

Fair Division Among Couples and Small Groups

Paul Gözl, Hannane Yaghoubzade

Cornell University, School of Operations Research and Information Engineering
136 Hoy Road, Ithaca, NY 14850, USA
paulgoelz@cornell.edu, hy465@cornell.edu

Abstract

We study the fair allocation of indivisible goods across groups of agents with additive valuations, where each agent fully enjoys all goods allocated to their group. We focus on groups of two (*couples*) and other groups of small size. For two couples, an EF1 allocation — one in which all agents find their group’s bundle no worse than the other group’s, up to one good — always exists and can be found efficiently. For three or more couples, EF1 allocations need not exist.

Turning to proportionality, we show that, whenever groups have size at most k , a PROP k allocation exists and can be found efficiently. In fact, our algorithm additionally guarantees (fractional) Pareto optimality, and PROP1 to the first agent in each group, PROP2 to the second, and so on, for an arbitrary agent ordering. In special cases, we show that there are PROP1 allocations for any number of couples.

Code — <https://github.com/hannayzade/fair-division-among-couples-and-small-groups>

Extended version — <https://arxiv.org/abs/2508.13432>

1 Introduction

Four siblings — Anna, Ben, Carmen, and Dave — jointly own a cottage on the coast and are currently deciding which sibling’s family will get to stay in the cottage in which weeks of the year. Each sibling i has a utility $u_i(\alpha) \geq 0$ for each week α ; for example, Anna prefers Spring weeks over Summer due to milder temperatures, and Ben would particularly value being at the cottage for July 4. Assume (as we will throughout the paper) that a sibling’s utility for a set B of weeks is *additive* (i.e., that $u_i(B) = \sum_{\alpha \in B} u_i(\alpha)$) and that the siblings are treating the weeks as *indivisible* goods, i.e., they don’t want to allocate fractions of a week or assign a week to several families.

To solve the siblings’ predicament as stated so far, the field of fair division offers allocation algorithms (e.g., Lipton et al. 2004; Caragiannis et al. 2019) with compelling axiomatic guarantees. In particular, algorithms like *envy-cycle elimination* and *Maximum Nash welfare* ensure *envy freeness up to one good* (EF1). EF1 means that each sibling i finds their assigned weeks B_i to be at least as valuable as the weeks B_j

assigned to any other sibling, at least when removing some week α from B_j : $u_i(B_i) \geq \min_{\text{good } \alpha} u_i(B_j \setminus \{\alpha\})$.¹ A second key axiom, *proportionality up to one good* (PROP1), states that each sibling i receives at least their proportional share of their utility for all weeks, at least when adding some week α : $\max_{\text{week } \alpha} u_i(B_i \cup \{\alpha\}) \geq \frac{u_i(M)}{n}$, where $n = 4$ is the number of siblings and M the set of all weeks.

Our scenario deviates from the classic fair division setting in that the siblings’ spouses also have utilities, which need not align with their partners’. An allocation of weeks that Anna finds fair may still make her husband Alex envy another sibling or perceive their family’s assignment as falling short of proportionality. Is it possible to allocate the weeks over the families so that axioms such as EF1 and PROP1 hold from the perspectives of all four siblings and their respective spouses? Equivalent fair division problems may arise in splitting an inheritance between families or dissolving business partnerships, whenever the entities receiving allocations consist of two persons whose preferences should be satisfied.

The question we raise above fits into the model of group fair division (Manurangsi and Suksompong 2017; Kyropoulou, Suksompong, and Voudouris 2020), but little is known for small groups such as couples. Kyropoulou, Suksompong, and Voudouris (2020) show that EF1 allocations need not exist for two groups with three members each and show that an (as of yet, unproven) graph conjecture by Jafari and Alipour (2017) would imply EF1 existence for two couples. With the basic question of EF1 existence for two couples unresolved, Kyropoulou et al. instead focus on how much axioms like EF1 must be relaxed to be satisfiable for two large groups. Meanwhile, the question of whether EF1 might exist for arbitrarily many couples has remained open.

In this paper, we aim to answer the following question:

For fair allocation over couples, and groups with few members more broadly, can we always guarantee the existence of EF1 or PROP1 allocations? If no, can we guarantee slight approximations?

A second motivation is that this question touches on the fundamental combinatorics of fair allocations to *individuals*.

¹The strengthening of this axiom without the removal of a good, *envy freeness*, is not always satisfiable for indivisible goods. E.g., if all siblings only have positive utility for a single week, the siblings who do not receive this week will always be envious.

Indeed, fix a number n of agents and m of indivisible goods, and consider the set \mathcal{A} of all n^m allocations of goods across agents 1 through n . Each vector of utility functions $u = (u_1, \dots, u_n)$ determines a subset $F_u \subseteq \mathcal{A}$ of allocations that are EF1 (or PROP1 etc.) for these utilities; call the family of all such sets \mathcal{F} . An allocation to n couples is EF1 iff it lies in F_u for the utilities u of the first partners in each couple and in $F_{u'}$ for the second partners' utilities u' . That is, EF1 allocations always exist for couples iff any two sets $F_u \in \mathcal{F}$ intersect, i.e., if \mathcal{F} is an *intersecting family* (Jukna 2011).

1.1 Our Techniques and Results

We begin by proving in Section 3.1 that EF1 allocations always exist for two couples, proving a case left open by Kyropoulou, Suksompong, and Voudouris (2020) without relying on the graph conjecture mentioned above. Our proof takes a different path, by rounding a fractional allocation computed via an unconventional linear program, which also yields a polynomial-time algorithm. Though there are only a few ways of rounding this LP, our proof that one of these roundings must satisfy EF1 requires a non-trivial combinatorial argument. In Section 3.2, we prove that EF1 cannot be guaranteed for three (or more) couples.

