Partitioning Friends Fairly

Lily Li[1] , **Evi Micha**[1] , **Aleksandar Nikolov**[1] and **Nisarg Shah**[1]
[1]University of Toronto
{xinyuan,emicha,anikolov,nisarg}@cs.toronto.edu

Abstract

We consider the problem of partitioning n agents in an undirected social network into k almost equal in size (differing by at most one) groups, where the utility of an agent for a group is the number of her neighbors in the group. The core and envy-freeness are two compelling axiomatic fairness guarantees in such settings. The former demands that there be no coalition of agents such that each agent in the coalition has more utility for that coalition than for her own group, while the latter demands that no agent envy another agent for the group they are in. We provide (often tight) approximations to both fairness guarantees, and many of our positive results are obtained via efficient algorithms.

1 Introduction

The computer science department at University X is organizing a visit day for its newly admitted students. One of the most anticipated activity is the campus tour, during which the admitted students get to see the department they might one day join. Due to COVID-19 related capacity restrictions, the admitted students are divided into k separate tours. But more tours means the need for more volunteers. Luckily, n current graduate students have volunteered to help lead the tours. We want to partition them almost equally between the k tours so that all the admitted students have equal opportunity to socialize with the current students. However, the current students have developed friendships during their time at the university. We would like to ensure that each volunteer is assigned to a tour with as many of their friends as possible, so they have a good experience and will want to volunteer again next year.

In this paper, we introduce and study a model that captures such real-life applications. Specifically, we consider the problem of partitioning n agents into k almost equal-sized (either $\lfloor n/k \rfloor$ or $\lceil n/k \rceil$) groups, when the agents are connected via an undirected social network indicating *friendships*. An agent's utility for being part of a group is the number of her friends who are in that group.

Formally, this model sits within the *hedonic games* formalism in cooperative game theory with nontransferable utilities [Aziz and Savani, 2016]. Two compelling axiomatic guarantees that have received significant attention in this literature are *the core* [Gillies, 1953], which informally requires that there be no *deviating coalition* of agents such that each agent in the coalition has strictly more utility for the coalition than for her group in the given partition, and *envy-freeness* [George and Marvin, 1958], which informally requires that no agent receive strictly more utility when swapping places with another agent in the given partition. However, this literature typically does not impose any restriction on the partition (including on the number of groups it has). This would make our problem trivial because the grand coalition — a single group containing all agents — would trivially satisfy both the core and envy-freeness requirements. To study the core and envy-freeness meaningfully, this literature allows agents to have negative utility for other agents.

We are interested in the case where the utilities are non-negative. But we require there to be exactly k groups and the groups to be of approximately equal sizes.[1] This can be viewed as a multi-dimensional generalization of the *stable roommates problem* [Irving, 1985], in which the goal is to partition $2n$ agents between n rooms of capacity 2 each when agents have preferences over who they wish to have as a roommate. The core becomes a notion of stability: if a pair of agents prefer each other to their assigned roommates, they may actually deviate and rent a room by themselves. But the core is also applicable to contexts where groups cannot really deviate; in such contexts, it is often viewed as a group fairness notion [Fain *et al.*, 2018; Conitzer *et al.*, 2019], demanding that each group be treated at least as well as what it deserves, where deserve is defined based on how happy it can be on its own.

1.1 Our Results

We consider general graphs, and also focus on the special case where the social network forms a tree for both fairness guarantees. For the core, we study bicriteria approximations of the form (α, β)-core, where a deviating coalition must improve the utility of each of its members by more than a multiplicative factor of α and an additive factor of β.[2]

[1]In appendix, we consider imposing only the former restriction, allowing k arbitrarily-sized non-empty groups.

[2]That is, an agent with utility u must receive utility more than $\alpha u + \beta$ after deviating.

