

Group Fairness for the Allocation of Indivisible Goods

Vincent Conitzer

Duke University

conitzer@cs.duke.edu

Nisarg Shah

University of Toronto

nisarg@cs.toronto.edu

Rupert Freeman

Microsoft Research

rupert.freeman@microsoft.com

Jennifer Wortman Vaughan

Microsoft Research

jenn@microsoft.com

Abstract

We consider the problem of fairly dividing a collection of indivisible goods among a set of players. Much of the existing literature on fair division focuses on notions of individual fairness. For instance, envy-freeness requires that no player prefer the set of goods allocated to another player to her own allocation. We observe that an algorithm satisfying such individual fairness notions can still treat groups of players unfairly, with one group desiring the goods allocated to another. Our main contribution is a notion of group fairness, which implies most existing notions of individual fairness. Group fairness (like individual fairness) cannot be satisfied exactly with indivisible goods. Thus, we introduce two “up to one good” style relaxations. We show that, somewhat surprisingly, certain local optima of the Nash welfare function satisfy both relaxations and can be computed in pseudo-polynomial time by local search. Our experiments reveal faster computation and stronger fairness guarantees in practice.

1 Introduction

Algorithms have come to play an increasingly prominent role in our everyday lives, augmenting, or even replacing, traditional human decision making. As our dependence on algorithms in high-stakes domains has increased, the spotlight has been placed on the potential for algorithms to exacerbate inequalities, highlighting the need to design algorithms with fairness in mind to ensure that some segments of the population are not treated differently than others.

While fairness is a relatively new design criterion in many areas of algorithmic decision making, it has a long history of study in the literature on resource allocation, in which a set of goods or resources must be divided among players with competing needs. In the context of resource allocation—often referred to as *fair division*—fairness is usually considered at an individual level. For instance, the classic definition of envy-freeness (Foley 1967) requires that no player should prefer another’s allocation to her own. When goods are indivisible, envy-freeness cannot always be guaranteed; consider a single good that must be given to one of two players. Therefore, it is often relaxed to envy-freeness “up to one good,” which allows for a player to envy another as long as

this envy can be eliminated by removing a single good from the envied player’s bundle.

In this paper, we ask whether such individual-level guarantees can be strengthened to ensure that algorithmically generated allocations are fair with respect to arbitrary segments of the population. We consider a setting in which a set of indivisible goods must be divided among players with heterogeneous, additive preferences. As an example, consider a manager in a corporate setting who needs to allocate resources (interns, conference rooms, time slots for shared machines, equipment, etc.) to employees. She may want to simultaneously ensure that no business team envies another team, that the women do not envy the men, that people in one location do not envy those in another, that people in one role do not envy those in another, and so on. Envy-freeness alone is not enough in this setting, in the sense that an allocation that is envy-free up to one good may still yield significant levels of inequality and envy between groups of players.

To address this problem, we introduce a notion of *group fairness*. Loosely speaking, an allocation is said to be group fair if no group of players would prefer to receive and redistribute among themselves the goods allocated to any other group in place of the goods that they were originally allocated, modulo some scaling to account for possibly different group sizes. Group fairness is a stronger property than envy-freeness, so it is not satisfiable in general. We therefore relax group fairness by requiring that any unfairness can be eliminated by removing a single good *per player* from the envied allocation. We obtain two distinct relaxations by distinguishing between removing one good per player in the more favored group *before* redistributing goods, and removing one good per player in the less favored group *after* goods have been redistributed among them.

Our main theoretical result is algorithmic. We first show that certain local optima of the Nash welfare function (the product of players’ utilities) satisfy both relaxations of group fairness. In particular, we show this for *locally Nash-optimal* allocations, in which transferring a single good from one player to another does not increase the product of those players’ utilities. Thus although local Nash optimality only imposes a requirement on pairs of individuals, it is strong enough to guarantee fairness to groups of arbitrary size.

We then show that a locally Nash-optimal allocation, and therefore an allocation satisfying both group fairness relax-

ations, can be computed in pseudo-polynomial time via a local search algorithm. In contrast, we show that the problem of *checking* whether an arbitrary given allocation satisfies either relaxation is coNP-hard. In experiments on real and synthetic data, we show that the local search algorithm converges quickly, and is likely to output an efficient allocation.

Related Work. Several definitions of fairness at a group level have been considered in the resource allocation literature. Most closely related to ours is the work of Berliant, Thomson, and Dunz (1992) (and the later work of Huseinovic (2011)), who defined group envy-freeness, an extension of envy-freeness to pairs of equal-sized groups, in a setting with a single divisible good. Our notion of group fairness extends group envy-freeness to cover groups of different sizes, but since we consider indivisible goods, our results are technically not comparable to theirs even if we restrict attention to groups of equal size. Aleksandrov and Walsh (2018) defined an alternative notion of group envy-freeness between groups of possibly unequal sizes. They extended individual preferences to group preferences by taking the arithmetic mean of utilities of group members, which requires interpersonal comparison of utilities. We avoid such comparisons by working only with individual utilities. Todo et al. (2011) also extended envy-freeness to groups, but considered mechanisms with monetary transfers.

