

Almost Envy-Freeness for Groups: Improved Bounds via Discrepancy Theory

Pasin Manurangsi¹ and Warut Suksompong²

¹Google Research, USA

²School of Computing, National University of Singapore, Singapore

pasin@google.com, warut@comp.nus.edu.sg

Abstract

We study the allocation of indivisible goods among groups of agents using well-known fairness notions such as envy-freeness and proportionality. While these notions cannot always be satisfied, we provide several bounds on the optimal relaxations that can be guaranteed. For instance, our bounds imply that when the number of groups is constant and the n agents are divided into groups arbitrarily, there exists an allocation that is envy-free up to $\Theta(\sqrt{n})$ goods, and this bound is tight. Moreover, we show that while such an allocation can be found efficiently, it is NP-hard to compute an allocation that is envy-free up to $o(\sqrt{n})$ goods even when a fully envy-free allocation exists. Our proofs make extensive use of tools from discrepancy theory.

1 Introduction

Resource allocation problems arise in numerous facets of modern society, from allotting supplies to neighborhoods in a city to distributing personnel among governmental organizations. A principal consideration when allocating resources is *fairness*: the society is better off when all parties involved feel that they receive a fair share of the resource. It therefore comes as no surprise that the study of how to allocate resources fairly—commonly referred to as *fair division*—has received substantial attention in economics and, as societies become more interconnected and applications grow in scale, in computer science [Brams and Taylor, 1996; Moulin, 2003; Thomson, 2016; Moulin, 2019]

The vast majority of the fair division literature assumes that each involved party consists of a single agent. Yet, in many resource allocation scenarios, especially large-scale ones, resources are allocated to *groups* of agents—even though the agents in each group share the same set of goods, they may have varying preferences over different goods in the set. Indeed, some citizens of a neighborhood may benefit from new books allotted to the public library, while others would rather have additional fitness equipment in their local park. Similarly, members of an organization may have diverse opinions about the new personnel that they would like to have in their organization. These scenarios cannot be captured by the tra-

ditional fair division setting, in which each recipient of a bundle of goods is represented by a single preference.

The group aspect of fair division has been addressed in a number of recent papers [Manurangsi and Suksompong, 2017; Ghodsi *et al.*, 2018; Suksompong, 2018; Segal-Halevi and Nitzan, 2019; Segal-Halevi and Suksompong, 2019; Kyropoulou *et al.*, 2020]. Most of these papers studied the important fairness notion of *envy-freeness*: an agent is said to be *envy-free* if she values the goods allocated to her group at least as much as those allocated to any other group. When goods are discrete—books, personnel, fitness equipment, and many other common supplies fall into this category—envy-freeness cannot always be satisfied even when allocating the goods among individual agents; indeed, this can be easily seen when there is a single valuable good and at least two agents. This observation has motivated relaxing the envy-freeness criterion to *envy-freeness up to c goods (EF c)*, which means that any agent’s envy toward another group can be eliminated by removing at most c goods from that group’s bundle, where $c \geq 1$ is an integer parameter.

When allocating goods among individual agents, an EF1 allocation can be found regardless of the number of agents [Lipton *et al.*, 2004]. However, the picture for group allocation is much less clear, even in the simplest case of two groups. Segal-Halevi and Suksompong [2019] showed that if the two groups contain n agents in total and the agents have additive valuations, then an EF n allocation is guaranteed to exist. Their result follows by applying a classic theorem on *consensus halving*, i.e., a partition of a set of *divisible* goods into two parts such that every agent values both parts equally. Since there is always a consensus halving in which at most n goods are divided [Alon, 1987; Simmons and Su, 2003], rounding such a consensus halving yields an EF n allocation.¹ On the other hand, Kyropoulou *et al.* [2020, Prop. 3.5] gave a simple example showing that it is impossible to ensure EF c for $c \in o(\log n)$, thereby leaving an exponential gap in this fundamental question. Can we always achieve an impressive fairness guarantee of mere logarithmic envy, or does the envy scale linearly with the number of agents in the worst case?

¹In fact, Segal-Halevi and Suksompong [2019] gave a slightly better guarantee of EF $(n - 1)$. This guarantee was obtained by finding a consensus halving for $n - 1$ of the agents, and letting the remaining agent choose the part that she prefers.

1.1 Our Results

In this paper, we give a precise answer to the above question, and much more. We consider a general setting with $n = n_1 + \dots + n_k$ agents distributed into $k \geq 2$ groups consisting of $n_1, \dots, n_k \geq 1$ agents, respectively. As is common in fair division, we assume that the agents have additive utilities over the goods. Besides EFC, we investigate relaxations of two other important fairness notions: *proportionality*—every agent believes that the share allocated to her group is worth at least $1/k$ of the entire set of goods—and *consensus $1/k$ -division*—each agent finds all k bundles to be of equal value.² The precise definitions can be found in **Section 2.1**.

For each fairness notion and each n_1, \dots, n_k , we are interested in the smallest positive value c such that an allocation satisfying that notion up to c goods always exists for agents with arbitrary additive utilities. For envy-freeness and proportionality, we denote this value of c by $c^{\text{EF}}(n_1, \dots, n_k)$ and $c^{\text{PROP}}(n_1, \dots, n_k)$, respectively. On the other hand, for consensus $1/k$ -division, the partition of agents into groups is inconsequential, so we use the notation $c_k^{\text{CD}}(n)$. Our main results provide bounds on these values:

Theorem 1.1. For any $k, n_1, \dots, n_k \in \mathbb{N}$,

$$\begin{aligned} O(\sqrt{n}) &\geq c^{\text{EF}}(n_1, \dots, n_k) \\ &\geq \Omega(\sqrt{\max\{n_1, \dots, n_k\}/k^3}). \end{aligned}$$

