Given that an EF1 allocation for chores can be computed through a picking sequence [4], an immediate implication of Proposition 1 is a polynomial-time algorithm for computing an EF1+PO allocation of chores under lexicographic preferences.

Corollary 1 (EF1+PO for chores). *Under lexicographic preferences an* EF1+PO *allocation of chores can be computed in polynomial time.*

However, when dealing with *mixed* items, sequencibility is no longer a sufficient condition for guaranteeing PO, even in the lexicographic domain.

Proposition 2 (PO and sequencibility for mixed items). *For mixed items under lexicographic preferences, Pareto optimality implies sequencibility, but the converse is not true even for objective mixed items.*

Proof. (sketch) To see why sequencibility does not imply PO, consider the objective mixed items instance with three items $\{o_1^+, o_2^+, o_3^-\}$ and two agents where agent 1's importance ordering is $o_1^+ \rhd o_2^+ \rhd o_3^-$, and agent 2's ordering is $o_3^- \rhd o_1^+ \rhd o_2^+$. The picking sequence $\langle 1, 2, 2\rangle$ allocates $\{o_1^+\}$ to agent 1 and $\{o_3^-, o_2^+\}$ to agent 2. However, this allocation is Pareto dominated by the allocation that gives all items to agent 1.

Given an instance with mixed items, there always exists a Pareto optimal allocation (since there are only finitely many allocations and Pareto domination is a transitive relation). Furthermore, one such allocation can be computed in polynomial time; in particular, the *rank-maximal* allocation is Pareto optimal [35]. □

4 RESULTS

We start our investigation by considering the strongest fairness notion—*envy-freeness*. As we will see, this notion will provide us our first point of distinction between goods and chores.

4.1 Envy-Freeness

With indivisible items, a complete and envy-free allocation may not always exist. Thus, it is of interest to ask whether one can efficiently determine the existence of such solutions. This problem admits a polynomial-time algorithm in case of goods under lexicographic preferences [34], but turns out to be NP-complete for chores, and by extension, for mixed items (Theorem 1).

Theorem 1 (EF for chores). *Determining whether a chores-only instance with lexicographic preferences admits an envy-free allocation is* NP-complete.

To understand the reason behind the sharp contrast in the complexity of the goods and chores problems, notice that for goods under lexicographic preferences, an allocation is envy-free if and only if each agent gets its top-ranked item. One can efficiently check whether there exists a partial allocation satisfying this property via a straightforward matching computation (by considering a bipartite graph whose vertex sets are the agents and the items and an edge between each agent and its top-ranked item). Furthermore, if such a partial allocation exists, *any* completion of it is also envy-free.

By contrast, envy-freeness for chores entails that for every agent, the *worst* or least-preferred chore (i.e., highest-ranked in the importance ordering) in its own bundle is strictly preferred over the worst chore in any other agent's bundle. Thus, given an envy-free partial allocation, its completion may no longer be envy-free since, upon receiving more items, a different chore could become the worst.

We note that the allocation constructed in the forward direction in the proof of Theorem 1 is sequencible. Due to the equivalence between sequencibility and PO (Proposition 1), this implies that NP-hardness also holds for EF+PO. Furthermore, the EF+PO problem is actually NP-complete because EF and PO are both efficiently checkable properties; the latter because of its equivalence with sequencibility which can be checked in polynomial time.[4]

4.2 Envy-Freeness up to any Item (EFX)

Let us now turn our attention to a relaxation of envy-freeness called envy-freeness up to any item (EFX). Prior work has shown that an EFX and Pareto optimal allocation always exists for goods under lexicographic preferences [34]. In the full version [35], we show that a similar positive result can be achieved for the chores-only problem via the following simple procedure: Fix a priority ordering σ over agents. Let the first agent in σ pick its most preferred $m - n$ chores. Then, all agents (including the first agent) pick one chore each according to σ from the remaining items.

Our main result in this section is that the above positive results for goods-only and chores-only models fail to extend to the mixed items setting: We show that an EFX allocation may not exist even for *objective* mixed items, i.e., when each item is either a common good or a common chore (Theorem 2).

Theorem 2 (Non-existence of EFX). *There exists an instance with objective mixed items and lexicographic preferences that does not admit any* EFX *allocation.*

Since lexicographic preferences are a subclass of additive valuations, our counterexample also shows that an EFX allocation fails to exist under non-monotone and additive valuations (Corollary 2).[5] Our result complements that of Bérczi et al. [13] who showed that an EFX allocation could fail to exist for two agents with non-monotone, *non-additive*, and identical utility functions.