Barman et al. (2018) defined a notion of groupwise maximin share, which strengthens the maximin share guarantee (Budish 2011). Their definition is of a different flavor than ours. In particular, they provide an individual-level fairness guarantee relative to all subgroups of players, while we provide guarantees between groups. Finally, several papers considered the problem of fair division among specific groups that are fixed in advance (Segal-Halevi and Suksompong 2018; Segal-Halevi and Nitzan 2015; Manurangsi and Suksompong 2017; Suksompong 2018) considering different notions of fairness, again of a different flavor than ours.

This line of work serves as a complement and point of contrast to research on algorithmic fairness in other areas of AI, including the burgeoning subfield of fairness in machine learning. Our definition of group fairness is *stronger* than individual fairness, as it requires fair treatment for groups of all sizes, individuals included. In contrast, in machine learning, individual fairness is too strong as decisions that are unfavorable to some individuals are typically unavoidable. Thus, it is common to introduce group-level parity metrics as *weaker* notions than individual fairness. These metrics ensure that outcomes are fair on average across pre-specified demographic categories, often defined by protected attributes (Kleinberg, Mullainathan, and Raghavan 2017; Calders and Verwer 2010; Hardt, Price, and Srebro 2016; Kamiran and Calders 2012; Agarwal et al. 2018).

Kearns et al. (2018) argue that algorithms that are fair only to pre-specified groups are at risk of “fairness gerrymandering,” in which the algorithm appears to be fair to each group while violating fairness constraints on particular structured subgroups. They and others (Hebert-Johnson et al. 2018; Zhang and Neill 2016) have proposed algorithms that pro-

vide or audit for fairness with respect to exponentially many groups. In a similar spirit, our group fairness definition requires fairness with respect to *all possible groups* at once.

2 Preliminaries

Throughout the paper, we use the notation $[K]$ to denote the set $\{1, \dots, K\}$. For vectors x and y of length K , we say that x Pareto dominates y if $x_k \geq y_k$ for all $k \in [K]$, with at least one inequality strict, and we say that x strictly dominates y if $x_k > y_k$ for all $k \in [K]$. For set X and element t , we use $X + t$ to denote $X \cup \{t\}$ and $X - t$ to denote $X \setminus \{t\}$.

Let M be a set of m goods, and N a set of n players. Each player i has a valuation $v_i : 2^M \rightarrow \mathbb{R}_+ \cup \{0\}$ over subsets of goods. For a single good $g \in M$, we slightly abuse notation and let $v_i(g) = v_i(\{g\})$. We assume that players have additive valuations, so that $v_i(Z) = \sum_{g \in Z} v_i(g)$ for all $Z \subseteq M$ and $v_i(\emptyset) = 0$. Without loss of generality, we assume that each good is positively valued by at least one player, and each player positively values at least one good.

An allocation A is a partition of the goods in M into (possibly empty) bundles A_i for each player i . An allocation A is non-wasteful if $g \in A_i$ implies $v_i(g) > 0$ for all g .

Much of the literature on fair division is concerned with finding allocations that satisfy particular notions of fairness. One basic notion, proportionality, requires that each player receive a set of goods that she values at least $1/n$ as much as she values the entire set of goods (Steinhaus 1948).

Definition 1 (Proportionality). An allocation A is *proportional* if for all $i \in N$, $v_i(A_i) \geq (1/n)v_i(M)$.

Envy-freeness (Foley 1967), a stronger notion, says that no player should prefer another’s allocation to her own.

Definition 2 (Envy-freeness). An allocation A is *envy-free* if for all $i, j \in N$, $v_i(A_i) \geq v_i(A_j)$.

Since envy-freeness cannot always be satisfied, relaxations have been proposed. Envy-freeness up to one good allows a player i to envy a player j , but only if the removal of a single good from j ’s bundle would remove the envy (Lipton et al. 2004). Such an allocation is guaranteed to exist.

Definition 3 (Envy-freeness up to one good). An allocation A is *envy-free up to one good (EF1)* if for all $i, j \in N$ such that $A_j \neq \emptyset$, there exists a good $g \in A_j$ such that $v_i(A_i) \geq v_i(A_j - g)$.

Finally, envy-freeness up to the least valued good says that if i envies j , the removal of *any* good that i values positively from j ’s bundle should eliminate the envy (Caragiannis et al. 2016). It is an open question whether such an allocation always exists.

Definition 4 (Envy-freeness up to the least valued good). An allocation A is *envy-free up to the least valued good (EFX)* if for all $i, j \in N$ such that $A_j \neq \emptyset$, and all $g \in A_j$ with $v_i(g) > 0$, $v_i(A_i) \geq v_i(A_j - g)$.

In addition to fairness, it is desirable to produce economically efficient allocations. The standard notion of efficiency is Pareto optimality, which says that it should not be possible to improve a player’s utility without harming someone else.

Definition 5 (Pareto optimality). An allocation A is *Pareto optimal* if for all allocations A' such that $v_i(A'_i) > v_i(A_i)$ for some $i \in N$, $v_j(A'_j) < v_j(A_j)$ for some $j \in N$.

