

Equitable Allocations of Indivisible Goods

Rupert Freeman¹, Sujoy Sikdar², Rohit Vaish^{2*} and Lirong Xia²

¹Microsoft Research New York City

²Rensselaer Polytechnic Institute

rupert.freeman@microsoft.com, sikdas@rpi.edu, vaishr2@rpi.edu, xial@cs.rpi.edu

Abstract

In fair division, *equitability* dictates that each participant receives the same level of utility. In this work, we study equitable allocations of indivisible goods among agents with additive valuations. While prior work has studied (approximate) equitability in isolation, we consider equitability in conjunction with other well-studied notions of fairness and economic efficiency. We show that the Leximin algorithm produces an allocation that satisfies equitability up to any good and Pareto optimality. We also give a novel algorithm that guarantees Pareto optimality and equitability up to one good in pseudopolynomial time. Our experiments on real-world preference data reveal that approximate envy-freeness, approximate equitability, and Pareto optimality can often be achieved simultaneously.

1 Introduction

We consider fair division problems that require a central planner to divide a set of *goods* among a group of *agents*—each with their own individual preferences over the goods—such that the resulting allocation is fair. How exactly one can certify that an allocation is “fair” remains a subject of debate, but the literature suggests two distinct viewpoints. In the first viewpoint, an agent should prefer her bundle of goods to some comparison bundle. The gold standard of fairness here is *envy-freeness*, which says that each agent should prefer her bundle of goods to any other agents’ bundle.

In this work, we consider the second viewpoint, in which agents compare their happiness levels, or *utilities*. Here, an allocation is considered fair if the planner is able to make all agents equally well-off. A central fairness notion in this context is *equitability*: An equitable allocation is one where agents derive equal utilities from their assigned shares. Stated differently, an equitable allocation seeks to minimize the disparity between the best-off and the worst-off agents.

Both perspectives have merit, but the practical importance of equitability as a fairness criterion has been highlighted in an experimental study conducted by [2009]. They asked human subjects to deliberate over an assignment of indivisible

goods subject to a time limit. It was found that the chosen outcomes were equitable (and Pareto optimal) far more often than they were envy-free. They concluded that equitability is a significant predictor of the perceived fairness of an allocation, often more so than envy-freeness.

Like many other fairness notions, equitability has been traditionally studied for *divisible* goods (i.e., cake-cutting). In this setting, it is known that an equitable allocation always exists [Dubins and Spanier, 1961; Alon, 1987]. On the computability side, it is known that no finite procedure can find an (exact) equitable division [Procaccia and Wang, 2017], though an ϵ -equitable division can be computed in a finite number of steps [Cechlárová and Pillárová, 2012a; Cechlárová and Pillárová, 2012b].

For *indivisible* goods, an equitable (EQ) allocation might fail to exist even with two agents and a single good, motivating the need for approximations. To this end, [2014] proposed the notion of *near jealousy-freeness*, under which for any pair of agents, the disparity can be reversed by removing any good from the bundle of the agent with higher utility. We refer to this notion as *equitability up to any good* (EQx) in keeping with the nomenclature for a similar relaxation of envy-freeness [Caragiannis *et al.*, 2016]. We also study *equitability up to one good* (EQ1), requiring only that inequity can be eliminated by removing *some* good from the higher-utility-agent’s bundle. [2014] showed that for additive valuations, an EQx (hence, EQ1) allocation always exists and can be computed in polynomial time. However, they did not study Pareto optimality (PO), a fundamental and often desirable notion of economic efficiency that may still be violated by an (approximately) equitable allocation.

Our work takes a deeper dive into the study of (approximately) equitable allocations of indivisible goods—in conjunction with Pareto optimality as well as other well-studied notions of fairness (envy-freeness and its relaxations)—and considers a host of existence and computational questions. Table 1 provides a comprehensive summary of our results. Some of the highlights are:

- We strengthen the aforementioned result of [2014] to show that an EQx *and* PO allocation always exists for strictly positive valuations (Proposition 3). Without the positivity assumption, even an EQ1+PO allocation might fail to exist (Example 1), and finding an EQ/EQx/EQ1+PO allocation becomes strongly NP-

*Contact Author

Fairness + Efficiency	Existence Results		Computational Results	
	general	special case	general	special case
EQ EQx EQ1	✗ even for two agents and one good		strongly NP-c even for id (Proposition 1)	
	✓ (Proposition 2)		Poly-time (Proposition 2)	
PO + { EQ EQx EQ1	✗ even for two agents and one good		strongly NP-h (Remark 1)	Poly-time for bin (Theorem 2)
	✗ (Example 1)	✓ for pos (Proposition 3)		Poly-time for bin (Theorem 4)
EF + PO + { EQ EQx EQ1	✗ even for two agents and one good		strongly NP-h (Theorem 1)	Pseudopoly-time for pos (Theorem 3)
	EFx + PO + { EQ EQx EQ1	✗ even for bin (Example 1)		Poly-time for bin (Remark 3)
EF1 + PO + { EQ EQx EQ1			✗ even for pos (Proposition 4)	NP-c even for bin (Remark 4)
			strongly NP-h (Corollary 1)	Poly-time for bin (Remark 3)
			Poly-time for bin (Theorem 4)	
			Poly-time for bin (Remark 3)	
			Poly-time for bin (Theorem 4)	

Table 1: Summary of results. For “Existence Results,” a ✓ denotes guaranteed existence while a ✗ indicates that existence might fail for some instance. For “Computational Results,” NP-c/NP-h refers to NP-complete/NP-hard. The shorthands bin, id, and pos refer to binary, identical, and strictly positive valuations, respectively.

hard (Theorem 1 and Remark 1).