Corollary 2. *An* EFX *allocation can fail to exist for instances with non-monotone and additive valuations.*

The counterexample in the proof of Theorem 2 (given below) uses only *four* agents and *seven* items. Interestingly, for the said number of agents and items, an EFX allocation is guaranteed to exist for *goods-only* instances even under *monotone* valuations [41], which is significantly more general than additive (or lexicographic) preferences. It is also known that when agents belong to one of two given "types", an EFX allocation is guaranteed to exist for *goods-only* instances under *monotone* valuations [41]. Our result

[4] To verify sequencibility of a given allocation, consider the following procedure: Identify all chores that are allocated to agents who "prefer them the most" (i.e., chores that are lowest ranked in the importance orderings of their owners). Add these chores to the sequence and remove them from further consideration. Now, among the remaining chores, again identify the ones allocated in the "most preferred" manner (i.e., lowest ranked in the owner's importance ordering among the remaining chores). Again, add these chores to the sequence and remove them from further consideration. Repeat this process for as long as possible. It can be observed that the given allocation is sequencible if and only if the sequence constructed above includes all chores.

[5] A valuation function $v_i : 2^M \rightarrow \mathbb{R}$ is *non-monotone* if for some subsets $T \subset S \subseteq M$, we have $v_i(T) > v_i(S)$ and for some (possibly different) subsets $T' \subset S' \subseteq M$, we have $v_i(T') < v_i(S')$.

in Theorem 2, which also has two types of agents, demonstrates a barrier to extending this result in the non-monotone setting, even under lexicographic preferences.

Proof. (of Theorem 2) Consider an objective mixed items instance with four agents. Agents 1 and 2 have the same importance ordering, and so do agents 3 and 4, as shown below:

$$1,2: \quad o_2^- \rhd o_3^- \rhd o_4^- \rhd o_1^+ \rhd o_5^- \rhd o_6^- \rhd o_7^-$$
$$3,4: \quad o_5^- \rhd o_6^- \rhd o_7^- \rhd o_1^+ \rhd o_2^- \rhd o_3^- \rhd o_4^-$$

Since the items are objective, we will find it convenient to use the phrases 'the good o_1' and 'the chore o_2' instead of just calling them 'items'.

Suppose, for contradiction, that an EFX allocation exists. Without loss of generality, suppose agent 1 gets the good o_1^+. Let A_i denote the bundle allocated to agent i. We will show a contradiction via case analysis, depending on the chores allocated to agent 1.

Case 1: Suppose $A_1 \cap \{o_2^-, o_3^-, o_4^-\} = \emptyset$. That is, agent 1's allocated chores are a (possibly empty) subset of $\{o_5^-, o_6^-, o_7^-\}$, which are all ranked below o_1^+ according to agent 1's importance ordering.

This means that regardless of what agent 2 gets, it prefers the bundle A_1 to its own bundle A_2. Therefore, A_2 must be empty, as otherwise agent 2 will prefer A_1 even when some chore is removed from A_2. Thus, the chores o_2^-, o_3^-, and o_4^- must be allocated to agents 3 and 4, which means that one of these agents must get at least two of these chores. Suppose, without loss of generality, that agent 3 gets at least two chores. Then, agent 3 would prefer the empty bundle A_2 after any chore is removed from A_3, a contradiction to EFX.

Case 2: Suppose $A_1 \cap \{o_5^-, o_6^-, o_7^-\} = \emptyset$. That is, agent 1's allocated chores are a subset of $\{o_2^-, o_3^-, o_4^-\}$, which are all ranked above o_1^+ according to agent 1's importance ordering.

This means that regardless of how the remaining chores are assigned, both agents 3 and 4 will strictly prefer A_1 over their respective bundles (because their most important item in A_1 is o_1^+). Now, if agent 3 or 4 is assigned any item, which must be a chore, then even after removing this chore, it would still envy A_1. Therefore, agents 3 and 4 cannot be allocated any item. This means that agent 2 gets at least $\{o_5^-, o_6^-, o_7^-\}$, which implies that after any item (which must be a chore) is removed from agent 2's bundle, agent 2 envies agent 3 (who is not allocated any item). This contradicts EFX.

Case 3: If $A_1 \cap \{o_2^-, o_3^-, o_4^-\} \neq \emptyset$ and $A_1 \cap \{o_5^-, o_6^-, o_7^-\} \neq \emptyset$. That is, agent 1 gets at least one chore above good o_1^+ and at least one chore below o_1^+ according to its importance ordering.