Since EF1 is not achievable for more than two couples, we turn our focus on PROP1 and its variants in Section 4. In the setting of arbitrarily many couples, we can guarantee the slightly weaker guarantee of PROP2, in addition to fractional Pareto optimality, an axiom of allocation efficiency. This follows from our main result — an efficient, iterative-rounding algorithm that works for groups of arbitrary sizes, and guarantees PROP1 to the first member of each group, PROP2 to the second, etc., according to an arbitrary ordering of group members. In various special cases, we can show that PROP1 allocations exist for arbitrarily many couples, for example if utilities are dichotomous and all agents value the same number of goods. PROP1 allocations need not exist for groups of three agents.

In Section 5, we study fair division among couples empirically, by taking utilities from real-world allocation problems submitted to Spliddit (Goldman and Procaccia 2014) and pairing up the agents. EF1 and PROP1 allocations exist for every single fair division problem we study, suggesting that they are ubiquitous in practice. We also study the iterative rounding algorithm and find that it almost always provides PROP1 for all agents, and even EF1 in most cases.

1.2 Related Work

Fair division among groups was independently introduced by Manurangsi and Suksompong (2017), and by Segal-Halevi and Nitzan (2019); the former studying indivisible goods and the latter divisible goods. This line of work has since been expanded, with several works exploring fair allocation in both divisible and indivisible settings (Suksompong 2018; Ghodsi et al. 2018; Kyropoulou, Suksompong, and Voudouris 2020; Segal-Halevi and Suksompong 2020; Manurangsi and Suksompong 2022; Segal-Halevi and Suksompong 2023; Bu et al. 2024; Caragiannis, Larsen, and Shyam 2025; Manurangsi and Meka 2025). In contrast to our work, most research on group

fair division with indivisible goods focuses on asymptotic analysis for large groups.

The paper by Bu et al. (2025) is closely related to our work. While their motivation for studying two sets of utilities does not make reference to groups, their setting is equivalent to fair division among couples. Concurrently to us,² Bu et al. (2025) also show that EF1 allocations exist for two couples. Their argument makes heavy use of combinatorial results to show EF1 existence even for general monotone valuations. They study proportionality as well, proving the existence of PROP- $O(\log(n))$ allocations for additive functions, where n is the number of couples. Our iterative rounding algorithm improves the gap in proportionality to a constant (PROP2) and naturally extends to arbitrary group sizes.

2 Preliminaries

A group fair division instance consists of a set of goods $M = [m] = \{1, \dots, m\}$, a set \mathcal{G} of $n \geq 2$ groups of agents, and the agents' valuations.

For a group $g \in \mathcal{G}$, we use $|g|$ to denote the number of agents in that group. We refer to the i -th agent in g as (g, i) for $1 \leq i \leq |g|$. Each agent (g, i) has a value $u_{gi}(\alpha)$ for each good $\alpha \in M$, which induce an additive valuation function $u_{gi}(B) := \sum_{\alpha \in B} u_{gi}(\alpha)$ over sets of goods B . We say that the agent's valuation is *binary* if $u_{gi}(\alpha) \in \{0, 1\}$ for all $\alpha \in M$.

A *fractional allocation* is a vector $x \in [0, 1]^{M \times \mathcal{G}}$, where $x_{\alpha g}$ denotes the fraction of good α assigned to group g , such that $\sum_{g \in \mathcal{G}} x_{\alpha g} = 1$ for each $\alpha \in M$. If $x \in \{0, 1\}^{M \times \mathcal{G}}$, we call it a (*discrete*) *allocation*. Equivalently, we represent an allocation as a partition of the goods into *bundles* $\{B_g\}_{g \in \mathcal{G}}$, where B_g is the bundle group g receives. An allocation is *balanced* if $||B_g| - |B_{g'}|| \leq 1$ for all $g, g' \in \mathcal{G}$. Since all agents in g fully enjoy the goods in their bundle B_g , agent (g, i) 's utility for such an allocation is $u_{gi}(B_g)$. We linearly extend utilities to fractional allocations so that (g, i) 's utility for allocation x is $\sum_{\alpha \in M} x_{\alpha, g} u_i(\alpha)$.

For any $k \geq 0$, an allocation $\{B_g\}_{g \in \mathcal{G}}$ is

- *envy-free up to k goods (EF k)* for an agent (g, i) if, for any $g' \in \mathcal{G}$, there is a set $B \subseteq B_{g'}$ such that $|B| \leq k$ and $u_{gi}(B_g) \geq u_{gi}(B_{g'} \setminus B)$.
- *proportional up to k goods (PROP k)* for an agent (g, i) if there exists a set $B \subseteq M \setminus B_g$ such that $|B| \leq k$ and $u_{gi}(B_g \cup B) \geq u_{gi}(M)/n$.

An allocation is *envy-free (EF)* for an agent if it is EF0 and *proportional (PROP)* if its PROP0. We say an allocation is EF k if it is EF k for every agent, and analogously for EF, PROP k , and PROP. EF and PROP also naturally extend to fractional allocations.

A fractional allocation x *Pareto dominates* another fractional allocation x' if, for all agents (g, i) , $u_{gi}(x) \geq u_{gi}(x')$ and if this inequality is strict for at least one agent. A fractional allocation x is *fractionally Pareto optimal (fPO)* if it

²An earlier version (Bu et al. 2024) relied on the conjecture by Jafari and Alipour (2017) to prove EF1 for two couples. The very recently revised preprint (Bu et al. 2025) removes this assumption and includes an algorithm for additive utilities.

is not Pareto dominated by any other fractional allocation. Note that, for a discrete allocation, being fPO implies the more classic axiom of Pareto optimality (i.e., not being Pareto dominated by any discrete allocation).

Linear Programming. We recall some notions from linear programming. A set $P = \{x \in \mathbb{R}^n : Ax \leq b\}$ for some $A \in \mathbb{R}^{m \times n}$, and $b \in \mathbb{R}^m$ is a *polyhedron*, and a bounded polyhedron is called a *polytope*. A point $z \in P$ is called a *Basic Feasible Solution (BFS)* if there are n linearly independent rows of A such that $Az \leq b$ holds with equality in these n rows. *Linear Programming (LP)* consists of maximizing (or minimizing) a linear function over a polyhedron, i.e., $\max\{c^T x : Ax \leq b, x \in \mathbb{R}^n\}$. We repeatedly use:

Proposition 2.1 (Bertsimas and Tsitsiklis 1997, Thm. 2.8). *If P is a non-empty polytope, there exists a BFS that achieves $\max\{c^T x : x \in P\}$ for a given $c \in \mathbb{R}^n$.*

Such an optimal BFS can be found in (weakly) polynomial time using the ellipsoid method (Khachiyan 1980).

