

Democratic Fair Allocation of Indivisible Goods

Erel Segal-Halevi and Warut Suksompong

Abstract

We study the problem of fairly allocating indivisible goods to groups of agents. Agents in the same group share the same set of goods even though they may have different preferences. Previous work has focused on *unanimous fairness*, in which all agents in each group must agree that their group’s share is fair. Under this strict requirement, fair allocations exist only for small groups. We introduce the concept of *democratic fairness*, which aims to satisfy a certain fraction of the agents in each group. This concept is better suited to large groups such as cities or countries. We present protocols for democratic fair allocation among two or more arbitrarily large groups of agents with monotonic, additive, or binary valuations. Our protocols approximate both envy-freeness and maximin-share fairness. As an example, for two groups of agents with additive valuations, our protocol yields an allocation that is envy-free up to one good and gives at least half of the maximin share to at least half of the agents in each group.

1 Introduction

Fair division is the study of how to allocate resources among agents with different preferences so that agents perceive the resulting allocation as fair. This problem occurs in a wide range of situations, from negotiating over international interests and reaching divorce settlements [4] to dividing household tasks and sharing apartment rent [7].

Two kinds of fairness criteria are common in the literature. The first, *envy-freeness (EF)*, means that each agent finds her share at least as good as the share of any other agent. When allocating indivisible goods, envy-freeness is sometimes unattainable (consider two agents quarreling over a single good), so it is often relaxed to *envy-freeness up to one good (EF1)*, which is always attainable [11, 5]. The second kind, *maximin share fairness*, means that each agent finds his share at least as good as his *maximin share* (MMS), which is the best share he can secure by dividing the goods into n parts and getting the worst part, where n is the number of agents. Again, with indivisible goods MMS fairness cannot be guaranteed, but at least a constant fraction of the MMS can [15].

Most works on fair division involve individual agents, each of whom has individual preferences. But in reality, resources often have to be allocated among *groups of agents*, such as families or states. A good allocated to a group is shared among the group members and all of them derive full utility from the good. For example, when dividing real estate among families, all members of a family enjoy their allocated house and backyard. In international negotiations, the divided rights and settled outcomes are enjoyed by all citizens of a country. When resources are allocated between different buildings of a university, all occupants of a building benefit from the whiteboards and open space allocated to their building. However, different group members may have different preferences. The same share can be perceived as fair by one member and unfair by another member of the same group. Ideally, we would like to find an allocation considered fair by *all* agents in *all* groups. However, two recent works show that this “unanimous fairness” might be too strong to be practical.

(a) Suksompong [23] shows that when allocating indivisible goods among groups, there might be no allocation that is unanimously EF1. Moreover, there might be no division that

gives all agents a positive fraction of their MMS. This impossibility occurs even for two groups of three agents.

(b) Segal-Halevi and Nitzan [17] show that when allocating a divisible good (“cake”) among groups, there might be no division that is unanimously envy-free and gives each group a single connected piece, or even a constant number of connected pieces. In contrast, with individual agents a connected envy-free division always exists [18].

What do groups do when they cannot attain unanimity? In democratic societies, they use some kind of voting. The premise of voting is that it is impossible to satisfy everyone, so we should try to satisfy as many members as possible. Based on this observation, we say that a division is *h-democratic fair*, for some fairness notion and for some $h \in [0, 1]$, if at least a fraction h of the agents in each group believe it is fair. In this paper we focus on allocating indivisible goods. We would like h , the fraction of *happy* agents, to be as large as possible. We thus pose the following question:

Given a fairness notion, what is the largest h such that an h -democratic fair allocation of indivisible goods can always be found?

We study democratic fairness under three different assumptions on the agents’ valuations. In the most general case, the agents can have arbitrary *monotonic* valuations on bundles of goods. A more common assumption in the literature is that agents’ valuations are *additive* (the value of a bundle is the sum of the value of its goods). We also study a special case of additive valuations in which agents’ valuations are *binary* (each agent has a set of desired goods and her utility equals the number of desired goods allocated to her group).

1.1 Overview of our results

Initially (**Section 3**) we consider two groups with binary agents. We study a relaxation of envy-freeness that we call *envy-freeness up to c goods* (EF c), a generalization of EF1. One might expect to have a trade-off curve where a larger c corresponds to a larger h . However, we find that the actual trade-off curve is degenerate: for every constant c , it is possible to guarantee 1/2-democratic EF c and the 1/2 is tight. The same holds for MMS fairness. Moreover, a positive fraction of the MMS can be guaranteed to at least 3/5 and at most 2/3 of the agents in both groups. To get a more flexible trade-off curve, we study a generalization of MMS called *1-out-of- c -MMS*, which is the best share an agent can secure by dividing the goods into c subsets and receiving the worst one. We prove that for every integer $c \geq 2$, 1-out-of- c MMS can be guaranteed to at least $1 - 1/2^{c-1}$ and at most $1 - 1/c^2 2^{c+1}$ of the agents in both groups.

Our positive results are attained by an efficient round-robin protocol where each group in turn picks a good using *weighted approval voting* with carefully calculated weights. We believe this weighted voting scheme can be interesting in its own right as a way to make fair group decisions.

Next (**Section 4**) we consider two groups whose agents have arbitrary monotonic valuations. We present an efficient protocol that guarantees EF1 to at least 1/2 of the agents in each group (which is tight even for binary agents). When all agents are additive, this protocol guarantees 1/2 of the MMS to 1/2 of the agents. This is tight: one cannot guarantee more than 1/2 of the MMS to more than 1/3 of the agents.

