

Yankee Swap: A Fast and Simple Fair Allocation Mechanism for Matroid Rank Valuations

Vignesh Viswanathan

University of Massachusetts, Amherst
Amherst, USA
vviswanathan@umass.edu

Yair Zick

University of Massachusetts, Amherst
Amherst, USA
yzick@umass.edu

ABSTRACT

We study fair allocation of indivisible goods when agent valuations are matroid rank functions (MRFs). Our main contribution is a simple algorithm based on the colloquial Yankee Swap procedure that computes provably fair and efficient Lorenz dominating allocations. While there exist polynomial time algorithms to compute fair and efficient allocations for MRF valuations, we improve on them in two ways: (a) Our approach is easy to understand and does not use complex matroid optimization algorithms as subroutines. (b) Our approach is scalable; it is provably faster than all known algorithms to compute Lorenz dominating allocations. These two properties are key to the adoption of algorithms in any real fair allocation setting; our contribution brings us one step closer to this goal.

KEYWORDS

Fair Allocation, Indivisible Goods, Matroid Rank Valuations

ACM Reference Format:

Vignesh Viswanathan and Yair Zick. 2023. Yankee Swap: A Fast and Simple Fair Allocation Mechanism for Matroid Rank Valuations. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), London, United Kingdom, May 29 – June 2, 2023*, IFAAMAS, 9 pages.

“It Sounds Mean”

– Angela Martin, “The Office” S2E10

1 INTRODUCTION

Fair allocation of indivisible goods is an extremely popular problem in the EconCS community. We would like to assign a set of indivisible *goods* to a set of *agents* who express *preferences* over the sets of goods (or *bundles*) they receive. For example, consider the problem of assigning course slots to students [11, 12]. Each student i has a preference over the set of classes they are assigned. Student preferences have implicit constraints on their structure: for example, a student may not take two classes if their schedules overlap (they will simply drop one if assigned conflicting classes). Similarly, students gain no value from receiving two seats in the same class. In addition, universities often impose bounds on the number of classes that a student may take in a given semester, or on the number of classes a student takes from a given major. We wish to identify an assignment of course seats to students that is both *efficient* – no student wants any additional available slots – and *fair* – no student would prefer another student’s assignment to their own. Course allocation can thus be naturally cast as an

instance of a fair allocation problem. To this end, one might wish to implement some fair allocation mechanism from the literature. However, the algorithms proposed in the fair division literature are becoming increasingly complex, which often precludes their consideration in university-wide applications. Consider for example the CourseMatch algorithm [12], used to assign MBA students to classes at the Wharton School of Business. Budish et al. state that “To find allocations, CourseMatch performs a massive parallel heuristic search that solves billions of mixed-integer programs to output an approximate competitive equilibrium in a fake-money economy for courses”. This framework has been applied to the course allocation system at the UPenn Wharton School of Business, which admits approximately 1700 students to roughly 350 courses. To understand the system, students are referred to nine instructional videos and a 12-page manual. While this system may be appropriate for a specialized MBA program, it may not be as effective for university-wide application, especially in settings with non-expert end-users. Spliddit [17] is another application of fair allocation mechanisms to real-world instances; however, it too does not scale well, as its underlying mechanism solves a mixed integer-linear program to find a Nash-welfare maximizing allocation [13].

Indeed, without any restriction on student preferences, computing a fair and efficient allocation is computationally intractable. However, under some reasonable assumptions on student preferences, we can compute fair and efficient allocations in polynomial time. If we assume that students simply want to take as many classes as they are allowed to, subject to scheduling constraints and course limits, then student preferences induce *binary submodular* valuations [8, 9]. Submodular valuations exhibit decreasing returns to scale: the larger the bundle agents already have, the less marginal gain they get from additional goods. Under binary submodular valuations, each agent values each additional good at either 1 or 0.

Since these valuations correspond to the rank function of some matroid, they are commonly referred to as *matroid rank functions* (MRFs) [24]. MRFs are highly structured, a structure that has been exploited in the optimization literature [19, 24] and more recently, in fair allocation [2, 6, 7, 9]. Most notably, Babaioff et al. [2] show that when agents have MRF valuations, there is a truthful poly-time algorithm to compute a leximin, envy free up to any good allocation that maximizes both the utilitarian social welfare and Nash social welfare. However, analysis of their algorithm places its runtime at roughly $O(n^6 m^{9/2})$ time (where n and m are the number of agents and goods respectively) which significantly hinders scalability. Moreover, their algorithm uses complex matroid optimization algorithms as subroutines; this prevents non-expert users from understanding the algorithm. Both issues are detrimental to the deployment of such an algorithm for course allocation. Thus, our main goal is to

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom. © 2023 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

develop a simple and fast algorithm to compute fair and efficient allocations under binary submodular valuations.

1.1 Our Contribution

Our main contribution is the introduction of a novel algorithm (known colloquially as *Yankee Swap*¹) for computing fair and efficient allocations for agents with MRF valuations.

More specifically, Yankee Swap computes prioritized Lorenz dominating allocations. When agents have binary submodular valuations, prioritized Lorenz dominating allocations are known to be leximin, envy free upto any good (EFX) and maximize both the utilitarian social welfare and Nash social welfare. In addition, randomizing over agent priorities can result in *ex-ante* fairness guarantees such as *ex-ante* envy-freeness and *ex-ante* proportionality. Finally, prioritized Lorenz dominating allocations can be computed in a strategyproof manner.

Yankee Swap is similar in spirit to the well known round robin algorithm. In the round robin algorithm, agents initially start with an empty bundle and proceed in rounds, picking a good in each round one by one from the pool of unallocated goods. Agents take sequential actions under Yankee Swap as well. Unlike the round robin algorithm, agents have the power to steal goods from other agents if they do not like any unassigned good. The agents who have goods stolen from them make up for their loss by either taking an unassigned good or by stealing a good from someone else. This procedure induces *transfer paths* — an agent steals a good from someone, who potentially steals a good from someone else and this goes on until someone takes an unassigned good. The main difference between our method and the colloquial Yankee Swap is that we only allow an agent to steal a good when a transfer path exists. This means that the utility of every agent on the transfer path remains the same (except for the agent that initiates it, whose utility increases by 1).