3 EF1 Among Couples

In this section, we study EF1 allocations for *couples*, i.e., groups g that all have size 2. For $n = 2$ couples, we show that the “gold standard” (Freeman et al. 2019) of EF1 can be achieved even when we must satisfy twice as many agents per bundle, compared to the classic, individual setting. However, this positive result does not extend further, as we prove that EF1 allocations may not exist for $n \geq 3$ couples.

3.1 Existence of EF1 with Two Couples

In this part, we consider the case with two groups, which we call the *first* group f and the *second* group s . Kyropoulou, Suksompong, and Voudouris (2020) proved that a balanced EF1 allocation always exists when $(|f|, |s|) = (2, 1)$, and left the existence of such an allocation for $(|f|, |s|) = (2, 2)$ as an open question, which we answer in the affirmative.

Our high-level approach is to round an appropriate fractional allocation into an EF1 (discrete) allocation. Starting from a fractional allocation is promising because envy-freeness is always achievable in this domain (say, by splitting each good equally between groups). Broadly speaking, fractional allocations x are easiest to round if they are already “almost discrete” (i.e., most entries are 0 or 1). In such cases, we only need to round the few remaining fractional entries to 0 and 1, which yields a limited number of discrete allocations to reason over, all of which are still close to x and might therefore be “almost” envy-free.

The most direct attempt at pursuing this approach would be to round a BFS from the polytope of envy-free allocations, which is defined as follows, where x_α is the share of good α given to the first group:

$$\begin{aligned} \sum_{\alpha \in M} x_\alpha u_{fi}(\alpha) &\geq \sum_{\alpha \in M} (1 - x_\alpha) u_{fi}(\alpha) & i = 1, 2 \\ \sum_{\alpha \in M} (1 - x_\alpha) u_{si}(\alpha) &\geq \sum_{\alpha \in M} x_\alpha u_{si}(\alpha) & i = 1, 2 \\ 0 \leq x_\alpha &\leq 1 & \forall \alpha \in M. \end{aligned}$$

Using a BFS of this polytope allows us to obtain an almost discrete allocation. Indeed, since there are m variables, m

constraints must be tight in a BFS, and hence at least $m - 4$ constraints of the shape $0 \leq x_\alpha$ or $x_\alpha \leq 1$ are tight. This implies that at most four goods α are allocated *fractionally* (i.e., $0 < x_\alpha < 1$), leaving $2^4 = 16$ ways of rounding.

Unfortunately, a BFS of this polytope may not have a way to be rounded into an EF1 allocation. For example, consider the following valuations over goods $\{1, 2, 3, 4\}$:

valuation	1	2	3	4
u_{f1}	1	0	0.1	0.1
u_{f2}	0	1	0.1	0.1
u_{s1}	0.5	0.5	0.1	0.1
u_{s2}	0.2	0	0.5	0.5

In this case, one BFS³ allocates a $3/5$ fraction of goods 1 and 2 to group f , and the rest of 1 and 2 plus the entirety of 3 and 4 to group s . But no way of rounding the fractional goods leaves all agents EF1: agent $(f, 1)$ requires good 1 to be given to f , $(f, 2)$ requires the same for good 2, but giving both to f leaves $(s, 1)$ envious.

To obtain a working rounding argument, we devise an alternative polytope whose BFS has even fewer fractional variables, and whose roundings are all EF1 for agent $(f, 1)$, leaving us with one fewer agent to worry about. For this, assume w.l.o.g. that the number of goods m is even (otherwise, we add a dummy good with value 0 for every agent) and that the goods are ordered according to $(f, 1)$ ’s valuation, i.e., $u_{f1}(1) \geq \dots \geq u_{f1}(m)$.

We restrict ourselves to allocations in which each group receives exactly one out of the goods $\{1, 2\}$, one out of $\{3, 4\}$, ..., and one out of $\{m - 1, m\}$. This structure is inspired by Kyropoulou, Suksompong, and Voudouris (2020), who observe that any allocation with this structure is EF1 for $(f, 1)$, allowing us to focus on satisfying EF1 for the remaining three agents. The structure also ensures balancedness.

This structure naturally generalizes to fractional allocations by ensuring that each group receives a total of one unit from each pair $\{2j - 1, 2j\}$ (for $1 \leq j \leq m/2$). Specifically, if the first group receives a fraction y_j of good $2j - 1$, it receives $1 - y_j$ of good $2j$. The polytope of fractional allocations of this structure, which are furthermore EF for the three remaining agents, can be written as follows:

$$\begin{aligned} \sum_{j \in [m/2]} (2y_j - 1) u_{f2}(2j - 1) + (1 - 2y_j) u_{f2}(2j) &\geq 0 \\ \sum_{j \in [m/2]} (1 - 2y_j) u_{si}(2j - 1) + (2y_j - 1) u_{si}(2j) &\geq 0 \quad i = 1, 2 \\ 0 \leq y_j &\leq 1 \quad j = 1, \dots, m/2. \end{aligned}$$

A BFS of this polytope has at most three fractional values. We avoid one more fractional value and simplify the argument by not only requiring EF, but maximizing the minimum gap d by which any agent prefers their bundle over the other. We believe that this trick for eliminating one more fractional

³Specifically, the EF fractional allocation with maximum utilitarian welfare.

