

Fair and Efficient Allocations under Lexicographic Preferences

Hadi Hosseini,¹ Sujoy Sikdar,² Rohit Vaish,³ Lirong Xia⁴

¹ Pennsylvania State University, University Park, Pennsylvania, USA 16802

² Binghamton University, Binghamton, New York 13902-6000

³ Tata Institute of Fundamental Research, Mumbai, India 400005

⁴ Rensselaer Polytechnic Institute, Troy, New York 12180

hadi@psu.edu, ssikdar@binghamton.edu, rohit.vaish@tifr.res.in, xial@cs.rpi.edu

Abstract

Envy-freeness up to any good (EFX) provides a strong and intuitive guarantee of fairness in the allocation of indivisible goods. But whether such allocations always exist or whether they can be efficiently computed remains an important open question. We study the existence and computation of EFX in conjunction with various other economic properties under *lexicographic preferences*—a well-studied preference model in artificial intelligence and economics. In sharp contrast to the known results for additive valuations, we not only prove the existence of EFX and Pareto optimal allocations, but in fact provide an algorithmic characterization of these two properties. We also characterize the mechanisms that are, in addition, strategyproof, non-bossy, and neutral. When the efficiency notion is strengthened to rank-maximality, we obtain non-existence and computational hardness results, and show that tractability can be restored when EFX is relaxed to another well-studied fairness notion called maximin share guarantee (MMS).

Introduction

Fair and efficient allocation of scarce resources is a fundamental problem in economics and computer science. The quintessential fairness notion—*envy-freeness*—enjoys strong existential and computational guarantees for *divisible* resources (Varian 1974). However, in notable applications such as course allocation (Budish 2011) and property division (Pruhs and Woeginger 2012) that involve *indivisible* resources, (exact) envy-freeness could be too restrictive. In these settings, it is natural to consider notions of approximate fairness such as *envy-freeness up to any good* (EFX) wherein pairwise envy can be eliminated by the removal of any good in the envied bundle (Caragiannis et al. 2019).

EFX is arguably the closest analog of envy-freeness in the indivisible setting, and, as a result, has been actively studied especially for the domain of additive valuations. However, it also suffers from a number of limitations: First, barring a few special cases, the existence and computation of EFX allocations remains an open problem. Second, for additive valuations, EFX can be incompatible with *Pareto optimality* (PO)—a fundamental notion of economic efficiency (Plaut and Roughgarden 2020). Finally, EFX could also be at odds

with *strategyproofness* (Amanatidis et al. 2017), which is another desirable property in the economic analysis of allocation problems.

The aforementioned limitations of EFX prompt us to explore the *domain restriction* approach in search of positive results (Elkind, Lackner, and Peters 2016). Specifically, we deviate from the framework of cardinal preferences for which EFX allocations have been most extensively studied, and instead focus on the purely ordinal domain of *lexicographic preferences*.

Lexicographic preferences have been widely studied in psychology (Gigerenzer and Goldstein 1996), machine learning (Schmitt and Martignon 2006), and social choice (Taylor 1970) as a model of human decision-making. Several real-world settings such as evaluating job candidates and the desirability of a product involve lexicographic preferences over the set of features. In the context of fair division, too, lexicographic preferences can arise naturally. For example, when dividing an inheritance consisting of a house, a car, and some home appliances, a stakeholder might prefer any division in which she gets the house over one where she doesn't (possibly because of its sentimental value), subject to which she might prefer any outcome that includes the car over one that doesn't, and so on.

On the computational side, lexicographic preferences provide a succinct language for representing preferences over combinatorial domains (Saban and Sethuraman 2014; Lang, Mengin, and Xia 2018), and have led to numerous positive results at the intersection of artificial intelligence and economics (Fujita et al. 2018; Hosseini and Larson 2019). Motivated by these considerations, our work examines the existence and computation of *fair* (i.e., EFX) and *efficient* allocations from the lens of lexicographic preferences.

Our Contributions. Figure 1 summarizes our theoretical contributions.

- **EFX+PO:** Our first result provides a family of polynomial-time algorithms for computing EFX+PO allocations under lexicographic preferences. Furthermore, we show that *any* EFX+PO allocation can be computed by some algorithm in this family, thus providing an algorithmic characterization of such allocations (Theorem 1). This result establishes a sharp contrast with the additive valuations domain where the two properties are incompatible in general.




Figure 1: Summary of our theoretical results.

- **EFX+PO+strategyproofness:** The positive result for EFX+PO motivates us to investigate a more demanding property combination of EFX, PO, and strategyproofness. Once again, we obtain an algorithmic characterization (Theorem 2): Subject to some common axioms (non-bossiness and neutrality), any mechanism satisfying EFX, PO, and strategyproofness is characterized by a special class of *quota-based serial dictatorship mechanisms* (Pápai 2000b; Hosseini and Larson 2019).

- **EFX+rank-maximality:** When the efficiency notion is strengthened to *rank-maximality*, we encounter incompatibility with strategyproofness (Example 1) as well as with EFX (Example 2). Furthermore, checking the existence of EFX and rank-maximal allocations turns out to be NP-complete (Theorem 3), suggesting that our algorithmic results are, in a certain sense, ‘maximal’. The intractability persists even when EFX is relaxed to envy-freeness up to k goods (EF k) (Theorem 4), but efficient computation is possible if EFX is relaxed to another well-studied fairness notion called *maximin share guarantee* or MMS (Theorem 5).

Related Work. Envy-free solutions may not always exist for indivisible goods. As a result, the literature has focused on notions of approximate fairness, most notably *envy-freeness up to one good* (EF1) and its strengthening called *envy-freeness up to any good* (EFX). The former enjoys strong existential and algorithmic support, as an EF1 allocation always exists for general monotone valuations and can be efficiently computed. However, achieving EF1 together with economic efficiency seems non-trivial: For additive valuations, EF1+PO allocations always exist (Caragiannis et al. 2019; Barman, Krishnamurthy, and Vaish 2018) but no polynomial-time algorithm is known for computing such allocations.

The stronger notion of EFX has proven to be more challenging. As mentioned previously, the existence of EFX for additive valuations remains an open problem. Additionally, EFX and Pareto optimality are known to be incompatible for non-negative additive valuations (Plaut and Roughgarden 2020).