Finally (**Section 5**), we present two generalizations of our results to $k \geq 3$ groups. The first generalization has stronger fairness guarantees: when all valuations are binary, it guarantees to $1/k$ of the agents in all groups both EF1 and MMS (the $1/k$ is tight for EF1). When valuations are additive, it guarantees an additive approximation to EF and MMS. However, the run-time of the protocol might be exponential. The second generalization uses a polynomial-time protocol but has weaker guarantees: when all valuations are binary, it

Happy $h \downarrow$ Share $q \rightarrow$	Positive	$(0, 1/2]$	$(1/2, 1]$
$(0, 1/3]$	Yes (Cor. 3.8)	Yes (Cor. 4.2)	Bin: Yes (Thm. 3.6), Add: ?
$(1/3, 1/2]$			Bin: Yes (\uparrow), Add: No (\downarrow)
$(1/2, 3/5]$?	Bin: ?, Add: No (Prop. 4.3)
$(3/5, 2/3]$			
$(2/3, 1]$	No (Prop. 3.1)		

Happy h	EFc for any constant $c \geq 1$
$(0, 1/2]$	Yes (Thm. 4.1)
$(1/2, 1]$	No (Prop. 3.4)

Table 1: Summary of results for two groups with **Binary** and **Additive** agents. For each range of $h, q \in (0, 1]$, the table shows whether there always exists an allocation that gives at least a fraction h of the agents in each group at least a fraction q of their maximin share. For EFc, the results hold for monotonic valuations too. The arrows refer to the directions pointed to in the table.

guarantees MMS to $1/k$ of the agents, and when valuations are additive, it guarantees an additive approximation to MMS. Some of our results and open questions are summarized in Table 1.

1.2 Related work

The group resource allocation problem is relatively new. We already mentioned the impossibility result of Suksompong [23], which is for worst-case agents' utilities. On the other hand, if the agents' utilities are drawn at random from probability distributions, Manurangsi and Suksompong [12] showed that a unanimously envy-free allocation exists with high probability as the number of agents and goods grows. In our terminology, unanimous fairness is called *1-democratic-fairness*.

The term democratic fairness already appears in [16, 17]; however, they use it in the narrower sense that at least $1/2$ of the agents in each group must be satisfied. In our terminology this is called *1/2-democratic fairness*. Hence, our democratic fairness notion generalizes existing notions of group fairness.

A related model, in which a subset of public goods is allocated to a single group of agents but the rest of the goods remain unallocated, has also been studied [13, 21].

MMS fairness was introduced by Budish [5] based on earlier concepts by Moulin [14]. Budish also considered its relaxation to 1-out-of- $(n+1)$ MMS. The notion 1-out-of- c MMS is a special case of l -out-of- d MMS, recently defined by Babaioff et al. [1].

Group preferences are important in matching markets, too. For example, when matching doctors to hospitals, usually a husband and a wife want to be matched to the same hospital. This issue poses a substantial challenge to stable-matching mechanisms [8, 9, 10].

2 Preliminaries

There is a set $G = \{g_1, \dots, g_m\}$ of goods. A *bundle* is a subset of G . There is a set A of agents. The agents are partitioned into k groups A_1, \dots, A_k with n_1, \dots, n_k agents, respectively. Let a_{ij} denote the j th agent in group A_i . Each agent a_{ij} has a nonnegative utility $u_{ij}(G')$ for each $G' \subseteq G$. For any agent a_{ij} , denote by $u_{ij, \max} := \max_{l=1, \dots, m} u_{ij}(g_l)$ the maximum utility of the agent for any single good. Denote by $\mathbf{u}_{ij} = (u_{ij}(g_1), \dots, u_{ij}(g_m))$ the utility vector of agent a_{ij} for single goods. The agents' utility functions are *monotonic*, i.e., $u_{ij}(G'') \leq u_{ij}(G')$ for every $G'' \subseteq G' \subseteq G$ and every agent a_{ij} . A subclass of monotonic

utilities is the class of *additive* utilities, i.e., for every bundle $G' \subseteq G$ and every agent $a_{ij} \in A$, we have $u_{ij}(G') = \sum_{g \in G'} u_{ij}(g)$.

Sometimes we will study a special case of additive utilities in which utilities are *binary*, i.e., each agent either approves or disapproves each good. Since we will not engage in interpersonal comparison of utilities, we may assume without loss of generality that in this case $u_{ij}(g) \in \{0, 1\}$ for each i, j, g .

We allocate a bundle $G_i \subseteq G$ to each group A_i . All goods should be allocated. The goods are treated as public goods within each group, i.e., for every group i , the utility of every agent a_{ij} is $u_{ij}(G_i)$. We refer to a setting with agents partitioned into groups, goods and utility functions as an *instance*.

We now define the fairness notions considered in this paper. We start by defining what it means for an allocation to be fair for a *specific* agent. We start with envy-freeness.

Definition 2.1. Given an agent a_{ij} and an integer $c \geq 0$, an allocation is called *envy-free up to c goods (EF c)* for a_{ij} if for every i' there is a set $C_{i'} \subseteq G_{i'}$ with $|C_{i'}| \leq c$ such that

$$u_{ij}(G_i) \geq u_{ij}(G_{i'} \setminus C_{i'}).$$

In other words, one can remove the envy of a_{ij} toward group i' by removing at most c goods from the group's bundle.

An EF0 allocation is also known as *envy-free*.

Next, we define the maximin share concepts.

Definition 2.2. Given an agent a_{ij} and an integer $c \geq 2$, the *1-out-of- c maximin share (MMS)* of a_{ij} is defined as the maximum, over all partitions of G into c sets, of the minimum of the agent's utilities for the sets in the partition:

$$\text{MMS}_{ij}^c(G) := \max_{G'_1, \dots, G'_c} \min(u_{ij}(G'_1), \dots, u_{ij}(G'_c)),$$

where the maximum ranges over all partitions (G'_1, \dots, G'_c) of G . When $c = k$ (the number of groups), the 1-out-of- k MMS of an agent is simply called his *MMS* and denoted $\text{MMS}_{ij}(G)$. An allocation (G_1, \dots, G_k) is said to be:

- *1-out-of- c MMS-fair* for a_{ij} , if $u_{ij}(G_i) \geq \text{MMS}_{ij}^c(G)$.
- *MMS-fair* for a_{ij} , if $u_{ij}(G_i) \geq \text{MMS}_{ij}(G)$.
- *q -MMS-fair* for a_{ij} , for some fraction $q \in (0, 1)$, if $u_{ij}(G_i) \geq q \cdot \text{MMS}_{ij}(G)$.
- *positive-MMS-fair* if every agent with positive MMS gets positive value: $\text{MMS}_{ij}(G) > 0 \rightarrow u_{ij}(G_i) > 0$.