Theorem 1.2. For any $k, n_1, \dots, n_k \in \mathbb{N}$,

$$\begin{aligned} O(\sqrt{n}) &\geq c^{\text{PROP}}(n_1, \dots, n_k) \\ &\geq \Omega(\sqrt{\max\{n_1, \dots, n_k\}/k^3}). \end{aligned}$$

Theorem 1.3. For any $n, k \in \mathbb{N}$,

$$O(\sqrt{n}) \geq c_k^{\text{CD}}(n) \geq \Omega(\sqrt{n/k}).$$

Note that since $\max\{n_1, \dots, n_k\} \geq n/k$, all three bounds are asymptotically tight when k is constant. In particular, Theorem 1.1 answers the question that we posed earlier: taking $k = 2$, we find that the optimal envy-freeness guarantee for two groups is EFC where $c \in \Theta(\sqrt{n})$. This significantly improves upon the lower bound of $\Omega(\log n)$ and upper bound of $O(n)$ from prior work, and implies that a decent, though not outstanding, fairness guarantee can be obtained. We establish Theorem 1.3 along with the upper bounds of Theorems 1.1 and 1.2 in **Section 3**, and the lower bounds of Theorems 1.1 and 1.2 in **Section 4**.

Our main tools and techniques throughout this work come from *discrepancy theory*, an area of mathematics that studies how much deviation from the desired state is necessary in various settings—we provide the relevant background in **Section 2.2**. The tools that we use imply that for each of the fairness notions, an allocation satisfying the corresponding upper bound in Theorems 1.1–1.3 can be found efficiently. In light of this, a natural question is whether we can compute an allocation improving upon these bounds *if* such allocations are known to exist in a given instance. In **Section 5**, we provide a strong negative answer to this question: for example,

²When $k = 2$, consensus $1/k$ -division is better known as *consensus halving* [Simmons and Su, 2003].

we show that even if a fully envy-free allocation is known to exist for a certain instance, it is still NP-hard to find an allocation that is envy-free up to $o(\sqrt{n})$ goods for that instance.

1.2 Further Related Work

While fair division has a long and storied history, several fairness notions for the indivisible goods setting, including envy-freeness relaxations, have only been proposed and studied in the past few years [Bouveret *et al.*, 2016; Markakis, 2017]. In the group setting, Kyropoulou *et al.* [2020] showed that EF1 can be guaranteed for all agents only when the groups are small—for instance, with two groups, an EF1 allocation does not always exist when both groups have size at least three. Segal-Halevi and Suksompong [2019] investigated *democratic fairness*, where the goal is to satisfy a certain fraction of the agents in each group. They showed that for two groups with any number of agents, there exists an allocation that is EF1 for at least half of the agents in each group—this ratio is tight in the worst case, and continues to be tight even if we relax EF1 to EFC for any constant c .

Besides the model that we consider, a number of papers have studied related models and notions. Ghodsi *et al.* [2018] addressed *rent division* among groups, where in addition to deciding the allocation of the rooms, the agents must determine how to split the rent of their apartment. Benabbou *et al.* [2019] examined a group setting where the goods allocated to each group are further divided among the members of the group, so in contrast to our setting, each agent does not derive full utility from the bundle of her group. Several authors studied individual resource allocation using fairness notions relating different groups of agents, for example notions aiming to minimize envy that arises between groups [Berliant *et al.*, 1992; Husseinov, 2011; Todo *et al.*, 2011; Aleksandrov and Walsh, 2018; Conitzer *et al.*, 2019; Aziz and Rey, 2020].

Like fair division in general, consensus $1/k$ -division and consensus halving have been studied by mathematicians and economists for several decades [Hobby and Rice, 1965; Alon, 1987; Simmons and Su, 2003], and attracted recent interest from computer scientists in light of new computational complexity results [Filos-Ratsikas and Goldberg, 2018; Filos-Ratsikas *et al.*, 2020; Goldberg *et al.*, 2020; Deligkas *et al.*, 2021]. In particular, Filos-Ratsikas and Goldberg [2018] proved that approximate consensus halving of a one-dimensional heterogeneous divisible resource is PPA-complete—this constituted the first PPA-completeness result for a problem that is “natural” in the sense that its description does not involve a polynomial-sized circuit.

2 Preliminaries

Let $G = [m]$ be the set of goods, where $[r] := \{1, 2, \dots, r\}$ for any positive integer r . There are $n = n_1 + \dots + n_k$ agents divided into $k \geq 2$ groups, where group i contains $n_i \geq 1$ agents. Denote by $a^{(i,j)}$ the j th agent in group i . The utility of $a^{(i,j)}$ for good ℓ is given by $u^{(i,j)}(\ell)$. We assume that the agents’ utilities are additive, that is, $u^{(i,j)}(G') = \sum_{\ell \in G'} u^{(i,j)}(\ell)$ for every $G' \subseteq G$. An *allocation* (A_1, \dots, A_k) is an ordered partition of the goods into

k bundles, where bundle A_i is allocated to group i . In particular, $A_1 \cup \dots \cup A_k = G$ and $A_i \cap A_j = \emptyset$ for $i \neq j$.

2.1 Fairness Notions

We are interested in the following fairness notions:

Definition 2.1. Let c be a nonnegative integer. An allocation (A_1, \dots, A_k) is said to be

- *envy-free up to c goods (EF c)* if, for every agent $a^{(i,j)}$ and every $i' \neq i$, there exists a set $B \subseteq A_{i'}$ with $|B| \leq c$ such that $u^{(i,j)}(A_i) \geq u^{(i,j)}(A_{i'} \setminus B)$.
- *proportional up to c goods (PROP c)* if, for every agent $a^{(i,j)}$, there exists a set $B \subseteq G \setminus A_i$ with $|B| \leq c$ such that $u^{(i,j)}(A_i) \geq u^{(i,j)}(G/k - u^{(i,j)}(B))$.
- *a consensus $1/k$ -division up to c goods* if, for every agent a and every pair of bundles $A_i, A_{i'}$, there exists $B \subseteq A_{i'}$ with $|B| \leq c$ such that a values A_i no less than $A_{i'} \setminus B$.