The aforementioned limitations of EFX have motivated the study of further relaxations or special cases in search of positive results. Some recent results establish the existence of partial allocations that satisfy EFX after discarding a small number of goods while also fulfilling certain efficiency criteria (Caragiannis, Gravin, and Huang 2019; Chaudhury et al. 2020). Similarly, EFX allocations have been shown to exist for the special case of three agents with

additive valuations (Chaudhury, Garg, and Mehlhorn 2020), or when the agents can be partitioned into two *types* (Mahara 2020), or when agents have dichotomous preferences (Amanatidis et al. 2020). For cardinal utilities, various multiplicative approximations of EFX (and its variant that involves removing an *average* good) have been considered (Plaut and Roughgarden 2020; Amanatidis, Markakis, and Ntokos 2020; Chaudhury, Garg, and Mehta 2020; Farhadi et al. 2020). Another emerging line of work studies EFX for *non-monotone* valuations, i.e., when the resources consist of both goods and chores (Chen and Liu 2020; Bérczi et al. 2020).

The interaction between fairness and efficiency is further complicated with the addition of *strategyproofness* due to several fundamental impossibility results both in deterministic (Zhou 1990) as well as randomized settings (Bogomolnaia and Moulin 2001; Kojima 2009). Indeed, while ordinal efficiency is compatible with envy-freeness, such outcomes cannot, in general, be achieved via (weakly) strategyproof mechanisms even under strict preferences (Kojima 2009). Moreover, sd-efficiency and sd-strategyproofness (here, *sd* stands for stochastic dominance) are incompatible even with a weak notion of stochastic fairness called equal treatment of equals (Aziz and Kasajima 2017). In a similar vein, for deterministic mechanisms, any strategyproof mechanism could fail to satisfy EF1 even for two agents under additive valuations (Amanatidis et al. 2017).

Lexicographic preferences have also been successfully used as a domain restriction to circumvent impossibility results in mechanism design (Sikdar, Adali, and Xia 2017; Fujita et al. 2018). In fair division of indivisible goods, lexicographic (sub)additive utilities have facilitated constant-factor approximation algorithms for egalitarian and Nash social welfare objectives (Baumeister et al. 2017; Nguyen 2020). Hosseini and Larson (2019) showed that under lexicographic preferences, a mechanism is Pareto optimal, strategyproof, non-bossy, and neutral if and only if it is a serial dictatorship quota mechanism. In randomized settings, too, lexicographic preferences have led to the design of mechanisms that simultaneously satisfy stochastic efficiency, envy-freeness, and strategyproofness (Schulman and Vazirani 2015; Hosseini and Larson 2019).

Preliminaries

Model For any $k \in \mathbb{N}$, define $[k] := \{1, \dots, k\}$. An *instance* of the allocation problem is a tuple $\langle N, M, \succ \rangle$, where $N := [n]$ is a set of n agents, M is a set of m goods, and $\succ := (\succ_1, \dots, \succ_n)$ is a *preference profile* that specifies the ordinal preference of each agent $i \in N$ as a linear order $\succ_i \in \mathcal{L}$ over the set of goods; here, \mathcal{L} denotes the set of all (strict and complete) linear orders over M .

Allocation and bundles A *bundle* is any subset $X \subseteq M$ of the set of goods. An *allocation* $A = (A_1, \dots, A_n)$ is an n -partition of M , where $A_i \subseteq M$ is the bundle assigned to agent i . We will write Π to denote the set of all n -partitions of M . We say that allocation A is *partial* if $\bigcup_{i \in N} A_i \subset M$, and *complete* if $\bigcup_{i \in N} A_i = M$.

Lexicographic preferences We will assume that agents’ preferences over the bundles are given by the lexicographic

extension of their preferences over individual goods. Informally, this means that if an agent ranks the goods in the order $a \succ b \succ c \succ \dots$, then it prefers a bundle containing a over any other bundle that doesn't, subject to that, it prefers a bundle containing b over any other bundle that doesn't, and so on. Formally, given any pair of bundles $X, Y \subseteq M$ and any linear order $\succ_i \in \mathcal{L}$, we have $X \succ_i Y$ if and only if there exists a good $g \in X \setminus Y$ such that $\{g' \in Y : g' \succ_i g\} \subseteq X$. Notice that since \succ_i is a linear order over M , the corresponding lexicographic extension is a linear order over 2^M .

For any agent $i \in N$ and any pair of bundles $X, Y \in M$, we will write $X \succeq_i Y$ if either $X \succ_i Y$ or $X = Y$.

Envy-freeness Given a preference profile \succ , an allocation A is said to be (a) *envy-free* (EF) if for every pair of agents $i, h \in N$, we have $A_i \succeq_i A_h$; (b) *envy-free up to any good* (EFX) if for every pair of agents $i, h \in N$ such that $A_h \neq \emptyset$ and every good $j \in A_h$, we have $A_i \succeq_i A_h \setminus \{j\}$, and (c) *envy-free up to k goods* (EF k) if for every pair of agents $i, h \in N$ such that $A_h \neq \emptyset$, there exists a set $S \subseteq A_h$ such that $|S| \leq k$ and $A_i \succeq_i A_h \setminus S$. Clearly, $\text{EFX} \Rightarrow \text{EF1} \Rightarrow \text{EF2} \Rightarrow \dots$.

Maximin Share An agent's maximin share is its most preferred bundle that it can guarantee itself as a divider in an n -person cut-and-choose procedure against adversarial opponents (Budish 2011). Formally, the maximin share of agent i is given by $\text{MMS}_i := \max_{A \in \Pi} \min_i \{A_1, \dots, A_n\}$, where $\min\{\cdot\}$ and $\max\{\cdot\}$ denote the least-preferred and most-preferred bundles with respect to \succ_i . An allocation A satisfies *maximin share guarantee* (MMS) if each agent receives a bundle that it weakly prefers to its maximin share. That is, the allocation A is MMS if for every $i \in N$, $A_i \succeq_i \text{MMS}_i$. It is easy to see that $\text{EF} \Rightarrow \text{MMS}$. Additionally, for lexicographic preferences, we have that $\text{EFX} \Rightarrow \text{MMS}$ (the converse is not true) while EF1 and MMS can be incomparable (see the full version (Hosseini et al. 2020) for details).