Choose any $x \in A_1 \cap \{o_2^-, o_3^-, o_4^-\}$ and $y \in A_1 \cap \{o_5^-, o_6^-, o_7^-\}$. Then, because of EFX, agent 1 should not prefer any other agent's bundle after y is removed from A_1. This means that for any $i \in \{2,3,4\}$, A_i must contain a chore that is ranked higher than x according to agent 1's importance order. However, there are at most two chores perceived to be ranked higher than x by agent 1, which contradicts EFX. □

ALGORITHM 1: Finding an EFX+PO allocation when there is an agent whose top-ranked item is a good.

Input: A lexicographic mixed instance $\langle N, M, G, C, \rhd \rangle$
Output: An EFX+PO allocation A

Select an arbitrary agent $i \in N$ such that $\rhd_i(1) \in G_i$
Let $C' := \{o \in M : \forall j \in N \setminus \{i\}, o \in C_j\}$ `// The set of all common chores for the remaining agents.`
$A_i \leftarrow \rhd_i(1) \cup C'$
$N \leftarrow N \setminus \{i\}$
$M \leftarrow M \setminus A_i$
▷ `The remaining instance has no common chore.`
while *there exists an unallocated item* **do**
  if $|N| = 1$ **then**
    Assign all items to the remaining agent
  else
    Find the smallest $k \in \{1, 2, \ldots, |M|\}$ such that the set $S^k := \{i \in N : \rhd_i(k) \in G_i\}$ is non-empty `// set of agents whose` k^{th}`-ranked item is a good.`
    Select any agent $j \in S^k$
    $C' := \{o \in M : \forall i \in N \setminus \{j\}, o \in C_i\}$
    $A_j \leftarrow \{\rhd_j(k)\} \cup C'$
    $N \leftarrow N \setminus \{j\}$
    $M \leftarrow M \setminus A_j$
return A

4.3 EFX and Pareto optimality

We have seen that an EFX allocation may not exist for mixed items. This negative result prompts us to identify a subclass of lexicographic instances with subjective mixed items for which an EFX and Pareto optimal allocation is guaranteed to exist. Specifically, we will now require that there be an agent whose *top-ranked* item in its importance ordering is a good (Theorem 3).

Theorem 3 (EFX+PO when some agent has a top-ranked good). *Given a lexicographic mixed instance where some agent's top-ranked item is a good, an* EFX+PO *allocation always exists and can be computed in polynomial time.*

Proof. (sketch) Let us start by discussing why the allocation returned by our algorithm is EFX, followed by a similar discussion for PO.

Description of Algorithm and EFX Guarantee. Intuitively, the assumption about some agent's top-ranked item being a good allows us to deal with the common chores without violating EFX as follows (see Algorithm 1): An agent whose top-ranked item is a good can be assigned that item together with all items that are common chores for the rest of the $n-1$ agents. Since the preferences are lexicographic, this agent will not envy any other agent regardless of how the remaining items are allocated.

The first agent is now eliminated from the instance along with its assigned bundle. Observe that the reduced instance (with $n-1$ agents) has *no common chore*, that is, each item is considered as a good by at least one agent. The algorithm now uses the following strategy iteratively: It identifies an agent with the highest-ranking good (say agent j and good g), gives good g to agent j together with the common chores of the remaining $n-2$ agents, and then eliminates agent j.

Note that since agent j receives its highest-ranked good among the remaining items, it will not envy any agent that is eliminated *after* it, regardless of how the remaining items are assigned. Furthermore, by the 'no common chores' property, any item that is a chore for the rest of the agents must be a good for agent j. This means that agent j only receives those items that it considers to be goods. Thus, when evaluating EFX from agent j's perspective, we only need to look at the items in other agents' bundles that agent j considers to be goods. For any agent that was eliminated *before* j, there can be at most one such item (by virtue of assigning common chores), and thus EFX is maintained.

Guaranteeing PO. Suppose, for contradiction, that the allocation A returned by Algorithm 1 is Pareto dominated by the allocation B. We will argue by induction that for every agent i, we must have $A_i \subseteq B_i$, which would contradict Pareto optimality since A and B must be distinct. For ease of discussion, let us name the agents according to the order in which they are eliminated by Algorithm 1.