Note that unlike the first two notions, the third notion does not depend on how the agents are distributed across groups. For each k, n_1, \dots, n_k , let $c^{\text{EF}}(n_1, \dots, n_k)$ (resp., $c^{\text{PROP}}(n_1, \dots, n_k)$) denote the smallest value of c such that an EF c (resp., PROP c) allocation is guaranteed to exist for agents with additive utilities. Similarly, let $c_k^{\text{CD}}(n)$ denote the analogous value for consensus $1/k$ -division up to c goods when there are n agents and k bundles. We have the following relations between these values.

Proposition 2.2. For any $k, n', n_1, \dots, n_k \in \mathbb{N}$, we have

- $c^{\text{EF}}(n_1, \dots, n_k) \leq c_k^{\text{CD}}(n_1 + \dots + n_k)$;
- $c^{\text{PROP}}(n_1, \dots, n_k) \leq c^{\text{EF}}(n_1, \dots, n_k)$;
- $c^{\text{EF}}(n', \dots, n') \geq c_k^{\text{CD}}(n')$, where there are k copies of n' on the left-hand side.

Proof. We prove the three relations in turn.

- This follows immediately from the observation that a consensus $1/k$ -division up to c goods for $n_1 + \dots + n_k$ agents is also envy-free up to c goods for these agents regardless of how the agents are distributed into groups.
- It suffices to show that for any c , every EF c allocation is also PROP c . Let (A_1, \dots, A_k) be an EF c allocation, and consider agent $a^{(i,j)}$. By definition of EF c , for each $i' \neq i$, there exists $B_{i'} \subseteq A_{i'}$ with $|B_{i'}| \leq c$ such that $u^{(i,j)}(A_i) \geq u^{(i,j)}(A_{i'} \setminus B_{i'})$.

Let B denote the set of the c most valuable goods for $a^{(i,j)}$ outside of A_i , breaking ties arbitrarily. We have $u^{(i,j)}(B) \geq u^{(i,j)}(B_{i'})$ for all $i' \neq i$. Letting $B_i = \emptyset$, we have

$$\begin{aligned} u^{(i,j)}(A_i) &\geq \frac{1}{k} \sum_{i' \in [k]} u^{(i,j)}(A_{i'} \setminus B_{i'}) \\ &\geq \frac{1}{k} \sum_{i' \in [k]} \left(u^{(i,j)}(A_{i'}) - u^{(i,j)}(B) \right) \\ &= u^{(i,j)}(G)/k - u^{(i,j)}(B). \end{aligned}$$

Hence, the allocation is PROP c , as desired.

- Given n' agents, we make a copy of each agent in each of the k groups. Any EF c allocation with respect to these groups is also a consensus $1/k$ -division for the original agents. The conclusion follows. \square

2.2 Discrepancy Theory

In this section, we outline the tools from discrepancy theory that we will use in this work. Intuitively, a basic connection between discrepancy theory and our group fair division setting is the following: Discrepancy theory considers a setting where there is a collection of subsets (also known as a *set system*) and we want to color the elements of the ground set in two colors so that each subset contains roughly the same number of elements of each color. The elements of the ground set correspond to the goods in our setting, while each subset represents an agent and its elements correspond to the goods that the agent values. The goal of discrepancy theory is therefore similar to that of dividing the goods into two sets so that each agent values the two sets almost equally.

In this work, we view a vector $\mathbf{v} \in \mathbb{R}^k$ also as a column matrix $\mathbf{v} \in \mathbb{R}^{k \times 1}$. We use $\mathbf{1}_k$ to denote the k -dimensional all-1 vector; when the dimension is clear from context, we may drop the subscript and simply write $\mathbf{1}$. Furthermore, for a set $S \subseteq [k]$, we write $\mathbf{1}(S) \in \{0, 1\}^k$ to denote the indicator vector of S , i.e., $(\mathbf{1}(S))_i = 1$ if and only if $i \in S$.

For every $p \in [1, \infty)$, we use $\|\mathbf{v}\|_p$ to denote the ℓ_p norm of \mathbf{v} , defined by $\left(\sum_{i \in [k]} |\mathbf{v}_i|^p \right)^{1/p}$. The ℓ_∞ norm of \mathbf{v} , denoted by $\|\mathbf{v}\|_\infty$, is $\max_{i \in [k]} |\mathbf{v}_i|$. The Hamming norm of \mathbf{v} , denoted by $\|\mathbf{v}\|_0$, is the number of non-zero coordinates of \mathbf{v} .

2-Color Discrepancy

As mentioned earlier, a classic scenario in discrepancy theory is when there is a set system and the goal is to color the elements in two colors in such a way that each subset contains roughly the same number of elements of each color. As Lovász et al. [1986] noted, this notion generalizes naturally to any matrix. Specifically, the *discrepancy* of a matrix $\mathbf{A} \in \mathbb{R}^{n \times m}$ is defined as

$$\text{disc}(\mathbf{A}) := \min_{\mathbf{x} \in \{0, 1\}^m} \|\mathbf{A}(0.5 \cdot \mathbf{1} - \mathbf{x})\|_\infty.$$

Here \mathbf{x} can be thought of as a 2-coloring and the quantity $\|\mathbf{A}(0.5 \cdot \mathbf{1} - \mathbf{x})\|_\infty$ measures how “unbalanced” it is.

Let

$$\text{disc}^{\max}(n) := \sup_{m \in \mathbb{N}} \sup_{\mathbf{A} \in [0, 1]^{n \times m}} \text{disc}(\mathbf{A}).$$

A priori, $\text{disc}^{\max}(n)$ might even be infinite since we allow \mathbf{A} to have an arbitrary number of columns. Remarkably, however, it is known^{3,4} that $\text{disc}^{\max}(n)$ is bounded by $O(\sqrt{n})$:

Lemma 2.3 ([Alon and Spencer, 2000, Corollary 12.3.4]). For any $n \in \mathbb{N}$, $\text{disc}^{\max}(n) \leq O(\sqrt{n})$.