The final notion we require is local Nash optimality. An allocation is locally Nash-optimal if it is non-wasteful and transferring a single good from one player to another does not increase the product of their utilities. Note that local Nash optimality does not imply Pareto optimality.

Definition 6 (Locally Nash-optimal allocation). An allocation A is *locally Nash-optimal* if for all $i, j \in N$ and $g \in A_j$, $v_j(g) > 0$ and $v_i(A_i) \cdot v_j(A_j) \geq v_i(A_i + g) \cdot v_j(A_j - g)$.

3 Group Fair Allocations

In this section, we move beyond the standard fairness notions that operate on individuals or pairs of players and introduce a new definition of group fairness. Our definition is modeled on envy-freeness. It requires that no group of players S envy another group T , where S envies T if the players in S could redistribute the goods allocated to T among themselves in a way that yields a Pareto improvement, adjusting appropriately for any difference in the group sizes. Note that our definition does not require S and T to be disjoint.

Definition 7 (Group Fairness). An allocation A is *group fair* if for every non-empty $S, T \subseteq N$ and every partition $(B_i)_{i \in S}$ of $\cup_{j \in T} A_j$, $(|S|/|T|) \cdot (v_i(B_i))_{i \in S}$ does not Pareto dominate $(v_i(A_i))_{i \in S}$.

Group fairness is a strengthening of several properties from the fair division literature. Group envy-freeness (Berliant, Thomson, and Dunz 1992) requires the no-envy condition in the definition to hold when $|S| = |T|$, while envy-freeness requires it to hold only when $|S| = |T| = 1$. The core (Foley 1967; Fain, Goel, and Munagala 2016; Fain, Munagala, and Shah 2018) requires that it hold when $T = N$, while proportionality requires that it hold when $|S| = 1$ and $T = N$. Finally, Pareto optimality requires the condition when $S = T = N$.

When goods are divisible, it is easy to check that the globally Nash-optimal allocation, which coincides with a strong form of competitive equilibrium from equal incomes (Segal-Halevi and Sziklai 2018), satisfies group fairness. However, when goods are indivisible, it cannot be guaranteed; this is easy to see from the simple example with a single good and two competing players, one of whom necessarily gets nothing. We therefore turn to relaxed notions.

3.1 Two Relaxations of Group Fairness

Before presenting the relaxations, let us step back to consider what an “approximately group fair” allocation should look like. Consider the example in Figure 1. Here there are five players: one of type “circle” who values only circle goods (with zero value for squares), two of type “square” who value only square goods (with zero value for circles), and two of type “flex” who are more flexible and value both squares and circles equally. There are four goods: two circles and two squares. Because it is impossible to give a good to every player, there is no envy-free allocation, and therefore no group fair allocation. However, the allocation A shown in




Figure 1: An arguably unfair allocation satisfying EF1/EFX.

the figure, which gives one circle and one square to each of the flex players, satisfies EF1 and EFX. According to these criteria, we would thus call this allocation fair.

However, we argue that allocation A is not fair to all groups of players, in a way that we will soon make precise. Suppose group S consists of the circle player and one square player, and let T consist of both flex players. Collectively, players in S have demand for all of the goods that have been allocated to players in T . In fact, if T 's goods were transferred to S and distributed to the players who value them most, each player in S could be made significantly (that is, more than “up to one good”) happier than they are under allocation A . We argue that an (approximately) fair solution should split the goods more evenly between sets S and T to rectify this asymmetry, and we would like our “up to one good” relaxation of group fairness to capture this idea.

As a first attempt, one might hope to require that no set S envy another set T (modulo rescaling for size) once a single good has been removed from T 's allocation. However, it is easy to see that any relaxation that removes only a single good is still too strong to be satisfiable in general. Suppose that there are n identical players (for any even n) and $3n/2$ identical goods. Intuitively, the fairest allocation would give half of the players one good each (call these players S), and the other half two each (call these T). Even if we remove a single good from a player in T , the remaining $n - 1$ goods allocated to T can still be distributed among S in a way that yields a Pareto improvement. Indeed, the same problem arises if we remove any fewer than $n/2 = |T| = |S|$ goods. Therefore, minimal relaxations of group fairness must remove one good *per player*.

There are two natural ways to do this: remove one good from each player in T before the set of goods is handed over to S (“before”), or remove one good from each player in S after the goods have been redistributed among them (“after”). We consider both in turn.

Group Fairness up to One Good (After). We first present the version of our relaxation in which goods are removed from each player in S after redistribution occurs. To motivate our specific choice of definition, consider the example shown in Figure 2 (left) with two circle players, four square players, one circle good, and three square goods. We would argue that the allocation A that is pictured is the unique fair allocation, up to permutations of identical players; all other non-wasteful allocations involve one player receiving multiple goods while another player of the same type receives none. Thus, if we want our relaxed notion of group fairness to be satisfiable, it must be satisfied by this allocation.