Recall from the above discussion on EFX that for each agent i, the most important item in its bundle under A, namely $\rhd_i(1, A_i)$, must be a good. We will first show by induction (over i) that every agent i must retain the item $\rhd_i(1, A_i)$ in B_i. Indeed, agent 1 must retain $\rhd_1(1, A_1)$ in B_1 because it is agent 1's most important item in M and is a good for agent 1. Suppose each agent $h \in \{1, \dots, i-1\}$ retains $\rhd_h(1, A_h)$ in B_h. Then, by virtue of choosing an agent with the highest-ranking good (see Lines 11–14 in Algorithm 1), agent i's most important item in A_i, namely $\rhd_i(1, A_i)$, is also its most important "achievable" good, i.e., the most important item in the set $G_i \setminus \{\rhd_1(1, A_1), \rhd_2(2, A_2) \dots, \rhd_{i-1}(1, A_{i-1})\}$. Therefore, due to lexicographic preferences, $\rhd_i(1, A_i)$ must be retained in B_i, implying the induction hypothesis.

A similar inductive argument in the reverse direction (i.e., $n, n-1, \dots, 2, 1$) implies that for every agent i, we have $A_i \subseteq B_i$. Indeed, the last agent, namely agent n, must retain all of its items since every item in A_n is a good for agent n, and every item in $M \setminus \big(A_n \cup \{\rhd_1(1, A_1), \dots, \rhd_{n-1}(1, A_{n-1})\}\big)$ is a chore. Thus, $A_n \subseteq B_n$.

Next, suppose $A_k \subseteq B_k$ for all agents $k \in \{n, n-1, \dots, i+1\}$, where $i > 1$. We want to show that $A_i \subseteq B_i$. Since $i > 1$, all items in A_i are goods for agent i. If $A_i \setminus B_i \neq \emptyset$, then in order for B_i to be more preferable than A_i, there must be a good $g \in B_i \setminus A_i$ such that g has a higher importance than any item in $A_i \setminus B_i$. This, however, is not possible, since agent i gets its most important "achievable" good in A_i, i.e., the most important good in the set $G_i \setminus \{\rhd_1(1, A_1), \rhd_2(2, A_2) \dots, \rhd_{i-1}(1, A_{i-1})\}$. Thus, $A_i \subseteq B_i$ for all $i \in \{2, 3, \dots, n\}$. This, in turn, implies that $A_1 \subseteq B_1$, thereby finishing the induction and giving the desired contradiction. □

Another subclass of lexicographic instances where an EFX+PO allocation is guaranteed to exist is when every item is considered a good by at least one agent, i.e., there are no *common chores*.

Corollary 3 (EFX+PO for mixed instances without common chores). *Given a lexicographic mixed instance without any common chore, an* EFX+PO *allocation always exists and can be computed in polynomial time.*

Corollary 3 and our counterexample for EFX in Theorem 2 together raise an interesting question: Under lexicographic preferences, EFX allocation always exists with *zero* common chores (Corollary 3), but fails to exist with *six* common chores (Theorem 2). What happens for intermediate values 1, 2, 3, 4 or 5 common chores?

4.4 Maximin Share (MMS)

In light of the failure in guaranteeing EFX even for objective mixed items, we investigate the existence of MMS allocations for mixed items. We show that not only does an MMS allocation exist for *subjective* mixed items under lexicographic preferences, but also that such an allocation can be computed efficiently.

We start by characterizing MMS bundles by examining the structure of an agent's maximin share. Given a lexicographic mixed instance, an agent's maximin share is identified by its top-ranked item: If agent i's top-ranked item is a good, MMS_i is an empty set if the number of goods is less than the number of agents (i.e., $|G_i| < n$), or else it is the set of the least-preferred $|G_i| - n + 1$ goods. Otherwise, when agent i's top-ranked item is a chore, then MMS_i is uniquely defined by the union of the top-ranked item (worst chore) and all the goods.

Proposition 3 (Characterizing MMS for mixed items). *Given an instance $\langle N, M, G, C, \rhd \rangle$ with lexicographic mixed items, the maximin share of agent i can be defined based on whether its top-ranked item is a good or a chore, as follows:*

$$MMS_i = \begin{cases} G_i \setminus \rhd_i([n-1], G_i), & \text{if } \rhd_i(1) \in G_i \wedge |G_i| \geq n \\ \emptyset, & \text{if } \rhd_i(1) \in G_i \wedge |G_i| < n \\ \rhd_i(1, C_i) \cup G_i, & \text{if } \rhd_i(1) \in C_i. \end{cases}$$

Proof. The MMS partition of any agent $i \in N$ is uniquely defined based on whether its top-ranked item $\rhd_i(1)$ is a good or a chore:

Case 1. Top ranked item is a good, that is, $\rhd_i(1) \in G_i$: There are two cases according to the size of G_i.