- As a step towards making the above existence result constructive, we design a pseudopolynomial-time algorithm that always returns an EQ1+PO allocation for strictly positive valuations (Theorem 3).
- We construct an instance in which no allocation can be EQ1+EF1+PO (Proposition 4).¹ We show that determining whether such an allocation exists is, in general, strongly NP-hard (Corollary 1), but the special case of binary valuations is efficiently solvable (Theorem 4).
- We validate our theoretical results via experiments on the data from the popular fair division website *Spliddit*² as well as on synthetically generated instances (Section 4).

Related Work

For *divisible* goods (i.e., cake-cutting), [1961] showed that an equitable division always exists (without providing a bound on the number of cuts). Subsequent work has established the existence of equitable divisions where each agent gets a contiguous piece [Cechlárová *et al.*, 2013; Aumann and Dömbb, 2015; Chèze, 2017].

Equitability has also been studied in combination with other fairness and efficiency notions. It is known that there always exists a cake division that is simultaneously equitable and envy-free [Alon, 1987]. However, existence might fail if, in addition, one also requires Pareto optimality [Brams *et al.*, 2013] or contiguous pieces [Brams *et al.*, 2006]. Connections between Pareto optimality and social welfare maximizing equitable divisions have also been studied [Brams *et al.*, 2012].

For *indivisible* goods, in addition to the work of [2014] discussed above, [2019] studies equitable and connected allocations of indivisible goods (i.e., when the goods constitute the vertices of a graph and a feasible allocation assigns every agent a connected subgraph).

¹EF1 stands for *envy-freeness up to one good*, which is a (necessary) relaxation of envy-freeness defined for indivisible goods; see Section 2 for the relevant definitions.

²<http://www.spliddit.org/>

2 Preliminaries

Problem instance. An instance $\langle [n], [m], \mathcal{V} \rangle$ of the fair division problem is defined by a set of $n \in \mathbb{N}$ agents $[n] = \{1, 2, \dots, n\}$, a set of $m \in \mathbb{N}$ goods $[m] = \{1, 2, \dots, m\}$, and a valuation profile $\mathcal{V} = \{v_1, v_2, \dots, v_n\}$ that specifies the preferences of every agent $i \in [n]$ over each subset of the goods in $[m]$ via a valuation function $v_i : 2^{[m]} \rightarrow \mathbb{N} \cup \{0\}$.³ We will assume that the valuation functions are *additive*, i.e., for any agent $i \in [n]$ and any set of goods $S \subseteq [m]$, $v_i(S) := \sum_{j \in S} v_i(\{j\})$, where $v_i(\emptyset) = 0$. For a singleton good $j \in [m]$, we will write $v_{i,j}$ instead of $v_i(\{j\})$.

Allocation. An allocation $A := (A_1, \dots, A_n)$ is an n -partition of the set of goods $[m]$, where $A_i \subseteq [m]$ is the *bundle* allocated to the agent i . Given an allocation A , the *utility* of an agent $i \in [n]$ for the bundle A_i is $v_i(A_i) = \sum_{j \in A_i} v_{i,j}$.

Equitable allocations. An allocation A is said to be *equitable* (EQ) if for every pair of agents $i, k \in [n]$, we have $v_i(A_i) = v_k(A_k)$. An allocation A is *equitable up to one good* (EQ1) if for every pair of agents $i, k \in [n]$ such that $A_k \neq \emptyset$, there exists some good $j \in A_k$ such that $v_i(A_i) \geq v_k(A_k \setminus \{j\})$. An allocation A is *equitable up to any good* (EQx) if for every pair of agents $i, k \in [n]$ such that $A_k \neq \emptyset$ and for every good $j \in A_k$ such that $v_{k,j} > 0$, we have $v_i(A_i) \geq v_k(A_k \setminus \{j\})$.⁴

Envy-free allocations. An allocation A is *envy-free* (EF) if for every pair of agents $i, k \in [n]$, we have $v_i(A_i) \geq v_i(A_k)$. An allocation A is *envy-free up to one good* (EF1) if for every pair of agents $i, k \in [n]$ such that $A_k \neq \emptyset$, there exists some good $j \in A_k$ such that $v_i(A_i) \geq v_i(A_k \setminus \{j\})$. An allocation A is *envy-free up to any good* (EFx) if for every pair of agents $i, k \in [n]$ such that $A_k \neq \emptyset$ and for every good $j \in A_k$ such that $v_{i,j} > 0$, we have $v_i(A_i) \geq v_i(A_k \setminus \{j\})$. The notions

³Integrality of valuations is required only for Theorem 3. Our positive results (i.e., existence and algorithmic results) hold even in the absence of this assumption, and negative results (i.e., non-existence and hardness results) hold even for integral valuations.