We show that a balanced k-partition in the $(2k-1, 0)$-core always exists when $n \geqslant k^2 + k$, and a k-partition in the $(1, k)$-core always exists when $n < k^2 + k$. We achieve all these upper bounds via efficient algorithms, and prove lower bounds showing that these guarantees are asymptotically the best possible. For trees, we show that we can find a balanced 2-partition in the core in polynomial time for $k = 2$, while the best approximation guarantee for larger k is $(1, 1)$-core.

Similarly, we consider an additive approximation of envy-freeness, EF-r, where an agent's utility cannot increase by more than r when swapping places with another agent. We make a connection to discrepancy theory [Chen *et al.*, 2014] to show that a balanced EF-$O(\sqrt{\frac{n}{k} \cdot \log k})$ partition always exists for any k, and it can be computed efficiently. We conjecture that even a balanced EF-2 partition always exists for any k. For trees, we show that a balanced EF partition does not always exist, and it is NP-complete to decide whether a instance admits a balanced EF partition, while we can always find a balanced EF-1 partition in polynomial time.

1.2 Related Work

Our work can be viewed as a hedonic game with symmetric, binary, additively separable preferences and with the restriction that the partition produced have exactly k almost equal-sized parts. As noted in the introduction, this is a generalization of the stable roommates problem of partitioning $2n$ agents into n pairs, where the widely studied notion of stability coincides with the core. In this problem, with asymmetric preferences a solution in the core does not always exist — unlike in the bipartite version, referred to as the *stable marriage problem*, in which it is guaranteed to exist [Gale and Shapley, 1962] — but can be found in polynomial time when it does [Irving, 1985]. When preferences are symmetric, however, a solution in the core always exists and can be found efficiently; for instance, one can repeatedly match and remove a pair of agents with the highest utility. The three-dimensional version of this problem — partitioning $3n$ agents into groups of size 3 each — has also received significant attention. In this case, even with symmetric additive preferences, a solution in the core may not exist [Arkin *et al.*, 2009], and checking whether it does is NP-hard [Chen and Roy, 2021]. However, if we further restrict the preferences to be binary, then McKay and Manlove [2021] show that a solution in the core always exists and can be found efficiently. Our problem can be seen as a multidimensional generalization of the roommate problem with symmetric binary additive preferences.

Envy-freeness has been studied recently in the hedonic games literature [Peters, 2016; Barrot and Yokoo, 2019], again with possibly negative utilities. Another concept similar to envy-freeness is *Nash-stability* [Bogomolnaia and Jackson, 2002; Olsen *et al.*, 2012], which requires that no agent be happier by *joining* another part (rather than by swapping places with an agent in another part).[3] In our graph theoretic framework, this is equivalent to asking that each node have at least as many neighbors in its own part as in any other part. This has been studied extensively in graph theory using terms such as satisfactory partitions [Bazgan *et al.*, 2010], friendly partitions [Aharoni *et al.*, 1990], and internal partitions [Ban and Linial, 2016], but under only the restriction that each part is non-empty. This problem is also studied in the case, where the parts are required to be of almost the same size [Bazgan *et al.*, 2010]. However, since such partitions do not always exist, this literature primarily focuses on the computational complexity of checking the existence of such partitions and approximating the most satisfactory partitions.

Instead, our focus is on providing worst-case guarantees on the necessary violation of envy-freeness, as is commonly done in the literature on fair resource allocation [Lipton *et al.*, 2004; Caragiannis *et al.*, 2019; Aziz *et al.*, 2019]. We make a connection to discrepancy theory [Chen *et al.*, 2014] to establish an $O(\sqrt{n})$ bound. In discrepancy theory, the goal is to distribute each agent's friends as evenly as possible between the parts, so that not only does an agent not have many more friends in another part than her own part, she also does not have many more friends in her own part than in any other part. The latter restriction, a flipped version of the satisfactory partition problem, has also been studied separately as the co-satisfactory or unfriendly partition problem [Aharoni *et al.*, 1990]. Manurangsi and Suksompong [2021] use discrepancy theory in a similar problem with n agents partitioned into k groups, but with the agents having utilities over goods being allocated to the groups, not over the other agents.