Pareto optimality Given a preference profile \succ , an allocation A is said to be *Pareto optimal* (PO) if there is no other allocation B such that $B_i \succeq_i A_i$ for every agent $i \in N$ and $B_k \succ_k A_k$ for some agent $k \in N$.

Rank-maximality A *rank-maximal* (RM) allocation is one that maximizes the number of agents who receive their favorite good, subject to which it maximizes the number of agents who receive their second favorite good, and so on (Irving et al. 2006; Paluch 2013). Given an allocation A , its *signature* refers to a tuple (n_1, n_2, \dots, n_m) where n_i is the number of agents who receive their i^{th} favorite good (note that an agent can contribute to multiple n_i 's). All rank-maximal allocations for a given instance have the same signature. Computing *some* rank-maximal allocation for a given instance is easy: Assign each good to an agent that ranks it the highest among all agents (tiebreak arbitrarily). This procedure provides a computationally efficient way of computing the signature of a rank-maximal allocation as well as verifying whether a given allocation is rank-maximal. Notice that rank-maximality is a strictly stronger requirement than Pareto optimality.

Mechanism A mechanism $f : \mathcal{L}^n \rightarrow \Pi$ is a mapping from preference profiles to allocations. For any preference profile $\succ \in \mathcal{L}^n$, we use $f(\succ)$ to denote the allocation returned by

f , and $f_i(\succ)$ to denote the bundle assigned to agent i .

Properties of mechanisms A mechanism $f : \mathcal{L}^n \rightarrow \Pi$ is said to satisfy EF/EFX/EF k /PO/RM if for every preference profile $\succ \in \mathcal{L}^n$, the allocation $f(\succ)$ has that property. In addition, a mechanism f satisfies

- *strategyproofness* (SP) if no agent can improve by misreporting its preferences. That is, for every preference profile $\succ \in \mathcal{L}^n$, every agent $i \in N$, and every (misreported) linear order $\succ'_i \in \mathcal{L}$, we have $f_i(\succ) \succeq_i f_i(\succ')$, where $\succ' := (\succ_1, \dots, \succ_{i-1}, \succ'_i, \succ_{i+1}, \dots, \succ_n)$.
- *non-bossiness* if no agent can modify the allocation of another agent by misreporting its preferences without changing its own allocation. That is, for every profile $\succ \in \mathcal{L}^n$, every agent $i \in N$, and every (misreported) linear order $\succ'_i \in \mathcal{L}$, we have $f_i(\succ') = f_i(\succ) \Rightarrow f(\succ') = f(\succ)$, where $\succ' := (\succ_1, \dots, \succ_{i-1}, \succ'_i, \succ_{i+1}, \dots, \succ_n)$.
- *neutrality* if relabeling the goods results in a consistent change in the allocation. That is, for every preference profile $\succ \in \mathcal{L}^n$ and every relabeling of the goods $\pi : M \rightarrow M$, it holds that $f(\pi(\succ)) = \pi(f(\succ))$, where $\pi(\succ) := (\pi(\succ_1), \dots, \pi(\succ_n))$ and $\pi(A) := (\pi(A_1), \dots, \pi(A_n))$ for any allocation $A = (A_1, \dots, A_n)$.

EFX and Pareto Optimality

Recall that for additive valuations, establishing the existence of EFX allocations remains an open problem, and there exist instances where no allocation is simultaneously EFX and PO (Plaut and Roughgarden 2020). Our first result (Theorem 1) shows that there is no conflict between fairness and efficiency for lexicographic preferences: Not only does there exist a family of polynomial-time algorithms that always return EFX+PO allocations, but *every* EFX+PO allocation can be computed by some algorithm in this family. We will start with an easy observation concerning EFX allocations.

Proposition 1. *An allocation A is EFX if and only if each envied agent in A gets exactly one good.*

Description of algorithm Each algorithm in this family (Algorithm 1) is specified by an ordering σ over the agents, and consists of two phases. Phase 1 involves a single round of serial dictatorship according to σ . Phase 2 assigns the remaining goods among the *unenvied agents* according to a picking sequence τ . Note that the set of unenvied agents after Phase 1 must be nonempty; in particular, the last agent in σ belongs to this set since every other agent prefers its good picked in Phase 1 to any good in the last agent's bundle.

Theorem 1. *For any ordering σ of the agents, the allocation computed by Algorithm 1 satisfies EFX and PO. Conversely, any EFX+PO allocation can be computed by Algorithm 1 for some choice of σ .*

Proof. We will start by showing that the allocation A returned by Algorithm 1 satisfies EFX. From Proposition 1, it suffices to show that any envied agent gets exactly one good in A . Notice that any agent that is envied at the end of Phase 1 does not receive any good in Phase 2. Furthermore, the pairwise envy relations remain unchanged during Phase 2 since each agent has already picked its favorite available

ALGORITHM 1: EFX+PO

Input: An instance $\langle N, M, \succ \rangle$ with lexicographic preferences**Parameters:** A permutation $\sigma : N \rightarrow N$ of the agents**Output:** An allocation A $A \leftarrow (\emptyset, \dots, \emptyset)$ ▷ Phase 1: Serial dictatorship for assigning n goodsAgents arrive according to σ , and each picks a favorite good from the set of remaining goods. Update partial allocation A .

▷ Phase 2: Allocate leftover goods via picking sequence

if the set of remaining goods is nonempty **then** $U \leftarrow \{i \in N : i \text{ is not envied by any agent under } A\}$. Fix any picking sequence τ of length $m - n$ consisting only of the agents in U (i.e., the *unenvied* agents). Assign remaining goods according to τ and update A .**return** A

good in Phase 1, and because of lexicographic preferences, any goods assigned in Phase 2 are strictly less preferred. Thus, A is EFX.

To prove Pareto optimality (PO), suppose, for contradiction, that A is Pareto dominated by an allocation B . Then, there must exist some agent, say i , who receives a good under A that it does not receive under B (i.e., $A_i \setminus B_i \neq \emptyset$); we will call any such item a *difference* good. Observe that the execution of Algorithm 1 can be described in terms of a combined picking sequence $\langle \sigma, \tau \rangle$. Thus, without loss of generality, we can define i to be the *first* agent according to $\langle \sigma, \tau \rangle$ to receive a difference good. Let g denote the corresponding difference good picked by i , and note that $g \in A_i \setminus B_i$ by assumption.