⁴Our results hold analogously for the following variant of EQx due to [2014]: For every pair of agents $i, k \in [n]$ such that $A_k \neq \emptyset$, $v_i(A_i) \geq v_k(A_k \setminus \{j\})$ for every good $j \in A_k$.

of EF, EF1, and EFX are due to [1967], [2011],⁵ and [2016], respectively.

Pareto optimality. An allocation A is *Pareto optimal* if there is no other allocation B such that $v_k(B_k) \geq v_k(A_k)$ for every agent $k \in [n]$ with at least one of the inequalities being strict.

Nash social welfare. Given an instance $\langle [n], [m], \mathcal{V} \rangle$, the *Nash social welfare* of an allocation A is defined as $\text{NSW}(A) := \left(\prod_{i \in [n]} v_i(A_i) \right)^{1/n}$. An allocation A^* is called *Nash optimal* or *MNW* (Maximum Nash Welfare) if it maximizes the Nash social welfare among all allocations.⁶

Leximin-optimal allocations. A Leximin-optimal allocation [Dubins and Spanier, 1961] is one that maximizes the minimum utility that any agent achieves, subject to which the second-minimum utility is maximized, and so on. The utilities induced by a Leximin-optimal allocation are unique, although there may exist more than one such allocation.

3 Results

This section presents our theoretical results, summarized in Table 1. We first consider equitability and its relaxations, then consider them in conjunction with Pareto optimality, before finally adding envy-freeness (and its relaxations) to the mix.

3.1 Existence and Computation of EQ, EQ1, EQx

We will start by observing that envy-freeness and equitability (and their corresponding relaxations) become equivalent when the valuations are *identical* (i.e., when, for every good $j \in [m]$, $v_{i,j} = v_{k,j}$ for all $i, k \in [n]$).

Proposition 1. *For identical valuations, an allocation is EF/EF1/EFx if and only if it is EQ/EQ1/EQx.*

It is known that determining whether a given instance has an envy-free (EF) allocation is NP-complete even for identical valuations (via a straightforward reduction from PARTITION) [Lipton *et al.*, 2004].⁷ Proposition 1 implies that the same holds for equitable (EQ) allocations. By contrast, an EQx (and therefore EQ1) allocation always exists and can be efficiently computed (Proposition 2) even for non-identical valuations.

Proposition 2 ([Gourvès *et al.*, 2014]). *An EQx allocation always exists and can be computed in polynomial time.*

Briefly, [2014] prove Proposition 2 using a greedy algorithm. In each round, the algorithm assigns a least-happy agent its favorite good from among the remaining goods. Thus, at any stage, the most recent good assigned to an agent is also its least-favorite good in its own bundle. Since each new good is assigned to an agent with the least utility, an allocation that is EQx prior to the assignment continues to be so

⁵ [2004] previously defined a slightly weaker notion than EF1, but their algorithm can, in fact, compute an EF1 allocation.

⁶[2016] define a Nash optimal allocation as one that provides positive utility to the largest set of agents, and subject to that, maximizes the geometric mean of valuations. Our results hold even under this extended definition.

⁷In fact, the problem is strongly NP-complete due to a similar reduction from 3-PARTITION.

after it (up to the removal of the most recently assigned good). The claim now follows by induction over the rounds.

Proposition 2 presents an interesting contrast between the notions of EQx and EFX: An EQx allocation is guaranteed to exist and can be efficiently computed, whereas for EFX, even the question of guaranteed existence is an open problem.

3.2 Equitability and Pareto Optimality

We now turn our attention to computing an allocation that is both equitable up to one good and Pareto optimal (we use the shorthand EQ1+PO for such allocations). Unfortunately, such allocations might fail to exist when the valuations are allowed to be zero-valued (Example 1). This provides an interesting contrast with the analogous relaxation of envy-freeness; it is known that an allocation satisfying EF1 and PO always exists [Caragiannis *et al.*, 2016; Barman *et al.*, 2018a].

Example 1 (Non-existence of EQ1+PO). Consider an instance with three agents a_1, a_2, a_3 and six goods g_1, \dots, g_6 . The goods g_1, g_2, g_3 are valued at 1 by a_1 and at 0 by a_2 and a_3 . The goods g_4, g_5, g_6 are valued at 1 by a_2 and a_3 and at 0 by a_1 . Any PO allocation must assign g_1, g_2, g_3 to a_1 (giving it a utility of 3) and allocate g_4, g_5, g_6 between a_2 and a_3 . Either a_2 or a_3 receives at most one good, creating an EQ1 violation with a_1 . Thus, an EQ1 and PO allocation might fail to exist even under *binary* valuations. \square

Worse still, when the valuations can be zero-valued, determining whether there exists an EQ1+PO allocation is strongly NP-hard. Similar hardness results hold for EQx+PO and EQ+PO allocations as well (Remark 1).