(a) If $|G_i| \geq n$: the MMS partition for i is defined as

$$\{\{\rhd_i(1, G_i) \cup C_i\}, \{\rhd_i(2, G_i)\}, \dots, \{\rhd_i(n-1, G_i)\}, G_i \setminus \{\rhd_i([n-1], G_i)\}\}.$$

The MMS partition for i is the least-preferred bundle. Since preferences are lexicographic, we have $\text{MMS}_i = G_i \setminus \bigcup_{l \in [n-1]} \{\rhd_i(l, G_i)\} = G_i \setminus \rhd_i([n-1], G_i)$.

(b) If $|G_i| < n$: the MMS partition for i is uniquely defined as

$$\{\{\rhd_i(1, G_i) \cup C_i\}, \{\rhd_i(2, G_i)\}, \dots, \{\rhd_i(|G_i|, G_i)\}, \{\}, \dots, \{\}\}.$$

Therefore, $\text{MMS}_i = \emptyset$.

Case 2. Top ranked item is a chore, that is, $\rhd_i(1) \in C_i$: The MMS partition is uniquely defined as

$$\{\{\rhd_i(1, C_i) \cup G_i\}, \{\rhd_i(2, C_i)\}, \dots, \{\rhd_i(n, C_i)\}\}.$$

Note that if $|C_i| < n$, then $\{\rhd_i(k, C_i)\} = \emptyset$ for all $k < |C_i|$. The MMS for agent i is the least-preferred partition above. Since preferences are lexicographic, $\text{MMS}_i = \{\rhd_i(1, C_i) \cup G_i\}$. □

Although EFX may not always exist for mixed items (Theorem 2), we show that whenever such an allocation exists, it also satisfies

ALGORITHM 2: Algorithm for finding an MMS allocation for mixed items.

Input: A lexicographic mixed instance $\langle N, M, G, C, \rhd \rangle$
Output: An MMS allocation A

1 Let $C' := \{o \in M : \forall i \in N, o \in C_i\}$
 ▷ `Step 1: Assign chores according to top-ranked items`
2 **if** $\exists i \in N$ *such that* $\rhd_i(1) \in G_i$ **then**
3 Run Algorithm 1
4 **else** // `Else if` $\forall i \in N, \rhd_i(1) \in C_i$
5 Fix a priority ordering σ over n agents
6 **if** $|C'| \geq n$ **then**
7   Run a serial dictatorship where σ_1 picks its favorite (lowest ranked) $|C'| - n + 1$ chores
8   All remaining agents pick one chore each (lowest ranked chore among remaining chores)
9 **else**
10   Agents pick one chore (lowest ranked chore that remains) according to σ, and none if no chore is remaining
11 If there exists an agent who picked its worst chore (first in importance ordering), give that agent its remaining goods
 ▷ `Step 2: Serial dictatorship to assign remaining items`
12 Run a serial dictatorship according to any priority ordering; agents pick any number of goods among remaining items or nothing (if no item is a good for them).
13 **return** A

MMS. Note that the converse does not hold, that is, even for chores-only instances (where EFX always exists), MMS does not imply EFX (refer to the full version [35]).

Proposition 4 (EFX $\implies$ MMS for mixed items). *For mixed items under lexicographic preferences, an* EFX *allocation (whenever it exists) satisfies* MMS*, but the converse is not true.*

We develop an algorithm that computes an MMS allocation for *any* lexicographic instance—even with subjective mixed items—in polynomial time.

Description of algorithm. Our algorithm (Algorithm 2) first identifies the set C' of all common chores and proceeds in two steps: In **Step 1**, all common chores are allocated without violating MMS, and in **Step 2**, all remaining items are allocated as goods.

Step 1. *If there exists an agent whose top-ranked item is a good,* then run Algorithm 1 to achieve an EFX allocation. By Proposition 4, EFX implies MMS for mixed items.

Otherwise, if every agent's top item is a chore, a priority ordering σ over agents is fixed, and a serial dictatorship is run where agent σ_1 picks its most preferred (least important) $|C'| - n + 1$ chores, from the set of all common chores, C', and the remaining agents each pick one remaining chore from C'. Note that if $|C'| < n$, the first $n - |C'|$ agents pick one chore and the rest receive nothing. If an agent k receives its worst chore from C_k, it is given its remaining goods in G_k.

Step 2. All remaining items are allocated through a serial dictatorship. In each turn, an agent picks all remaining items it considers as goods, or picks nothing. All remaining items are only allocated as goods, and thus, do not violate MMS.

Theorem 4 (MMS for mixed items). *Given a lexicographic mixed instance, there is a polynomial-time algorithm that computes an* MMS *allocation even for subjective items.*

Proof. Algorithm 2 guarantees MMS for *any* lexicographic instance with mixed items. Let A be the output of the algorithm.