While the algorithm itself is remarkably simple, the analysis is decidedly non-trivial. First, we show via a combinatorial argument that transfer paths must exist (Lemma 3.1). Next, we show that Yankee Swap outputs balanced allocations, in the sense that once an agent cannot initiate transfer paths, the agents that hold goods they want have bundles of approximately the same size (Lemma 3.4); using these facts and careful analysis, we prove Theorem 3.5. We then turn to analyzing the most complex part of Yankee Swap: computing transfer paths. To do so, we construct a *good exchange graph*, and find shortest paths from the goods owned by the least utility agent and an unassigned good (Algorithm 2); we show that such paths on the good exchange graph correspond to valid path transfers (Theorem 3.9). We conclude our analysis by comparing the worst-case runtime of Yankee Swap to the current state of the art. Yankee Swap runs in $O((n+m)(n+\tau)m^2)$ time; τ is the time it takes to compute the valuation $v_i(S)$ for any bundle $S \subseteq G$ and agent $i \in N$. This is a significant speedup compared to Babaioff et al.’s runtime of $O(n^6 m^{7/2} (m+\tau) \log nm)$ (Section 3.5). We further note that, when $m = \Theta(n)$, our algorithm computes a MAX-USW allocation faster than the matroid intersection based method used by Benabbou et al. [9] and Babaioff et al. [2].

¹Yankee swap is also known as “Nasty Christmas” or “White Elephant”. See <https://youtu.be/19ulSNSRKYU> for a discussion.

We are currently in the process of implementing Yankee Swap for assigning classes at the University of Massachusetts, Amherst, working with the university registrar’s office, as well as the computer science department at the University of Massachusetts, Amherst.

1.2 Related Work

Binary valuation functions (otherwise called *dichotomous preferences*) have been studied in various contexts in the economics and computer science literature. More specifically, binary valuations have been studied in mechanism design [10, 23], auctions [3, 22], and exchange problems [1, 27].

Binary valuations have also been extensively studied in fair allocation. Halpern et al. [18] and Suksompong and Teh [29] study fair allocation in the restricted setting of binary additive valuations. Barman and Verma [6] study the computation of the maximin share under binary submodular valuations. Barman et al. [4], Darmann and Schauer [16] and Barman and Verma [5] study the computation of a max Nash welfare allocation under various binary valuation classes. Our work is not the first to explore transfer paths: Barman et al. [4], Suksompong and Teh [29] and Barman and Verma [6] also utilize transfer path techniques in their algorithm design.

Lastly, Benabbou et al. [9] and Babaioff et al. [2] study the computation of fair and efficient allocations under matroid rank valuations. Benabbou et al. [9] present the first positive algorithmic result for the valuation class by showing that a utilitarian social welfare maximizing and envy free up to one good allocation can be computed in polynomial time. This result was later significantly improved on by Babaioff et al. [2] whose work we discuss in detail in Section 2.1.

2 PRELIMINARIES

We use $[t]$ to denote the set $\{1, 2, \dots, t\}$. For the sake of readability, for a set A and a good g , we replace $A \setminus \{g\}$ (resp. $A \cup \{g\}$) with $A - g$ (resp. $A + g$).

We have a set of n agents $N = [n]$ and a set of m goods $G = \{g_1, g_2, \dots, g_m\}$. Each agent i has a *valuation function* $v_i : 2^G \mapsto \mathbb{R}_+$ — valuation function $v_i(S)$ corresponds to the value agent i has for the bundle of goods S . We let $\Delta_i(S, g) \triangleq v_i(S + g) - v_i(S)$ be the marginal utility of agent i from receiving the good g , given that they already own the bundle S . Unless otherwise stated, we assume that v_i is a *matroid rank function* (MRF). Due to their equivalence, we use binary submodular valuation and matroid rank function interchangeably. More formally, a function v_i is a matroid rank function if (a) $v_i(\emptyset) = 0$, (b) for every $S \subseteq G$ and every $g \in G$, $\Delta_i(S, g) \in \{0, 1\}$, and (c) v_i is submodular: for every $S \subseteq T \subseteq G$ and every $g \in G \setminus T$, $\Delta_i(S, g) \geq \Delta_i(T, g)$. Since there may not be a polynomial space representation of these valuation functions, we assume oracle access to each v_i : given a bundle of goods $S \subseteq G$, we can compute $v_i(S)$ in at most time τ .

An *allocation* is a partition of the set of goods $X = (X_0, X_1, \dots, X_n)$ where each agent i receives the bundle X_i , and X_0 consists of the unallocated goods. An allocation is *non-redundant* (or *clean*) if for every agent $i \in N$, and every good $g \in X_i$, $v_i(X_i) > v_i(X_i - g)$. Benabbou et al. [9] show that for MRF valuations, this is equivalent to having $v_i(X_i) = |X_i|$ for every $i \in N$. We sometimes refer to $v_i(X_i)$ as the *utility* (or *value*) of i under the allocation X . For ease of analysis, we treat 0 as an agent whose valuation function

is $v_0(S) = |S|$; this valuation function is trivially an MRF. Due to the choice of v_0 , any clean allocation for the set of agents N is also trivially clean for the set of agents $N + 0$. However, none of the fairness notions we discuss consider the (dummy) agent 0. Several fairness desiderata have been proposed and studied in the literature; three of them stand out:

Envy-Freeness: An allocation is *envy free* if no agent prefers another agent’s bundle to their own. This is impossible to guarantee when all goods are allocated – e.g. when there are two agents and only one item. Due to this impossibility, several relaxations have been studied in the literature. The most popular relaxation of envy freeness is *envy freeness up to one good* (EF1) [11, 21]. An allocation X is EF1 if no agent envies another agent after dropping some good for the latter agent’s bundle, i.e. if for every $i, j \in N$, if $v_i(X_i) < v_i(X_j)$ there exists some $g \in X_j$ such that $v_i(X_i) \geq v_i(X_j - g)$. An EF1 allocation can be computed in polynomial time for most realistic valuation classes [21]. More recently, a stronger relaxation called *envy free up to any good* (EFX) [13] has gained popularity: an allocation is EFX if no agent envies another agent after dropping *any* good from the latter agent’s bundle, i.e. if $v_i(X_i) < v_i(X_j)$ then for every $g \in X_j$ $v_i(X_i) \geq v_i(X_j - g)$. In contrast to EF1 allocations, the existence of EFX allocations is still an open question for several valuation classes [25].

Maximin Share: An agent’s *maximin share* (MMS) is defined as the value they would obtain had they divided the goods into n bundles themselves and picked the worst of these bundles. More formally,

$$\text{MMS}_i = \max_{X=(X_1, X_2, \dots, X_n)} \min_{j \in [n]} v_i(X_j)$$

Procaccia and Wang [26] show that agents cannot always be guaranteed their maximin share; however, past works [20] guarantee that every agent receives a fraction of their maximin share. For some $c \in (0, 1]$, an allocation X is c -MMS if for every agent $i \in N$, $v_i(X_i) \geq c \cdot \text{MMS}_i$.

Leximin: an allocation is leximin if it maximizes the value provided to the agent with least value and conditioned on this, maximizes the value provided to the agent with the second least value and so on. While leximin allocations are computationally intractable when agent valuations are unrestricted [9, Theorem 4.2], they can be computed in polynomial time under MRF valuations [2].