Theorem 1 (Hardness of EQ1 + PO). *Given any fair division instance with additive valuations, determining whether there exists an allocation that is equitable up to one good (EQ1) and Pareto optimal (PO) is strongly NP-hard.*

Proof. We will show a reduction from 3-PARTITION, which is known to be strongly NP-hard. An instance of 3-PARTITION consists of a set of $3r$ numbers $S = \{b_1, \dots, b_{3r}\}$ where $r \in \mathbb{N}$, and the goal is to find a partition of S into r subsets S_1, \dots, S_r such that the sum of numbers in each subset is T , where $T := \frac{1}{r} \sum_{a_i \in S} b_i$.⁸

We will construct a fair division instance as follows: There are $r+1$ agents a_1, \dots, a_{r+1} and $3r+2$ goods g_1, \dots, g_{3r+2} . For every $i \in [r]$ and $j \in [3r]$, agent a_i values the good g_j at b_j . The agents a_1, \dots, a_r all value the goods g_{3r+1} and g_{3r+2} at 0. Finally, the agent a_{r+1} values g_{3r+1} and g_{3r+2} at T each, and all other goods at 0.

(\Rightarrow) Suppose S_1, \dots, S_r is a solution of 3-PARTITION. Then, an EQ1 and PO allocation $A = (A_1, \dots, A_{r+1})$ can be constructed as follows: For every $i \in [r]$, $A_i := \{g_j : b_j \in S_i\}$, and $A_{r+1} := \{g_{3r+1}, g_{3r+2}\}$. Notice that A is EQ1 because each of the agents a_1, \dots, a_r has utility T , and the utility of the agent a_{r+1} exceeds T only by a single good g_{3r+2} . Furthermore, A is PO because each good is assigned to an agent with the highest valuation for it.

(\Leftarrow) Now suppose that $A = (A_1, \dots, A_{r+1})$ is an EQ1 and PO allocation. Since A is PO, it must assign g_{3r+1} and g_{3r+2}

⁸Note that we do not require S_1, \dots, S_r to be of size three each; 3-PARTITION remains strongly NP-hard even without this constraint.

to a_{r+1} . Furthermore, since A is EQ1, each of the agents a_1, \dots, a_r should have a utility of at least T under A , i.e., for every $i \in [r]$, $v_i(A_i) \geq v_{r+1}(A_{r+1} \setminus \{g_{3r+2}\}) = T$. This induces a solution of the 3-PARTITION instance. \square

Remark 1 (Hardness of EQx+PO/EQ+PO). The reduction in Theorem 1 can also be used to prove strong NP-hardness of finding an EQx+PO allocation (same construction works) or an EQ+PO allocation (if a_{r+1} values g_{3r+2} at 0).

Our next result shows that for the special case of *binary* valuations (i.e., for all $i \in [n], j \in [m], v_{i,j} \in \{0, 1\}$), an EQ+PO allocation, if it exists, can be computed in polynomial time. Later, we will show similar tractability results for EQ1+PO and EQx+PO allocations (Theorem 4).

Theorem 2 (Algorithm for EQ+PO for binary valuations). *There is a polynomial-time algorithm that given as input any fair division instance with additive and binary valuations, returns an allocation that is equitable (EQ) and Pareto optimal (PO) whenever such an allocation exists.*

Proof. We will use a maximum flow algorithm. For binary valuations, an allocation is PO if and only if it assigns each good to an agent that approves it. For an EQ allocation A , we have $v_i(A_i) = v_k(A_k) = c$ (say) for every $i, k \in [n]$. Consider a bipartite graph $G = ([n] \cup [m], E)$ over the set of agents and goods with an edge $(i, j) \in E$ for every $i \in [n]$ and $j \in [m]$ such that $v_{i,j} = 1$. For any fixed $c \in \mathbb{N}$, construct a flow network where the source node S is connected to each agent node in $[n]$ with an edge of capacity c . Each node corresponding to a good in $[m]$ is connected to the sink node T with an edge of capacity 1. The edges between agents and goods are of capacity 1. It is straightforward to check that there exists an EQ+PO allocation in the fair division instance (with common utility c) if and only if the above network admits a feasible flow of value $n \cdot c$. The desired algorithm simply iterates over all integral values of c between 1 and $\lfloor m/n \rfloor$. \square

On the other hand, when all valuations are *strictly positive* (i.e., $v_{i,j} > 0$ for all i, j), there always exists an allocation that is both equitable up to *any* good and Pareto optimal.

Proposition 3 (Existence of EQx+PO for positive valuations). *Given any fair division instance with additive and strictly positive valuations, an allocation that is equitable up to any good (EQx) and Pareto optimal (PO) always exists.*

Proof. (Sketch.) We will show that any Leximin-optimal allocation, say A , satisfies EQx (Pareto optimality is easy to verify). Suppose, for contradiction, that there exist agents $i, k \in [n]$ and some good $j \in A_k$ such that $v_i(A_i) < v_k(A_k \setminus \{j\})$. Let B be an allocation derived from A by transferring the good j from agent k to agent i . Notice that under B , both agents i and k have strictly greater utility than $v_i(A_i)$, while all other agents have exactly the same utility as under A . Thus, B is a ‘Leximin improvement’ over A —a contradiction. \square

Although Proposition 3 offers a strong existence result, it does not automatically provide a constructive procedure for finding such allocations. Indeed, computing a Leximin-optimal allocation is known to be intractable [Bezáková and Dani, 2005; Plaut and Roughgarden, 2018]. Our

next result (Theorem 3) addresses this gap by providing a pseudopolynomial-time algorithm for finding an EQ1 and PO allocation when the valuations are strictly positive.