Since B Pareto dominates A and $A_i \neq B_i$, we must have $B_i \succ_i A_i$. For lexicographic preferences, this means that there exists a good $g' \in B_i \setminus A_i$ such that $g' \succ_i g$. Since all agents preceding i in $\langle \sigma, \tau \rangle$ pick the goods that they also own under B , the good g' must be available (along with g) when it is i 's turn to pick. Thus, i would not pick g , which is a contradiction. Hence, A satisfies PO.

To prove the converse, note that any PO allocation can be induced by a picking sequence.¹ Given any EFX+PO allocation A , let S denote the corresponding picking sequence. We claim that without loss of generality, the first n positions in S belong to n different agents. Indeed, if some agent i appears more than once in the n -prefix of S , then $|A_i| > 1$. By Proposition 1, i must not be envied by any other agent. Because of the no-envy condition, a modified sequence where the repeated appearances of unenvied agents are “pushed behind” the first appearance of any other agent will also result in allocation A . We can now instantiate Algorithm 1 with σ as the n -prefix of S and τ as $S \setminus \sigma$ to compute A . \square

Characterizing Strategyproof Mechanisms

In addition to fairness and efficiency, an important desideratum for allocation mechanisms is strategyproofness. For ad-

¹Indeed, in any PO allocation, some agent must receive its favorite good (otherwise a cyclic exchange of the top-ranked goods gives a Pareto improvement). Add this agent to the picking sequence, and repeat the procedure for the remaining goods.

ALGORITHM 2:

Input: An instance $\langle N, M, \succ \rangle$ with lexicographic preferences**Parameters:** A permutation $\sigma : N \rightarrow N$ of the agents**Output:** An allocation A $A \leftarrow (\emptyset, \dots, \emptyset)$ Execute one round of serial dictatorship according to σ .Assign all remaining goods to the last agent in σ .**return** A

ditve valuations, strategyproofness is known to be incompatible even with EF1 (Amanatidis et al. 2017). By contrast, for lexicographic preferences, we will show that strategyproofness can be achieved in conjunction with a stronger fairness guarantee (EFX) as well as Pareto optimality, non-bossiness, and neutrality (Theorem 2). Indeed, Algorithm 2, a special case of Algorithm 1 where the last agent gets all the remaining goods characterizes these properties.

Our characterization result builds upon an existing result of Hosseini and Larson (2019, Theorem 5.6) (see Proposition 2) that characterizes four out of the five properties mentioned above (excluding EFX) in terms of *Serial Dictatorship Quota Mechanisms* (SDQ), as defined below.

Definition 1. The *Serial Dictatorship Quota* (SDQ) mechanism is specified by a permutation $\sigma : N \rightarrow N$ of the agents and a set of quotas (q_1, \dots, q_n) such that $\sum_{i=1}^n q_i = m$. Given a lexicographic instance $\langle N, M, \succ \rangle$ as input, the SDQ mechanism considers agents in the order σ , and assigns the i^{th} agent its most preferred bundle of size q_i from the remaining goods. The resulting allocation is returned as output.

Proposition 2 (Hosseini and Larson 2019). For lexicographic preferences, a mechanism is Pareto optimal, strategyproof, non-bossy, and neutral if and only if it is SDQ.

The next result (Theorem 2) provides an algorithmic characterization of EFX, PO, strategyproofness, non-bossiness, and neutrality for lexicographic preferences.

Theorem 2. For any ordering σ of the agents, Algorithm 2 is EFX, PO, strategyproof, non-bossy, and neutral. Conversely, any mechanism satisfying these properties can be implemented by Algorithm 2 for some σ .

Proof. Note that Algorithm 2 is a special case of SDQ for the quotas $q_i = 1$ for all $i \in [n - 1]$ and $q_n = m - (n - 1)$. Therefore, from Proposition 2, it is PO, strategyproof, non-bossy, and neutral. Furthermore, Algorithm 2 is also a special case of Algorithm 1 and is therefore EFX (Theorem 1).

To prove the converse, let f be an arbitrary mechanism satisfying the desired properties. From Proposition 2, f must be an SDQ mechanism for some ordering σ and some set of quotas (q_1, \dots, q_n) such that $\sum_{i=1}^n q_i = m$. If $m < n$, the claim follows easily from Theorem 1, so we can assume, without loss of generality, that $m \geq n$. Then, by Proposition 1, we must have that $q_i \geq 1$ for all $i \in [n]$. Therefore, it suffices to show that $q_i = 1$ for all $i \in [n - 1]$.

Assume, without loss of generality, that $\sigma = (1, 2, \dots, n)$. Consider a preference profile \succ with identical preferences, i.e., $\succ_i = \succ_k$ for all $i, k \in [n]$. Let $g_1 \succ_i g_2 \succ_i \dots \succ_i g_m$ for any $i \in [n]$. Suppose, for

contradiction, that $q_i > 1$ for some $i \in [n - 1]$, and let $k \in [n - 1]$ be the smallest index for which this happens. Since f is an SDQ mechanism, we have that $g_k \in f_k(\succ)$ and $|f_k(\succ)| = q_k > 1$. Then, for every $\ell > k$, agent ℓ envies agent k . By Proposition 1, f violates EFX, which is a contradiction. Therefore, f must be identical to Algorithm 2 for the ordering σ , as desired. \square

We note that in this context any deterministic strategyproof and non-bossy mechanism is also group-strategyproof (Pápai 2000a). Therefore, Algorithm 2 also characterizes the set of EFX, PO, group-strategyproof, non-bossy, and neutral mechanisms under lexicographic preferences. In addition, we show that the set of properties considered in Theorem 2 is *minimal*. That is, dropping any property from the characterization necessarily allows for feasible mechanisms beyond those in Algorithm 2. Details and the proof are relegated to the full version (Hosseini et al. 2020).

Proposition 3. *The set {EFX, PO, strategyproofness, non-bossiness, neutrality} is a minimal set of properties for characterizing the family of mechanisms in Algorithm 2.*

The efficiency guarantee in Theorem 2 cannot be strengthened much further, as there exists an instance where any *rank-maximal* (RM) mechanism violates strategyproofness (Example 1).