Note that MMS-fairness implies q -MMS-fairness (for any q), which implies positive-MMS-fairness. The next lemma shows an interesting link between EF1 and MMS-fairness; the proof can be found in the appendix.

Lemma 2.3. *If an allocation is EF1 for an agent with an additive utility function, then (a) it is also 1/ k -MMS-fair for that agent—the 1/ k is tight; (b) if the agent's utility function is binary, then the allocation is also MMS-fair for that agent.*

Now we are ready to define our main group fairness notion:

Definition 2.4. For any given fairness notion, an allocation (G_1, \dots, G_k) is said to be *h -democratic fair* if it is fair for at least $h \cdot n_i$ agents in group A_i , for all $i \in \{1, \dots, k\}$.

3 Two Groups with Binary Valuations

This section considers the special case in which there are two groups, and each agent either *desires* a good (in which case her utility for the good is 1) or does not desire it (in which case her utility is 0). Even in this special case, some fairness guarantees are unattainable.

Proposition 3.1. *For any $h > 2/3$, there is a binary instance in which no allocation is h -democratic positive-MMS-fair.*

Proof. Consider an instance with three goods where in each group, each of the three agents desires a unique subset of two goods. Each agent has a positive MMS (1), but no allocation gives all agents a positive utility. \square

We next leverage a combinatorial construction of Erdős to show the limitations of 1-out-of- c MMS-fairness.

Lemma 3.2. *For any integer $c \geq 2$, there is an instance with two groups consisting of $c^2 2^{c+1}$ agents with binary valuations in each group, such that each agent desires c goods but no allocation gives all agents a positive utility.*

Proof. Erdős [6] proved that for any positive integer c , there exists a collection \mathcal{C} of $c^2 2^{c+1}$ subsets of size c of a base set G that does not have “property B”.¹ This means that, no matter how we partition G into two subsets G_1 and G_2 , some subset in \mathcal{C} has an empty intersection with G_1 or G_2 .

Take the elements of G to be our goods. Each group consists of $c^2 2^{c+1}$ agents, each of whom desires a unique subset of goods in \mathcal{C} . Then every agent desires c goods, but no allocation gives all agents a positive utility. \square

Proposition 3.3. *For any integer $c \geq 2$ and $h > 1 - 1/c^2 2^{c+1}$, there is a binary instance with two groups in which no allocation is h -democratic 1-out-of- c MMS-fair.*

Proof. Consider the instance from Lemma 3.2. The 1-out-of- c MMS of an agent who desires c goods is positive (1), but no allocation gives all agents a positive utility. \square

If we change the fairness requirement to EF1, then we can satisfy no more than half of the agents in each group.

Proposition 3.4. *For any constant integer $c \geq 1$ and $h > 1/2$, there is an instance with two groups with binary agents in which no allocation is h -democratic EFC.*

Proof. Consider an instance with $m = 4l$ goods and $\binom{4l}{2l}$ agents in each group, for some $l \geq 1$. Each agent desires a unique subset of $2l$ goods. An allocation is EFC for an agent iff her group receives at least $l - \lfloor c/2 \rfloor$ of her $2l$ desired goods.

The symmetry between the groups implies that the best fairness guarantee can be attained by giving exactly $2l$ goods to each group; the symmetry between the goods implies that it does not matter which $2l$ goods are given to which group. In each such allocation, the number of a group’s members who receive exactly j desired goods is $\binom{2l}{j} \cdot \binom{4l-2l}{2l-j} = \binom{2l}{j}^2$. Therefore, the number of a group’s members who receive at least $l - \lfloor c/2 \rfloor$ desired goods is: $\sum_{j=l-\lfloor c/2 \rfloor}^{2l} \binom{2l}{j}^2 = \sum_{j=l-\lfloor c/2 \rfloor}^{l-1} \binom{2l}{j}^2 + \frac{1}{2} \binom{2l}{l}^2 + \frac{1}{2} \binom{4l}{2l}$, where the equality follows from expanding the central binomial coefficient $\binom{4l}{2l}$. The fraction of a group’s members who think the division is EFC is attained by dividing this expression by $\binom{4l}{2l}$. This fraction is:

$$\frac{1}{2} + \frac{1}{2} \cdot \frac{\binom{2l}{l}^2}{\binom{4l}{2l}} + \frac{\sum_{j=l-\lfloor c/2 \rfloor}^{l-1} \binom{2l}{j}^2}{\binom{4l}{2l}} \leq \frac{1}{2} \left(1 + (c+1) \frac{\binom{2l}{l}^2}{\binom{4l}{2l}} \right).$$

¹We are grateful to Fedor Petrov for suggesting the connection with Property B.

Using Stirling's approximation, we find that $\binom{2l}{l} \sim \frac{4^l}{\sqrt{\pi l}}$, so

$$\frac{\binom{2l}{l}^2}{\binom{4l}{2l}} \sim \frac{\frac{4^{2l}}{\pi l}}{\frac{4^{2l}}{\sqrt{2\pi l}}} = \sqrt{\frac{2}{\pi l}}.$$

As $l \rightarrow \infty$, the fraction of agents in each group who think that the allocation is EFC approaches 1/2, as claimed. \square

We now move on to some positive results. These are attained with a protocol we call Round-robin with Weighted Approval Voting (RWAV). Both groups take turn picking a single good until all goods are taken. Each group picks its good using the following weighted-approval-voting scheme:

- (a) Initially, each member pays to the "group account" some amount (to be calculated later) of fiat money.
- (b) Whenever it is the group's turn to pick, each member is assigned a positive *weight* (to be calculated later).
- (c) For each good, the *total weight* is calculated as the sum of the weights of the members who desire this good. The group picks a good with a maximal total weight (breaking ties arbitrarily).
- (d) Every member whose desired good was picked by the group pays his weight to the group account.
- (e) When it is the other group's turn to pick, each member whose desired good was picked by the other group receives his weight from the group account. This rule has one exception: it is not executed for the second group in the first turn of the first group (so that each execution of step (e) is preceded by an execution of step (d)).