Theorem 3 (Algorithm for EQ1+PO for positive valuations). *Given any fair division instance $\mathcal{I} = \langle [n], [m], \mathcal{V} \rangle$ with additive and strictly positive valuations, an allocation that is equitable up to one good (EQ1) and Pareto optimal (PO) always exists and can be computed in $O(\text{poly}(m, n, v_{\max}))$ time, where $v_{\max} = \max_{i,j} v_{i,j}$.*

In particular, when the valuations are polynomially bounded (i.e., for every $i \in [n]$ and $j \in [m]$, $v_{i,j} \leq \text{poly}(m, n)$), our algorithm runs in *polynomial time*. In contrast, computing a Leximin-optimal allocation remains NP-hard even under this restriction [Bezáková and Dani, 2005].

The proof of Theorem 3 is deferred to the full version [Freeman *et al.*, 2019] but a brief idea is as follows: Our algorithm uses the framework of *Fisher markets*, which are well-studied models of a set of buyers spending their budgets of virtual money on utility-maximizing bundles of goods. Standard welfare theorems in economics guarantee that equilibrium (i.e., market clearing) outcomes in these markets are economically efficient. However, such outcomes could, in general, lead to *fractional* allocations and be highly *inequitable*. Our algorithm addresses the first challenge by starting with (and always maintaining) an *integral* equilibrium of *some* Fisher market. To meet the second challenge, our algorithm uses a combination of local search and price-rise routines to gradually move towards an *approximately* equitable equilibrium. The analysis for achieving the desired running time and correctness guarantees is intricate, and involves a number of structural observations and potential function arguments.

Our techniques are inspired from a similar recent algorithm of [2018a] for finding allocations that are envy-free up to one good (EF1) and Pareto optimal (PO). A key difference between the two algorithms lies in the way a local improvement is defined: For [2018a], a local improvement is defined in terms of equalizing the agents’ *spendings*, whereas for us, it pertains to equalizing the agents’ *utilities*. We believe that the latter approach is more direct, and leads to a simpler algorithm and analysis. This distinction is also *necessary*, because as we will show in Proposition 4, an EQ1+EF1+PO allocation might fail to exist even with strictly positive valuations. Therefore, any algorithm that is tailored to return an EF1 outcome—including the algorithm of [2018a]—will invariably fail to find the desired EQ1+PO allocation, motivating the need for an alternative approach.

Given the success of market-based algorithms in finding EQ1+PO allocations, it is natural to ask whether these techniques can be extended to find an EQx+PO allocation. Unfortunately, this is where these techniques hit a roadblock. The problem stems from the fact that the market-based algorithm always outputs a *fractionally Pareto optimal* (fPO) allocation, but there exist instances where no EQx allocation satisfies fPO [Freeman *et al.*, 2019]. Whether an EQx+PO allocation can be computed in (pseudo-)polynomial time with strictly positive valuations is an intriguing question for future research.

3.3 Equitability, Envy-Freeness and Pareto Optimality

We will now consider all three notions—equitability, envy-freeness, and Pareto optimality—together. Recall from Proposition 3 that for strictly positive valuations, an EQ1+PO (in fact, an EQx+PO) allocation is guaranteed to exist. It is also known that an EF1+PO allocation always exists. One might therefore ask whether an EQ1+EF1+PO allocation also always exists. Our next result (Proposition 4) rules it out.

Proposition 4 (Non-existence of EQ1+EF1+PO). *There exists an instance with strictly positive valuations in which no allocation is simultaneously equitable up to one good (EQ1), envy-free up to one good (EF1) and Pareto optimal (PO).*

Proof. Fix some $n \geq 2$ and $0 < \varepsilon < \frac{1}{2n+2}$. Consider an instance with $n + 1$ agents a_1, \dots, a_{n+1} and $3n + 1$ goods g_1, \dots, g_{3n+1} . Each of a_1, \dots, a_n values each of g_1, \dots, g_{n-1} at 2 and each of g_n, \dots, g_{3n+1} at ε . Agent a_{n+1} values every good at 1. By the pigeonhole principle for the goods g_1, \dots, g_{n-1} , some agent among a_1, \dots, a_n must have utility at most $(2n + 2)\varepsilon < 1$. This means that a_{n+1} can be assigned at most one good (otherwise EQ1 is violated). Therefore, if all the goods are allocated (which is a necessary condition for a PO allocation), at least $3n$ goods must be assigned among a_1, \dots, a_n . This means that one of these agents gets at least three goods, creating an EF1 violation with a_{n+1} . \square

Remark 2. Proposition 4 has several interesting implications. First, it shows that a Nash optimal allocation—which is guaranteed to be EF1 and PO [Caragiannis *et al.*, 2016]—need not satisfy EQ1. Similarly, the algorithm of [2018a] for computing an EF1 and PO allocation could also fail to return an EQ1 allocation. By contrast, our algorithm in Theorem 3 is guaranteed to find an EQ1 and PO allocation. Finally, it shows that the Leximin-optimal allocation—which is guaranteed to be EQx and PO for strictly positive valuations (Proposition 3)—need not be EF1.