³We note here that some of the bounds we refer to are stated only for 0-1 matrices \mathbf{A} . However, one can check that they also hold for any matrix $\mathbf{A} \in [0, 1]^{n \times m}$. See also the full version of our paper where we sketch how the constructive versions of these bounds can be derived [Manurangsi and Suksompong, 2021].

⁴See also [Spencer, 1985] on which this bound is based.

The above bound is also known to be asymptotically tight:

Lemma 2.4 ([Spencer, 1985]). *For any $n \in \mathbb{N}$, it holds that $\text{disc}^{\max}(n) \geq \Omega(\sqrt{n})$.*

Weighted Discrepancy. The p -weighted discrepancy [Doerr and Srivastav, 2003] is a generalization of discrepancy where 0.5 is replaced by some $p \in [0, 1]$:

$$\text{wdisc}_p(\mathbf{A}) := \min_{\mathbf{x} \in \{0,1\}^m} \|\mathbf{A}(p \cdot \mathbf{1} - \mathbf{x})\|_{\infty}.$$

Similarly to above, let

$$\text{wdisc}_p^{\max}(n) := \sup_{m \in \mathbb{N}} \sup_{\mathbf{A} \in [0,1]^{n \times m}} \text{wdisc}_p(\mathbf{A}).$$

Using standard techniques in discrepancy theory, we can prove the following lower bound on $\text{wdisc}_p^{\max}(n)$. (This bound was also implicit in the work of Doerr and Srivastav.)

Proposition 2.5. *For any $p \in (0, 1/2]$ and any $n \in \mathbb{N}$ such that $n \geq 16/p$, we have $\text{wdisc}_p^{\max}(n) \geq \Omega(\sqrt{pn})$.*

The proof of Proposition 2.5, as well as all omitted proofs, can be found in the full version of our paper [Manurangsi and Suksompong, 2021].

Multi-Color Discrepancy

We will also use the extension of the 2-color definition to multi-color cases due to Doerr and Srivastav [2003]. Recall that a k -coloring of $[m]$ is a function $\chi : [m] \rightarrow [k]$. The k -color discrepancy of \mathbf{A} is defined as

$$\text{disc}(\mathbf{A}, k) := \min_{\chi: [m] \rightarrow [k]} \max_{s \in [k]} \left\| \mathbf{A} \left(\frac{1}{k} \cdot \mathbf{1} - \mathbf{1}(\chi^{-1}(s)) \right) \right\|_{\infty}.$$

Note that $\text{disc}(\mathbf{A}, 2)$ coincides with $\text{disc}(\mathbf{A})$ defined earlier.

Similarly to above, we let

$$\text{disc}^{\max}(n, k) := \sup_{m \in \mathbb{N}} \sup_{\mathbf{A} \in [0,1]^{n \times m}} \text{disc}(\mathbf{A}, k).$$

The following lemma is a consequence of Corollary 3.5 of Doerr and Srivastav [2003] and Lemma 2.3.

Lemma 2.6 ([Doerr and Srivastav, 2003]). *For any $n, k \in \mathbb{N}$, $\text{disc}^{\max}(n, k) \leq O(\sqrt{nk})$.*

Furthermore, Doerr and Srivastav also proved the following lower bound:⁵

Lemma 2.7 ([Doerr and Srivastav, 2003, Theorem 5.2]). *For any $n, k \in \mathbb{N}$ such that $k \geq 2$, $\text{disc}^{\max}(n, k) \geq \Omega(\sqrt{nk})$.*

3 Approximate Fair Division From Multi-Color Discrepancy

In this section, we derive generic upper and lower bounds for the value c_k^{CD} based on the multi-color discrepancy bounds $\text{disc}^{\max}(n, k)$. Our results are stated formally below.

Theorem 3.1. *For any $n, k \in \mathbb{N}$, we have*

$$c_k^{\text{CD}}(n) \geq \lceil \text{disc}^{\max}(n, k) \rceil.$$

⁵Doerr and Srivastav [2003] only stated their lower bound for n such that a Hadamard matrix of order n exists. However, as explained in our proof of Proposition 2.5, this implies the same asymptotic bound for all $n \in \mathbb{N}$.

Theorem 3.2. *For any $n, k \in \mathbb{N}$, we have*

$$c_k^{\text{CD}}(n) \leq 4 \cdot \lceil \text{disc}^{\max}(2n, k) \rceil.$$

By the known bounds for disc^{\max} (Lemmas 2.6 and 2.7), the two results above yield Theorem 1.3.

Using the relationships between c_k^{CD} , c^{EF} , and c^{PROP} established in Proposition 2.2, we get the following corollary:

Corollary 3.3. *For any $k, n', n_1, \dots, n_k \in \mathbb{N}$, we have*

- (a) $c^{\text{EF}}(n_1, \dots, n_k) \leq 4 \cdot \lceil \text{disc}^{\max}(2(n_1 + \dots + n_k), k) \rceil$;
- (b) $c^{\text{PROP}}(n_1, \dots, n_k) \leq 4 \cdot \lceil \text{disc}^{\max}(2(n_1 + \dots + n_k), k) \rceil$;
- (c) $c^{\text{EF}}(n', \dots, n') \geq \lceil \text{disc}^{\max}(n', k) \rceil$, where there are k copies of n' on the left-hand side.

Parts (a) and (b), together with the known upper bound for disc^{\max} (Lemma 2.6), give us the upper bounds in Theorems 1.1 and 1.2. On the other hand, part (c) is not sufficient for the lower bounds in these theorems yet, since the above corollary does not apply to the ‘‘unbalanced groups’’ case where some groups are small, e.g., when $n_1 = \dots = n_{k-1} = 1$ and $n_k = n'$. Indeed, the lower bound parts of Theorems 1.1 and 1.2 will be handled in Section 4.