Case 1. There exists an agent i with top-ranked item as a good. That is, $\rhd_i(1) \in G_i$. Then, run Algorithm 1 that satisfies EFX (and PO). By Proposition 4, any EFX allocation is also MMS, thus, Algorithm 2 is MMS. In this case, the algorithm does not allocate any item in 'Step 2', thus, the allocation vacuously remains MMS.

Case 2. Every agent's top-ranked item is a chore. That is, $\forall i \in N, \rhd_i(1) \in C_i$. The proof relies on allocating items that are considered as chores by all agents, i.e., $C' := \{o \in M : \forall i \in N, o \in C_i\}$. All remaining items in $M \setminus C'$ by construction are considered goods by at least one agent. Algorithm 2 proceeds to first allocate items in C'—via a serial dictatorship specified by σ—such that the first agent σ_1 either receives its least important chore (if $|C'| < n$) or its $|C'| - n + 1$ least important chores (if $|C'| \geq n$). All other agents pick a single chore from $C' \setminus A_{\sigma_1}$ or an empty set, which satisfies MMS.

Suppose agent h receives its most disliked (i.e., rank 1 in importance ordering) chore in C_h. Then, since C' did not contain any item that is considered good by any agent, agent h receives all goods in G_h (Line 11) and $A_h = \rhd_h(1, C_h) \cup G_h$. Notice that only the last agent to pick a chore from C' (according from σ) can receive its worst (top-ranked) chore. If an agent h does not receive $\rhd_h(1, C_h)$, then $A_h \succeq_h \rhd_h(1, C_h) \cup G_h$ by lexicographic preferences. Together, this implies implies that A satisfies MMS by Proposition 3. The allocation of remaining items only improves the outcome for all agents since all remaining items are assigned as goods in a serial dictatorship by Algorithm 2 (Line 12). Thus, all agents' allocations weakly improves. □

The significance of Theorem 4 stems from providing an efficient algorithm for computing an MMS allocation for any lexicographic mixed instance (including subjective instances). Yet, the problem of computing an MMS+PO allocation remains open even for objective lexicographic instances.

5 CONCLUDING REMARKS

We studied the interaction between fairness and efficiency for a mixture of indivisible goods and chores under lexicographic preferences. We showed that an EFX allocation may not always exist for mixed items. Nonetheless, we identified natural classes of lexicographic instances for which an EFX+PO allocation exists and can always be computed efficiently. We further proved that an MMS allocation always exists and can be computed efficiently even for subjective mixed instances.

Going forward, it will be interesting to resolve the computational complexity of checking the existence of EFX allocations for mixed items. Another relevant direction will be to explore the space of strategyproof mechanisms satisfying desirable fairness and efficiency guarantees.

ACKNOWLEDGMENTS

We thank the anonymous reviewers for helpful comments and suggestions. HH acknowledges support from NSF grants #2144413, #2052488, and #2107173. RV acknowledges support from DST INSPIRE grant no. DST/INSPIRE/04/2020/000107 and SERB grant no. CRG/2022/002621. LX acknowledges support from NSF #1453542 and #2107173, and a Google gift fund.

REFERENCES

[1] Georgios Amanatidis, Georgios Birmpas, George Christodoulou, and Evangelos Markakis. 2017. Truthful Allocation Mechanisms without Payments: Characterization and Implications on Fairness. In *Proceedings of the 18th ACM Conference on Economics and Computation*. 545–562.

[2] Georgios Amanatidis, Georgios Birmpas, Aris Filos-Ratsikas, Alexandros Hollender, and Alexandros A Voudouris. 2021. Maximum Nash Welfare and Other Stories about EFX. *Theoretical Computer Science* 863 (2021), 69–85.

[3] Haris Aziz and Florian Brandl. 2021. Efficient, Fair, and Incentive-Compatible Healthcare Rationing. In *Proceedings of the 22nd ACM Conference on Economics and Computation*. 103–104.

[4] Haris Aziz, Ioannis Caragiannis, Ayumi Igarashi, and Toby Walsh. 2022. Fair Allocation of Indivisible Goods and Chores. *Autonomous Agents and Multi-Agent Systems* 36, 1 (2022), 1–21.

[5] Haris Aziz, Gerhard Rauchecker, Guido Schryen, and Toby Walsh. 2017. Algorithms for Max-Min Share Fair Allocation of Indivisible Chores. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*. 335–341.