Consider sets S_1 and T . These sets are witness to a vi-




Figure 2: Examples illustrating the necessary technical conditions in the definitions of GF1A (left) and GF1B (right).

olation of group fairness, because T 's goods can be reallocated among S_1 such that, even after scaling by a factor of $|S_1|/|T| = 1/3$, S_1 has an allocation that Pareto improves over the original. In fact, even if we remove a good from the player in S_1 , she would receive two goods that she values, still a Pareto improvement. Thus relaxing the group fairness definition by removing a good from each player in S_1 is not sufficient to guarantee existence. To get around this technicality, which is due to the way in which scaling occurs, we consider a slight variant of the same idea: instead of removing a single good from the bundle B_i received by player i in group S_1 and then comparing the (scaled) value of the remaining bundle to the (unscaled) value of the original allocation, we add this good to the original allocation and compare its (unscaled) value to the (scaled) value of the whole B_i .

There is one other technicality our definition must account for. Consider sets S_2 and T . When we partition T 's goods among players in S_2 as pictured, we have $(|S|/|T|) \cdot (v_i(B_i))_{i \in S_2} = (0, 4/3)$, which Pareto dominates the utilities under the original allocation to S_2 even if each player in S_2 were given a single good from group T . This problem arises from the fact that the circle player is essentially serving as a dummy player; she does not value T 's goods at all, yet still inflates the size of the set S_2 , changing the scaling factor without meaningfully changing the fairness constraint that we want to capture. We can avoid this issue by requiring that the partition B must give positive value to all players in the set S , which rules out sets with dummy players included.

We are now ready to formally define our first relaxation of group fairness.

Definition 8 (GF1A). An allocation A satisfies *GF1A* if for every non-empty $S, T \subseteq N$ and every partition $(B_i)_{i \in S}$ of $\cup_{j \in T} A_j$ such that $v_i(B_i) > 0$ for all $i \in S$, there exists a good $g_i \in B_i$ for each $i \in S$ such that $(|S|/|T|) \cdot (v_i(B_i))_{i \in S}$ does not Pareto dominate $(v_i(A_i + g_i))_{i \in S}$.

Returning to the example from Figure 1, we see that, as desired, the pictured allocation fails to satisfy GF1A, as witnessed by the set S consisting of the circle player and one square player and the set T consisting of the two flex players. To provide more intuition for what this definition does and does not allow, we point out that the set S' consisting of the two square players does not serve as a witness with the same set T . Loosely speaking, this is because the set S' collectively has no demand for circle goods, and so the allo-

cation of the circle goods to players in T does not preclude these players from also receiving the square goods.

Group Fairness up to One Good (Before). In our second relaxation of group fairness, we consider removing one good from the bundle of each player in set T before T 's goods are redistributed among players in S , requiring that the (scaled) values of the resulting bundles do not provide a Pareto improvement for S .

Once again, the most straightforward definition would be susceptible to dummy players in S inflating the scale factor without impacting the underlying fairness of the allocation, as illustrated in Figure 2 (right). Here the allocation A is the only intuitively fair and non-wasteful allocation, up to permutations of identical players, so it must satisfy our relaxation if we want the relaxation to be satisfiable in general. If we remove a single good from the one player in T and reallocate her remaining good to S as pictured, both players in S would get the same value as they would under allocation A , but since $|S|/|T| = 2$, their scaled values under partition B would Pareto dominate their values under A . Like before, we avoid this problem by considering only pairs S and T for which it is possible to partition T 's goods among S so that all players in S receive positive value.

Definition 9 (GF1B). An allocation A satisfies *GF1B* if for every non-empty $S, T \subseteq N$ for which there exists a partition $(C_i)_{i \in S}$ of $\cup_{j \in T} A_j$ with $v_i(C_i) > 0$ for all $i \in S$, there exists a good $g_j \in A_j$ for every $j \in T$ with $A_j \neq \emptyset$ such that for every partition $(B_i)_{i \in S}$ of $\cup_{j \in T} A_j \setminus \cup_{j \in T: A_j \neq \emptyset} \{g_j\}$, $(|S|/|T|) \cdot (v_i(B_i))_{i \in S}$ does not Pareto dominate $(v_i(A_i))_{i \in S}$.

Once again it is easy to verify that the allocation pictured in Figure 1 fails to satisfy GF1B, as witnessed by the same sets S and T as before. And just as it was with GF1A, the set S' consisting of the two square players does not serve as a witness with the same set T .

Notice that in the definition of group fairness, the no-envy condition is agnostic about the exact allocation A_j for each $j \in T$; only $\cup_{j \in T} A_j$ is relevant. While this is also true for the GF1A relaxation, it is not true for GF1B since we require that only a single good be removed from each player in T . An alternative, weaker definition of GF1B would be to remove $|T|$ goods in total from players in T , without the requirement that one is removed from each player. (See the

full version of the paper¹ for an example where the two definitions differ.) We present the stronger definition here, but note that all of our results hold for the weaker version also.

3.2 A Comparison of GF1A and GF1B.

To gain further intuition, we briefly discuss examples of cases in which GF1A and GF1B differ, as shown in Figure 3. The players in these examples are again of type circle, square, or flex, but the goods in Figure 3 (left) are more general. A circle (respectively, square) with a label v is valued v by players who value circles (respectively, squares), and 0 by other players. A pentagon labeled v is valued v by all.