We now calculate the weights such that, when the protocol ends, the net payment paid by each happy agent is 1 and by each unhappy agent is 0 (i.e., each unhappy agent got all his money back). An agent's weight will be a function of the number r of his desired goods that remain untaken, and the number s of desired goods that he is missing to be happy. Both r and s of an agent weakly decrease as the protocol runs. An agent becomes happy when $s = 0$ and unhappy when $r < s$. Let $w(r, s)$ be the weight of such an agent and $B(r, s)$ the net amount paid by such an agent. Then, $\forall r \geq s > 0$:

$$\begin{aligned} (*) \quad B(r, s) + w(r, s) &= B(r - 1, s - 1) && \text{by step (d)} \\ B(r, s) - w(r, s) &= B(r - 1, s) && \text{by step (e)} \end{aligned}$$

This implies the following recurrence relation for $B(r, s)$:

$$\begin{aligned} \forall r \geq s > 0: \quad B(r, s) &= \frac{B(r - 1, s) + B(r - 1, s - 1)}{2} \\ \forall r \geq 0: \quad B(r, 0) &= 1 \quad (\text{happy agents}) \\ \forall r < s: \quad B(r, s) &= 0 \quad (\text{unhappy agents}) \end{aligned}$$

$$\begin{aligned} \text{Its solution is (**):} \quad \forall r \geq s \geq 0: \quad B(r, s) &= \frac{1}{2^r} \sum_{i=s}^r \binom{r}{i} \\ \text{and (***)}: \quad w(r, s) &= B(r, s) - B(r - 1, s) \end{aligned}$$

We make several observations. First, when r is fixed, $B(r, s)$ decreases when s increases, since the sum in (***) has fewer elements. Second, when s is fixed, $B(r, s)$ increases when r increases, since $B(r, s)$ is an average of $B(r - 1, s)$ with the larger term $B(r - 1, s - 1)$. Now, (***) implies that all weights are positive, as required by the protocol.

To complete the specification of the protocol, we have to calculate the initial payments in step (a). We define a generalized fairness criterion called $f(d)$ -fairness, where f is some integer function. An allocation is $f(d)$ -fair for agent j if the agent's group receives at least $f(d_j)$ desired goods whenever the agent has d_j desired goods. Note that EF1 and MMS-fairness are both equivalent to $\lfloor \frac{d}{2} \rfloor$ -fairness. Suppose we are interested in $f(d)$ -fairness for some function f . Then, in the first group, an agent with d desired goods has initially $r = d$ and $s = f(d)$. Hence his initial payment should be $B(d, f(d))$. In the second group, such an agent might have lost a desired good in the first turn of the first group, so $r = d - 1$ and the initial payment is only $B(d - 1, f(d))$, which by (*) is smaller than $B(d, f(d))$. We are now ready to state the main lemma:

Lemma 3.5. *For every integer function $f(d)$, it is possible to select weights such that the RWAV protocol attains h -democratic $f(d)$ -fairness, where:*

$$h = \inf_{d=1,2,\dots} B(d-1, f(d)) = \inf_{d=1,2,\dots} \frac{1}{2^{d-1}} \sum_{i=f(d)}^{d-1} \binom{d-1}{i}$$

Proof. In step (a), each agent pays at least h . Therefore, the initial balance of group i is at least $h \cdot n_i$. In step (e) the group pays the total weight of a good, while in step (d) the group receives the total weight of another good. Since the good picked in step (c) has a maximal total weight, the balance of the group weakly increases, so its final balance is at least $h \cdot n_i$. By (*), the final balance of each unhappy agent is 0 and of each happy agent -1 . Since the total balance of the group plus the agents is 0, there are at least $h \cdot n_i$ happy agents. \square

Lemma 3.5 implies the following positive result.

Theorem 3.6. *For every integer $c \geq 2$, it is possible to select weights such that RWAV attains $(1 - 1/2^{c-1})$ -democratic 1-out-of- c MMS-fairness.*

Proof. Suppose an agent wants $d = l \cdot c + l'$ goods, for some integers $l \geq 1$ and $l' < c$. So her 1-out-of- c -MMS is l . We can arbitrarily ignore l' of her desired goods and aim to give her a utility of l . Applying Lemma 3.5 with $f(lc) = l$ gives $h = \inf_{l=1,2,\dots} B(lc - 1, l)$. It remains to prove that for every l , $B(lc - 1, l) \geq B(c - 1, 1) = 1 - 1/2^{c-1}$. That is:

$$\begin{aligned} & \frac{1}{2^{lc-1}} \sum_{i=l}^{lc-1} \binom{lc-1}{i} \geq 1 - \frac{1}{2^{c-1}} \\ \iff & \frac{1}{2^{lc-1}} \sum_{i=0}^{l-1} \binom{lc-1}{i} \leq \frac{1}{2^{c-1}} \\ \iff & \sum_{i=0}^{l-1} \binom{lc-1}{i} \leq 2^{lc-c} \end{aligned}$$

We prove the latter inequality using combinatorial arguments.² Consider a row of $lc - c$ light-bulbs, each of which can be either on or off. The bulbs are divided to c groups with $l - 1$ bulbs in each group. There are $lc - 1$ switches. Of these, $lc - c$ are "normal" switches,

²We are grateful to Y. Forman for suggesting these arguments in <https://math.stackexchange.com/a/2598269/29780>.

each of which toggles a single distinct bulb. The remaining $c - 1$ are “special” switches: special switch number i toggles all the bulbs in groups i and c . Initially all lights are off. We are allowed to push at most $l - 1$ switches. How many light patterns can we make?

Clearly, the left-hand side is the number of choices of at most $l - 1$ switches, and the right-hand side is the number of different light patterns. Therefore, to prove the inequality it is sufficient to prove that any two distinct switch choices generate different light patterns.