Example 1 (Strategyproofness and RM). *Consider the instance with $k + 2$ goods g_1, \dots, g_{k+2} and three agents:*

$$\begin{aligned} a_1 : & g_1 \succ g_2 \succ g_3 \succ \dots \succ g_{k+1} \succ g_{k+2} \\ a_2 : & g_1 \succ g_2 \succ g_3 \succ \dots \succ g_{k+1} \succ g_{k+2} \\ a_3 : & g_2 \succ g_3 \succ g_4 \succ \dots \succ g_{k+2} \succ g_1. \end{aligned}$$

Each of the goods g_2, \dots, g_{k+2} is ranked higher by a_3 than by a_1 or a_2 , and therefore must be assigned to a_3 in any rank-maximal allocation. Suppose, under truthful reporting, g_1 is assigned to a_1 , and a_2 gets an empty bundle. Then, a_2 could falsely report g_3 as its favorite good. By rank-maximality, g_3 is now assigned to a_2 , resulting in a strict improvement. \square

The non-existence result in Example 1 prompts us to forego strategyproofness (as well as non-bossiness and neutrality) and focus only on (approximate) envy-freeness and rank-maximality.

Envy-Freeness and Rank-Maximality

For lexicographic preferences, it is easy to see that a complete allocation is envy-free if and only if each agent receives its favorite good. Checking the existence of an envy-free allocation therefore boils down to computing a (left-)perfect matching in a bipartite graph where the left and the right vertex sets correspond to the agents and the goods, respectively, and the edges denote the top-ranked good of each agent. If an envy-free partial allocation of the top-ranked goods exists, then it can be extended to a complete rank-maximal allocation by assigning each remaining good to an agent that has the highest rank for it (note that the assignment of the remaining goods does not introduce any envy). Thus, the existence of an envy-free and rank-maximal allocation can

be checked efficiently for lexicographic preferences (Proposition 4).

Proposition 4. *There is a polynomial-time algorithm that, given a lexicographic instance as input, computes an envy-free and rank-maximal allocation, whenever one exists.*

Since an envy-free allocation is not guaranteed to exist, one could ask whether rank-maximality can always be achieved alongside *approximate* envy-freeness; in particular, EFX and EFX. Example 2 shows that both of these notions could conflict with rank-maximality. Specifically, for any fixed $k \in \mathbb{N}$, an EFX+RM allocation could fail to exist. Since EFX implies EF1, a similar incompatibility holds for EFX+RM as well.

Example 2 (EFX and RM). *Consider again the instance in Example 1. Any rank-maximal allocation assigns the goods g_2, \dots, g_{k+2} to a_3 . If g_1 is assigned to a_1 , then a_2 gets an empty bundle and the pair $\{a_2, a_3\}$ violates EFX.* \square

Given the non-existence result in Example 2, a natural question is whether there exists an efficient algorithm for checking the existence of an approximately envy-free and rank-maximal allocation. Unfortunately, the news here is also negative, as the problem turns out to be NP-complete (Theorem 3). Thus, while EFX can always be achieved in conjunction with Pareto optimality (Theorems 1 and 2), strengthening the efficiency notion to rank-maximality results in non-existence and computational hardness.

Theorem 3. *Determining whether a given instance admits an EFX and rank-maximal allocation is NP-complete.*

Proof. Membership in NP follows from the fact that both EFX and rank-maximality can be checked in polynomial time. To prove NP-hardness, we will show a reduction from a restricted version of 3-SAT called (2/2/3)-SAT, which is known to be NP-complete (Ahadi and Dehghan 2019). An instance of (2/2/3)-SAT consists of a collection of r variables X_1, \dots, X_r and s clauses C_1, \dots, C_s , where each clause is specified as a disjunction of three literals, and each variable occurs in exactly four clauses, twice negated and twice non-negated. The goal is to determine if there is a truth assignment that satisfies all clauses.

Construction of the reduced instance: We will construct a fair division instance with $n = 4r$ agents and $m = 4r + s$ goods. The set of agents consists of $2r$ *literal* agents $\{x_i, \bar{x}_i\}_{i \in [r]}$, and $2r$ *dummy* agents $\{d_i, \bar{d}_i\}_{i \in [r]}$. The set of goods consists of $2r$ *signature* goods $\{S_i, \bar{S}_i\}_{i \in [r]}$, s *clause* goods $\{C_j\}_{j \in [s]}$, and $2r$ *dummy* goods $\{T_i, B_i\}_{i \in [r]}$; here T_i and B_i denote the *top* and the *bottom* dummy goods associated with the variable X_i , respectively.

Preferences: Table 1 shows the preferences of the agents. Let \triangleright define a *reference* ordering on the set of goods. For every $i \in [r]$, if C_j and C_k denote the two clauses containing the positive literal x_i , then the literal agent x_i ranks S_i at the top, and the clause goods C_j and C_k at ranks $j + 1$ and $k + 1$, respectively. The missing positions consist of remaining goods ranked according to \triangleright (we write \triangleright_ℓ to denote the top ℓ goods in \triangleright that have not been ranked so far). The symbol $*$ indicates rest of the goods ordered according to \triangleright . The

\triangleright :	$S_1 \succ \bar{S}_1 \succ \dots \succ S_r \succ \bar{S}_r \succ T_1 \succ \dots \succ T_r$ $\succ C_1 \succ \dots \succ C_s \succ B_1 \succ \dots \succ B_r$
x_i :	$S_i \succ \triangleright_{(j-1)} \succ C_j \succ \triangleright_{(k-j-1)} \succ C_k \succ *$
\bar{x}_i :	$\bar{S}_i \succ \triangleright_{(p-1)} \succ C_p \succ \triangleright_{(q-p-1)} \succ C_q \succ *$
d_i :	$T_i \succ S_i \succ B_i \succ *$ and $\bar{d}_i : T_i \succ \bar{S}_i \succ B_i \succ *$.

Table 1: Preferences of agents in the proof of Theorem 3.

preferences of the (negative) literal agent \bar{x}_i and the dummy agents d_i, \bar{d}_i are defined similarly as shown in Table 1. This completes the construction of the reduced instance.