[6] Haris Aziz, Toby Walsh, and Lirong Xia. 2015. Possible and Necessary Allocations via Sequential Mechanisms. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence*. 468–474.

[7] 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*, Vol. 35. 5119–5126.

[8] Moshe Babaioff, Noam Nisan, and Inbal Talgam-Cohen. 2019. Fair Allocation through Competitive Equilibrium from Generic Incomes. In *Proceedings of the 2019 ACM Conference on Fairness, Accountability and Transparency*. 180–180.

[9] Ana Babus, Sanmay Das, and SangMok Lee. 2020. The Optimal Allocation of Covid-19 Vaccines. *medRxiv* (2020).

[10] Siddharth Barman, Sanath Kumar Krishnamurthy, and Rohit Vaish. 2018. Finding Fair and Efficient Allocations. In *Proceedings of the 2018 ACM Conference on Economics and Computation*. 557–574.

[11] Xiaohui Bei, Zihao Li, Jinyan Liu, Shengxin Liu, and Xinhang Lu. 2021. Fair Division of Mixed Divisible and Indivisible Goods. *Artificial Intelligence* 293 (2021), 103436.

[12] Nawal Benabbou, Mithun Chakraborty, Xuan-Vinh Ho, Jakub Sliwinski, and Yair Zick. 2020. The Price of Quota-based Diversity in Assignment Problems. *ACM Transactions on Economics and Computation* 8, 3 (2020), 1–32.

[13] Kristóf Bérczi, Erika R Bérczi-Kovács, Endre Boros, Fekadu Tolessa Gedefa, Naoyuki Kamiyama, Telikepalli Kavitha, Yusuke Kobayashi, and Kazuhisa Makino. 2020. Envy-Free Relaxations for Goods, Chores, and Mixed Items. *arXiv preprint arXiv:2006.04428* (2020).

[14] Aurélie Beynier, Sylvain Bouveret, Michel Lemaître, Nicolas Maudet, Simon Rey, and Parham Shams. 2019. Efficiency, Sequenceability and Deal-Optimality in Fair Division of Indivisible Goods. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*. 900–908.

[15] Umang Bhaskar, AR Sricharan, and Rohit Vaish. 2021. On Approximate Envy-Freeness for Indivisible Chores and Mixed Resources. In *Proceedings of the 24th International Conference on Approximation Algorithms for Combinatorial Optimization Problems*.

[16] Anna Bogomolnaia, Hervé Moulin, Fedor Sandomirskiy, and Elena Yanovskaya. 2017. Competitive Division of a Mixed Manna. *Econometrica* 85, 6 (2017), 1847–1871.

[17] Sylvain Bouveret and Jérôme Lang. 2011. A General Elicitation-Free Protocol for Allocating Indivisible Goods. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence*. 73–78.

[18] Steven J Brams and Alan D Taylor. 1996. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press.

[19] Eric Budish. 2011. The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes. *Journal of Political Economy* 119, 6 (2011), 1061–1103.

[20] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D Procaccia, Nisarg Shah, and Junxing Wang. 2019. The Unreasonable Fairness of Maximum Nash Welfare. *ACM Transactions on Economics and Computation* 7, 3 (2019), 12.

[21] Bhaskar Ray Chaudhury, Jugal Garg, and Kurt Mehlhorn. 2020. EFX Exists for Three Agents. In *Proceedings of the 21st ACM Conference on Economics and Computation*. 1–19.

[22] Soroush Ebadian, Dominik Peters, and Nisarg Shah. 2022. How to Fairly Allocate Easy and Difficult Chores. In *Proceedings of the 21st International Conference on Autonomous Agents and MultiAgent Systems*. 372–380.

[23] Edith Elkind, Martin Lackner, and Dominik Peters. 2016. Preference Restrictions in Computational Social Choice: Recent Progress. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence*. 4062–4065.

[24] Duncan K. Foley. 1967. Resource Allocation and the Public Sector. *Yale Economic Essays* 7 (1967), 45–98.

[25] Etsushi Fujita, Julien Lesca, Akihisa Sonoda, Taiki Todo, and Makoto Yokoo. 2018. A Complexity Approach for Core-Selecting Exchange under Conditionally Lexicographic Preferences. *Journal of Artificial Intelligence Research* 63 (2018), 515–555.

[26] George Gamow and Marvin Stern. 1958. Puzzle-Math. , 45–46 pages.

[27] Jugal Garg and Aniket Murhekar. 2021. Computing Fair and Efficient Allocations with Few Utility Values. In *Proceedings of the 14th International Symposium on Algorithmic Game Theory*. Springer, 345–359.