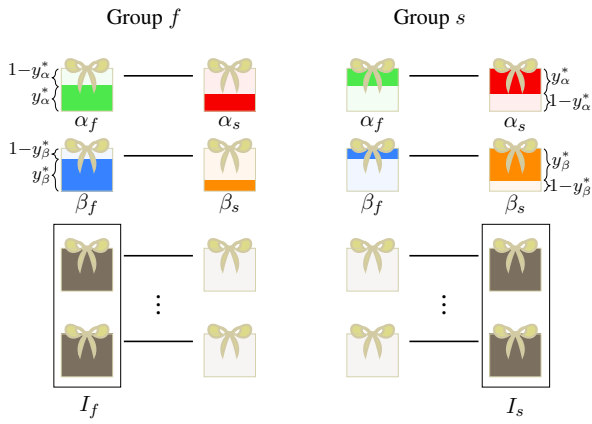


Figure 1: Illustration of y^* . Each good is paired with the good next to it. The bright, solid portion of a box represents the fraction of the good a group receives, while the faded portion represents the fraction allocated to the other group.

variable can be useful for other settings as well.

$$\begin{aligned} & \max d \\ & \text{s.t. } \sum_{j \in [m/2]} (2y_j - 1) u_{f2}(2j - 1) + (1 - 2y_j) u_{f2}(2j) \geq d \\ & \sum_{j \in [m/2]} (1 - 2y_j) u_{s1}(2j - 1) + (2y_j - 1) u_{s1}(2j) \geq d \quad i = 1, 2 \\ & 0 \leq y_j \leq 1 \quad j = 1, \dots, m/2. \end{aligned}$$

Because this formulation has one more variable, one more constraint is binding at a BFS, which ensures that there are at most two fractional y_j . Let (y^*, d^*) be an optimal BFS for the LP, which can be found efficiently (Proposition 2.1). Since setting all $y_j = 1/2$ and $d = 0$ is feasible, we know that d^* is nonnegative and that y^* describes a fractional allocation that is EF for the three agents.

We are now ready to prove EF1 existence, by rounding the fractional solution y^* into an EF1 allocation, in which each group receives exactly one good among $\{2j - 1, 2j\}$ for each $1 \leq j \leq m/2$. Since the rounding argument is fairly complex, we outline the main ideas here and defer the proof to Appendix A of the full version.

Theorem 3.1. *In the case of two groups with two agents each, a balanced EF1 allocation always exists and can be found in (weakly) polynomial time.*

Proof sketch. Since y^* has at most two fractional values, we can fix $1 \leq \alpha, \beta \leq m/2$ such that all variables except for y_α^*, y_β^* are integral. Let I_f and I_s be the set of remaining goods that are allocated entirely to the first and second group, respectively. We can assume w.l.o.g. that $y_\alpha^*, y_\beta^* \geq 1/2$.⁴ For convenience, set $\alpha_f := 2\alpha - 1$ to be the good of which group f receives a y_α^* fraction and group s receives a $1 - y_\alpha^*$ fraction; $\alpha_s := 2\alpha$ to be the good of which s receives a y_α^* fraction and f receives a $1 - y_\alpha^*$ fraction; and analogously for β_f, β_s . This allocation is illustrated in Fig. 1.

⁴Otherwise, one can swap the roles of, say, goods $2\alpha - 1$ and 2α , which keeps $(f, 1)$ EF1.

Case $y_\alpha^* + y_\beta^* \geq 3/2$. If $y_\alpha^* + y_\beta^* \geq 3/2$, the fractional allocation is already very close to being integral. In this case, allocating $\{\alpha_f, \beta_f\}$ to f and $\{\alpha_s, \beta_s\}$ to s turns out to be EF1. Since this allocation gives each good to the group that had the larger fraction of it in y^* , we refer to this as the *natural rounding*. To see that the natural rounding is EF1, observe that the natural rounding is reached by starting from y^* , and transferring goods as follows:

$$\begin{array}{ccc} & \xrightarrow{(1 - y_\alpha^*) \times \alpha_s, (1 - y_\beta^*) \times \beta_s} & \\ \text{group } f & & \text{group } s \\ & \xleftarrow{(1 - y_\alpha^*) \times \alpha_f, (1 - y_\beta^*) \times \beta_f} & \end{array}$$

Taking the perspective of, say, $(f, 1)$, they start from the envy-free allocation y^* and receive some fraction of α_f and β_f , which only reduces their envy. Then, f hands some fraction of α_s, β_s to s , but the amount of this transfer is $1 - y_\alpha^* + 1 - y_\beta^* \leq 1/2$ goods. This transfer increase $(f, 1)$'s envy twofold because f 's allocation shrinks and that of s grows. But $(f, 1)$'s envy is now at most $\max(u_{f1}(\alpha_s), u_{f1}(\beta_s))$, which can be eliminated by removing the higher-valued good from s 's bundle.

Case $y_\alpha^* + y_\beta^* < 3/2$. In the remaining case, in which $3/2 > y_\alpha^* + y_\beta^* \geq 1$, the fractional allocation is further from the natural rounding. As a consequence, we have to reason about which of the four rounding options (in which f receives $\{\alpha_f, \beta_f\}$, $\{\alpha_f, \beta_s\}$, $\{\alpha_s, \beta_f\}$, or $\{\alpha_s, \beta_s\}$, respectively) are EF1 for each agent to find a rounding option that works for everyone. We call agent (g, i) *unhappy* with α_g if they prefer $\alpha_{g'}$ over it (where g' is the other group) and unhappy with β_g if they prefer $\beta_{g'}$. The following two observations follow from arguments similar to the one of the previous case:

- (A) If an agent (g, i) is unhappy with both α_g and β_g , they are EF1 for all rounding options except the natural one.
- (B) If an agent is happy with at least one of α_g or β_g , they are EF1 for the natural rounding and at least one other rounding option.

Since any two agents have at least one EF1 option in common, we successfully find an EF1 allocation whenever at least one of the three agents is EF1 for all four rounding options.

If, finally, all three agents have some rounding option that is not EF1, none of the four rounding options discussed so far may work, but we have one more ace up our sleeve:

- (C) If an agent is not EF1 under some rounding option, then they become EF1 under the other three options if we *swap the integral parts* I_f and I_s .