Consider two distinct switch choices and let S_1, S_2 be their set of special switches. If $S_1 = S_2$ then the two choices must contain different sets of normal switches and therefore generate different light patterns. Suppose $S_1 \neq S_2$. Any set of special switches leaves an even number of groups of lights on, so the light patterns resulting only from the special switches differ in at least two groups, i.e., $2l - 2$ positions. However, both choices must contain at most $l - 1$ switches overall, and at least one of them already contains at least one special switch so only $l - 2$ normal switches remain. Therefore the total number of normal switches in both groups is at most $2l - 3$. The difference cannot be canceled using normal switches, so the patterns must be different. \square

Our results raise the following open question: what is the maximum fraction of agents that can be guaranteed 1-out-of- c -fairness? By Theorem 3.6 it is at least $1 - 1/2^{c-1}$; by Proposition 3.3 it is at most $1 - 1/c^2 2^{c+1}$.

An interesting special case of Theorem 3.6 is when $c = 2$, since with two groups 1-out-of-2 MMS-fairness is equivalent to MMS-fairness, and with binary agents this is also equivalent to EF1. Hence Theorem 3.6 implies that 1/2-democratic EF1 and MMS-fairness are attainable. The fraction 1/2 exactly matches the upper bound in Proposition 3.4.

Note that the guarantee in Theorem 3.6 (as in Lemma 3.5) is for the second group; the guarantee for the first group is even higher—it is $1 - 1/2^c$. In particular, in the first group at least 3/4 of all agents are guaranteed EF1 and MMS.

What happens if we are only interested in positive-MMS-fairness? The upper bound on h in Proposition 3.1 is 2/3, but the RWAV protocol only guarantees 1/2. This is because each agent with two desired goods must be given at least one good, and $B(2 - 1, 1) = 1/2$. However, a variant of RWAV can do better even for additive valuations. We say that an allocation is *one-of-best-two* for an agent with additive valuations if he receives at least one of his two most valuable goods. The proof of the following result can be found in the appendix.

Theorem 3.7. *For two groups with additive agents, there exists a 3/5-democratic one-of-best-two allocation.*

Corollary 3.8. *For two groups with additive agents, there exists a 3/5-democratic positive-MMS allocation.*

4 Two Groups with General Valuations

In this section we assume that there are two groups and each agent can have an arbitrary monotonic utility function.

We start with a positive result: it is always possible to efficiently allocate goods so that at least half of the agents in each group believe the division is EF1. Despite the simplicity of the protocol, we find the result important since, unlike previous results in this setting [12, 23], our result holds for worst-case instances with any number of agents in the groups and very general utility functions.

Theorem 4.1. *For two groups with agents with monotonic valuations, 1/2-democratic EF1 is attainable.*

Proof. We arrange the goods in a line and process them from left to right. Starting from an empty block, we add one good at a time until the current block is EF1 for at least half of the agents in at least one group. We allocate the current block to one such group, and the remaining goods to the other group.

Since the whole set of goods is EF1 for both groups, the protocol terminates. Assume without loss of generality that the left block G_1 is allocated to the first group A_1 , and the right block G_2 to the second group A_2 . By the description of the protocol, the allocation is EF1 for at least half of the agents in A_1 , so it remains to show that the same holds for A_2 . Let g be the last good added to the left block. More than half of the agents in A_2 think that $G_1 \setminus \{g\}$ is not EF1, so for these agents, $G_1 \setminus \{g\}$ is worth less than $G_2 \cup \{g\} \setminus \{g'\}$ for any $g' \in G_2 \cup \{g\}$. Taking $g' = g$, we find that these agents value $G_1 \setminus \{g\}$ less than G_2 . But this implies that the agents find G_2 to be EF1, completing the proof. \square

Theorem 4.1 shows that if the goods lie on a line, we can find a $1/2$ -democratic EF1 allocation that moreover gives each group a contiguous block on the line. This may be important, for example, if the goods are houses on a street and each group wants to have all its houses in a contiguous block [3, 22].

If agents have additive valuations, Lemma 2.3 implies:

Corollary 4.2. *For two groups with additive agents, $1/2$ -democratic $1/2$ -MMS-fairness is attainable.*

For EF1, the factor $1/2$ in Theorem 4.1 is tight even for binary valuations, as shown in Proposition 3.4. For $1/2$ -MMS-fairness, the factor $1/2$ in Corollary 4.2 is “almost” tight:

Proposition 4.3. *For any $h > 1/3$ and $q > 1/2$, there is an additive instance with two groups in which no allocation is h -democratic q -MMS-fair.*

Proof. Consider an instance with $m = 3$ goods and $n_1 = n_2 = 3$ agents in each group, with utility vectors: $\mathbf{u}_{i1} = (2, 1, 1)$, $\mathbf{u}_{i2} = (1, 2, 1)$, and $\mathbf{u}_{i3} = (1, 1, 2)$ for $i = 1, 2$. The MMS of every agent is 2. In any allocation, one group receives at most one good, so at most one of its three agents receives utility more than 1. So in that group, at most $1/3$ of the agents receive more than $1/2$ of their MMS. \square

A corollary of Proposition 4.3 is that, for every $h \in (1/3, 1/2]$, the maximum fraction q such that there always exists an h -democratic q -MMS-fair allocation is $q = 1/2$.

5 Three or More Groups

In this section we study the most general setting where we allocate goods among any number of groups. When there are two groups, the protocol in Theorem 4.1 is efficient and yields an allocation that is both approximately envy-free and approximately MMS-fair. We present two ways of generalizing the result to multiple groups: one keeps the approximate envy-freeness guarantee but loses computational efficiency, while the other keeps only the approximate MMS-fairness guarantee but also retains computational efficiency.

5.1 Approximate envy-freeness

The main theorem in this subsection is:

Theorem 5.1. *When all agents have binary valuations, there exists an allocation that is $1/k$ -democratic EF1 and $1/k$ -democratic MMS-fair. The factor $1/k$ is tight for EF1.*

To establish this theorem, we prove two lemmas that may be of independent interest—one on cake-cutting and the other on group allocation for agents with additive valuations.

The result on cake-cutting generalizes the theorems of Stromquist [18] and Su [20], who prove the existence of contiguous envy-free cake allocations for individual agents. Since these results are well-known, we present the model and proof quite briefly, focusing on the changes required to generalize from individuals to groups.