2 Preliminaries

For $t \in \mathbb{N}$, let $[t] = \{0, \ldots, t-1\}$. We consider a set $V = [n]$ of agents who are members of a social network. The network is represented by an undirected graph $G(V, E)$, where the agents are the nodes and an edge $(i, i') \in E$ indicates friendship between agents i and i'. This induces the utility function of agent i, denoted $u_i : V \to \{0, 1\}$, where $u_i(i') = 1$ if $(i, i') \in E$ and 0 otherwise. Let $N_G(i)$ denote the set of neighbors of agent i in G, i.e., $N_G(i) = \{i' \in V : (i, i') \in E\}$. We refer to $d_G(i) = |N_G(i)|$ as the degree of agent i. We omit G when it is clear from the context.

A *k-partition* of V is given by $X = (X_0, \ldots, X_{k-1})$, where $X_j \cap X_{j'} = \emptyset$ for all distinct $j, j' \in [k]$; $X_j \neq \emptyset$ for all $j \in [k]$; and $\cup_{j \in [k]} X_j = V$. We refer to an individual group X_j as a *part*. With slight abuse of notation, we denote by $X(i)$ the part X_j to which agent i belongs (i.e., $i \in X_j$). We assume that $n \geqslant k$, so a k-partition exists. A k-partition is called *balanced* if $\lfloor n/k \rfloor \leqslant |X_j| \leqslant \lceil n/k \rceil$ for all $j \in [k]$. The utility of agent i for $S \subseteq V$ is denoted by, with slight abuse of notation, $u_i(S)$. We assume that utilities are additive, i.e., $u_i(S) = \sum_{i' \in S} u_i(i') = |S \cap N(i)|$.

In this work, we focus on two fairness criteria. The first one is the *core* which, informally, requires that there be no group of agents (coalition) of size $\lfloor n/k \rfloor \leqslant |S| \leqslant \lceil n/k \rceil$ such that every agent in the coalition prefers to be in that coalition than in her own part; such a coalition is called "blocking".

Definition 1. Fix $\alpha \geqslant 1$ and $\beta \geqslant 0$. A coalition $S \subseteq V$ is called (α, β)-blocking for a balanced k-partition X if

$$u_i(S) > \alpha \cdot u_i(X(i)) + \beta$$

for every $i \in S$. A balanced k-partition X is said to be in the (α, β)-*core* if there is no (α, β)-blocking coalition S with

[3] The two differ only when the other part consists entirely of the agent's friends.

$\lfloor n/k \rfloor \leqslant |S| \leqslant \lceil n/k \rceil$. When $\alpha = 1$ and $\beta = 0$, we simply use the terms blocking coalition, and core.

Another fairness criterion we focus on is envy-freeness. Because we will often be able to provide approximate envy-freeness guarantee with a small additive error, we only focus on additive approximations in this case.

Definition 2. For $r \geqslant 0$, a balanced k-partition X is called *envy-free up to* r, denoted EFr or EF-r, if, for every pair of agents $i, i' \in V$, $u_i(X(i)) \geqslant u_i(X(i') \cup \{i\} \setminus \{i'\}) - r$. When $r = 0$, we simply refer to this as envy-freeness (EF).

For the proof techniques we plan to use, we need the following additional terminology. The *cut size* of a k-partition X, denoted $\mathrm{cut}(X)$, is the number of edges between its different parts, i.e., $\mathrm{cut}(X) = |\{(i, i') \in E : X(i) \neq X(i')\}|$. A balanced k-partition with the smallest cut size is called a *balanced min k-cut*. Note that

$$\mathrm{cut}(X) = \sum_{i \in V} (|N(i)| - u_i(X(i))) = 2|E| - \sum_{i \in V} u_i(X(i)).$$

Hence, balanced min k-cut also maximizes the social welfare among all balanced k-partitions. Some of our results show that such solutions also satisfy good approximations of the core. Give disjoint sets of nodes A and B, $E(A, B)$ denotes the set of edges with one endpoint in A and the other in B.

