2018a] S. Barman, S. K. Krishnamurthy, and R. Vaish. Finding fair and efficient allocations. In *Proc. EC*, pages 557–574. ACM, 2018.
- [Barman *et al.*, 2018b] S. Barman, S. K. Krishnamurthy, and R. Vaish. Greedy algorithms for maximizing Nash social welfare. In *Proc. AAMAS*, pages 7–13, 2018.
- [Benabbou *et al.*, 2018] N. Benabbou, M. Chakraborty, V. H. Xuan, J. Sliwinski, and Y. Zick. Diversity constraints in public housing allocation. *Proc. AAMAS*, pages 973–981, 2018.
- [Benabbou *et al.*, 2019] N. Benabbou, M. Chakraborty, E. Elkind, and Y. Zick. Fairness towards groups of agents in the allocation of indivisible items. In *Proc. IJCAI*, pages 95–101, 2019.
- [Benabbou *et al.*, 2020] N. Benabbou, M. Chakraborty, A. Igarashi, and Y. Zick. Finding fair and efficient allocations when valuations don’t add up. *arXiv preprint arXiv:2003.07060*, 2020.
- [Bouveret *et al.*, 2016] S. Bouveret, Y. Chevaleyre, and N. Maudet. Fair allocation of indivisible goods. In Brandt, Felix and Conitzer, Vincent and Endriss, Ulle and Lang, Jérôme and Procaccia, Ariel D., editor, *Handbook of Computational Social Choice*, chapter 12, pages 284–310. Cambridge University Press, 2016.
- [Budish, 2011] E. Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103, 2011.
- [Caragiannis *et al.*, 2016] I. Caragiannis, D. Kurokawa, H. Moulin, A. D. Procaccia, N. Shah, and J. Wang. The unreasonable fairness of maximum Nash welfare. In *Proc. EC*, pages 305–322. ACM, 2016.
- [Darmann and Schauer, 2015] A. Darmann and J. Schauer. Maximizing Nash product social welfare in allocating indivisible goods. *European Journal of Operational Research*, 247(2):548–559, 2015.
- [Edmonds, 1979] J. Edmonds. Matroid intersection. *Annals of Discrete Mathematics*, 4:39–49, 1979. Discrete Optimization I.
- [Foley, 1967] D. Foley. Resource allocation and the public sector. *Yale Economics Essays* 7, pages 45–98, 1967.
- [Frank and Murota, 2019] A. Frank and K. Murota. Discrete Decreasing Minimization, Part III: Network Flows. *arXiv e-prints*, Sep 2019.
- [Halpern *et al.*, 2020] D. Halpern, A. D. Procaccia, A. Psomas, and N. Shah. Fair division with binary valuations: One rule to rule them all. Technical report, Tech. rep, 2020.
- [Korte and Vygen, 2006] B. Korte and J. Vygen. *Combinatorial Optimization: Polyhedra and Efficiency*. Algorithms and Combinatorics. Springer, 2006.
- [Lehmann *et al.*, 2006] B. Lehmann, D. Lehmann, and N. Nisan. Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behavior*, 55(2):270–296, 2006.
- [Leme, 2017] R. P. Leme. Gross substitutability: An algorithmic survey. *Games and Economic Behavior*, 106:294–316, 2017.
- [Lipton *et al.*, 2004] R. J. Lipton, E. Markakis, E. Mossel, and A. Saberi. On approximately fair allocations of indivisible goods. In *Proc. EC*, pages 125–131. ACM, 2004.
- [Munkres, 1957] J. Munkres. Algorithms for the assignment and transportation problems. *Journal of the Society for Industrial and Applied Mathematics*, 5(1):32–38, 1957.
- [Oxley, 2011] J. Oxley. *Matroid Theory*. Oxford Graduate Texts in Mathematics. OUP Oxford, 2011.
- [Sen, 1970] A. Sen. *Collective choice and social welfare*. Holden Day, San Francisco, 1970.
- [Shapley, 1958] L. S. Shapley. Complements and substitutes in the optimal assignment problem. *ASTIA Document*, 1958.