We consider a “cake” modeled as the interval $[0, 1]$. Each agent a_{ij} has a value-density function $v_{ij} : [0, 1] \rightarrow \mathbb{R}_{\geq 0}$. The value of an agent for a piece X is $V_{ij}(X) = \int_{x \in X} v_{ij}(x) dx$. Denoting by X_i the allocation to group i , an allocation is *envy-free* for an agent a_{ij} if $V_{ij}(X_i) \geq V_{ij}(X_{i'})$ for every group i' . A *contiguous allocation* is an allocation of the cake in which each group gets a contiguous interval.

Lemma 5.2. *There always exists a contiguous cake allocation that is $1/k$ -democratic envy-free. The factor $1/k$ is tight.*

Proof. The space of all contiguous partitions corresponds to the standard simplex in \mathbb{R}^k . Triangulate that simplex and assign each vertex of the triangulation to one of the groups. In each vertex, ask the group owning that vertex to select one of the k pieces using *plurality voting* among its members, breaking ties arbitrarily. Label that vertex with the group’s selection. The resulting labeling satisfies the conditions of *Sperner’s lemma* (see [20]). Therefore, the triangulation has a *Sperner subsimplex*—a subsimplex all of whose labels are different. We can repeat this process with finer and finer triangulations. This gives an infinite sequence of smaller and smaller Sperner subsimplices. This sequence has a subsequence that converges to a single point. By the continuity of preferences, this limit point corresponds to a partition in which each group selects a different piece. Since the selection is by plurality, at least $1/k$ of the agents in each group prefer their group’s piece over all other pieces.

The tightness of the $1/k$ factor follows from Lemma 6 of Segal-Halevi and Nitzan [17]. It shows an example with k groups and n' agents in each group with the property that in order to give a positive value to q out of n' agents in each group, we need to cut the cake into at least $k(kq - n')/(k - 1)$ intervals. In a contiguous partition there are exactly k intervals. Therefore, the fraction of agents in each group that can be guaranteed a positive value is $q/n' \leq 1/k + 1/n' - 1/kn'$. Since n' can be arbitrarily large, the largest fraction that can be guaranteed is $1/k$. \square

The next lemma presents a reduction from approximate envy-free allocation of indivisible goods to envy-free cake-cutting. We call this approximation “EF-minus-2”. An allocation is *EF-minus-2* for agent a_{ij} if for every group i' , $u_{ij}(G_i) > u_{ij}(G_{i'}) - 2u_{ij, \max}$. The reduction generalizes Theorem 5 of Suksompong [22]; a similar reduction was used in Theorem 3 of Barrera et al. [2].

Lemma 5.3. *When agents have additive valuations, there always exists a contiguous allocation of indivisible goods that is $1/k$ -democratic EF-minus-2.*

Proof. We create an instance of the cake-cutting problem in the following way. Given an instance of indivisible goods allocation, we create a cake-cutting instance where the cake is the half-open interval $(0, m]$. The value-density functions are piecewise constant: for every $l \in \{1, \dots, m\}$, the value-density v_{ij} in the half-open interval $(l - 1, l]$ equals $u_{ij}(g_l)$.

By Lemma 5.2, there exists a contiguous cake allocation that is envy-free for at least $1/k$ of the agents in each group. From this allocation we construct an allocation of goods as follows. If point g of the cake is in the interior of a piece, then good g is given to the group owning that piece. Else, point g is at the boundary between two pieces, and good g is given to the group owning the piece to its left. A group gets good g only if it owns a positive

fraction of the interval $(g - 1, g]$. Hence, in the allocation, each group loses strictly less than the value of a good and gains strictly less than the value of a good (relative to its value in the cake division). This means that every agent who believes that the cake allocation is envy-free also believes that the goods allocation is EF-minus-2. \square

Proof of Theorem 5.1. Suppose an allocation is EF-minus-2 for some agent. This means that the agent's envy towards any other group is less than $2u_{\max}$. Since the agent has binary valuations, the envy is at most 1, meaning that the allocation is EF1 for that agent. Hence any $1/k$ -democratic EF-minus-2 allocation, which is guaranteed to exist by Lemma 5.3, is also $1/k$ -democratic EF1. By Lemma 2.3 it is also $1/k$ -democratic MMS-fair.

We next show that the factor $1/k$ is tight. Assume that there are $m = km'$ goods for some large positive integer m' . Each group consists of 2^m agents, each of whom values a distinct combination of the goods. Consider first an allocation that gives exactly m' goods to each group, and fix a group. We claim that the fraction of the agents in the group whose utilities for some two bundles differ by at most 1 converges to 0 for large m' . Indeed, this follows from the central limit theorem: Fix two bundles and consider a random agent from the group; let X be the random variable denoting the (possibly negative) difference between the agent's utilities for the two bundles. Then X is a sum of m' independent and identically distributed random variables with mean 0. The central limit theorem implies that for any fixed $\epsilon > 0$, there exists a constant c such that $\Pr[|X| \leq 1] \leq \Pr[|X| \leq c\sqrt{m'}] \leq \epsilon$ for any sufficiently large m' . Taking the union bound over all pairs of bundles, we find that the fraction of agents in the group who value some two bundles within 1 of each other approaches 0 as m' goes to infinity. This means that all but a negligible fraction of the agents find only one bundle to be EF1. By symmetry, $1/k$ of these agents find the bundle allocated to the group to be EF1. It follows that the fraction of agents in the group for whom the allocation is EF1 converges to $1/k$.