Comparison with cake-cutting. It is worth comparing Proposition 4 with the corresponding results for *divisible* goods (i.e., cake-cutting). [2013] have shown that there might not exist a division of the cake that simultaneously satisfies EQ, EF, and PO. Our result in Proposition 4 shows an analogous impossibility for *indivisible* goods. Interestingly, the impossibility for cake-cutting goes away when PO is relaxed to *completeness* (i.e., only requiring that the entire cake is allocated). Under this relaxation, it is known that a *perfect* allocation of the cake exists [Alon, 1987].⁹ By contrast, for indivisible goods, the impossibility remains even when PO is relaxed to completeness and EF1 is relaxed to proportionality up to one good (Prop1).¹⁰ Indeed, the proof of Proposition 4 works even under these relaxations. Moreover, the proof can be easily extended to show the non-existence of EQk, Prop ℓ and complete allocations for any constants $k, \ell \in \mathbb{N}$.

⁹An allocation A is *perfect* if for every $i, k \in [n]$, $v_i(A_k) = \frac{1}{n}$.

¹⁰An allocation A is *proportional* if for every $i \in [n]$, we have $v_i(A_i) \geq \frac{1}{n} \sum_{k \in [n]} v_i(A_k)$. An allocation A is *proportional up to one good* [Conitzer *et al.*, 2017] if for every $i \in [n]$, there exists a good g such that $v_i(A_i \cup \{g\}) \geq \frac{1}{n} \sum_{k \in [n]} v_i(A_k)$.

We now turn to the computational aspects of allocations with all three properties. Note that the allocation constructed in the proof of Theorem 1 is envy-free. Therefore, from Theorem 1 and Remark 1, we obtain strong NP-hardness of all combinations of the three properties.

Corollary 1 (Hardness of EF+EQ+PO). *Let $X \in \{\text{EF}, \text{EFx}, \text{EF1}\}$, $Y \in \{\text{EQ}, \text{EQx}, \text{EQ1}\}$, and $Z = \text{PO}$. Then, determining whether a given instance admits an allocation that is simultaneously X , Y , and Z is strongly NP-hard.*

The intractability in Corollary 1 can, in certain cases, be alleviated when the valuations are restricted to be *binary*. We will start with an observation concerning EQ and PO allocations under this restriction.

Proposition 5. *For binary valuations, an allocation that is equitable (EQ) and Pareto optimal (PO) is also envy-free (EF).*

Proof. Suppose each agent gets a utility k under the said EQ allocation. For binary valuations, PO implies that each agent i approves all the goods in its bundle. Furthermore, any other agent j gets at most k goods approved by i (simply because agent j gets exactly k goods). Hence, the allocation is EF. \square

Remark 3. Proposition 5 shows that for binary valuations, an EQ+PO allocation (if it exists) is, in fact, EQ+PO+EF (hence also EQ+PO+EFx/EF1). From Theorem 2, we know that there is a polynomial-time algorithm for determining whether an instance with binary valuations admits an EQ+PO allocation. A similar implication therefore also holds for EQ+PO+EF/EF1/EFx allocations.

Theorem 4 shows that binary valuations are also useful when one considers the combination of EQ1, EF1, and PO.

Theorem 4 (Algorithm for EQ1+EF1+PO for binary valuations). *There is a polynomial-time algorithm that given as input any fair division instance with additive and binary valuations, returns an allocation that is equitable up to one good (EQ1), envy-free up to one good (EF1), and Pareto optimal (PO), whenever such an allocation exists.*

The proof of Theorem 4 is deferred to the full version [Freeman *et al.*, 2019]. The idea is to show that any EQ1+PO allocation, if it exists, is also Nash optimal. For binary valuations, all Nash optimal allocations induce identical utility profiles (up to renaming of agents). As a result, every Nash optimal allocation satisfies EQ1. It is known that every Nash optimal allocation satisfies EF1 and PO [Caragiannis *et al.*, 2016]. Moreover, for binary valuations, a Nash optimal allocation can be computed in polynomial time [Darmann and Schauer, 2015; Barman *et al.*, 2018b]. Therefore, determining the existence of an EQ1+EF1+PO allocation reduces to checking whether an arbitrary Nash optimal allocation satisfies EQ1, which can be done in polynomial time.

Notice that for binary valuations, a Pareto optimal allocation is EF1 if and only if it is EFx, and is EQ1 if and only if it is EQx. Therefore, when the valuations are binary, the above algorithm works for all combinations of $X + Y + \text{PO}$, where $X \in \{\text{EFx}, \text{EF1}\}$ and $Y \in \{\text{EQx}, \text{EQ1}\}$.

We conclude this section by observing that some of the problems discussed in Corollary 1 continue to be intractable even for binary valuations. This follows from a result of




Figure 1: Experimental results for Spliddit and synthetic datasets.

[2008], who showed that finding an envy-free (EF) and Pareto optimal (PO) allocation under binary valuations is NP-complete (refer to Proposition 21 in their paper).