3.1 Lower Bound

We prove Theorem 3.1 via a simple reduction that views each row of a matrix \mathbf{A} as a vector of utilities for the goods. The existence of a consensus $1/k$ -division up to a small number of goods would imply a strong upper bound on the discrepancy of \mathbf{A} . This is formalized below; since the distribution of agents into groups is irrelevant for consensus $1/k$ -division, we use the notation a^j and u^j instead of $a^{(i,j)}$ and $u^{(i,j)}$.

Proof of Theorem 3.1. Let $\Delta := \lceil \text{disc}^{\max}(n, k) \rceil$. Note that $\Delta - 1 < \text{disc}^{\max}(n, k)$. Thus, there exists $m \in \mathbb{N}$ and $\mathbf{A} \in [0, 1]^{n \times m}$ such that $\text{disc}(\mathbf{A}, k) > \Delta - 1$. We define the agents’ utilities by $u^j(\ell) = \mathbf{A}_{j,\ell}$ for all $j \in [n]$ and $\ell \in [m]$. It suffices to show that there is no consensus $1/k$ -division up to $\Delta - 1$ goods with respect to these utilities.

Suppose for the sake of contradiction that there is a consensus $1/k$ -division up to $\Delta - 1$ goods, (A_1, \dots, A_k) . Let $\chi : [m] \rightarrow [k]$ denote the coloring corresponding to this allocation, i.e., $\chi^{-1}(s) = A_s$.

Consider any agent a^j . Since the allocation is a consensus $1/k$ -division up to $\Delta - 1$ goods and each good has value at most 1 to each agent, the following holds for all $i, s \in [k]$:

$$\Delta - 1 \geq |u^j(A_i) - u^j(A_s)|. \quad (1)$$

Fix any $s \in [k]$. For each $j \in [n]$, this inequality allows us to bound the j th entry of $\mathbf{A} \left(\frac{1}{k} \cdot \mathbf{1} - \mathbf{1}(\chi^{-1}(s)) \right)$ as

$$\begin{aligned} & \left| \left(\mathbf{A} \left(\frac{1}{k} \cdot \mathbf{1} - \mathbf{1}(\chi^{-1}(s)) \right) \right)_j \right| = \left| \frac{u^j(G)}{k} - u^j(A_s) \right| \\ & = \left| \frac{1}{k} \sum_{i \in [k]} (u^j(A_i) - u^j(A_s)) \right| \leq \frac{1}{k} \sum_{i \in [k]} |u^j(A_i) - u^j(A_s)| \\ & \stackrel{(1)}{\leq} \Delta - 1, \end{aligned}$$

where the first inequality follows from the triangle inequality.

Applying this for all $j \in [n]$, we have $\text{disc}(\mathbf{A}, k) \leq \Delta - 1$, contradicting our assumption that $\text{disc}(\mathbf{A}, k) > \Delta - 1$. \square

3.2 Upper Bound

We next prove our upper bound (Theorem 3.2), which turns out to be more challenging than the lower bound. To demonstrate this, let us consider using the “reverse” of the reduction in the proof of Theorem 3.1; specifically, suppose we create one row for each agent corresponding to her utilities. The discrepancy bound ensures that we can divide the goods into k bundles so that each agent’s utilities for the k bundles are not too different. However, this does not translate into any bound on the number of goods necessary in the relaxation of any of the fairness notions, since it is possible that an agent has a tiny utility for every good in some bundle.

To tackle this issue, we must also ensure that each agent has some “large” (i.e., valuable) goods in every bundle. To this end, we divide the set of goods (with respect to each agent) into the set of large goods and the set of “small” goods. We create one row as before, but only for the small goods; this is to ensure that the utilities of the agent for the small goods do not differ by much between bundles. Additionally, we create a row corresponding to the large goods, which ensures that the agent has a non-trivial number of large goods in each bundle. When choosing the size of the set of large goods appropriately, this gives us the desired bound. We formalize this intuition in the proof below.

Proof of Theorem 3.2. Let $T := \lceil \text{disc}^{\max}(2n, k) \rceil$. For every agent a^j , let $S_{\text{large}}^j \subseteq G$ denote the set of $L := \min\{m, 3Tk\}$ goods that the agent values the most (ties broken arbitrarily), and let $p^j := \min_{\ell \in S_{\text{large}}^j} u^j(\ell)$. We define \mathbf{y}^j as the m -dimensional indicator vector of S_{large}^j , i.e., $\mathbf{y}^j = \mathbf{1}(S_{\text{large}}^j)$. We also define \mathbf{z}^j as the utility vector of the goods outside of S_{large}^j , scaled by $1/p^j$, i.e.,

$$\mathbf{z}_\ell^j = \begin{cases} u^j(\ell)/p^j & \text{if } \ell \notin S_{\text{large}}^j; \\ 0 & \text{otherwise.} \end{cases}$$

By our choice of S_{large}^j and p^j , we have $\mathbf{z}^j \in [0, 1]^m$. (We use the convention $0/0 = 0$, i.e., \mathbf{z}^j is the all-zero vector when $p^j = 0$.)

Finally, we define our matrix $\mathbf{A} = [\mathbf{y}^1 \dots \mathbf{y}^n \mathbf{z}^1 \dots \mathbf{z}^n]^T$; note that $\mathbf{A} \in [0, 1]^{2n \times m}$. From the definition of disc^{\max} , there exists a coloring $\chi : [m] \rightarrow [k]$ such that

$$\left\| \mathbf{A} \left(\frac{1}{k} \cdot \mathbf{1} - \mathbf{1}(\chi^{-1}(i)) \right) \right\|_\infty \leq T. \quad (2)$$

for all $i \in [k]$. We pick our allocation (A_1, \dots, A_k) according to χ , that is, $A_i = \chi^{-1}(i)$ for all $i \in [k]$.