It remains to consider the case where the allocation does not give the same number of goods to all groups. In this case, let \mathcal{G} denote the set of bundles with the smallest number of goods, which must be strictly smaller than m' goods. If we move goods from bundles with more than m' goods to bundles in \mathcal{G} in such a way that the number of goods in each bundle in \mathcal{G} increases by exactly one, the fraction of agents in an arbitrary group that receives a bundle in \mathcal{G} who finds the allocation to be EF1 can only increase. We can repeat this process, at each step possibly adding bundles to \mathcal{G} , until all bundles contain the same number of goods, which is the case we have already handled. Since the fraction of agents for whom the allocation is EF1 is bounded above by $1/k$ for large m' in the latter allocation, and this fraction only increases during our process of moving goods, the same is true for the original allocation. \square

The cake-cutting protocol of Lemma 5.2 might take infinitely many steps to converge. In fact, there is no finite protocol for contiguous envy-free cake-cutting even for individuals [19]. However, the division guaranteed by Lemma 5.3 and Theorem 5.1 can be found in finite time (exponential in the input size) by checking all possible allocations. It is interesting whether a faster algorithm exists.

5.2 Approximate MMS

In this subsection, we show that if we weaken our fairness requirement to approximate MMS, it is possible to compute a fair allocation in time polynomial in the input size.

Lemma 5.4. *When agents have additive valuations, there always exists an allocation such that at least $1/k$ of the agents a_{ij} in each group j receive utility at least $\frac{1}{k} \cdot u_{ij}(G) - \frac{k-1}{k} \cdot u_{ij, \max}$, and such an allocation can be computed efficiently.*

This lemma generalizes the corresponding result for the setting with one agent per group by Suksompong [22]. The factor $(k-1)/k$ is tight even for individual agents.

Proof. We arrange the goods in a line and process them from left to right. Starting from an empty block, we add one good at a time until the current block yields utility at least $\frac{1}{k} \cdot u_{ij}(G) - \frac{k-1}{k} \cdot u_{ij,\max}$ for at least $1/k$ of the agents in at least one group. We allocate the current block to one such group and repeat the process with the remaining $k-1$ groups. It is clear that this algorithm can be implemented efficiently. Any group that receives a block from the algorithm meets the requirement, so it suffices to show that the algorithm allocates a block to every group. We claim that if l groups are yet to receive a block, at least l/k of the agents a_{ij} in each of these groups have utility at least $\frac{l}{k} \cdot u_{ij}(G) - \frac{k-l}{k} \cdot u_{ij,\max}$ for the remaining goods. This would imply the desired result because for the last group, at least $1/k$ of the agents have utility $\frac{1}{k} \cdot u_{ij}(G) - \frac{k-1}{k} \cdot u_{ij,\max}$, which is exactly our requirement.

To show the claim, we proceed by backward induction on l . The claim trivially holds when $l = k$. Suppose that the statement holds when there are $l+1$ groups left, and consider a group j that is not the next one to receive a block. At least $(l+1)/k$ of the agents a_{ij} in the group have utility at least $\frac{l+1}{k} \cdot u_{ij}(G) - \frac{k-l-1}{k} \cdot u_{ij,\max}$ for the remaining goods. Since the group does not receive the next block, less than $1/k$ of the agents in the group have utility at least $\frac{1}{k} \cdot u_{ij}(G) - \frac{k-1}{k} \cdot u_{ij,\max}$ for the block excluding the last good. Hence, less than $1/k$ of the agents have utility at least $\frac{1}{k} \cdot u_{ij}(G) + \frac{1}{k} \cdot u_{ij,\max}$ for the whole block. This means that at least l/k of the agents have utility at least $(\frac{l+1}{k} \cdot u_{ij}(G) - \frac{k-l-1}{k} \cdot u_{ij,\max}) - (\frac{1}{k} \cdot u_{ij}(G) + \frac{1}{k} \cdot u_{ij,\max}) = \frac{l}{k} \cdot u_{ij}(G) - \frac{k-l}{k} \cdot u_{ij,\max}$, completing the induction. \square

It is clear by definition that the MMS of any agent a_{ij} is at most $\frac{1}{k} \cdot u_{ij}(G)$. Lemma 5.4 therefore implies the following:

Theorem 5.5. *When agents have additive valuations, there always exists an allocation such that at least $1/k$ of the agents a_{ij} in each group j receive utility at least $\text{MMS}_{ij}(G) - \frac{k-1}{k} \cdot u_{ij,\max}$, and such an allocation can be computed efficiently.*

For binary valuations, if we change the stopping condition in Lemma 5.4 to be when the current block yields the MMS for at least $1/k$ of the agents in some group, we get:

Theorem 5.6. *When agents have binary valuations, there always exists a $1/k$ -democratic MMS-fair allocation, and such an allocation can be computed efficiently.*

6 Conclusion and Future Work

For two groups, we have a rather comprehensive understanding of possible democratic-fairness guarantees. We have a complete characterization of possible envy-freeness approximations, and upper and lower bounds for maximin-share-fairness approximations. Some remaining gaps are shown in Table 1; closing them raises interesting combinatorial challenges.

For $k \geq 3$ groups, the challenges are much greater. Currently all our fairness guarantees are to $1/k$ of the agents in all groups. From a practical perspective, it may be important in some settings to give fairness guarantees to at least half of the agents in all groups. Preliminary numerical calculations indicate that the RWAV protocol can be modified to provide such guarantees. We plan to check this further in future work. From an algorithmic perspective, it is interesting whether there exists a polynomial-time algorithm that guarantees EF1 to any positive fraction of the agents.

Acknowledgments

We are grateful to Fedor Petrov, Arnaud Mortier, Darij Grinberg, Michael Korn, Kevin P. Costello, Nick Gill, Jack D’Aurizio and Leon Bloy for helpful discussions.