Let us also introduce standard graph theory terminology. We denote by K_n, K_{n_1,n_2} and K_{n_1,n_2,n_3} the complete undirected graph of n vertices; the complete bipartite graph with n_1 and n_2 vertices on the two sides; and the complete tripartite graph with n_1, n_2, and n_3 vertices on the three sides, respectively. We refer to $K_{1,n-1}$ as a *star*. Finally, P_n denotes a path graph with n vertices.

3 Core

In this section, we study balanced k-partitions in the (approximate) core. While for $k \geqslant 3$ below we show that the core is not always non-empty, we start by pointing it an interesting open question for $k = 2$:

Open Question 1. *Does every graph admit a balanced* 2*-partition in the core?*

Now, we show that the core can be empty when $k \geqslant 3$.

Theorem 1. When $k \geqslant 3$, there exists an instance in which no k-partition is in the $(\alpha, 0)$-core for any $\alpha \geqslant 1$, and there also exists an instance in which no k-partition is in the $(1, \beta)$-core for any $\beta < k - 2$.

Proof. Fix $k > 2$. For the first claim, consider a cycle with $n = k + 1 \geqslant 4$ nodes. Fix an arbitrary k-partition X. Note that X must consist of one part with two nodes and $k - 1$ parts with a single node each. Without loss of generality, let X_0 be the part with $|X_0| = 2$. Note that in a cycle of length at least 4, the size of the smallest maximal matching is at least 2. Hence, there must exist agents $i, i' \notin X_0$ that are connected by an edge. Since the coalition $\{i, i'\}$ is allowed to deviate, they can both go from receiving utility 0 to receiving utility 1, implying that X is not in the $(\alpha, 0)$-core for any $\alpha > 0$.

For the second claim, consider the complete graph K_n with $n \geqslant k \cdot (k - 1)$. Let X be any k-partition of this graph. Due to the pigeonhole principle, there exists $r^* \in [k]$ such that $|X_{r^*}| \geqslant n/k \geqslant k - 1$. Hence, the coalition $S = \cup_{r \in [k] \setminus \{r^*\}} X_j$ is allowed to deviate as $|S| \leqslant n - k + 1$. Since each X_r part of this coalition is non-empty, we have $u_i(S) \geqslant u_i(X(i)) + k - 2$ for each $i \in S$, implying that X is not in the $(1, \beta)$-core for any $\beta < k - 2$. ☐

While above we show that that when $k \geqslant 3$ core can be empty, these examples are somewhat unsatisfactory as they crucially rely on n not being divisible by k, which leads to another interesting open question:

Open Question 2. *Does every graph with n nodes admit a balanced k-partition in the core, if k divides n?*

Now, we provide algorithms for finding balanced k-partitions which are in the approximate core. We begin with the case of $k = 2$. We show that the $(2, 0)$-core is always non-empty, and in particular, contains every balanced min 2-cut.

Theorem 2. For $k = 2$, a balanced min 2-cut is in the $(2, 0)$-core.

Proof. Let $X = (X_0, X_1)$ be a balanced min 2-cut. Suppose for contradiction that there exists a $(2, 0)$-blocking coalition S of size $\lceil n/2 \rceil$ or $\lfloor n/2 \rfloor$. Let $X_0^* = X_0 \cap S$ and $X_1^* = X_1 \cap S$.