When envy is the main consideration, an allocation where no agent gets any good is envy free. While this is fair, it is very inefficient. Therefore, coupled with fairness metrics, algorithms usually guarantee some *efficiency* criterion as well. We consider two popular notions of efficiency:

Utilitarian Social Welfare: The *utilitarian social welfare* of an allocation is defined as the *sum* of the value obtained by each agent i.e. $\text{USW}(X) = \sum_{i \in N} v_i(X_i)$.

Nash Social Welfare: The *Nash social welfare* of an allocation is defined as the *product* of the value obtained by each agent i.e. $\text{NSW}(X) = \prod_{i \in N} v_i(X_i)$.

Allocations which maximize utilitarian social welfare and Nash social welfare are referred to as MAX-USW and MNW respectively. Since allocations have $\text{NSW}(X) = 0$ when some agent receives no items, we adopt the same convention as Caragiannis et al. [13], and

first minimize the number of agents with zero utility; subject to that, we maximize the product of positive utilities.

Before we proceed, we present a simple useful result about matroid rank valuations – if an agent values the bundle Y more than the bundle X , there must be a good $g \in Y$ such that $\Delta_i(X, g) = 1$. Variants of this result have also been shown by Babaioff et al. [2] and Benabbou et al. [9] and therefore, we omit the proof.

Observation 2.1 (Benabbou et al. [9]). Suppose that agents have binary submodular valuations. If X and Y are two allocations and $v_i(X_i) < v_i(Y_i)$ for some $i \in N + 0$, there exists a good $g \in Y_i \setminus X_i$ such that $\Delta_i(X_i, g) = 1$.

2.1 Prioritized Lorenz Dominating Allocations

We define the *sorted utility vector* of an allocation X as $\vec{u}^X = (u_1^X, u_2^X, \dots, u_n^X)$ which corresponds to the vector of agent valuations $(v_1(X_1), v_2(X_2), \dots, v_n(X_n))$ sorted in ascending order (ties broken arbitrarily). An allocation X *Lorenz dominates* the allocation Y (denoted $X \succeq_{\text{lorenz}} Y$) if for all $k \in [n]$, $\sum_{j=1}^k u_j^X \geq \sum_{j=1}^k u_j^Y$. An allocation X is *Lorenz dominating* if for every allocation Y , we have $X \succeq_{\text{lorenz}} Y$.

Babaioff et al. [2] show that when agents have MRF valuations, a non-redundant Lorenz dominating allocation always exists and satisfies several desirable fairness and efficiency guarantees. We formalize this result below.

THEOREM 2.2 (BABAIOFF ET AL. [2]). *When agents have MRF valuations, a non-redundant Lorenz dominating allocation always exists and is MNW, MAX-USW, EFX, leximin and $\frac{1}{2}$ -MMS.*

The one minor drawback of the ordering defined above is that it does not distinguish between two allocations with the same sorted utility vector, even when agents receive a different utility in the two allocations.

Example 2.3. Consider a problem instance with two agents $\{1, 2\}$ and three goods $\{g_1, g_2, g_3\}$. The valuation function for each agent $v_i(S) = |S|$ for $i \in \{1, 2\}$. Any allocation that gives two goods to one agent and one good to the other is Lorenz dominating, but under one agent 1 may receive two goods, and under another they may receive one.

It is desirable to distinguish between the two allocations described above; if we can create a solution concept where agent 1 always gets two goods whereas under another agent 2 gets two goods, randomizing between the two allocations will give us a (random) allocation which is arguably more fair since both agents have the same expected utility.

To this end, Babaioff et al. [2] introduce a priority order over the set of agents. The priority ordering is modelled as a permutation over the set of agents $\pi : N \mapsto [n]$ where agents with a lower value of π have a higher priority. This ordering is enforced by perturbing the valuations in the original instance. More formally, to enforce a priority ordering, Babaioff et al. [2] create a new fair allocation instance, where agent i has the valuation function v'_i :

$$v'_i(S) = v_i(S) + \frac{\pi(i)}{n^2}$$

We refer to this problem instance (resp. valuation function v') as the *augmented* problem instance (resp. augmented valuation function) with the priority order π . When π is clear from context, we simply refer to this problem instance as the augmented problem instance. We sometimes refer to the value $\frac{\pi(i)}{n^2}$ as the perturbation. Babaioff et al. [2] show that any Lorenz dominating allocation for the augmented instance is also a Lorenz dominating allocation for the original fair allocation instance [2, Theorem 4] and therefore retain all the desirable fairness properties in Theorem 2.2. They refer to this allocation as a Lorenz dominating allocation w.r.t. the priority order π .

This priority order can be used to guarantee additional fairness properties — by randomly choosing this priority order, we can generate allocations that are ex-ante envy-free and ex-ante proportional. A random allocation X is *ex-ante envy-free* if, in expectation, no agent envies another agent i.e. $\mathbb{E}_X[v_i(X_i)] \geq \mathbb{E}_X[v_i(X_j)]$ for all $i, j \in N$. Similarly, a random allocation X is *ex-ante proportional* if each agent, in expectation, receives a utility greater than the n -th fraction of their value for the entire bundle of goods i.e. $\mathbb{E}_X[v_i(X_i)] \geq \frac{v_i(G)}{n}$.

Babaioff et al. [2] define the randomized prioritized egalitarian (RPE) mechanism which performs the following steps:

- (1) Choose a priority order π uniformly at random.
- (2) Elicit the preferences of each agent. If the agent's valuation function is not an MRF, set their value for all bundles to be equal to 0.
- (3) Compute a Lorenz dominating allocation with respect to the ordering π for the above elicited preferences.

Babaioff et al. [2] show that the RPE mechanism computes an ex-ante envy-free and ex-ante proportional allocation. In addition, the RPE mechanism is *strategyproof*: no agent can get a better outcome by misreporting their valuation function. This strategyproofness result is independent of the algorithm used to compute Lorenz dominating allocations; therefore, it applies to Yankee Swap.

3 YANKEE SWAP

We now present a simple, fast algorithm for computing prioritized Lorenz dominating allocations. The algorithm we propose is known colloquially as a *Yankee swap*. Our objective is to ensure that at every round, the least utility agent receives a *useful* good, i.e. a good for which they have a positive marginal utility. We proceed in rounds, and at every round agents have a choice of either taking an unallocated useful good or stealing a useful good from another agent, who then either takes an unassigned useful good or steals one from another agent, and so on. Thus, the first agent increases their utility by 1, and the other agents' utilities remain the same: if a good was stolen from them, they must have recovered their utility by either stealing a good from another agent or by taking an unassigned good.