Note that for any fixed $i \in [r]$, the signature good S_i (or \bar{S}_i) is ranked at the top position by the literal agent x_i (or \bar{x}_i), and at a lower position by all other agents. Therefore, any rank-maximal allocation must assign S_i to x_i and \bar{S}_i to \bar{x}_i . For a similar reason, a rank-maximal allocation must assign the clause good C_j to a literal agent corresponding to a literal contained in the clause C_j , and the dummy goods T_i, B_i to the dummy agents d_i, \bar{d}_i . The aforementioned *necessary* conditions for rank-maximality are also *sufficient* since each clause good C_j is ranked at the same position by all literal agents corresponding to the literals contained in clause C_j , and the goods T_i and B_i are ranked identically by d_i and \bar{d}_i .

We will now argue the equivalence of solutions.

(\Rightarrow) Given a satisfying truth assignment, the desired allocation, say A , can be constructed as follows: For every $i \in [r]$, assign the signature goods S_i and \bar{S}_i to the literal agents x_i and \bar{x}_i , respectively. If $x_i = 1$, then assign T_i to d_i and B_i to \bar{d}_i , otherwise, if $x_i = 0$, then assign T_i to \bar{d}_i and B_i to d_i . For every $j \in [s]$, the clause good C_j is assigned to a literal agent x_i (or \bar{x}_i) if the literal x_i (or \bar{x}_i) is contained in the clause C_j and the clause is satisfied by the literal, i.e., $x_i = 1$ (or $\bar{x}_i = 1$). Note that under a satisfying assignment, each clause must have at least one such literal.

Observe that allocation A satisfies the aforementioned sufficient condition for rank-maximality. Furthermore, any envied agent in A receives exactly one good; in particular, if d_i receives a bottom dummy good B_i , then we have $x_i = 0$ in which case the literal agent x_i , who is envied by d_i , does not receive any clause goods. By Proposition 1, A is EFX.

(\Leftarrow) Now suppose there exists an EFX and rank-maximal allocation A . Then, A must satisfy the aforementioned necessary condition for rank-maximality. That is, for every $i \in [r]$, the signature goods S_i and \bar{S}_i are assigned to the literal agents x_i and \bar{x}_i , respectively (i.e., $S_i \in A_{x_i}$ and $\bar{S}_i \in A_{\bar{x}_i}$), and the dummy goods T_i and B_i are allocated between the dummy agents d_i and \bar{d}_i (i.e., $\{T_i, B_i\} \subseteq A_{d_i} \cup A_{\bar{d}_i}$). In addition, for every $j \in [s]$, the clause good C_j is assigned to a literal agent x_i (or \bar{x}_i) such that the literal x_i (or \bar{x}_i) is contained in the clause C_j . Also, by Proposition 1, each dummy agent must get exactly one dummy good.

We will construct a truth assignment for the (2/2/3)-SAT instance as follows: For every $i \in [r]$, if $T_i \in A_{d_i}$, then set $x_i = 1$, otherwise set $x_i = 0$. Note that the assignment is feasible as no literal is assigned conflicting values. To see why this is a satisfying assignment, consider any clause C_j .

Suppose the clause good C_j is assigned to a literal agent x_i (an analogous argument works when \bar{x}_i gets C_j). Then, due to rank-maximality, we know that the literal x_i must be contained in the clause C_j . Furthermore, since agent x_i gets more than one good ($S_i, C_j \in A_{x_i}$), it cannot be envied under A (Proposition 1). Thus, the dummy agent d_i must get the top good T_i . Recall that in this case we set $x_i = 1$. Since clause C_j contains x_i , it must be satisfied, as desired. \square

The intractability in Theorem 3 persists even when we relax the fairness requirement from EFX to EFk.

Theorem 4. *For any fixed $k \in \mathbb{N}$, determining the existence of an EFk and rank-maximal allocation is NP-complete.*

We note that the proof of Theorem 4 (see the full version (Hosseini et al. 2020) for details) differs considerably from that of Theorem 3 as neither result is an immediate consequence of the other. Indeed, a YES instance of EFX+RM is also a YES instance of EFk+RM, but the same is not true for a NO instance.

A corollary of Theorem 4 is that checking the existence of EF1+RM allocations for *additive valuations* is also NP-complete.² For this setting, Aziz et al. (2019) have shown NP-completeness even for three agents. By contrast, we will show that for lexicographic preferences, the problem is efficiently solvable when $n = 3$ (Proposition 5). The proof of this result is deferred to the full version of the paper (Hosseini et al. 2020).

Proposition 5. *There is a polynomial-time algorithm that, given as input a lexicographic instance with three agents, computes an EF1 and rank-maximal allocation, whenever one exists.*

Another avenue for circumventing the intractability in Theorem 3 is provided by *maximin share guarantee* (MMS). For additive valuations, EFX and MMS are incomparable notions (Amanatidis, Birmpas, and Markakis 2018). However, for lexicographic preferences, MMS is strictly weaker than EFX (see the full version (Hosseini et al. 2020) for details). This relaxation of EFX turns out to be computationally useful, as the existence of an MMS and RM allocations can be checked in polynomial time.

Theorem 5. *There is a polynomial-time algorithm that, given as input a lexicographic instance, computes an MMS and rank-maximal allocation, whenever one exists.*

We defer the proof of Theorem 5 to the full version (Hosseini et al. 2020), and briefly outline the algorithm below.

Fix any agent $i \in N$, and suppose its preference is given by $\succ_i := g_1 \succ g_2 \succ \dots \succ g_m$. Under lexicographic preferences, the MMS partition of agent $i \in N$ is uniquely defined as $\{\{g_1\}, \{g_2\}, \dots, \{g_{n-1}\}, \{g_n, g_{n+1}, \dots, g_m\}\}$. This observation gives a characterization of MMS allocations: An allocation is MMS if and only if each agent either receives one or more of its top- $(n-1)$ goods, or it receives all of its bottom- $(m-n+1)$ goods.

Construct a bipartite graph $G = (N \cup M, E)$ between agents and goods where an edge $(i, j) \in E$ exists if agent i

²An additive valuations instance in which agent i values its j^{th} favorite good at 2^{m-j+1} is equivalent to the lexicographic instance.