Since this observation applies to all three agents, each of them rules out at most one of the four rounding options after swapping the integral parts, which leaves one that is EF1 for all of them.⁵ Since this allocation is also still EF1 for the set-aside agent $(f, 1)$, this establishes the claim. \square

⁵For the straight-forward EF polytope, swapping integral allocation parts does not overcome the rounding counterexample.

As the counting arguments do not refer to which group each agent is in, the same argument also shows the existence of EF1 allocations for two groups of sizes $(|f|, |s|) = (3, 1)$, left open by Kyropoulou, Suksompong, and Voudouris (2020), for the natural adaptation of the LP. We conclude:

Corollary 3.2. *When there are two groups with a total of four agents, a balanced EF1 allocation exists and can be computed in (weakly) polynomial time.*

3.2 EF1 Impossibility for Three or More Couples

Since we were able to guarantee EF1 existence for two couples, one may hope that this existence extends to any number of couples, an audacious hope that has not been contradicted by earlier papers. In the special case where the agents $(g, 1)$ across all groups g have identical valuations, Bu et al. (2024) show that EF1 allocations do exist, using a variant of envy-cycle elimination due to Barman and Biswas (2020).

In general, however, we find that EF1 allocation need no longer exist for three or more couples:

Theorem 3.3. *For $n \geq 3$ couples, there exist some instances that admit no EF1 allocation.*

Proof. We prove the claim for $n = 3$ and generalize to $n > 3$ in Appendix B of the full version. Consider the following instance with 3 couples called f, s, t and goods $\{1, 2, 3, 4, 5\}$:

valuation	1	2	3	4	5
u_{f1}	2	2	0	0	1
u_{f2}	0	0	2	2	1
u_{s1}	0	2	0	2	1
u_{s2}	2	0	2	0	1
u_{t1}	2	0	0	2	1
u_{t2}	0	2	2	0	1

Each agent has positive valuation for three goods: two with value 2 and one with value 1. The agent must receive at least one such good since, otherwise, some other group receives two or more of those goods, violating EF1.

For the sake of contradiction, suppose that an EF1 allocation exists. Since there are five goods and three groups, one group must receive a single good. Let this be the case for group f , w.l.o.g. by symmetry. Since this good must have positive value for both agents in the group, it must be good 5.

Because the remaining four goods are never liked by both agents in a group, the other two groups must receive two goods each. Consider the two goods given to group s . By construction, there must be some agent (g, i) (not necessarily in group s) for whom both of these goods have value 2, and for whose partner (g, i') both goods have value 0. If $g = s$, then (g, i') receives value 0 and must be envious. Otherwise, (g, i) envies group s by more than one good. \square

4 Proportionality

Since an EF1 allocation among couples may not exist, it is natural to ask whether the weaker axiom of PROP1 can be guaranteed instead. For n couples, Bu et al. (2024) establish the existence of a PROP- $O(\log n)$ allocation by iteratively bipartitioning the agents and applying, in each step, a rounding

argument for a fair allocation among two groups.⁶ In this section, we get much closer to the standard axiom of PROP1; for n couples, we achieve PROP1 for the first agent in each group and PROP2 for each second agent.

4.1 Almost PROP Allocations for Small Groups

In fact, the claim for couples follows from a general result for groups of arbitrary sizes, which shows the existence of an FPO allocation in which every agent (g, i) is PROP1.

We prove this existence using an algorithm based on the *iterative rounding* method (Jain 2001). This method has been widely used in combinatorial optimization, including fair allocation (Chakrabarty, Chuzhoy, and Khanna 2009; Nguyen, Peivandi, and Vohra 2016; Cembrano, Moraga, and Verdugo 2025). Our rounding is inspired by the algorithm of Shmoys and Tardos (1993) for the Generalized Assignment Problem.

Our algorithm maintains a sequence of fractional allocations, and iteratively freezes coordinates at 0 and 1 until it reaches a discrete allocation. The steps of the algorithm are most easily explained by considering a bipartite graph, whose nodes on one side are the goods $M' \subseteq M$ not yet discretely allocated and on the other side are a set of n groups \mathcal{G}' , obtained from \mathcal{G} by removing some agents. The set of edges $E \subseteq M' \times \mathcal{G}'$ denotes the *allowed assignments* by specifying, for each good, the groups that the good may still be allocated to. Initially, we have $M' := M, \mathcal{G}' := \mathcal{G}, E := M \times \mathcal{G}$, meaning that no goods have been allocated, no agents eliminated, and all possible assignments are allowed. Based on M', \mathcal{G}', E , and the bundles B_g of goods already discretely allocated to group g , we consider the following polytope, which describes the currently allowed fractional allocations that are PROP for all agents remaining in \mathcal{G}' :

$$\begin{aligned} \sum_{\alpha: (\alpha, g) \in E} x_{\alpha g} u_{gi}(\alpha) &\geq \frac{u_{gi}(M)}{n} - u_{gi}(B_g) \quad \forall g \in \mathcal{G}', i \in [|g|] \\ \sum_{g: (\alpha, g) \in E} x_{\alpha g} &= 1 \quad \forall \alpha \in M' \\ 0 \leq x_{\alpha g} &\leq 1 \quad \forall (\alpha, g) \in E. \end{aligned}$$

We refer to the three types of constraints, in order from top to bottom, as *agent* constraints, *good* constraints, and *edge* constraints. Our algorithm is based on the following lemma, which will guide the rounding procedure.

Lemma 4.1. *If $M' \neq \emptyset$, every BFS x^* of the polytope above satisfies at least one of the following two conditions:*

- (i) $x_{\alpha g}^* \in \{0, 1\}$ for some $(\alpha, g) \in E$, or
- (ii) $\sum_{\alpha: (\alpha, g) \in E} x_{\alpha g}^* \leq |g|$ for some nonempty group $g \in \mathcal{G}'$.