Next, we argue that for every pair $i, i' \in [k]$ and every agent a^j , there exists $B \subseteq A_{i'}$ of size at most $4T$ such that $u^j(A_i) \geq u^j(A_{i'} \setminus B)$; this suffices to finish the proof. From (2) and the definition of \mathbf{A} , we have

$$T \geq \left\langle \mathbf{y}^j, \left(\frac{1}{k} \cdot \mathbf{1} - \mathbf{1}(A_i) \right) \right\rangle = \left| \frac{L}{k} - |A_i \cap S_{\text{large}}^j| \right|$$

and

$$T \geq \left\langle \mathbf{z}^j, \left(\frac{1}{k} \cdot \mathbf{1} - \mathbf{1}(A_i) \right) \right\rangle$$

$$= \frac{1}{p^j} \cdot \left| \frac{1}{k} \cdot u^j(G \setminus S_{\text{large}}^j) - u^j(A_i \setminus S_{\text{large}}^j) \right|.$$

Rearranging these, we have

$$\left| \frac{L}{k} - |A_i \cap S_{\text{large}}^j| \right| \leq T \quad (3)$$

and

$$\left| \frac{1}{k} \cdot u^j(G \setminus S_{\text{large}}^j) - u^j(A_i \setminus S_{\text{large}}^j) \right| \leq p^j \cdot T. \quad (4)$$

A similar argument on bundle $A_{i'}$ implies

$$\left| \frac{L}{k} - |A_{i'} \cap S_{\text{large}}^j| \right| \leq T \quad (5)$$

and

$$\left| \frac{1}{k} \cdot u^j(G \setminus S_{\text{large}}^j) - u^j(A_{i'} \setminus S_{\text{large}}^j) \right| \leq p^j \cdot T. \quad (6)$$

Let $B := A_{i'} \cap S_{\text{large}}^j$. By (5), we have $|B| \leq L/k + T \leq 4T$. Now, if $m \leq 3Tk$, then we have $L = m$ and $A_{i'} \setminus B = \emptyset$. Thus, $u^j(A_i) \geq u^j(A_{i'} \setminus B)$ trivially holds in this case.

Next, consider the case $m > 3Tk$, so $L = 3Tk$. In this case, we may bound $u^j(A_i) - u^j(A_{i'} \setminus B)$ as follows:

$$\begin{aligned} & u^j(A_i) - u^j(A_{i'} \setminus B) \\ &= u^j(A_i \cap S_{\text{large}}^j) + \left(u^j(A_i \setminus S_{\text{large}}^j) - u^j(A_{i'} \setminus S_{\text{large}}^j) \right) \\ &\geq p^j \cdot |A_i \cap S_{\text{large}}^j| + \left(u^j(A_i \setminus S_{\text{large}}^j) - u^j(A_{i'} \setminus S_{\text{large}}^j) \right) \\ &\stackrel{(3)}{\geq} p^j(L/k - T) + \left(u^j(A_i \setminus S_{\text{large}}^j) - u^j(A_{i'} \setminus S_{\text{large}}^j) \right) \\ &\stackrel{(4),(6)}{\geq} p^j(L/k - T) - 2p^j \cdot T \geq 0, \end{aligned}$$

where the first inequality follows from our definition of p^j and the last inequality follows from $L = 3Tk$. This concludes the proof. \square

4 Lower Bounds From Weighted Discrepancy

In this section, we prove a lower bound on c^{PROP} for k groups via weighted discrepancy:

Theorem 4.1. *For any $n', k \in \mathbb{N}$, we have*

$$c^{\text{PROP}}(2n', 1, \dots, 1) \geq \lceil \text{wdisc}_{1/k}^{\max}(n')/k \rceil.$$

Combined with Proposition 2.2, this gives a similar lower bound for c^{EF} (again, the left-hand side contains $k - 1$ 1’s):

Corollary 4.2. *For any $n', k \in \mathbb{N}$, we have*

$$c^{\text{EF}}(2n', 1, \dots, 1) \geq \lceil \text{wdisc}_{1/k}^{\max}(n')/k \rceil.$$

These two results, together with the lower bound on wdisc_p^{\max} (Proposition 2.5) and the observation that removing agents does not increase the value of c^{EF} or c^{PROP} , yield the lower bound parts of Theorems 1.1 and 1.2. More specifically, we claim that for any n_1, \dots, n_k , we have

$$c^{\text{EF}}(\max\{n_1, \dots, n_k\}, 1, 1, \dots, 1) \leq c^{\text{EF}}(n_1, \dots, n_k). \quad (7)$$

To see why this claim holds, first note that due to symmetry, we may assume that $n_1 = \max\{n_1, \dots, n_k\}$. Note also that if an allocation satisfies a fairness notion for a certain instance, it still satisfies the same fairness notion when we arbitrarily remove some agent(s) from some group(s). Hence, given any instance with group sizes $n_1, 1, \dots, 1$, we may add agents with arbitrary utilities so that the group sizes become n_1, \dots, n_k and apply the bound for the latter case. This implies (7), whereupon we can apply Corollary 4.2 and Proposition 2.5 to derive Theorem 1.1. A similar argument can be used to derive Theorem 1.2 from Theorem 4.1.

The proof of Theorem 4.1 can be found in the full version of our paper; we describe here the intuition behind it. We construct the utility functions of the agents so that if (A_1, \dots, A_k) is proportional up to a small number of goods, then $\mathbf{x} := \mathbf{1}(A_1)$ gives us a small $1/k$ -weighted discrepancy. Similarly to the proof of Theorem 3.1, we start by creating n' agents in the first group where each agent’s utilities correspond to a row of \mathbf{A} . This yields a lower bound on each entry of $\mathbf{A} \cdot \mathbf{1}(A_1)$. To get an upper bound, we simply create a “conjugate” of each of these agents in the first group, i.e., the conjugate utility of each good is simply 1 minus the original utility. However, this construction alone is not sufficient—for example, it is still possible to assign all goods to A_1 . To avoid this, we create one agent in each of the remaining groups with the same utility for all goods. This ensures that A_1 has size roughly m/k , which turns out to be sufficient for bounding the $1/k$ -weighted discrepancy (using $\mathbf{x} := \mathbf{1}(A_1)$).