In Figure 3 (left), S and T are witness to a violation of GF1B. After removing any good from the player in T , it is still possible to give one of the players in S a value of at least 4 and the other a value of 4.1. Since $|S|/|T| = 2$, this violates GF1B. However, allocation A does satisfy GF1A.

In Figure 3 (right), groups S and T are witness to a violation of GF1A. When T 's goods are redistributed to the players in S who value them most, both players in S end up better off even with a single good removed. However, it can be verified that A satisfies GF1B.

GF1A and GF1B both imply “up to one good” style variants of the core and group envy-freeness. They additionally both imply proportionality up to one good (Conitzer, Freeman, and Shah 2017) and envy-freeness up to one good (Budish 2011; Caragiannis et al. 2016).

In the special case in which all players have identical valuations, stronger implications hold. In this case, GF1A is stronger than GF1B, in the sense that any GF1A allocation satisfies GF1B but the converse does not hold. In fact, GF1B becomes equivalent to EF1 in this special case. To further complete the picture, the relationship between our group fairness relaxations and local Nash optimality explored in the next section allows us to show that all three properties are implied by EFX. The proof appears in the full version.

Theorem 1. *When all players have identical valuations, $EFX \Rightarrow GF1A \Rightarrow GF1B$, and $GF1B \Leftrightarrow EF1$, where \Rightarrow is strict logical implication and \Leftrightarrow is logical equivalence.*

4 Local Nash Optimality Implies GF1A/B

Our desire to relax the notion of group fairness stemmed from the fact that group fair allocations may not exist in general when goods are indivisible. In this section, we show that both GF1A and GF1B allocations are always guaranteed to exist. In particular, every locally Nash-optimal allocation is guaranteed to satisfy both GF1A and GF1B. This result is surprising given that local Nash optimality is a local property, involving only pairs of players, while GF1A and GF1B are global properties involving arbitrary player groups.

Theorem 2. *Every locally Nash-optimal allocation A satisfies GF1A and GF1B.*

The proof follows a similar structure to the proof due to Caragiannis et al. (2016) that Nash optimality implies EF1.² We observe that instead of removing a good from player j

that depends on the identity of the envying player i , the same good $g_j \in A_j$ can be removed irrespective of i . This observation is what allows us to extend the proof to groups. (It also implies some slightly stronger results for individual fairness, which we discuss in Section 7.)

Proof of Theorem 2. Here we provide the proof for GF1A. We refer the reader to the full version of the paper for the GF1B proof, which follows a similar outline.

Let A be a locally Nash-optimal allocation. Assume for contradiction that A does not satisfy GF1A, and let (S, T) be groups with smallest $|T|$ that are witness to the violation of GF1A. Note that this implies $|A_j| \geq 1$ for all $j \in T$, which in turn implies that $v_j(A_j) > 0$ by non-wastefulness; if $|A_j| = 0$ for some $j \in T$, $(S, T - j)$ would also be witness to the violation of GF1A.

Fix a partition $(B_i)_{i \in S}$ of $\cup_{j \in T} A_j$ for which the GF1A constraint is violated. For the constraint to be violated, it must be the case that $v_i(B_i) > 0$ for all $i \in S$, which implies that $B_i \neq \emptyset$ for all $i \in S$. For all $i \in S$, let $g_i^* \in \arg \max_{g \in B_i} v_i(g)$. Then, we have $v_i(g_i^*) > 0$, and hence, $v_i(A_i + g_i^*) > 0$.

With a little algebraic simplification, we can rewrite the final condition from Definition 6 as $v_i(A_i + g) \cdot v_j(g) \geq v_i(g) \cdot v_j(A_j)$. Then for all $i \in S$, $j \in T$, and $g \in B_i \cap A_j$, $v_i(g) \cdot v_j(A_j) \leq v_i(A_i + g) \cdot v_j(g) \leq v_i(A_i + g_i^*) \cdot v_j(g)$, where the second transition follows from the definition of g_i^* . Rearranging, we have

$$\frac{v_i(g)}{v_i(A_i + g_i^*)} \leq \frac{v_j(g)}{v_j(A_j)}.$$

Summing over $i \in S$, $j \in T$, and $g \in B_i \cap A_j$, we obtain

$$\sum_{i \in S} \frac{v_i(B_i)}{v_i(A_i + g_i^*)} \leq |T|.$$

Since the partition B violates the constraint, $(|S|/|T|) \cdot (v_i(B_i))_{i \in S}$ Pareto dominates $(v_i(A_i + g_i^*))_{i \in S}$, and so $v_i(B_i)/v_i(A_i + g_i^*) \geq |T|/|S|$ for each $i \in S$, with at least one inequality strict. This implies that $\sum_{i \in S} v_i(B_i)/v_i(A_i + g_i^*) > |T|$, a contradiction. \square

Since Nash-optimal allocations always exist, this immediately implies the existence of allocations that satisfy both GF1A and GF1B. In the next section, we provide an algorithm for computing such an allocation.

Corollary 3. *An allocation A satisfying both GF1A and GF1B always exists.*

5 Complexity Results

We have shown that any locally Nash-optimal allocation satisfies GF1A and GF1B. We now show that such an allocation can be computed in pseudo-polynomial time.