More formally, we define transfer paths recursively as follows: a *transfer path* in an allocation X is a sequence of agents in $N \cup \{0\}$, (p_1, p_2, \dots, p_r) , such that for some good $g \in X_{p_2}$: (a) $\Delta_{p_1}(X_{p_1}, g) = 1$. (b) If X' is the allocation that results from moving g from p_2 to p_1 in X ; then there exists a transfer path (p_2, \dots, p_r) in X' that does not involve the transfer of the good g . In other words, there is a set of goods that can be transferred along the path such that agent p_1 's

utility increases by 1, agents p_2, \dots, p_{r-1} 's utilities are unchanged, and agent p_r 's utility decreases by 1. While paths can be cyclic i.e. an agent can be present multiple times in a transfer path sequence, every good may be transferred at most once.

Since goods are transferred at most once, paths can be characterized by a sequence of goods $(g_{i_1}, g_{i_2}, \dots, g_{i_k})$ and an agent i where g_{i_k} gets transferred to the agent that has $g_{i_{k-1}}$, $g_{i_{k-1}}$ gets transferred to the agent that has $g_{i_{k-2}}$ and so on until finally, g_{i_1} gets transferred to agent i . Depending on the context, we use both notations of transfer paths in our algorithms and analysis.

3.1 The Algorithm

The Yankee Swap algorithm takes as input a fair allocation instance $(N, G, \{v_i\}_{i \in N})$ and a priority ordering $\pi : N \mapsto [n]$ over the set of agents. The algorithm first places all goods in X_0 — all goods are initially unassigned — and has all players playing (denoted by having them in the set P). We pick an agent $i \in P$ (who's not agent 0) with the least utility so far and check if there is a transfer path starting from them and ending at 0; ties are broken in favor of agents with higher priority. If a path exists, we transfer goods backwards along the path giving the agent an additional unit of value. Otherwise, we remove them from P , at which point their utility can no longer increase. Once all agents are removed from P , Algorithm 1 terminates.

Algorithm 1 Yankee Swap

Require: The set of agents $N = [n]$, the set of goods G , oracle access to valuation functions $\{v_i\}_{i \in N}$ and a priority order over the agents $\pi : N \mapsto [n]$

Ensure: A prioritized Lorenz dominating allocation X

- 1: $X = (X_0, X_1, \dots, X_n) \leftarrow (G, \emptyset, \dots, \emptyset)$
 - 2: $P \leftarrow N$
 - 3: **while** $P \neq \emptyset$ **do**
 - 4: Let $P' = \arg \min\{|X_i| : i \in [n]\}$
 - 5: Let i be the highest priority agent in P' according to π
 - 6: Check if there exists a transfer path in X starting at i which ends at 0
 - 7: **if** a path $(g_{i_1}, g_{i_2}, \dots, g_{i_k})$ exists **then**
 - 8: Transfer goods along the path and update X
 - 9: **else**
 - 10: $P \leftarrow P - i$
 - 11: **end if**
 - 12: **end while**
 - 13: **return** X
-

3.2 Analysis

Algorithm 1 computes prioritized Lorenz dominating allocations. Before we go into the technical details of the proof, we remark that our algorithm is very similar to the round robin algorithm in two ways. First, agents in P have roughly the same bundle size at any iteration of the algorithm. More formally, all agents in P have bundles that differ in size by at most 1, with higher priority agents receiving the slightly larger bundles. Second, much like how the round robin algorithm is EF1 at every iteration, the Yankee Swap

algorithm is Lorenz dominating at every iteration (ignoring the utility of agent 0 who controls the unassigned items).

Before we dive into any result about the algorithm, it is important to show sufficient conditions for a transfer path to exist. A similar version of the following lemma appears in Babaioff et al. [2, Lemma 17] and Benabbou et al. [9, Lemma 3.12].

Lemma 3.1. *Let X and Y be two non-redundant allocations for the set of agents $N + 0$. Let S^- be the set of all agents $i \in N + 0$ where $|X_i| < |Y_i|$, $S^=$ be the set of all agents $i \in N + 0$ where $|X_i| = |Y_i|$ and S^+ be the set of all agents $i \in N + 0$ where $|X_i| > |Y_i|$. For any agent $i \in S^-$, there exists a transfer path from i to some agent $k \in S^+$ in X .*

PROOF. Our proof is constructive. We construct the transfer path using the following recursive loop which we denote by $\text{Loop}(X, Y, i)$:

Since X and Y are non-redundant, using Observation 2.1, there is some good $g \in Y_j \setminus X_j$ such that $\Delta_j(X_j, g) = 1$. This good must belong to some other agent, say $j \in N + 0$. If $j \in S^+$, we are done. Otherwise, move the good from j to i to create a new allocation X' . We now compare X' and Y and define S'^- , $S'^=$ and S'^+ analogously to S^- , $S^=$ and S^+ . Since $j \in S^- \cup S^=$ and j lost a good in X' , we must have that $j \in S'^-$. Further, $i \in S'^= \cup S'^-$ which means $S'^+ = S^+$. We then repeat this process with the allocations X' and Y and the agent j i.e. we run $\text{Loop}(X', Y, j)$.

Note that the above loop must terminate since $\sum_{j \in N+0} |X'_j \cap Y_j|$ increases by 1 at every iteration, and is upper bounded by $|G|$. Since S^+ does not change, at some iteration, we take a good from an agent in S^+ and the above loop terminates.

Let $(p_1, p_2, \dots, p_r, k)$ be the sequence of agents we take a good from in the above loop — we take a good from p_1 and give it to i , take a good from p_2 and give it to p_1 , and so on. In order to show that (i, p_1, \dots, p_r, k) forms a transfer path, the last thing we need to show is that g never gets transferred out of i which would imply that goods get transferred at most once. This is easy to see: the only goods that can potentially get transferred in X' are the goods in $Y_j \setminus X'_j$ for any $j \in N + 0$. Since $g \in Y_i \cap X'_i$, it will never get transferred out. Similarly, no good that is transferred to some agent ever gets transferred out. \square

When X is the allocation computed by Yankee Swap and Y is the prioritized Lorenz dominating allocation, the above lemma shows that we can, in a way, move closer to Y from X using path transfers. While repeated applications of this lemma to any non-redundant allocation X will ultimately terminate at the prioritized Lorenz dominating allocation, this is not a very efficient process.

Our next three lemmata establish some important properties of the algorithm. The first Lemma shows that the allocation maintained by Yankee Swap is always non-redundant; thereby, allowing us to use Lemma 3.1.

Lemma 3.2. *At the beginning of every iteration, the allocation X of the Yankee Swap algorithm is non-redundant for the set of agents $N + 0$.*

PROOF. By the definition of transfer paths, agents only ever take unassigned goods/steal goods that they have a positive marginal gain for, i.e. useful goods. Thus, after executing a transfer path, the allocation remains non-redundant. Since at the beginning of the

first iteration the allocation is non-redundant (by our choice of v_0), it follows that non-redundancy is maintained throughout. \square

The second Lemma shows that picking the least utility agent with highest priority is equivalent to picking the agent with least value according to the augmented valuation v' . This will help us when analyzing the properties of the allocation output by Yankee Swap with respect to the augmented valuations v' .