5 Computational Complexity

Since efficient algorithms matching the bound in Lemma 2.6 are known [Bansal, 2010; Lovett and Meka, 2015]⁶ and all of our upper bounds are obtained by polynomial-time reductions to this bound, it immediately follows that given the goods and the agents’ utilities for them, we can efficiently find an allocation matching the upper bounds in our main theorems (Theorems 1.1–1.3). In summary, we have:

Corollary 5.1. *There exists a randomized polynomial-time algorithm that can compute a consensus $1/k$ -division (or an envy-free/proportional allocation) up to $O(\sqrt{n})$ goods.*

In light of the above corollary, a natural question is whether we can improve on this $O(\sqrt{n})$ bound if we know that, say, an unknown “fully fair” division exists for a given instance. For example, provided that there is a consensus $1/k$ -division in that instance, can we efficiently find an allocation that beats the upper bounds in Corollary 5.1?

A similar question has been asked in the context of discrepancy theory; for the bound in Lemma 2.3, the answer was shown to be negative [Charikar *et al.*, 2011], i.e., even when \mathbf{A} has discrepancy zero, it is NP-hard to find \mathbf{x} achieving discrepancy $o(\sqrt{n})$. In this section, we extend this hardness to the setting of fair division, as stated below.

Theorem 5.2. *For any constant $k \in \mathbb{N} \setminus \{1\}$, there exists a constant $\varepsilon_k > 0$ such that it is NP-hard, given m goods and*

⁶Please refer to the full version of our paper for more details [Manurangsi and Suksompong, 2021].

k groups, each containing n' agents with additive utilities, to distinguish between the following two cases:

- (YES) *There exists a consensus $1/k$ -division;*
- (NO) *No allocation is proportional up to $\varepsilon_k \sqrt{n'}$ goods.*

In other words, when k is constant, we cannot asymptotically improve upon the bound in Corollary 5.1 even when we are promised that a consensus $1/k$ -division exists. Note that since consensus $1/k$ -division is the strongest notion and proportionality the weakest (see Proposition 2.2), this theorem is the “strongest possible”.

6 Conclusion and Future Work

In this paper, we have studied the allocation of indivisible goods to groups of agents using the standard fairness notions of envy-freeness, proportionality, and consensus $1/k$ -division. We presented bounds on the optimal relaxations of these notions that can be guaranteed for agents with additive valuations; all of the bounds are asymptotically tight when the number of groups is constant. Our results imply that relatively strong fairness guarantees can be provided for all agents even when agents in the same group, who share the same set of resources, have highly differing preferences. Moreover, we showed that computing allocations that improve upon these bounds is NP-hard even in instances where such allocations are known to exist.

Besides closing the gaps left by our work, an interesting direction for future work is to consider agents with arbitrary monotonic utilities. Indeed, the techniques from discrepancy theory that we used crucially rely on the additivity assumption; so does the result of Alon [1987] that established the existence of a consensus $1/k$ -division for divisible goods. Even in the case of prime numbers k , where a consensus $1/k$ -division can be guaranteed for non-additive utilities [Filos-Ratsikas *et al.*, 2021],⁷ it is unclear whether such a division can be rounded into a discrete allocation with a loss that is bounded only in terms of n . Beyond the setting of our paper, one could also consider allocating a mixture of indivisible and divisible goods [Bei *et al.*, 2021] or allowing groups to have different entitlements which can correspond to the group sizes [Chakraborty *et al.*, 2020] as well.

Acknowledgments

This work was partially supported by an NUS Start-up Grant. We would like to thank Paul Goldberg, Alexandros Hollender, and Ayumi Igarashi for interesting discussions and the anonymous reviewers for valuable comments.

References

- [Aleksandrov and Walsh, 2018] Martin Aleksandrov and Toby Walsh. Group envy freeness and group Pareto efficiency in fair division with indivisible items. In *KI*, pages 57–72, 2018.
- [Alon and Spencer, 2000] Noga Alon and Joel H. Spencer. *The Probabilistic Method*. Wiley, 2nd edition, 2000.

⁷See Theorem 6.5 in their extended version. For $k = 2$, the existence of a consensus halving with non-additive utilities was shown by Simmons and Su [2003].