We consider a simple local search algorithm that works as follows. Begin with an arbitrary allocation A . At every step, check for a violation of local Nash optimality: that is,

Nash-optimal allocations, an identical proof holds for locally Nash-optimal allocations too.

¹The full version can be found on the authors' websites.

²While Caragiannis et al. (2016) state their result for globally




Figure 3: Allocations that satisfy GF1A but not GF1B (left) and GF1B but not GF1A (right).

find a pair of players $i, j \in N$ and a good $g \in A_j$ such that either $v_j(g) = 0$ and $v_i(g) > 0$, or transferring the good from A_j to A_i increases the product of utilities of i and j (i.e., $v_i(A_i + g) \cdot v_j(A_j - g) > v_i(A_i) \cdot v_j(A_j)$). If such a violation exists, transfer the good. Otherwise, terminate. We show that this algorithm terminates at a locally Nash-optimal allocation in a pseudo-polynomial number of steps.

Theorem 4. *A locally Nash-optimal allocation can be computed in pseudo-polynomial time.*

The proof appears in the full version. It proceeds by upper bounding the maximum possible Nash welfare, and lower bounding the multiplicative Nash welfare increase in a single step of the algorithm.

Corollary 5. *An allocation satisfying both GF1A and GF1B can be computed in pseudo-polynomial time.*

Whether an allocation satisfying GF1A or GF1B can be computed in polynomial time remains an interesting open question. We are able to show that the problem of *verifying* whether a given allocation satisfies GF1A or GF1B is strongly coNP-hard. The proofs are in the full version.

Theorem 6. *It is strongly coNP-hard to determine whether an allocation A satisfies GF1A.*

Theorem 7. *It is strongly coNP-hard to determine whether an allocation A satisfies GF1B.*

6 Simulations

In this section, we investigate the performance of the local search algorithm in practice, in terms of both its running time and the quality of the allocation it returns. Specifically, we measure the number of steps it takes to converge, how frequently it returns a Pareto optimal allocation, and how frequently it returns a *globally* Nash optimal, also known as *max Nash welfare* allocation (Caragiannis et al. 2016); the last number is guaranteed to be weakly lower than the former since all max Nash welfare allocations are Pareto optimal.

We first experiment with a dataset of fair division instances obtained from Spliddit.org, a not-for-profit website that allows its users to employ fair division algorithms for every-day problems, including allocation of (possibly indivisible) goods. The dataset contains 2754 division instances in which all goods are indivisible. These instances contain as many as 15 players (2.6 on average) and 93 goods (5.7 on average). The algorithm currently deployed on Spliddit




Figure 4: Running time (number of steps) and solution quality of local search varying K with $n = 5$ and $m = 15$.

computes a max Nash welfare (and thus also locally Nash-optimal) allocation (Caragiannis et al. 2016).

On this dataset, local search takes only 6.0 steps on average, and the maximum on any instance is 91. In over 88% of the instances, the algorithm returns a Pareto optimal allocation, while in over 68% of the instances, it returns a max Nash welfare allocation.

Since typical Spliddit instances are relatively small, we next explore the algorithm’s performance on larger simulated instances. We vary the number of players (n) from 3 to 10. For each n in this range, we vary the number of goods (m) from n to $5n$ in increments of n . To explore the effect of the magnitude of player valuations on the running time, we additionally vary a parameter K controlling this magnitude from 100 to 1000 in increments of 100. For each combination of n , m , and K , we generate 1000 instances in which the valuation v_i of each player i is sampled i.i.d. from the uniform distribution over all integral valuations that sum to K (i.e., uniformly at random subject to $v_i(M) = K$).

In Figure 4, we examine the effect of varying K . We note that it does not significantly affect the average number of steps until convergence (left figure), or the percentage of instances in which the algorithm finds a Pareto optimal or max Nash welfare allocation (right figure). For the remainder of this section, we report our findings for $K = 500$.

Figures 5a and 5b respectively show that the average number of steps until convergence appears to increase linearly with m (fixing $n = 5$) and increase linearly with n (fixing $m = 3n$). For the largest synthetic instances that we examined ($n = 10, m = 50$ and $K = 1000$), local search terminated in 220 steps on average. Instances of this size are close to the maximum size that the max Nash welfare algorithm can reliably handle, while they remain trivial for the

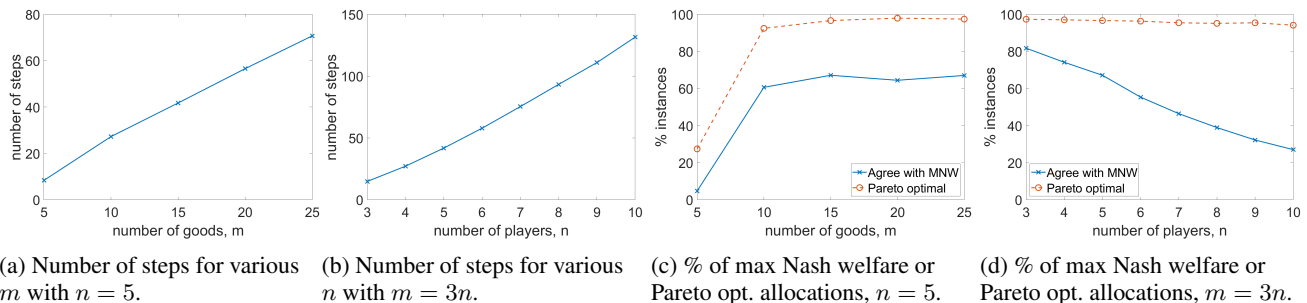


Figure 5: Running time (number of steps) and solution quality of the local search algorithm on synthetic data.

local search algorithm.