Figure 2: The plots show how the fraction of instances that admit $\{EF, EFX, EF1, MMS\} + RM$ allocations varies with the number of goods. The number of agents is fixed ($n = 5$). Preferences follow the Mallows model with dispersion parameter ϕ .

ranks good j within its top- $(n-1)$ goods, and good j can be ‘rank-maximally assigned’ to agent i . In other words, agent i ranks good j at least as high as any other agent.

If G admits a perfect matching (this can be checked in polynomial time), then, by the above characterization, we have a partial allocation that is MMS and rank-maximal. By assigning the unmatched goods in a rank-maximal manner (which can be done in polynomial time), we obtain a desired complete allocation.

Experiments

We now revisit the non-existence result in Example 2 by examining how frequently $\{EF, EFX, EF1, MMS\} + RM$ allocations exist in synthetically generated data. To that end, we consider a fixed number of agents ($n = 5$) whose preferences over a set of m goods (where $m \in \{5, \dots, 15\}$) are generated using the Mallows model (Mallows 1957). Given a reference ranking $\succ^* \in \mathcal{L}$ and a dispersion parameter $\phi \in [0, 1]$, the probability of generating a ranking $\succ_i \in \mathcal{L}$ under the Mallows model is given by $\frac{1}{Z} \phi^{d(\succ^*, \succ_i)}$, where Z is a normalization constant and $d(\cdot)$ is the Kendall’s Tau distance. Thus, $\phi = 0$ induces identical preferences (i.e., $\succ_i = \succ^*$) while $\phi = 1$ is the uniform distribution. For each combination of m , n , and $\phi \in \{0, 0.25, 0.5, 0.75, 1\}$, we sample 1000 preference profiles, and use an integer linear program to check the existence of $\{EF, EFX, EF1, MMS\} + RM$ allocations. Code and data for all our experiments is available at <https://github.com/sujoyksikdar/Envy-Freeness-With-Lexicographic-Preferences>.

Figure 2 presents our experimental results. For identical preferences ($\phi = 0$), every complete allocation is Pareto optimal as well as rank-maximal. Therefore, an EFX+RM (and hence $\{EF1, MMS\} + RM$) allocation always exists in this case, validating our theoretical result in Theorem 1. On the other hand, an EF+RM allocation fails to exist because of the conflict in top-ranked goods. At the other extreme for $\phi = 1$ (i.e., the uniform distribution), we note that the probability of existence of EF+RM outcomes grows steadily with m . This is because for (exact) envy-freeness, all five rankings should have distinct top goods, the probability of which is $(1 - \frac{1}{m}) \cdot (1 - \frac{2}{m}) \cdot (1 - \frac{3}{m}) \cdot (1 - \frac{4}{m})$. For $m = 100$, this value is roughly 0.9, suggesting that in the asymptotic regime, envy-free (and, by extension, EF+RM) allocations are increasingly likely to exist, and our algorithm in Proposition 4 will return EF+RM outcomes with high probability.

We observe that the gap between the fractions of instances that admit EFX+RM allocations and EF+RM allocations decreases with the number of goods. We conjecture that as the number of goods increases, the likelihood that every agent must be allocated more than one good in any RM allocation increases. Therefore, it is likely that envied agents receive more than one good, which is in direct conflict with EFX. In the full version (Hosseini et al. 2020), we test this conjecture through experiments for larger values of m , and observe an increasing trend in the fractions of instances that admit EF+RM allocations as well as those that admit EFX+RM allocations, while the gap between the two shrinks rapidly, and becomes negligible for larger number of goods. We also observe the general trend in our plots that $\{EF1, MMS\} + RM$ allocations tend to exist more frequently as the number of goods increases. Together, these observations suggest that the distributional approach could be a promising avenue for addressing the non-existence result in Example 2.

Concluding Remarks

We studied the interplay of fairness and efficiency under lexicographic preferences, obtaining strong algorithmic characterizations for EFX and Pareto optimality that addressed notable gaps in the additive valuations model, and outlining the computational limits of our approach for the stronger efficiency notion of rank-maximality. Going forward, it would be interesting to develop distribution-specific algorithms that can compute EFX+RM (or EF1+RM) allocations with, say, a constant probability. Extending our algorithmic characterization results to other fairness notions, specifically EF1 and MMS, will also be of interest.

Acknowledgments

HH acknowledges support from NSF grant #1850076. RV acknowledges support from ONR#N00014-171-2621 while he was affiliated with Rensselaer Polytechnic Institute, and is currently supported by project no. RTI4001 of the Department of Atomic Energy, Government of India. Part of this work was done while RV was supported by the Prof. R Narasimhan postdoctoral award. LX acknowledges support from NSF #1453542 and #1716333, ONR #N00014-171-2621, and a gift fund from Google. We thank the anonymous reviewers for their very helpful comments and suggestions.