Proposition 6 ([Bouveret and Lang, 2008]). *Given any fair division instance with additive and binary valuations, determining whether there exists an envy-free (EF) and Pareto optimal (PO) allocation is NP-complete.*

Remark 4. It is easy to verify that the allocation constructed in the reduction of [2008] is, without loss of generality, equitable up to one good (EQ1). Therefore, for binary valuations, determining whether there exists an allocation that is EF + EQ1/EQx + PO is NP-complete.

4 Experiments

In this section, we compare the proposed and existing algorithms (in particular, ALG-EQ1+PO, MNW, and Leximin) in terms of how frequently they satisfy various fairness and efficiency properties in the real-world and synthetic datasets.

For real-world preferences, we used data obtained from the popular fair division website *Spliddit* [Goldman and Procaccia, 2014]. Out of the 2212 instances in the Spliddit data, we used the 914 instances that had strictly positive valuations and $m \geq n$. The instances have between 3 and 9 agents, and between 3 and 29 goods.¹¹ Users are restricted to normalized, integral valuations. For synthetic data, we generated 1000 instances with $n = 5$, $m = 20$, and (strictly positive) valuations drawn i.i.d. from Dirichlet distribution. The concentration parameter for each item is set to 10 to generate normalized valuations.¹²

We consider the following fairness and efficiency properties: EQ+PO, EQ1+PO, EQx+PO, EQ1+EF1+PO, and EQx+EFx+PO. For each instance of the Spliddit and synthetic datasets, we check whether the property is satisfied by the output of ALG-EQ1+PO, MNW, and Leximin. Figure 1 presents the relevant histograms. Note that each of the algorithms we consider is Pareto optimal, so the histograms would be unaltered even if we did not assess PO.

Not surprisingly, we see that very few instances permit a solution that is Pareto optimal and exactly equitable. When-

¹¹More than 80% of the instances have three agents and six goods.

¹²We normalize the valuations in the synthetic data to allow for a fair comparison with the Spliddit data, which has normalized valuations by design. We remark that all algorithms studied in this paper work even in the absence of this assumption.

ever such a solution exists, it is provably achieved by Leximin, but this happens in only 1% of Spliddit instances and none of the synthetic instances. For the EQ1 relaxation, we see that not only do Leximin and ALG-EQ1+PO satisfy both EQ1 and PO, but so does MNW on over 94% of Spliddit instances (and over 88% of synthetic instances). However, this trend changes when we consider EQx. ALG-EQ1+PO, despite being guaranteed to satisfy EQ1, only satisfies EQx on 62% of Spliddit instances (and 52% of synthetic instances). A similar drop off is observed with MNW. Thus, for the purpose of achieving (approximately) equitable and Pareto optimal allocations, Leximin is a clear winner.

We observe little change when, in addition to approximate equitability and Pareto optimality, we also require approximate envy-freeness. Indeed, in most cases, an allocation that is EQ1+PO/EQx+PO is also EF1/EFx. It is interesting to note that while MNW—which is appealing from the perspective of achieving relaxed envy-freeness—quite often fails to satisfy EQx, Leximin provably satisfies relaxed equitability while also achieving EFx on a large fraction of instances.

5 Discussion

Our work reveals some intriguing similarities and differences between equitability and envy-freeness. In many places, our work parallels the existing literature on envy-freeness: We present Leximin as a canonical algorithm for EQ1+PO, just like MNW achieves EF1+PO. Also, our pseudopolynomial-time algorithm for EQ1+PO uses similar techniques to that of [2018a] for EF1+PO. However, in other places, the differences are more pronounced. Most notably, EQx comes with a universal existence guarantee (often in conjunction with PO), while the existence of EFx allocations remains an open problem. Finally, exact equitability is a knife-edge property often hard to achieve in practice, unlike envy-freeness which is often satisfiable [Dickerson *et al.*, 2014].

Going forward, it would be very interesting to extend our results to the public decisions model of [2017]. Extensions to models with additional feasibility constraints on the allocations [Bouveret *et al.*, 2017], or settings with both goods and chores [Aziz *et al.*, 2018] will also be interesting.

Acknowledgments

We are grateful to the anonymous IJCAI-19 reviewers for their helpful comments, and to Ariel Procaccia and Nisarg Shah for sharing with us the data from Spliddit. LX acknowledges NSF #1453542 and #1716333 for support.

References

[Alon, 1987] Noga Alon. Splitting Necklaces. *Advances in Mathematics*, 63(3):247–253, 1987.

[Aumann and Dombb, 2015] Yonatan Aumann and Yair Dombb. The Efficiency of Fair Division with Connected Pieces. *ACM Transactions on Economics and Computation*, 3(4):23, 2015.