Figures 5c and 5d show the percentage of instances in which the local search algorithm produces a Pareto optimal or max Nash welfare allocation, as a function of m (with $n = 5$) and n (with $m = 3n$), respectively. In Figure 5c, notice that when $m = n = 5$, only a very small percentage of allocations returned by local search are max Nash welfare allocations, or even Pareto optimal. This is because in almost all cases when $m = n$, any allocation in which every player receives a single good is locally Nash optimal, and local search might terminate at an arbitrary allocation of this form. However, with $m = 2n = 10$, local search returns a max Nash welfare allocation in nearly 60% of the instances, and almost always achieves Pareto optimality. Increasing m further only slightly improves performance.

Examining Figure 5d reveals a different story for the performance as a function of n . With $n = 3$, local search usually finds a Pareto optimal allocation and finds a max Nash welfare allocation in nearly 80% of the instances. As n increases, the allocation remains likely to be Pareto optimal, but quickly becomes unlikely to be globally Nash optimal.

We remark that for every combination of n , $m \geq 2n$, and K in our simulations, local search returns a Pareto optimal allocation in at least 85% of the instances.

7 Discussion

Our work opens up several avenues for future research on fair allocation and takes steps towards addressing existing open questions that go beyond group fairness.

Fairness with respect to fixed groups of players. We consider fairness guarantees that hold simultaneously for every pair of groups. In some applications, we may care only about fixed partitions of players into groups, for example, based on gender or race. An interesting open question is whether it is possible to provide stronger guarantees if we ask for fairness only with respect to fixed groups (potentially in conjunction with an individual fairness notion such as EF1). For instance, is it possible to provide “up to one good” guarantees with the removal of a single good *overall*, as opposed to a single good per player?

Strong envy-freeness up to one good. In the definition of GF1B, sets S and T are chosen before the selection of

the good g_j for each player $j \in T$. However, the proof of Theorem 2 establishes that locally Nash optimal allocations satisfy a slightly stronger version of the definition in which a single good g_j for each player j is chosen in advance (independent of sets S and T). This property is formally defined as strong GF1B in the full version. When restricting to sets S and T with $|S| = |T| = 1$, the original GF1B definition yields envy-freeness up to one good, while strong GF1B yields a slightly stronger property.

Definition 10 (Strong envy-freeness up to one good (s-EF1)). An allocation A is s-EF1 if for each $j \in N$ such that $A_j \neq \emptyset$ there exists a good $g_j \in A_j$ such that for all $i \in N$, $v_i(A_i) \geq v_i(A_j - g_j)$.

It follows that every locally Nash optimal allocation satisfies s-EF1. It is easy to check that the allocations produced by the round robin algorithm and the algorithm of Barman, Krishnamurthy, and Vaish (2018), which are both known to satisfy EF1, also satisfy the stronger s-EF1.

Locally Nash-optimal allocations and approximate market equilibria. When goods are divisible, it is known that globally Nash optimal allocations coincide with strong competitive equilibria with equal incomes (Segal-Halevi and Sziklai 2018), where (informally) each good is assigned a price, each player is given one unit of fake money (equal incomes), and each player purchasing her highest valued bundle of goods that she can afford perfectly partitions the set of goods (competitive equilibrium).

With indivisible goods, such an allocation may not exist. A recent line of work (Budish 2011; Babaioff, Nisan, and Talgam-Cohen 2017; Barman, Krishnamurthy, and Vaish 2018) proposes relaxations in which the competitive equilibrium condition is retained, but the equal incomes condition is relaxed to *almost* equal incomes. The relaxation due to Barman, Krishnamurthy, and Vaish (2018) is guaranteed to be satisfiable, and leads to allocations that are envy-free up to one good and Pareto optimal. However, Caragiannis et al. (2016) posed the open question of whether such relaxations retain any connection to the Nash welfare function.

In the full version of the paper, we explore a relaxation that is very different from the relaxation due to Barman, Krishnamurthy, and Vaish (2018). We retain exactly equal incomes, and instead relax the competitive equilibrium condition: each player now purchases an *almost* optimal bundle

of goods that she can afford. Our relaxation loses Pareto optimality while theirs guarantees it. However, our relaxation satisfies both GF1A and GF1B, while theirs can be shown to satisfy only GF1B and violate GF1A. Additionally, we recover an equivalence between approximate market equilibria and local Nash optimality, partially answering the open question by Caragiannis et al. (2016).

Acknowledgments

We thank David Pennock for helpful discussions. Part of this work was completed while Freeman was supported by a Facebook Ph.D. Fellowship. This work was also supported by NSF under awards IIS-1814056 and IIS-1527434 and by NSERC under the Discovery Grants program.