Proof. Fix a BFS x^* , and assume that Condition (i) does not hold. Since $|E|$ constraints must be tight, i.e., hold with equality, at a BFS, the number of agent and good constraints must be at least $|E|$: $\sum_{g \in \mathcal{G}'} |g| + |M'| \geq |E|$.

Furthermore, since all $x_{\alpha g}^*$ are fractional and each good $\alpha \in M'$ has a total incident weight of 1 by the good constraints, α must be incident to at least two edges. Hence, $\sum_{g \in \mathcal{G}'} |g| + |M'| \geq |E| \geq 2|M'|$, i.e., $\sum_{g \in \mathcal{G}'} |g| \geq |M'|$.

⁶This argument is much easier than our argument for EF1. Though EF and PROP are equivalent for two groups, PROP1 is weaker and easier to achieve through rounding than EF1.

By summing over all good constraints, we obtain that $\sum_{(\alpha,g) \in E} x_{\alpha g}^* = |M'|$, hence $\sum_{g \in \mathcal{G}'} |g| \geq \sum_{(\alpha,g) \in E} x_{\alpha g}^*$.

Suppose, for contradiction, that Condition (ii) were also violated. In this case, each group $g \in \mathcal{G}'$ would satisfy $\sum_{\alpha: (\alpha,g) \in E} x_{\alpha g}^* \geq |g|$ and this inequality would be strict for the nonempty groups.⁷ Summing up over all groups, we obtain $\sum_{(\alpha,g) \in E} x_{\alpha g}^* > \sum_{g \in \mathcal{G}'} |g|$, a contradiction. \square

In each iteration, our algorithm finds a BFS x^* for the polytope above, and then proceeds as below. In the first iteration, we specifically select a BFS representing an fPO allocation, say, by solving an LP that maximizes the sum of all agents' utilities over the polytope. Then:

1. We delete all edges (α, g) from E for which $x_{\alpha g}^* = 0$.
2. If $x_{\alpha g}^* = 1$ for some $(\alpha, g) \in E$, we discretely allocate α to g , remove α from M' and (α, g) from E .
3. We update \mathcal{G}' by removing the last agent of every group g for which Condition (ii) from Lemma 4.1 holds.

Since x^* (restricted to the remaining edges) remains feasible for the updated polytope, the polytope remains nonempty, so we can find a new BFS x^* and repeat the process from Step 1 until all goods are allocated. Since, by Lemma 4.1, each iteration removes either an agent or an edge, the algorithm terminates in polynomially many iterations.

We now state the main theorem and sketch its proof. We defer pseudocode for the algorithm and the formal proof to Appendix C of the full version.

Theorem 4.2. *In any group fair division instance with arbitrary group sizes, there exists an fPO allocation which is PROP i for every agent (g, i) where $g \in \mathcal{G}$ and $i \leq |g|$. This allocation can be computed in (weakly) polynomial time.*

Proof sketch. We have already argued that the algorithm makes progress and terminates in polynomially many iterations. Since each iteration involves solving a linear program and some polynomial computation, the total running time is polynomial. It remains to argue that the resulting allocation satisfies PROP i and fPO.

Proportionality. Fix an agent (g, i) . If the agent never gets eliminated, their agent constraint ensures that the allocation even satisfies PROP. Should the agent get eliminated in some iteration, it must hold that $\sum_{(\alpha,g) \in E} x_{\alpha g}^* \leq i$ for the BFS x^* of this iteration. Since x^* satisfies the agent's constraint, the bundle B_g already discretely allocated to g before this iteration satisfies $u_{gi}(B_g) + \sum_{\alpha: (\alpha,g) \in E} x_{\alpha g}^* u_{gi}(\alpha) \geq u_{gi}(M)/n$. Since $\sum_{\alpha: (\alpha,g) \in E} x_{\alpha g}^* u_{gi}(\alpha)$ is at most the value of the i most valuable goods outside of B_g , the agent is PROP i — even if the final allocation does not give their group any goods in addition to B_g .

Fractional Pareto Optimality. It is a classic result by Varian (1974) that a fractional allocation is fPO iff it maximizes a positively weighted sum of agent utilities. (We confirm that this equivalence persists in the group setting.) Observe that a fractional allocation maximizes the weighted sum of agent

⁷Some group is nonempty because $\sum_{g \in \mathcal{G}'} |g| \geq |M'| > 0$.

utilities with weights $w_{gi} > 0$ iff each good α is only allocated among groups $g \in \operatorname{argmax}_{g \in \mathcal{G}} \sum_{i \in [|g|]} w_{gi} u_{gi}(\alpha)$. It follows that, if the fractional allocation x is fPO and if, for another fractional allocation x' , $x_{g\alpha} = 0$ implies $x'_{g\alpha} = 0$ for all groups g and goods α , then x' is also fPO. This argument was previously used, for example, by Aziz, Moulin, and Sandmirskiy (2020) and Bai et al. (2022).

Since the first iteration of the algorithm starts with an fPO x^* , and the algorithm immediately removes all edges that were zero for x^* , the final allocation only allocates goods to groups who received a non-zero amount of this good in the initial fractional allocation. Hence, the final allocation found by the algorithm is fPO. \square

4.2 Possibility of PROP1 Allocations

Although we can only prove the existence of a PROP2 allocation among couples, we conjecture that PROP1 allocations exist for any number of couples.

Our conjecture is supported, in part, by a failure to find counter-examples by hand and with computer aid. More importantly, we were able to show the existence of PROP1 allocations for several special cases:

Theorem 4.3. *When each group $g \in \mathcal{G}$ has size 2, a PROP1 allocation is guaranteed and efficiently computable whenever one of the following conditions holds.*

- $m \leq 2n$.
- m divides n and $(g, 1)$ and $(g, 2)$ have opposite preference rankings over the goods for all $g \in \mathcal{G}$.
- All agents have binary valuations and approve the same number of goods.
- All agents have binary valuations and $n = 3$.