[28] Jugal Garg, Aniket Murhekar, and John Qin. 2022. Fair and Efficient Allocations of Chores under Bivalued Preferences. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (To Appear)*.

[29] Jugal Garg and Setareh Taki. 2021. An Improved Approximation Algorithm for Maximin Shares. *Artificial Intelligence* 300 (2021), 103547.

[30] Mohammad Ghodsi, Mohammad Taghi Hajiaghayi, Masoud Seddighin, Saeed Seddighin, and Hadi Yami. 2021. Fair Allocation of Indivisible Goods: Improvement. *Mathematics of Operations Research* 46, 3 (2021), 1038–1053.

[31] Gerd Gigerenzer and Daniel G Goldstein. 1996. Reasoning the Fast and Frugal Way: Models of Bounded Rationality. *Psychological Review* 103, 4 (1996), 650.

[32] Hadi Hosseini and Kate Larson. 2019. Multiple Assignment Problems under Lexicographic Preferences. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*. 837–845.

[33] Hadi Hosseini and Andrew Searns. 2021. Guaranteeing Maximin Shares: Some Agents Left Behind. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence*. 238–244.

[34] Hadi Hosseini, Sujoy Sikdar, Rohit Vaish, and Lirong Xia. 2021. Fair and Efficient Allocations under Lexicographic Preferences. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence*. 5472–5480.

[35] Hadi Hosseini, Sujoy Sikdar, Rohit Vaish, and Lirong Xia. 2022. Fairly Dividing Mixtures of Goods and Chores under Lexicographic Preferences. *arXiv preprint arXiv:2203.07279* (2022).

[36] Robert W Irving, Telikepalli Kavitha, Kurt Mehlhorn, Dimitrios Michail, and Katarzyna E Paluch. 2006. Rank-Maximal Matchings. *ACM Transactions on Algorithms* 2, 4 (2006), 602–610.

[37] Rucha Kulkarni, Ruta Mehta, and Setareh Taki. 2021. Indivisible Mixed Manna: On the Computability of MMS+PO Allocations. In *Proceedings of the 22nd ACM Conference on Economics and Computation*. 683–684.

[38] David Kurokawa, Ariel D Procaccia, and Junxing Wang. 2018. Fair Enough: Guaranteeing Approximate Maximin Shares. *J. ACM* 65, 2 (2018), 1–27.

[39] Jérôme Lang, Jérôme Mengin, and Lirong Xia. 2018. Voting on Multi-Issue Domains with Conditionally Lexicographic Preferences. *Artificial Intelligence* 265 (2018), 18–44.

[40] Richard J Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. 2004. On Approximately Fair Allocations of Indivisible Goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce*. 125–131.

[41] Ryoga Mahara. 2021. Extension of Additive Valuations to General Valuations on the Existence of EFX. In *Proceedings of the 29th Annual European Symposium on Algorithms*. 66:1–66:15.

[42] Katarzyna Paluch. 2013. Capacitated Rank-Maximal Matchings. In *Proceedings of the 8th International Conference on Algorithms and Complexity*. 324–335.

[43] Parag A Pathak, Tayfun Sönmez, M Utku Ünver, and M Bumin Yenmez. 2021. Fair Allocation of Vaccines, Ventilators and Antiviral Treatments: Leaving No Ethical Value Behind in Health Care Rationing. In *Proceedings of the 22nd ACM Conference on Economics and Computation*. 785–786.

[44] Benjamin Plaut and Tim Roughgarden. 2020. Almost Envy-Freeness with General Valuations. *SIAM Journal on Discrete Mathematics* 34, 2 (2020), 1039–1068.

[45] Daniela Saban and Jay Sethuraman. 2014. A Note on Object Allocation under Lexicographic Preferences. *Journal of Mathematical Economics* 50 (2014), 283–289.

[46] Michael Schmitt and Laura Martignon. 2006. On the Complexity of Learning Lexicographic Strategies. *Journal of Machine Learning Research* 7, Jan (2006), 55–83.

[47] Michael Taylor. 1970. The Problem of Salience in the Theory of Collective Decision-Making. *Behavioral Science* 15, 5 (1970), 415–430.

[48] Vignesh Viswanathan and Yair Zick. 2022. Yankee Swap: A Fast and Simple Fair Allocation Mechanism for Matroid Rank Valuations. *arXiv preprint arXiv:2206.08495* (2022).

[49] Shengwei Zhou and Xiaowei Wu. 2021. Approximately EFX Allocations for Indivisible Chores. *arXiv preprint arXiv:2109.07313* (2021).