For each agent $i \in X_0^*$, $i \in S$ implies $u_i(S) > 2 \cdot u_i(X_0)$, which in turn implies $|N(i) \cap X_1^*| > 2 \cdot |N(i) \cap X_0 \setminus X_0^*|$. Summing over all $i \in X_0^*$, we obtain

$$E(X_0^*, X_1^*) > 2 \cdot E(X_0^*, X_0 \setminus X_0^*).$$

Similarly, for each agent $i \in X_1^*$, we have $|N(i) \cap X_0^*| > 2 \cdot |N(i) \cap X_1 \setminus X_1^*|$. Summing over all $i \in X_1^*$, we get

$$E(X_0^*, X_1^*) > 2 \cdot E(X_1^*, X_1 \setminus X_1^*).$$

Combining the two equations, we have

$$\begin{aligned} E(X_0^*, X_1^*) &> 2 \cdot \max\{E(X_0^*, X_0 \setminus X_0^*), \\ &\qquad\qquad E(X_1^*, X_1 \setminus X_1^*)\} \\ &\geqslant E(X_0^*, X_0 \setminus X_0^*) + E(X_1^*, X_1 \setminus X_1^*). \quad (1) \end{aligned}$$

Now, consider the balanced 2-partition $X' = (S, V \setminus S)$. We will show that $\mathrm{cut}(X) > \mathrm{cut}(X')$, which will contradict X being a balanced min 2-cut. We have

$$\begin{aligned} \mathrm{cut}(X) &= E(X_0, X_1) \\ &= E(X_0^*, X_1^*) + E(X_0^*, X_1 \setminus X_1^*) \\ &\quad + E(X_1^*, X_0 \setminus X_0^*) + E(X_0 \setminus X_0^*, X_1 \setminus X_1^*) \\ &\geqslant E(X_0^*, X_1^*) + E(X_0^*, X_1 \setminus X_1^*) \\ &\quad + E(X_1^*, X_0 \setminus X_0^*) \\ &> E(X_0^*, X_0 \setminus X_0^*) + E(X_1^*, X_1 \setminus X_1^*) \\ &\quad + E(X_0^*, X_1 \setminus X_1^*) + E(X_1^*, X_0 \setminus X_0^*) \\ &= \mathrm{cut}(X'), \end{aligned}$$

where the strict inequality uses Equation (1). This is the desired contradiction. ☐

While Theorem 2 is a strong existential result, it does not come with an efficient algorithm as finding a balanced min 2-cut (also known as the minimum bisection problem) is NP-hard [Garey and Johnson, 1979]. This leads to our next open problem:

Algorithm 1 Local Min-Cut

1: $X \leftarrow$ an arbitrary balanced k-partition
2: **while** true **do**
3: Build a directed graph $G' = (V', E')$ with $V' = V$ and $E' = \{(i, i') : u_i(X(i')) > u_i(X(i)) + 1\}$
4: **if** there is a cycle $(i_0, i_1, \ldots, i_{s-1}, i_0)$ in G' **then** {Shift the nodes along the cycle}
5: **for** $\ell \in [s]$ **do**
6: $X(i_\ell) \leftarrow X(i_\ell) \setminus \{i_\ell\}$
7: $X(i_{\{\ell+1 \mod s\}}) \leftarrow X(i_{\{\ell+1 \mod s\}}) \cup \{i_\ell\}$
8: **end for**
9: **else if** $\exists (i, i')$ s.t. $u_{i'}(X(i')) = 0$ and $u_i(X(i')) > u_i(X(i))$ **then**
10: **if** $(i, i') \notin E$ or $u_i(X(i')) > u_i(X(i)) + 1$ **then**
11: $X(i) \leftarrow X(i) \cup \{i'\} \setminus \{i\}$
12: $X(i') \leftarrow X(i') \cup \{i\} \setminus \{i'\}$
13: **end if**
14: **else**
15: break
16: **end if**
17: **end while**
18: **return** X

Open Question 3. *Can a balanced* 2*-partition in the* $(2, 0)$*-core be computed in polynomial time?*

If our goal is efficient computation, the next result shows that we can find a balanced 2-partition in the $(3, 0)$-core (and more generally, a balanced k-partition in the $(2k-1, 0)$-core) in polynomial time, if $n \geqslant k^2 + k$. In this case, we in fact show that every balanced min k-cut is in the $(2k-1, 0)$-core, but we can also use an efficient local search, presented as Algorithm 1, to obtain the same approximation guarantee.