Since PROP1 is weaker than EF1, one might hope for the existence of PROP1 for even larger groups. However, PROP1 may fail to exist for groups of size three:

Theorem 4.4. *For $n \geq 5$ groups of three agents, PROP1 allocations need not exist, even when utilities are binary and the groups are all identical. Moreover, deciding whether a PROP1 (or EF1) allocation exists is NP-complete for groups of three agents, even for binary utilities.*

Proof sketch. We only give the counter-example for five groups of three agents here, and defer the rest to Appendix E of the full version. Consider an instance with goods $\{1, \dots, 9\}$ and five groups g , each of which has the following valuations:

valuation	1	2	3	4	5	6	7	8	9
u_{g1}	1	1	1	1	1	1	0	0	0
u_{g2}	1	1	1	0	0	0	1	1	1
u_{g3}	0	0	0	1	1	1	1	1	1

Each agent has a total valuation of 6 and her proportional share is $\frac{6}{5} > 1$. Hence, each agent must receive at least one good with value 1 to be PROP1. As there are 9 goods and 5 groups, one group receives only one good. By construction, this good has zero value for one of the agents in the group, implying a PROP1 allocation cannot be achieved. \square




Figure 2: Fraction of pairings for which fair allocations exist or are found by one of two algorithms, averaged over all Spliddit instances. Axioms imply axioms to their left. Error bars indicate 95% confidence intervals (bootstrapping).

5 Experiments

We now use real-world preference data to empirically examine how often fair allocations exist for practical allocation problems among couples and if our iterative-rounding algorithm exceeds its theoretical guarantees in practice. Our dataset consists of all allocation problems for indivisible goods (over individuals) submitted to the website *Spliddit* (Goldman and Procaccia 2014; Shah 2017) as of June 2025. To allow us to meaningfully group agents, we consider only Spliddit instances with at least four agents. The remaining data consists of 254 instances, whose number of agents ranges between 4 and 15 (median: 5) and whose number of goods ranges between 1 and 59 (median: 6). See Appendix F of the full version for more details on data and experiments.

We transform each Spliddit instance into fair allocation problems over couples by iterating over all partitions of agents into pairs (if the number of agents is odd, one agent remains on their own), considering 1000 random pairings if the number of pairings exceeds this number. Since the different pairings of the same Spliddit instance produce correlated observations, we do not treat them as independent datapoints. Instead, we calculate for each Spliddit instance the fraction of its pairings that satisfies some property (say, EF1), and report averages over these fractions of pairings. In Fig. 2, we display the average fractions of pairs for the existence of several fairness axioms, and for whether these axioms are achieved by two variants of our iterative rounding algorithm. In Appendix F of the full version, we show that the patterns remain similar when restricting to instances with many or few agents, or with many or few goods.

While our Theorem 3.3 shows that EF1 allocations do not exist for all fair allocation instances among couples, such allocations seem to exist for most practical problems. Strikingly, we find EF1 allocations (hence also PROPI allocations) for each of the over 13,000 instance–pairing combinations we study. We also tested the frequency of allocations satisfying EF and EFX,⁸ an axiom between EF and EF1, whose existence is a tantalizing open question (Caragiannis et al. 2019;

⁸An allocation $\{B_{g'}\}_{g' \in \mathcal{G}}$ is EFX for (g, i) if removing any good $\alpha \in B_{g'}$ with $u_{gi}(\alpha) > 0$ from $B_{g'}$ eliminates (g, i) 's envy.

Chaudhury, Garg, and Mehlhorn 2020) in the individual setting. As shown by the blue bars in Fig. 2, EFX exists for 96% of pairings on average, whereas EF is rarer at 44%.

In Section 4, we proposed a natural iterative-rounding algorithm with proportionality and efficiency guarantees. To test this algorithm's usefulness in practice, we apply it to the same datasets, and report the fraction of pairings for which the algorithm satisfies each fairness axiom. The orange bars in Fig. 2 (“remove all the agents”) represent a direct implementation of our algorithm. Though the algorithm only guarantees PROP2 for the second agents in the worst case, it satisfies PROPI almost always on our data (99% of pairings on average). The algorithm even finds EF1 (73% of pairings) and EFX (53%) reasonably often, though substantially less often than the existence of these axioms, which is to be expected since the algorithm does not avoid envy.

We also repeated the experiment with a variant of the iterative-rounding algorithm, in which we do not immediately eliminate the last agent from all groups satisfying condition (ii) of Lemma 4.1. Instead, we find the group with the lowest incident weight among eligible groups, and eliminate only a single agent, namely the remaining agent in this group with the largest utility from already discretely allocated goods. Heuristically, this might lead to fairer allocations by deferring when we drop the constraints of agents who have not yet reached proportionality.

The performance of this variant is shown in Fig. 2 by the green bar (“remove the best agent”). Eliminating only one agent per iteration leads to PROPI allocations on all our considered instances and pairings. For EF1, the variant increases the average fraction of pairings from 73% to 86%, an increase clearly beyond the confidence intervals of both estimates (see figure).

The change also moderately increases the fraction of EFX pairings from 53% to 60%, and has no discernible effect on EF.

In light of these improvements, it would be interesting to study in future work if starting from fPO allocations other than the one with maximum utilitarian welfare and using other heuristics for eliminating agents can lead to even better practical performance.

6 Conclusion

We studied the allocation of indivisible goods among small groups, a well-motivated setting that most prior works, due to their focus on asymptotic bounds in the group size, have left largely unexplored. For two couples and envy freeness, or for any number of small groups and proportionality, we showed that fairness axioms must not be relaxed by much more than in the individual setting to guarantee existence.

Though our hope of EF1 existence for arbitrary numbers of couples did not materialize, our work leaves open many possibilities for positive results. For example, we do not know if EF1 allocations exist for all allocation problems over couples with binary valuations, whether PROPI allocations exist for any number of couples with additive utilities (as we believe), or whether, say, it is possible to guarantee EF1 for one partner and PROPI for the other in each couple.

Acknowledgements

We thank Warut Suksompong for helpful discussions and for bringing the work of Bu et al. (2024) to our attention.