[Aziz *et al.*, 2018] Haris Aziz, Ioannis Caragiannis, and Ayumi Igarashi. Fair Allocation of Combinations

- of Indivisible Goods and Chores. *arXiv preprint arXiv:1807.10684*, 2018.
- [Barman *et al.*, 2018a] Siddharth Barman, Sanath Kumar Krishnamurthy, and Rohit Vaish. Finding Fair and Efficient Allocations. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 557–574, 2018.
- [Barman *et al.*, 2018b] Siddharth Barman, Sanath Kumar Krishnamurthy, and Rohit Vaish. Greedy Algorithms for Maximizing Nash Social Welfare. In *Proceedings of the 2018 International Conference on Autonomous Agents and Multiagent Systems*, pages 7–13, 2018.
- [Bezáková and Dani, 2005] Ivona Bezáková and Varsha Dani. Allocating Indivisible Goods. *ACM SIGecom Exchanges*, 5(3):11–18, 2005.
- [Bouveret and Lang, 2008] Sylvain Bouveret and Jérôme Lang. Efficiency and Envy-Freeness in Fair Division of Indivisible Goods: Logical Representation and Complexity. *Journal of Artificial Intelligence Research*, 32:525–564, 2008.
- [Bouveret *et al.*, 2017] Sylvain Bouveret, Katarína Cechlárová, Edith Elkind, Ayumi Igarashi, and Dominik Peters. Fair Division of a Graph. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, pages 135–141. AAAI Press, 2017.
- [Brams *et al.*, 2006] Steven J Brams, Michael A Jones, and Christian Klamler. Better Ways to Cut a Cake. *Notices of the AMS*, 53(11):1314–1321, 2006.
- [Brams *et al.*, 2012] Steven J Brams, Michal Feldman, John K Lai, Jamie Morgenstern, and Ariel D Procaccia. On Maxsum Fair Cake Divisions. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence*, pages 1285–1291, 2012.
- [Brams *et al.*, 2013] Steven J Brams, Michael A Jones, and Christian Klamler. N-Person Cake-Cutting: There May be No Perfect Division. *The American Mathematical Monthly*, 120(1):35–47, 2013.
- [Budish, 2011] Eric Budish. The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes. *Journal of Political Economy*, 119(6):1061–1103, 2011.
- [Caragiannis *et al.*, 2016] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D Procaccia, Nisarg Shah, and Junxing Wang. The Unreasonable Fairness of Maximum Nash Welfare. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 305–322, 2016.
- [Cechlárová and Pillárová, 2012a] Katarína Cechlárová and Eva Pillárová. On the Computability of Equitable Divisions. *Discrete Optimization*, 9(4):249–257, 2012.
- [Cechlárová and Pillárová, 2012b] Katarína Cechlárová and Eva Pillárová. A Near Equitable 2-Person Cake Cutting Algorithm. *Optimization*, 61(11):1321–1330, 2012.
- [Cechlárová *et al.*, 2013] Katarína Cechlárová, Jozef Doboš, and Eva Pillárová. On the Existence of Equitable Cake Divisions. *Information Sciences*, 228:239–245, 2013.
- [Chèze, 2017] Guillaume Chèze. Existence of a Simple and Equitable Fair Division: A Short Proof. *Mathematical Social Sciences*, 87:92–93, 2017.
- [Conitzer *et al.*, 2017] Vincent Conitzer, Rupert Freeman, and Nisarg Shah. Fair Public Decision Making. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 629–646, 2017.
- [Darmann and Schauer, 2015] Andreas Darmann and Joachim Schauer. Maximizing Nash Product Social Welfare in Allocating Indivisible Goods. *European Journal of Operational Research*, 247(2):548–559, 2015.
- [Dickerson *et al.*, 2014] John P Dickerson, Jonathan Goldman, Jeremy Karp, Ariel D Procaccia, and Tuomas Sandholm. The Computational Rise and Fall of Fairness. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence*, pages 1405–1411, 2014.
- [Dubins and Spanier, 1961] Lester E Dubins and Edwin H Spanier. How to Cut a Cake Fairly. *The American Mathematical Monthly*, 68(1P1):1–17, 1961.
- [Foley, 1967] Duncan Foley. Resource Allocation and the Public Sector. *Yale Economic Essays*, pages 45–98, 1967.
- [Freeman *et al.*, 2019] Rupert Freeman, Sujoy Sikdar, Rohit Vaish, and Lirong Xia. Equitable Allocations of Indivisible Goods. *arXiv preprint arXiv:1905.10656*, 2019.
- [Goldman and Procaccia, 2014] Jonathan Goldman and Ariel D Procaccia. Splidit: Unleashing Fair Division Algorithms. *ACM SIGecom Exchanges*, 13(2):41–46, 2014.
- [Gourvès *et al.*, 2014] Laurent Gourvès, Jérôme Monnot, and Lydia Tilane. Near Fairness in Matroids. In *Proceedings of the Twenty-First European Conference on Artificial Intelligence*, pages 393–398, 2014.
- [Herreiner and Puppe, 2009] Dorothea K Herreiner and Clemens D Puppe. Envy Freeness in Experimental Fair Division Problems. *Theory and Decision*, 67(1):65–100, 2009.
- [Lipton *et al.*, 2004] Richard J Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. On Approximately Fair Allocations of Indivisible Goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce*, pages 125–131, 2004.
- [Plaut and Roughgarden, 2018] Benjamin Plaut and Tim Roughgarden. Almost Envy-Freeness with General Valuations. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2584–2603, 2018.
- [Procaccia and Wang, 2017] Ariel D Procaccia and Junxing Wang. A Lower Bound for Equitable Cake Cutting. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 479–495, 2017.
- [Suksompong, 2019] Warut Suksompong. Fairly Allocating Contiguous Blocks of Indivisible Items. *Discrete Applied Mathematics*, 2019.