Lemma 3.3. *At the beginning of any iteration of Algorithm 1, let i be the highest priority agent with least utility such that $i \in P$ i.e. let i be the agent chosen by the algorithm to be the starting point of the transfer path. Let W be the allocation at the beginning of the iteration. Then i is the least valued agent in W with respect to the augmented valuation function among all the agents in P .*

PROOF. For any agent $u \in P$, if $|W_u| > |W_i|$, we have $v'_u(W_u) > v'_i(W_i)$ since the augmented valuations add a value smaller than 1 for both agents compared to the original valuation function and W is non-redundant (Lemma 3.2). If $|W_u| = |W_i|$, then we must have $v'_u(W_u) > v'_i(W_i)$ since we chose i as the agent with highest priority with a bundle of size $|W_i|$. Since our first constraint on i was that it needed to minimize $|W_i|$, we can never have $|W_u| < |W_i|$. \square

Finally, we formalize the round-robin balancedness notion described at the beginning of this subsection.

Lemma 3.4. *Let i be the agent chosen to initiate a transfer path in some iteration of Algorithm 1. Let W be the allocation at the beginning of the iteration. If $|W_i| = k$, then the following holds:*

- (1) *Agents with higher priority than i have a bundle of size at most $k + 1$.*
- (2) *Agents with a lower priority than i have a bundle of size at most k .*

PROOF. This result stems from the sequential nature of the allocation. Let t be the iteration of the algorithm being examined.

Let j be an agent with higher priority than i with a bundle of size at least $k + 2$. Consider the start of the iteration t' where W_j moved from a bundle of size $k + 1$ to a bundle of size $k + 2$. Since bundle sizes increase by at most 1 at every iteration, there is a unique iteration where this event occurred. It must also be that $t' < t$. To be selected on the iteration t' , j must have been the agent with least utility ($k + 1$) and highest priority among the agents in P . However, we also have that at iteration t' , i 's bundle had a size of at most k and it was in P as well. This is because at iteration $t > t'$, i 's bundle had a size of k and i was in P ; bundle sizes increase monotonically and agents removed from P never get added back. This is a contradiction since it implies j was not the agent with least utility among the agents in P at iteration t' ; therefore, j cannot have a bundle of size $\geq k + 2$.

The case where j has a lower priority than i is handled similarly. \square

We are now ready to prove our main result. This is done via a simple contradiction — if Yankee Swap fails and there exists some agent i with a higher value in the Lorenz dominating allocation, then there must have been a path from i to 0 that Yankee Swap somehow missed. However, showing that this path exists requires showing that Yankee Swap gets sequentially closer to the final solution. That is, at every round, the utility of every agent i is no

more than their utility under the prioritized Lorenz dominating allocation. This proof requires a careful combinatorial argument, and is shown in Lemma 3.6.

THEOREM 3.5. *When agents have MRF valuations, Yankee Swap computes a non-redundant Lorenz dominating allocation with respect to the priority order π .*

PROOF. It is easy to see that the algorithm always terminates: at every iteration, an agent is removed from P or $|X_0|$ reduces by 1. Since no item is ever deallocated, and no agents ever return to P once removed, the number of iterations is at most the size of P plus X_0 on initialization, i.e. $n + m$.

We now show that the allocation output by Algorithm 1 is Lorenz dominating with respect to the priority order π .

Assume for contradiction that the allocation output by Algorithm 1 is not Lorenz dominating with respect to the ordering π . Let X be the allocation output by Algorithm 1 and Y be a Lorenz dominating allocation with respect to the ordering π (recall that Y is shown to always exist by Babaioff et al. [2]). Note that Y is also a Lorenz dominating allocation with respect to the augmented problem instance with priority order π .

If for all $i \in N$ $v_i(X_i) \geq v_i(Y_i)$ then since Y is MAX-USW, this means that $v_i(X_i) = v_i(Y_i)$ for all i and X is Lorenz dominating as well with respect to the ordering π , contrary to our assumption. Thus, there must be at least one $i \in N$ such that $v_i(X_i) < v_i(Y_i)$. Let $i \in \arg \min\{v_i(X_i) \mid v_i(X_i) < v_i(Y_i)\}$. If there are multiple such agents, we pick the one with highest priority. Note that, by an argument similar to Lemma 3.3, this is equivalent to saying i is the least valued agent in the augmented problem instance such that $v'_i(X_i) < v'_i(Y_i)$.

Consider the iteration of the algorithm where i was removed from P . Let W be the non-redundant allocation at the beginning of the iteration (Lemma 3.2). We have the following Lemma.

Lemma 3.6. *For all $h \in N$, we have $|Y_h| \geq |W_h|$.*

PROOF. Assume for contradiction that this is not true. We show that Y does not Lorenz dominate W in the augmented problem instance, which contradicts the fact that Y is Lorenz dominating. Let j be the agent with lowest utility under Y such that $|Y_j| < |W_j|$; we break ties by choosing the agent with highest priority. Again, since W is non-redundant, this is equivalent to saying that j is the agent with least utility in Y with respect to the augmented problem instance such that $v'_j(Y_j) < v'_j(W_j)$.

Before we delve into the technical details, let us discuss the main idea of the proof. We examine the sorted utility vectors of W and Y . We show that for any agent h such that $v'_h(Y_h) < v'_j(Y_j)$, then $v'_h(Y_h) = v'_h(W_h)$. The ‘converse’ also holds: for any agent h , if $v'_h(W_h) < v'_j(Y_j)$ then $v'_h(W_h) = v'_h(Y_h)$. We can compare the first index where the sorted utility vectors of W and Y differ; we then show that this element is greater in the sorted utility vector of W than in Y , i.e. showing that Y does not Lorenz dominate W , and yielding a contradiction. This is summarized in Figure 1. Lastly, in this proof, we will be analyzing the augmented problem instance; unless specifically stated, any claims about preferences can be assumed to be with respect to the augmented valuations.




Figure 1: The idea behind the proof of Lemma 3.6: we plot the sorted utility vectors of W (blue) and Y (red). Each large block can be thought of as a utility of 1 derived from a good and each of the smaller light colored blocks can be thought of as the augmented utilities. We first show that all the values to the left of j are equal in both W and Y and correspond to the same agent. We next show that the bar corresponding to u is taller than that of j .

We first show that $v'_j(Y_j) < v'_i(W_i)$. If j has a higher priority than i , $|W_j| \leq |W_i| + 1$ (Lemma 3.4). By our choice of j , $|Y_j| \leq |W_j| - 1 \leq |W_i|$, which implies that $v'_j(Y_j) < v'_i(W_i)$. If j has a lower priority than i , $|W_j| \leq |W_i|$ (Lemma 3.4). By our choice of j , $|Y_j| \leq |W_j| - 1 \leq |W_i| - 1 < |W_i|$. Since $|W_i|$ is greater than $|Y_j|$ by at least 1, we have $v'_j(Y_j) < v'_i(W_i)$ irrespective of their priorities.