- [Alon, 1987] Noga Alon. Splitting necklaces. *Advances in Mathematics*, 63(3):247–253, 1987.
- [Aziz and Rey, 2020] Haris Aziz and Simon Rey. Almost group envy-free allocation of indivisible goods and chores. In *IJCAI*, pages 39–45, 2020.
- [Bansal, 2010] Nikhil Bansal. Constructive algorithms for discrepancy minimization. In *FOCS*, pages 3–10, 2010.
- [Bei *et al.*, 2021] Xiaohui Bei, Zihao Li, Jinyan Lu, Shengxin Liu, and Xinhang Lu. Fair division of mixed divisible and indivisible goods. *Artificial Intelligence*, 293:103436, 2021.
- [Benabbou *et al.*, 2019] Nawal Benabbou, Mithun Chakraborty, Edith Elkind, and Yair Zick. Fairness towards groups of agents in the allocation of indivisible items. In *IJCAI*, pages 95–101, 2019.
- [Berliant *et al.*, 1992] Marcus Berliant, William Thomson, and Karl Dunz. On the fair division of a heterogeneous commodity. *Journal of Mathematical Economics*, 21(3):201–216, 1992.
- [Bouveret *et al.*, 2016] Sylvain Bouveret, Yann Chevaleyre, and Nicolas Maudet. Fair allocation of indivisible goods. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 12, pages 284–310. Cambridge University Press, 2016.
- [Brams and Taylor, 1996] Steven J. Brams and Alan D. Taylor. A procedure for divorce settlements. *Mediation Quarterly*, 13(3):191–205, 1996.
- [Chakraborty *et al.*, 2020] Mithun Chakraborty, Ayumi Igarashi, Warut Suksompong, and Yair Zick. Weighted envy-freeness in indivisible item allocation. In *AAMAS*, pages 231–239, 2020.
- [Charikar *et al.*, 2011] Moses Charikar, Alantha Newman, and Aleksandar Nikolov. Tight hardness results for minimizing discrepancy. In *SODA*, pages 1607–1614, 2011.
- [Conitzer *et al.*, 2019] Vincent Conitzer, Rupert Freeman, Nisarg Shah, and Jennifer Wortman Vaughan. Group fairness for the allocation of indivisible goods. In *AAAI*, pages 1853–1860, 2019.
- [Deligkas *et al.*, 2021] Argyrios Deligkas, John Fearnley, Themistoklis Melissourgos, and Paul G. Spirakis. Computing exact solutions of consensus halving and the Borsuk-Ulam theorem. *Journal of Computer and System Sciences*, 117:75–98, 2021.
- [Doerr and Srivastav, 2003] Benjamin Doerr and Anand Srivastav. Multicolour discrepancies. *Combinatorics, Probability and Computing*, 12(4):365–399, 2003.
- [Filos-Ratsikas and Goldberg, 2018] Aris Filos-Ratsikas and Paul W. Goldberg. Consensus halving is PPA-complete. In *STOC*, pages 51–64, 2018.
- [Filos-Ratsikas *et al.*, 2020] Aris Filos-Ratsikas, Alexandros Hollender, Katerina Sotiraki, and Manolis Zampetakis. Consensus halving: Does it ever get easier? In *EC*, pages 381–399, 2020.
- [Filos-Ratsikas *et al.*, 2021] Aris Filos-Ratsikas, Alexandros Hollender, Katerina Sotiraki, and Manolis Zampetakis. A topological characterization of modulo- p arguments and implications for necklace splitting. In *SODA*, pages 2615–2634, 2021. Extended version available as arXiv:2003.11974.
- [Ghodsí *et al.*, 2018] Mohammad Ghodsí, Mohamad Latifian, Arman Mohammadi, Sadra Moradian, and Masoud Seddighin. Rent division among groups. In *COCOA*, pages 577–591, 2018.
- [Goldberg *et al.*, 2020] Paul W. Goldberg, Alexandros Hollender, Ayumi Igarashi, Pasin Manurangsi, and Warut Suksompong. Consensus halving for sets of items. In *WINE*, pages 384–397, 2020.
- [Hobby and Rice, 1965] Charles R. Hobby and John R. Rice. A moment problem in L1 approximation. *Proceedings of the American Mathematical Society*, 16(4):665–670, 1965.
- [Husseinov, 2011] Farhad Husseinov. A theory of a heterogeneous divisible commodity exchange economy. *Journal of Mathematical Economics*, 47(1):54–59, 2011.
- [Kyropoulou *et al.*, 2020] Maria Kyropoulou, Warut Suksompong, and Alexandros A. Voudouris. Almost envy-freeness in group resource allocation. *Theoretical Computer Science*, 841:110–123, 2020.
- [Lipton *et al.*, 2004] Richard J. Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. On approximately fair allocations of indivisible goods. In *EC*, pages 125–131, 2004.
- [Lovász *et al.*, 1986] László Lovász, Joel Spencer, and Katalin Vesztergombi. Discrepancy of set-systems and matrices. *European Journal of Combinatorics*, 7(2):151–160, 1986.
- [Lovett and Meka, 2015] Shachar Lovett and Raghu Meka. Constructive discrepancy minimization by walking on the edges. *SIAM Journal on Computing*, 44(5):1573–1582, 2015.
- [Manurangsi and Suksompong, 2017] Pasin Manurangsi and Warut Suksompong. Asymptotic existence of fair divisions for groups. *Mathematical Social Sciences*, 89:100–108, 2017.
- [Manurangsi and Suksompong, 2021] Pasin Manurangsi and Warut Suksompong. Almost envy-freeness for groups: Improved bounds via discrepancy theory. *CoRR*, abs/2105.01609, 2021.
- [Markakis, 2017] Evangelos Markakis. Approximation algorithms and hardness results for fair division. In Ulle Endriss, editor, *Trends in Computational Social Choice*, chapter 12, pages 231–247. AI Access, 2017.
- [Moulin, 2003] Hervé Moulin. *Fair Division and Collective Welfare*. MIT Press, 2003.
- [Moulin, 2019] Hervé Moulin. Fair division in the internet age. *Annual Review of Economics*, 11:407–441, 2019.
- [Segal-Halevi and Nitzan, 2019] Erel Segal-Halevi and Shmuel Nitzan. Envy-free cake-cutting among families. *Social Choice and Welfare*, 53(4):709–740, 2019.
- [Segal-Halevi and Suksompong, 2019] Erel Segal-Halevi and Warut Suksompong. Democratic fair allocation of indivisible goods. *Artificial Intelligence*, 277:103167, 2019.
- [Simmons and Su, 2003] Forest W. Simmons and Francis Edward Su. Consensus-halving via theorems of Borsuk-Ulam and Tucker. *Mathematical Social Sciences*, 45(1):15–25, 2003.
- [Spencer, 1985] Joel Spencer. Six standard deviations suffice. *Transactions of the American Mathematical Society*, 289(2):679–706, 1985.
- [Suksompong, 2018] Warut Suksompong. Approximate maximin shares for groups of agents. *Mathematical Social Sciences*, 92:40–47, 2018.
- [Thomson, 2016] William Thomson. Introduction to the theory of fair allocation. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 11, pages 261–283. Cambridge University Press, 2016.
- [Todo *et al.*, 2011] Taiki Todo, Runcong Li, Xuemei Hu, Takayuki Mouri, Atsushi Iwasaki, and Makoto Yokoo. Generalizing envy-freeness toward group of agents. In *IJCAI*, pages 386–392, 2011.