References

- Aziz, H.; Moulin, H.; and Sandomirskiy, F. 2020. A Polynomial-Time Algorithm for Computing a Pareto Optimal and Almost Proportional Allocation. *Operations Research Letters*, 48(5): 573–578.
- Bai, Y.; Feige, U.; Gözl, P.; and Procaccia, A. D. 2022. Fair Allocations for Smoothed Utilities. In *Proceedings of the ACM Conference on Economics and Computation (EC)*.
- Barman, S.; and Biswas, A. 2020. Fair Division Under Cardinality Constraints. arXiv:1804.09521.
- Bertsimas, D.; and Tsitsiklis, J. N. 1997. *Introduction to Linear Optimization*. Athena Scientific Series in Optimization and Neural Computation. Belmont, Mass: Athena Scientific.
- Bu, X.; Li, Z.; Liu, S.; Song, J.; and Tao, B. 2024. Fair Division with Allocator’s Preference. In Garg, J.; Klimm, M.; and Kong, Y., eds., *Web and Internet Economics*, 77–94. Cham: Springer.
- Bu, X.; Li, Z.; Liu, S.; Song, J.; and Tao, B. 2025. Fair Division with Allocator’s Preference. arXiv:2310.03475v2.
- Caragiannis, I.; Kurokawa, D.; Moulin, H.; Procaccia, A. D.; Shah, N.; and Wang, J. 2019. The Unreasonable Fairness of Maximum Nash Welfare. *ACM Transactions on Economics and Computation*, 7(3): 1–32.
- Caragiannis, I.; Larsen, K. G.; and Shyam, S. 2025. A new lower bound for multi-color discrepancy with applications to fair division. arXiv preprint arXiv:2502.10516.
- Cembrano, J.; Moraga, A.; and Verdugo, V. 2025. Near-feasible Fair Allocations in Two-sided Markets. In *Proceedings of the 26th ACM Conference on Economics and Computation*, EC ’25, 898–915. New York, NY, USA: Association for Computing Machinery. ISBN 9798400719431.
- Chakrabarty, D.; Chuzhoy, J.; and Khanna, S. 2009. On Allocating Goods to Maximize Fairness. In *2009 50th Annual IEEE Symposium on Foundations of Computer Science*, 107–116.
- Chaudhury, B. R.; Garg, J.; and Mehlhorn, K. 2020. EFX Exists for Three Agents. In *Proceedings of the ACM Conference on Economics and Computation (EC)*, 1–19.
- Freeman, R.; Sikdar, S.; Vaish, R.; and Xia, L. 2019. Equitable Allocations of Indivisible Goods. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, 280–286.
- Ghods, M.; Latifian, M.; Mohammadi, A.; Moradian, S.; and Seddighin, M. 2018. Rent Division Among Groups. In *Combinatorial Optimization and Applications: 12th International Conference, COCOA 2018, Atlanta, GA, USA, December 15-17, 2018, Proceedings*, 577–591. Berlin, Heidelberg: Springer.
- Goldman, J.; and Procaccia, A. D. 2014. Spliddit: Unleashing Fair Division Algorithms. *ACM SIGecom Exchanges*, 13(2): 41–46.
- Jafari, A.; and Alipour, S. 2017. On the Chromatic Number of Generalized Kneser Graphs. *Contributions to Discrete Mathematics*, 12(2): 69–76.
- Jain, K. 2001. A Factor 2 Approximation Algorithm for the Generalized Steiner Network Problem. *Combinatorica*, 21: 39–60.
- Jukna, S. 2011. *Intersecting Families*, 99–106. Berlin, Heidelberg: Springer.
- Khachiyan, L. G. 1980. Polynomial algorithms in linear programming. *USSR Computational Mathematics and Mathematical Physics*, 20: 53–72. English translation of: Zhurnal Vychislitel’noi Matematiki i Matematicheskoi Fiziki 20 (1980) 51–68.
- Kyropoulou, M.; Suksompong, W.; and Voudouris, A. A. 2020. Almost Envy-Freeness in Group Resource Allocation. *Theoretical Computer Science*, 841: 110–123.
- Lipton, R. J.; Markakis, E.; Mossel, E.; and Saberi, A. 2004. On Approximately Fair Allocations of Indivisible Goods. In *Proceedings of the ACM Conference on Economics and Computation (EC)*, 125–131.
- Manurangsi, P.; and Meka, R. 2025. Tight Lower Bound for Multicolor Discrepancy. arXiv preprint arXiv:2504.18489.
- Manurangsi, P.; and Suksompong, W. 2017. Asymptotic Existence of Fair Divisions for Groups. *Mathematical Social Sciences*, 89: 100–108.
- Manurangsi, P.; and Suksompong, W. 2022. Almost envy-freeness for groups: Improved bounds via discrepancy theory. *Theoretical Computer Science*, 930: 179–195.
- Nguyen, T.; Peivandi, A.; and Vohra, R. 2016. Assignment Problems with Complementarities. *Journal of Economic Theory*, 165: 209–241.
- Segal-Halevi, E.; and Nitzan, S. 2019. Fair cake-cutting among families. *Social Choice and Welfare*, 53(4): 709–740.
- Segal-Halevi, E.; and Suksompong, W. 2020. How to cut a cake fairly: A generalization to groups. *The American Mathematical Monthly*, 128(1): 79–83.
- Segal-Halevi, E.; and Suksompong, W. 2023. Cutting a cake fairly for groups revisited. *The American Mathematical Monthly*, 130(3): 203–213.
- Shah, N. 2017. Spliddit: Two Years of Making the World Fairer. *XRDS: Crossroads, The ACM Magazine for Students*, 24(1): 24–28.
- Shmoys, D. B.; and Tardos, E. 1993. An approximation algorithm for the generalized assignment problem. *Math. Program.*, 62(1–3): 461–474.
- Suksompong, W. 2018. Approximate maximin shares for groups of agents. *Mathematical Social Sciences*, 92: 40–47.
- Varian, H. R. 1974. Equity, envy, and efficiency. *Journal of Economic Theory*, 9(1): 63–91.