We observe that for any $h \in N$ whose utility is less than $v'_j(Y_j)$ under Y , we have $v'_h(Y_h) = v'_h(W_h)$. This is because we pick an agent j with minimum utility under Y for which $v'_j(Y_j) < v'_j(W_j)$. Therefore, for any agent h for whom $v'_h(Y_h) < v'_j(Y_j)$ we must have that $v'_h(Y_h) \geq v'_h(W_h)$. If the inequality is strict, we have $v'_h(W_h) < v'_h(Y_h) < v'_j(Y_j) < v'_i(W_i)$. This implies that h was removed from P before i was removed from P i.e. at the iteration where i was removed from P , h was already removed from P ; otherwise, according to Lemma 3.3, Algorithm 1 would have chosen h instead of i creating a contradiction. If h was not in P at the iteration in consideration, h has a utility of $|W_h|$ when Algorithm 1 terminates, i.e. $|W_h| = |X_h|$. This implies that $v'_h(X_h) = v'_h(W_h) < v'_i(W_i) \leq v'_i(X_i)$ and $v'_h(X_h) < v'_h(Y_h)$. However, we assumed that agent i is the lowest utility agent among those who have lower utility under X than under Y , a contradiction. We can similarly show that for any agent h with utility less than $v'_j(Y_j)$ in W , we have $v'_h(Y_h) = v'_h(W_h)$.

Note that in the valuations v' , no two agents have the same value in any allocation due to the perturbation. Let j be the ℓ -th least valued agent in Y . From our discussion above, the first $\ell - 1$ least valued agents in both W and Y have the same utility in both allocations. Let the ℓ -th least valued agent in W be u . From the definition of Lorenz dominance, if $v'_j(Y_j) < v'_u(W_u)$, Y does not Lorenz dominate W .

If $v'_j(Y_j) = v'_u(W_u)$, this implies that $j = u$ since the fractional part of $v'_h(Y_h)$ is unique for every agent $h \in N$. However, if $j = u$, then by assumption we have $v'_j(Y_j) < v'_j(W_j) = v'_u(W_u) = v'_j(Y_j)$, a contradiction. Therefore $v'_j(Y_j) \neq v'_u(W_u)$. If $v'_u(W_u) < v'_j(Y_j)$, then from our discussion, we must have $v'_u(W_u) = v'_u(Y_u) < v'_j(Y_j)$. This implies that there are ℓ agents (the first ℓ least valued agents in W) which have a lower utility than j in Y , contradicting our

assumption on ℓ . Therefore, we must have $v'_u(W_u) > v'_j(Y_j)$ proving that Y does not Lorenz dominate W ; the sum of the first ℓ elements in the sorted utility vector of W is greater than the sum of the first ℓ elements in the sorted utility vector of Y . Since Y is Lorenz dominating in the augmented problem instance, this results in a contradiction and completes the proof. \square

We construct a new allocation Z starting at Y and moving goods from agents in $N - i$ arbitrarily to 0 till $|Z_j| = |W_j|$ for all $j \in N - i$. In other words, let Z_j be a size $|W_j|$ subset of Y_j for all $j \in N - i$ and $Z_i = Y_i$. Due to Lemma 3.6, there will be no $j \in N - i$ where $|Z_j| < |W_j|$. Given the dummy agent 0's valuation function, this allocation is still non-redundant for the agents $N + 0$. Note that $|W_i| < |Z_i|$ and $|W_0| > |Z_0|$.

Invoking Lemma 3.1 on the allocations W and Z , there must be a transfer path from i to 0 in W . This stems from the fact that 0 is the only agent in S^+ . Since i has a valid transfer path, it could not have been removed from P at the current iteration, a contradiction. \square

In addition to proving that Yankee Swap outputs a Lorenz dominating allocation, Theorem 3.5 shows that the resulting allocation is non-redundant. This is highly desirable in the course allocation setting; indeed, while some works require that algorithms leave no item unallocated, it is *preferable* to have incomplete allocations in course allocation: if the resulting allocation is not non-redundant, then we have assigned students to classes that they either don't want to take, can't fit in their schedule, or they already have been signed up for. In any case, students will have to drop some of their assigned classes, creating additional administrative overheads for both management and the students themselves.

3.3 Computing Path Transfers

The most computationally intensive aspect of Algorithm 1 is finding path transfers. We now provide a simple algorithm to compute path transfers. A general version of this method can be used to decide if a set is independent in a matroid union [28]. A similar approach is also used by Barman and Verma [6] whose notation we follow.

We define the *exchange graph* of an allocation X as a directed graph $\mathcal{G}(X) = (G, E)$ on the set of goods G . If a good g is in X_j for some agent $j \in N + 0$, then an edge exists from g to some other good $g' \notin X_j$ if $v_j(X_j - g + g') = v_j(X_j)$. In other words, there is an edge from g to g' if the agent who owns g can replace it with g' with no loss to their utility. In particular, if $g \in X_0$ then there is a directed edge from g to any other good not in X_0 . This is because agent 0, who represents the unassigned items, has an additive utility over all items. We observe that if the good g is stolen from agent j , then that good can be replaced with g' iff the edge (g, g') exists.

In order to check if a transfer path exists from some agent i to 0, we construct the exchange graph $\mathcal{G}(X)$. We then compute the set of goods which have a marginal gain of 1 for the agent i under the allocation X i.e. $F_i(X) = \{g \in G \mid \Delta_i(X_i, g) = 1\}$. In the exchange graph, we find the shortest path (if there exists one) from $F_i(X)$ to goods in X_0 ; this can be done by adding a source node s in the exchange graph with edges to all the goods in $F_i(X)$ and then using breadth-first search (BFS) to find the shortest path from s to X_0 . From the path in the exchange graph, we can also determine exactly which goods to transfer along the path to update the allocation –

if we have a path $(s, g_1, g_2, \dots, g_k)$, we transfer the good g_k to the agent who has g_{k-1} , transfer g_{k-1} to the agent who has g_{k-2} and so on. Finally, we give g_1 to i ; see Algorithm 2.

Algorithm 2 Computing Transfer Paths

Require: An allocation X and two agents i and j
 Construct the exchange graph $\mathcal{G}(X)$
 Add a source node s to $\mathcal{G}(X)$ with an edge to all the goods in $F_i(X)$
 Find the shortest path from s to X_j using breadth-first search in $\mathcal{G}(X)$
if A path exists **then**
 Return the path $(s, g_{i_1}, g_{i_2}, \dots, g_{i_k})$
else
 Return false
end if

That the shortest path in the exchange graph from $F_i(X)$ to X_0 is a transfer path is well known in the matroid literature (albeit using different terminology). This result was first adapted by Barman and Verma [6, Lemma 1] to fair allocation. The missing proofs in this section can be found in the full version of the paper.