References

- [1] M. Babaioff, N. Nisan, and I. Talgam-Cohen. Competitive equilibria with indivisible goods and generic budgets. *CoRR*, abs/1703.08150, 2017.
- [2] R. Barrera, K. Nyman, A. Ruiz, F. E. Su, and Y. Zhang. Discrete envy-free division of necklaces and maps. *CoRR*, abs/1510.02132, 2015.
- [3] S. Bouveret, K. Cechlárová, E. Elkind, A. Igarashi, and D. Peters. Fair division of a graph. In *Proc. 26th International Joint Conference on Artificial Intelligence*, 135–141, 2017.
- [4] S. J. Brams and A. D. Taylor. A procedure for divorce settlements. *Meditation Quarterly*, 13(3):191–205, 1996.
- [5] E. Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103, 2011.
- [6] P. Erdős. On a combinatorial problem. II. *Acta Mathematica Academiae Scientiarum Hungarica*, 15(3–4):445–447, 1964.
- [7] J. Goldman and A. D. Procaccia. Spliddit: Unleashing fair division algorithms. *ACM SIGecom Exchanges*, 13(2):41–46, 2014.
- [8] B. Klaus and F. Klijn. Stable matchings and preferences of couples. *Journal of Economic Theory*, 121(1):75–106, 2005.
- [9] B. Klaus and F. Klijn. Paths to stability for matching markets with couples. *Games and Economic Behavior*, 58(1):154–171, 2007.
- [10] F. Kojima, P. A. Pathak, and A. E. Roth. Matching with couples: Stability and incentives in large markets. *The Quarterly Journal of Economics*, 128(4):1585–1632, 2013.
- [11] R. J. Lipton, E. Markakis, E. Mossel, and A. Saberi. On approximately fair allocations of indivisible goods. In *Proc. 5th ACM Conference on Economics and Computation*, 125–131, 2004.
- [12] P. Manurangsi and W. Suksompong. Asymptotic existence of fair divisions for groups. *Mathematical Social Sciences*, 89:100–108, 2017.
- [13] P. Manurangsi and W. Suksompong. Computing an approximately optimal agreeable set of items. In *Proc. 26th International Joint Conference on Artificial Intelligence*, 338–344, 2017.
- [14] H. Moulin. Uniform externalities: Two axioms for fair allocation. *Journal of Public Economics*, 43(3):305–326, 1990.
- [15] A. D. Procaccia and J. Wang. Fair enough: Guaranteeing approximate maximin shares. In *Proc. 15th ACM Conference on Economics and Computation*, 675–692, 2014.

- [16] E. Segal-Halevi and S. Nitzan. Proportional cake-cutting among families. *CoRR*, abs/1510.03903, 2015.
- [17] E. Segal-Halevi and S. Nitzan. Envy-free cake-cutting among families. *CoRR*, abs/1607.01517, 2016.
- [18] W. Stromquist. How to cut a cake fairly. *The American Mathematical Monthly*, 87(8):640–644, 1980.
- [19] W. Stromquist. Envy-free cake divisions cannot be found by finite protocols. *Electronic Journal of Combinatorics*, 15:#R11, 2008.
- [20] F. E. Su. Rental harmony: Sperner’s lemma in fair division. *The American Mathematical Monthly*, 106(10):930–942, 1999.
- [21] W. Suksompong. Assigning a small agreeable set of indivisible items to multiple players. In *Proc. 25th International Joint Conference on Artificial Intelligence*, 489–495, 2016.
- [22] W. Suksompong. Fairly allocating contiguous blocks of indivisible items. In *Proc. 10th International Symposium on Algorithmic Game Theory*, 333–344, 2017.
- [23] W. Suksompong. Approximate maximin shares for groups of agents. *Mathematical Social Sciences*, 2018, 92:40–47, 2018.

Erel Segal-Halevi
Department of Computer Science
Ariel University
Ariel, Israel
Email: erelsgl@gmail.com

Warut Suksompong
Department of Computer Science
Stanford University
Stanford, CA, USA
Email: warut@cs.stanford.edu

A Omitted proofs

A.1 Proof of Lemma 2.3

Denote by u the utility function of the agent and assume without loss of generality that the agent is in group A_1 .

(a) EF1 implies that in each bundle G_i (for $i \in \{1, \dots, k\}$) there exists a subset C_i with $|C_i| \leq 1$ such that $u(G_1) \geq u(G_i \setminus C_i)$. Summing over all groups gives that $k \cdot u(G_1) \geq u(G \setminus (C_2 \cup \dots \cup C_k))$. Now, in any partition of G into k bundles, there is at least one bundle that does not contain any good in $C_2 \cup \dots \cup C_k$. This bundle is contained in $G \setminus (C_2 \cup \dots \cup C_k)$. Therefore, the MMS is at most $u(G \setminus (C_2 \cup \dots \cup C_k))$ which is at most $k \cdot u(G_1)$. Therefore, $u(G_1)$ is at least $1/k$ of the MMS.

To show that the factor $1/k$ is tight, assume that there are $2k - 1$ goods with $u(g_1) = \dots = u(g_k) = 1$ and $u(g_{k+1}) = \dots = u(g_{2k-1}) = k$. If the agent's group (say, group A_1) gets g_1 and group $i \geq 2$ gets $\{g_i, g_{k+i-1}\}$, the agent gets utility 1 and finds the allocation EF1. However, the MMS is k , as can be seen from the partition $(\{g_1, \dots, g_k\}, \{g_{k+1}\}, \dots, \{g_{2k-1}\})$.

(b) Suppose the agent's group wins l of the agent's desired goods. EF1 implies that each of the other $k - 1$ groups wins at most $l + 1$ of the agent's desired goods. Hence the agent has at most $kl + k - 1$ desired goods. Hence the agent's MMS is at most l , so the allocation is MMS-fair for her.

A.2 Proof of Theorem 3.7

We first convert all valuations to binary by assuming each agent desires only his two most valuable goods (breaking ties arbitrarily).

If, in one of the groups, at least $3/5$ of the agents desire the same good g , then give g to that group and give all other goods to the other group. The allocation is obviously $3/5$ -democratic one-of-best-two.

Otherwise, run RWAV with the following initial payments. In the first group, all n_1 agents pay $B(2, 1) = 3/4$. In the second group, there are less than $\frac{3}{5}n_2$ agents whose desired good was taken; these agents pay $B(1, 1) = 1/2$ and the others pay $B(2, 1) = 3/4$. The initial balance of the first group is $\frac{3}{4}n_1$. Hence, as in the proof of Lemma 3.5, at the end at least $\frac{3}{4}n_1$ of its members are happy. The initial balance of the second group is at least $\frac{3}{5}n_2 \cdot \frac{1}{2} + \frac{2}{5}n_2 \cdot \frac{3}{4} = \frac{3}{5}n_2$. Hence, at the end at least $\frac{3}{5}n_2$ of its members are happy.