References

- Agarwal, A.; Beygelzimer, A.; Dudík, M.; Langford, J.; and Wallach, H. 2018. A reductions approach to fair classification. In *Proc. of the 35th International Conference on Machine Learning (ICML)*.
- Aleksandrov, M., and Walsh, T. 2018. Group envy freeness and group pareto efficiency in fair division with indivisible items. In *Proc. of the 41st German Conference on AI (KI)*.
- Babaioff, M.; Nisan, N.; and Talgam-Cohen, I. 2017. Competitive equilibria with indivisible goods and generic budgets. CoRR arXiv:1703.08150.
- Barman, S.; Biswas, A.; Krishnamurthy, S.; and Narahari, Y. 2018. Groupwise maximin fair allocation of indivisible goods. In *Proc. of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*.
- Barman, S.; Krishnamurthy, S. K.; and Vaish, R. 2018. Finding fair and efficient allocations. In *Proc. of the 19th ACM Conference on Economics and Computation (EC)*.
- Berliant, M.; Thomson, W.; and Dunz, K. 1992. On the fair division of a heterogeneous commodity. *Journal of Mathematical Economics* 21(3):201–216.
- Budish, E. 2011. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy* 119(6):1061–1103.
- Calders, T., and Verwer, S. 2010. Three naive Bayes approaches for discrimination-free classification. *Data Mining and Knowledge Discovery* 21(2):277–292.
- Caragiannis, I.; Kurokawa, D.; Moulin, H.; Procaccia, A. D.; Shah, N.; and Wang, J. 2016. The unreasonable fairness of maximum Nash welfare. In *Proc. of the 17th ACM Conference on Economics and Computation (EC)*.
- Conitzer, V.; Freeman, R.; and Shah, N. 2017. Fair public decision making. In *Proc. of the 18th ACM Conference on Economics and Computation (EC)*.
- Fain, B.; Goel, A.; and Munagala, K. 2016. The core of the participatory budgeting problem. In *Proc. of the 12th Conference on Web and Internet Economics (WINE)*.
- Fain, B.; Munagala, K.; and Shah, N. 2018. Fair allocation of indivisible public goods. In *Proc. of the 19th ACM Conference on Economics and Computation (EC)*.
- Foley, D. 1967. Resource allocation and the public sector. *Yale Economics Essays* 7:45–98.
- Hardt, M.; Price, E.; and Srebro, N. 2016. Equality of opportunity in supervised learning. In *Advances in Neural Information Processing Systems* 29.
- Hebert-Johnson, U.; Kim, M.; Reingold, O.; and Rothblum, G. 2018. Multicalibration: Calibration for the (computationally-identifiable) masses. In *Proc. of the 35th International Conference on Machine Learning (ICML)*.
- Husseinov, F. 2011. A theory of a heterogeneous divisible commodity exchange economy. *Journal of Mathematical Economics* 47(1):54–59.
- Kamiran, F., and Calders, T. 2012. Data preprocessing techniques for classification without discrimination. *Knowledge and Information Systems* 33(1):1–33.
- Kearns, M.; Neel, S.; Roth, A.; and Wu, Z. S. 2018. Preventing fairness gerrymandering: Auditing and learning for subgroup fairness. In *Proc. of the 35th International Conference on Machine Learning (ICML)*.
- Kleinberg, J.; Mullainathan, S.; and Raghavan, M. 2017. Inherent trade-offs in the fair determination of risk scores. In *Proc. of the 8th Innovations in Theoretical Computer Science Conference (ITCS)*.
- Lipton, R. J.; Markakis, E.; Mossel, E.; and Saberi, A. 2004. On approximately fair allocations of indivisible goods. In *Proc. of the 6th ACM Conf. on Electronic Commerce (EC)*.
- Manurangsi, P., and Suksompong, W. 2017. Asymptotic existence of fair divisions for groups. *Mathematical Social Sciences* 89:100–108.
- Segal-Halevi, E., and Nitzan, S. 2015. Fair cake-cutting among families. CoRR arXiv:1510.03903.
- Segal-Halevi, E., and Suksompong, W. 2018. Democratic fair allocation of indivisible goods. In *Proc. of the 27th International Joint Conf. on Artificial Intelligence (IJCAI)*.
- Segal-Halevi, E., and Sziklai, B. 2018. Monotonicity and competitive equilibrium in cake-cutting. *Economic Theory* 1–39.
- Steinhaus, H. 1948. The problem of fair division. *Econometrica* 16:101–104.
- Suksompong, W. 2018. Approximate maximin shares for groups of agents. *Mathematical Social Sciences* 92:40–47.
- Todo, T.; Li, R.; Hu, X.; Mouri, T.; Iwasaki, A.; and Yokoo, M. 2011. Generalizing envy-freeness toward group of agents. In *Proc. of the 22nd International Joint Conf. on Artificial Intelligence (IJCAI)*.
- Zhang, Z., and Neill, D. B. 2016. Identifying significant predictive bias in classifiers. CoRR arXiv:1611.08292.