Lemma 3.7. *Given a non-redundant allocation X , the shortest path from $F_i(X)$ to X_j in the exchange graph $\mathcal{G}(X)$ is a transfer path from i to j ($i \neq j$ and $i, j \in N + 0$) in X . i.e. transferring goods along the shortest path results in an allocation where i 's value for its bundle goes up by 1, j 's value for its bundle goes down by 1 and all the other agents see no change in the value of their bundles.*

Barman and Verma [6, Lemma 5] also show that a result similar to Lemma 3.1 is true for paths in the exchange graph as well.

THEOREM 3.8. *Let X and Y be two non-redundant allocations for the set of agents $N + 0$. Let S^- be the set of all agents $i \in N + 0$ where $|X_i| < |Y_i|$, $S^=$ be the set of all agents $i \in N + 0$ where $|X_i| = |Y_i|$ and S^+ be the set of all agents $i \in N + 0$ where $|X_i| > |Y_i|$. For any agent $i \in S^-$, there exists a path in the exchange graph from $F_i(X)$ to X_k for some $k \in S^+$.*

Armed with these two results, we are ready to show correctness i.e. a path exists from agent $F_i(X)$ to X_j in the exchange graph iff a transfer path exists from i to j in the allocation X .

THEOREM 3.9. *Given a non-redundant allocation X , a transfer path exists from agent i to agent j in X if and only if Algorithm 2 outputs a path. Furthermore, the path output by Algorithm 2 is a transfer path.*

PROOF. The second statement is implied by Lemma 3.7 so we only prove the first statement.

(\Rightarrow) Assume Algorithm 2 outputs a path. This implies there exists a path from $F_i(X)$ to X_j . From Lemma 3.7, this implies that there is a transfer path from i to j

(\Leftarrow) Assume there is a transfer path from i to j in X . Let Y be the non-redundant allocation that arises from transferring goods along the transfer path. Apply Theorem 3.8 to allocations X and Y and the agent i . Agent i is clearly in S^- and the only agent in S^+ is j from the definition of the transfer path. Therefore, there must exist

a path from $F_i(X)$ to X_j in $\mathcal{G}(X)$. This implies that Algorithm 2 outputs a path. \square

3.4 Time Complexity Analysis

We now analyze the time complexity of Algorithm 1. We represent the allocation X as a binary matrix where $X(i, g) = 1$ if and only if $g \in X_i$. Thus, checking if a good belongs to an agent, adding a good to a bundle, and removing a good from a bundle can be done in $O(1)$ time.

Our time complexity analysis is based on two simple observations. First, the loop in Algorithm 1 runs at most $n + m$ times. This is because at each round, we either reduce the size of X_0 by 1 or remove some agent from the game — the former happens at most m times and the latter happens at most n times; once an agent is removed, they do not return.

Second, Algorithm 2 runs in $O(m^2(n + \tau))$ time (where again, τ is the maximum time to compute $v_i(S)$). This can be seen by closely examining each step of Algorithm 2. We can construct the exchange graph by examining each possible pair of goods (g, g') , finding the agent i whose bundle contains g , and checking if $v_i(X_i - g + g') = v_i(X_i)$. This can trivially be done in $O(m^2(n + \tau))$ time. This is the most expensive step of the algorithm.

Adding a source node and the required edges takes $O(m\tau)$ time; we only need to check if $\Delta_i(X_i, g) = 1$ for each good g . Running BFS and finding the shortest path takes $O(m^2)$ time since there are $O(m)$ nodes in the graph. Executing path transfers takes $O(nm)$ time — we iterate through the path and transfer each good to the owner of the previous good; finding the owner takes $O(n)$ time.

Combining the two observations, we have the following result.

THEOREM 3.10. *Algorithm 1 runs in $O(m^2(n + \tau)(m + n))$ time.*

A practical application of Yankee Swap in course allocation can likely leverage structural properties of the problem. For example, while real-world instances contain thousands of students and courses, each individual student is typically assigned no more than 5–6 classes. This sparsity most likely allows for more compact representations and further reductions in running time. We leave this analysis to future work

3.5 Comparison to Babaioff et al. [2]

The algorithm to compute Lorenz dominating allocations by Babaioff et al. [2] works as follows: starting with a MAX-USW allocation, repeatedly check if there exists a path transfer that improves the objective function $\sum_{i \in N} \left(v_i(X_i) + \frac{\pi(i)}{n^2} \right)^2$. This is similar to the technique used by Benabbou et al. [9], who maximize $\sum_{i \in N} v_i(X_i)^2$ via single-item transfers, rather than path transfers. Babaioff et al. check for and compute path transfers via matroid intersection algorithms. Specifically, they use algorithms for the matroid intersection problem to compute a MAX-USW allocation for the modified problem instance where each agent i 's valuation function is upper bounded by some predefined value k_i i.e. $v_i^{\text{new}}(S) = \max\{v_i(S), k_i\}$. Note that computing the value of a bundle with respect to this modified valuation function still takes $O(\tau)$ time; therefore, from a time complexity perspective, computing a MAX-USW allocation for the valuation profile v is equivalent to computing a MAX-USW allocation for the valuation profile v^{new} .

We first analyze the time complexity of computing a MAX-USW allocation and then use it to analyze the time complexity of the algorithm by Babaioff et al. [2]. Our analysis uses the matroid intersection algorithm by Chakrabarty et al. [14], the state-of-the-art algorithm for the matroid intersection problem with access to a rank oracle. The proof can be found in the full version of the paper.

Lemma 3.11. *When agents have MRF valuations, computing a MAX-USW allocation takes $O(n^2 m^{3/2} (m + \tau) \log nm)$ time using the matroid intersection problem.*

Note immediately that when $m = \Theta(n)$, our algorithm computes a MAX-USW allocation faster than the matroid intersection based approach. We now present the runtime of Babaioff et al.'s algorithm.

THEOREM 3.12. *The algorithm by Babaioff et al. [2] computes Lorenz dominating allocations in $O(n^6 m^{7/2} (m + \tau) \log nm)$ time.*

PROOF. Babaioff et al. [2] show that their algorithm computes a MAX-USW solution at most $O(n^4 m^2)$ times. Combining this with Lemma 3.11, we get the required time complexity. \square

Indeed, our algorithm is significantly faster than that of Babaioff et al. [2]. This speedup mainly stems from two sources. First, even though Babaioff et al. [2] use transfer paths, our method of computing them is faster. Second, by carefully choosing which transfer paths to check for, we check for much fewer paths. Combining these two factors, Yankee Swap offers a significantly better worst-case runtime. In particular, when $m = \Theta(n)$, the worst-case runtime of our algorithm is faster by a factor of $O(n^{13/2} \log n)$.