References

- Ahadi, A.; and Dehghan, A. 2019. (2/2/3)-SAT Problem and its Applications in Dominating Set Problems. *Discrete Mathematics & Theoretical Computer Science* 21(4).
- Amanatidis, G.; Birmpas, G.; Christodoulou, G.; and Markakis, E. 2017. Truthful Allocation Mechanisms without Payments: Characterization and Implications on Fairness. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, 545–562.
- Amanatidis, G.; Birmpas, G.; Filos-Ratsikas, A.; Hollender, A.; and Voudouris, A. A. 2020. Maximum Nash Welfare and Other Stories about EFX. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence*, 24–30.
- Amanatidis, G.; Birmpas, G.; and Markakis, E. 2018. Comparing Approximate Relaxations of Envy-Freeness. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, 42–48.
- Amanatidis, G.; Markakis, E.; and Ntokos, A. 2020. Multiple Birds with One Stone: Beating $1/2$ for EFX and GMMS via Envy Cycle Elimination. *Theoretical Computer Science* 841: 94 – 109.
- Aziz, H.; Huang, X.; Mattei, N.; and Segal-Halevi, E. 2019. The Constrained Round Robin Algorithm for Fair and Efficient Allocation. *arXiv preprint arXiv:1908.00161*.
- Aziz, H.; and Kasajima, Y. 2017. Impossibilities for Probabilistic Assignment. *Social Choice and Welfare* 49(2): 255–275.
- Barman, S.; Krishnamurthy, S. K.; and Vaish, R. 2018. Finding Fair and Efficient Allocations. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, 557–574.
- Baumeister, D.; Bouveret, S.; Lang, J.; Nguyen, N.-T.; Nguyen, T. T.; Rothe, J.; and Saffidine, A. 2017. Positional Scoring-Based Allocation of Indivisible Goods. *Autonomous Agents and Multi-Agent Systems* 31(3): 628–655.
- Bérczi, K.; Bérczi-Kovács, E. R.; Boros, E.; Gedefa, F. T.; Kamiyama, N.; Kavitha, T.; Kobayashi, Y.; and Makino, K. 2020. Envy-free Relaxations for Goods, Chores, and Mixed Items. *arXiv preprint arXiv:2006.04428*.
- Bogomolnaia, A.; and Moulin, H. 2001. A New Solution to the Random Assignment Problem. *Journal of Economic Theory* 100(2): 295–328.
- Budish, E. 2011. The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes. *Journal of Political Economy* 119(6): 1061–1103.
- Caragiannis, I.; Gravin, N.; and Huang, X. 2019. Envy-Freeness Up to Any Item with High Nash Welfare: The Virtue of Donating Items. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, 527–545.
- Caragiannis, I.; Kurokawa, D.; Moulin, H.; Procaccia, A. D.; Shah, N.; and Wang, J. 2019. The Unreasonable Fairness of Maximum Nash Welfare. *ACM Transactions on Economics and Computation* 7(3): 12.
- Chaudhury, B. R.; Garg, J.; and Mehlhorn, K. 2020. EFX Exists for Three Agents. In *Proceedings of the Twenty-First ACM Conference on Economics and Computation*, 1–19.
- Chaudhury, B. R.; Garg, J.; and Mehta, R. 2020. Fair and Efficient Allocations under Subadditive Valuations. *arXiv preprint arXiv:2005.06511*.
- Chaudhury, B. R.; Kavitha, T.; Mehlhorn, K.; and Sgouritsa, A. 2020. A Little Charity Guarantees Almost Envy-Freeness. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, 2658–2672.
- Chen, X.; and Liu, Z. 2020. The Fairness of Leximin in Allocation of Indivisible Chores. *arXiv preprint arXiv:2005.04864*.
- Elkind, E.; Lackner, M.; and Peters, D. 2016. Preference Restrictions in Computational Social Choice: Recent Progress. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*, 4062–4065.
- Farhadi, A.; Hajiaghayi, M.; Latifian, M.; Seddighin, M.; and Yami, H. 2020. Almost Envy-Freeness, Envy-Rank, and Nash Social Welfare Matchings. *arXiv preprint arXiv:2007.07027*.
- Fujita, E.; Lesca, J.; Sonoda, A.; Todo, T.; and Yokoo, M. 2018. A Complexity Approach for Core-Selecting Exchange under Conditionally Lexicographic Preferences. *Journal of Artificial Intelligence Research* 63: 515–555.
- Gigerenzer, G.; and Goldstein, D. G. 1996. Reasoning the Fast and Frugal Way: Models of Bounded Rationality. *Psychological Review* 103(4): 650.
- Hosseini, H.; and Larson, K. 2019. Multiple Assignment Problems under Lexicographic Preferences. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, 837–845.
- Hosseini, H.; Sikdar, S.; Vaish, R.; and Xia, L. 2020. Fair and Efficient Allocations under Lexicographic Preferences. *arXiv preprint arXiv:2012.07680*.
- Irving, R. W.; Kavitha, T.; Mehlhorn, K.; Michail, D.; and Paluch, K. 2006. Rank-Maximal Matchings. *ACM Transactions on Algorithms* 2: 602–610.
- Kojima, F. 2009. Random Assignment of Multiple Indivisible Objects. *Mathematical Social Sciences* 57(1): 134–142.
- Lang, J.; Mengin, J.; and Xia, L. 2018. Voting on Multi-Issue Domains with Conditionally Lexicographic Preferences. *Artificial Intelligence* 265: 18–44.
- Mahara, R. 2020. Existence of EFX for Two Additive Valuations. *arXiv preprint arXiv:2008.08798*.
- Mallows, C. L. 1957. Non-Null Ranking Models. *Biometrika* 44(1/2): 114–130.
- Nguyen, T. T. 2020. How to Fairly Allocate Indivisible Resources Among Agents Having Lexicographic Subadditive Utilities. In *Frontiers in Intelligent Computing: Theory and Applications*, 156–166.
- Paluch, K. 2013. Capacitated Rank-Maximal Matchings. In *International Conference on Algorithms and Complexity*, 324–335. Springer.

- Pápai, S. 2000a. Strategyproof Assignment by Hierarchical Exchange. *Econometrica* 68(6): 1403–1433.
- Pápai, S. 2000b. Strategyproof Multiple Assignment Using Quotas. *Review of Economic Design* 5: 91–105.
- Plaut, B.; and Roughgarden, T. 2020. Almost Envy-Freeness with General Valuations. *SIAM Journal on Discrete Mathematics* 34(2): 1039–1068.
- Pruhs, K.; and Woeginger, G. J. 2012. Divorcing Made Easy. In *International Conference on Fun with Algorithms*, 305–314.
- Saban, D.; and Sethuraman, J. 2014. A Note on Object Allocation under Lexicographic Preferences. *Journal of Mathematical Economics* 50: 283–289.
- Schmitt, M.; and Martignon, L. 2006. On the Complexity of Learning Lexicographic Strategies. *Journal of Machine Learning Research* 7(Jan): 55–83.
- Schulman, L. J.; and Vazirani, V. V. 2015. Allocation of Divisible Goods Under Lexicographic Preferences. In *35th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science*, 543.
- Sikdar, S.; Adali, S.; and Xia, L. 2017. Mechanism Design for Multi-Type Housing Markets. In *Thirty-First AAAI Conference on Artificial Intelligence*, 684–690.
- Taylor, M. 1970. The Problem of Salience in the Theory of Collective Decision-Making. *Behavioral Science* 15(5): 415–430.
- Varian, H. R. 1974. Equity, Envy, and Efficiency. *Journal of Economic Theory* 9(1): 63–91.
- Zhou, L. 1990. On a Conjecture by Gale about One-Sided Matching Problems. *Journal of Economic Theory* 52(1): 123–135.