4 CONCLUSIONS AND FUTURE WORK

In this work, we show that when agents have binary submodular valuations, Yankee Swap offers a simple and fast method to output fair and efficient allocations. The entire algorithmic framework can be implemented using no more than a few lines of code (Algorithms 1 and 2), and offers a far better worst-case runtime guarantee than the current state of the art. The simplicity of Yankee Swap is its key strength: it is easy to understand (even by non-experts) and implement, and can easily be adapted to different settings. This is all achieved while offering the same strong fairness and efficiency guarantees as the current state of the art.

This work highlights the surprising power of combinatorial arguments in computing transfer paths. Unlike prior work in this space, we do not invoke complex matroid optimization algorithms, from which the proofs are directly derived, but rather utilize a *simple* approach, 'relegating' the complexity to our careful combinatorial analysis. We believe that Yankee Swap can be applied to compute justice criteria beyond leximin. More specifically, we conjecture that when agents have entitlements (or priority weights) [15], a modified version of Yankee Swap can be used to compute a weighted leximin allocation. We also believe that Yankee Swap can be applied in fair chore allocation problems, and with some adaptations, in settings where agents do not have binary valuations.

ACKNOWLEDGMENTS

The authors would like to thank anonymous reviewers at WINE 2022 and AAMAS 2023 for useful feedback.

REFERENCES

- [1] Haris Aziz. 2019. Strategyproof Multi-Item Exchange under Single-Minded Dichotomous Preferences. *Autonomous Agents and Multi-Agent Systems* 34, 1 (2019), 11.
- [2] Moshe Babaioff, Tomer Ezra, and Uriel Feige. 2021. Fair and Truthful Mechanisms for Dichotomous Valuations. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*. 5119–5126.
- [3] Moshe Babaioff, Ron Lavi, and Elan Pavlov. 2009. Single-Value Combinatorial Auctions and Algorithmic Implementation in Undominated Strategies. *J. ACM* 56, 1, Article 4 (2009), 32 pages.
- [4] Siddharth Barman, Sanath Kumar Krishnamurthy, and Rohit Vaish. 2018. Greedy Algorithms for Maximizing Nash Social Welfare. In *Proceedings of the 17th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*. 7–13.
- [5] Siddharth Barman and Paritosh Verma. 2021. Approximating Nash Social Welfare Under Binary XOS and Binary Subadditive Valuations. In *Proceedings of the 17th Conference on Web and Internet Economics (WINE)*. 373–390.
- [6] Siddharth Barman and Paritosh Verma. 2021. Existence and Computation of Maximin Fair Allocations Under Matroid-Rank Valuations. In *Proceedings of the 20th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*. 169–177.
- [7] Siddharth Barman and Paritosh Verma. 2022. Truthful and Fair Mechanisms for Matroid-Rank Valuations. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI)*. 4801–4808.
- [8] 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.
- [9] 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* 9, 4, Article 21 (2021), 41 pages.
- [10] Anna Bogomolnaia, Hervé Moulin, and Richard Stong. 2005. Collective choice under dichotomous preferences. *Journal of Economic Theory* 122, 2 (2005), 165–184.
- [11] Eric Budish. 2011. The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes. *Journal of Political Economy* 119, 6 (2011), 1061–1103.
- [12] Eric Budish, Gérard P. Cachon, Judd B. Kessler, and Abraham Othman. 2016. Course Match: A Large-Scale Implementation of Approximate Competitive Equilibrium from Equal Incomes for Combinatorial Allocation. *Operations Research* 65, 2 (2016), 314–336.
- [13] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. 2016. The Unreasonable Fairness of Maximum Nash Welfare. In *Proceedings of the 17th ACM Conference on Economics and Computation (EC)*. 305–322.
- [14] Deeparnab Chakraborty, Yin Tat Lee, Aaron Sidford, Sahil Singla, and Sam Chiu-wai Wong. 2019. Faster Matroid Intersection. In *Proceedings of the 60th Symposium on Foundations of Computer Science (FOCS)*. 1146–1168.
- [15] Mithun Chakraborty, Ayumi Igarashi, Warut Suksompong, and Yair Zick. 2021. Weighted Envy-Freeness in Indivisible Item Allocation. *ACM Transactions on Economics and Computation* 9, Article 18 (2021), 39 pages.
- [16] Andreas Darmann and Joachim Schauer. 2015. Maximizing Nash product social welfare in allocating indivisible goods. *European Journal Operations Research* 247 (2015), 548–559.
- [17] Jonathan Goldman and Ariel D. Procaccia. 2015. Spliddit: Unleashing Fair Division Algorithms. *SIGecom Exchanges* 13 (2015), 41–46.
- [18] 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)*. 370–383.
- [19] Andreas Krause and Daniel Golovin. 2014. Submodular Function Maximization. In *Tractability: Practical Approaches to Hard Problems*, Lucas Bordeaux, Youssef Hamadi, and Pushmeet Kohli (Eds.). Cambridge University Press, 71–104.
- [20] David Kurokawa, Ariel D. Procaccia, and Junxing Wang. 2018. Fair Enough: Guaranteeing Approximate Maximin Shares. *J. ACM* 65, 2 (2018).
- [21] R. J. Lipton, E. Markakis, E. Mossel, and A. Saberi. 2004. On Approximately Fair Allocations of Indivisible Goods. In *Proceedings of the 5th ACM Conference on Economics and Computation (EC)*. 125–131.
- [22] Debasis Mishra and Souvik Roy. 2013. Implementation in multidimensional dichotomous domains. *Theoretical Economics* 8, 2 (2013).
- [23] Josué Ortega. 2020. Multi-unit assignment under dichotomous preferences. *Mathematical Social Sciences* 103 (2020), 15–24.
- [24] James Oxley. 2011. *Matroid Theory* (2nd ed.). Number 21. Oxford University Press.
- [25] Benjamin Plaut and Tim Roughgarden. 2017. Almost Envy-Freeness with General Valuations. *ArXiv abs/1707.04769* (2017).
- [26] Ariel D. Procaccia and Junxing Wang. 2014. Fair Enough: Guaranteeing Approximate Maximin Shares. In *Proceedings of the 15th ACM Conference on Economics and Computation (EC)*. 675–692.
- [27] Alvin E. Roth, Tayfun Sönmez, and M. Utku Ünver. 2005. Pairwise kidney exchange. *Journal of Economic Theory* 125, 2 (2005), 151–188.
- [28] A. Schrijver. 2003. *Combinatorial Optimization - Polyhedra and Efficiency*. Springer.
- [29] Warut Suksompong and Nicholas Teh. 2022. On maximum weighted Nash welfare for binary valuations. *Mathematical Social Sciences* 117 (2022), 101–108.