Theorem 3. When $n \geqslant k^2 + k$, every balanced min k-cut is in the $(2k-1, 0)$-core, and Algorithm 1 returns a balanced k-partition in the $(2k-1, 0)$-core in polynomial time.

Proof. First, we show that Algorithm 1 terminates in polynomial time by arguing that $\mathrm{cut}(X)$ strictly decreases in every iteration of the while loop. If we find a cycle on Line 4, then during the cyclic shift of nodes along this cycle, each node gains at least 1 utility. Since the social welfare strictly increases, $\mathrm{cut}(X)$ strictly decreases. Similarly, if we find two agents i and i' such that i' has no neighbors in $X(i')$ but i has at least two more neighbors in $X(i')$ than in $X(i)$, then swapping i and i' also strictly decreases the cut size. Further, if i' is not a neighbor of i, then we only need i to have at least one more neighbor in $X(i')$ than in $X(i)$. Hence, in any case, $\mathrm{cut}(X)$ strictly reduces in every iteration of the while loop, resulting in termination in polynomial time.

Let X be either a balanced min k-cut or the output of Algorithm 1. Suppose for contradiction that there is a $(2k-1, 0)$-blocking coalition S of size $\lceil n/k \rceil$ or $\lfloor n/k \rfloor$. We first show the following lemma.

Lemma 1. *For* $i \in S$*, if* $u_i(S \cap X_j) \leqslant u_i(X(i)) + 1$ *for each* $j \in [k]$*, then* $u_i(X(i)) = 0$*.*

Suppose that there exists $i \in S$ such that $u_i(S \cap X_j) > u_i(X(i)) + 1$ for some $j \in [k]$. Let G' be the directed graph constructed from X according to Line 3 of Algorithm 1. Then, there must be an edge from i to every node in X_j in G', as $u_i(X(i)) + 1 < u_i(S \cap X_j) \leqslant u_i(X_j)$. Further, since $u_i(S \cap X_j) > 0$, $S \cap X_j \neq \emptyset$. Hence, i has an edge to some node in S in G'. Note that there can be no cycle in G': if X is the output of Algorithm 1, this would contradict the while loop terminating, and if X is a balanced min k-cut, a cyclic shift of nodes like in Algorithm 1 would reduce the cut size, which would be a contradiction. Since there is no cycle in G', consider the longest path in G' starting at i and only containing nodes in S. Suppose it is $(i, i_1, \ldots, i_t, i')$. Then, i' must satisfy the condition of Lemma 1, otherwise by the same reasoning as before, there would exist $j' \in [k]$ such that $S \cap X_{j'} \neq \emptyset$ and i' has edges to all nodes in $X_{j'}$ in G'. This would either lead to a cycle or a longer path in G' starting at i and only containing nodes in S, which is a contradiction. Since i' satisfies the condition of Lemma 1, we have $u_{i'}(X(i')) = 0$. We also have $u_{i_t}(X(i')) > u_{i_t}(X(i_t)) + 1$. If X is returned by Algorithm 1, we get a contradiction because Algorithm 1 would have continued by swapping i_t and i' in Line 9. If X is a balanced min k-cut, then swapping i_t and i' would reduce the cut size, which would also be a contradiction.

We have established that all $i \in S$ satisfy the condition from Lemma 1. Hence, $u_i(X(i)) = 0$ for all $i \in S$. However, since $n \geqslant k^2 + k$, we have $|S| \geqslant \lfloor n/k \rfloor \geqslant k + 1$, which implies that there exist $i_1, i_2 \in S$ with $X(i_1) = X(i_2)$, which contradicts $u_{i_1}(X(i_1)) = u_{i_2}(X(i_2)) = 0$. Hence, there is no such $(2k-1, 0)$-blocking coalition S. □

In the proof of Lemma 1, note that if we assumed the deviating coalition S to be a $(k, k-1)$-blocking coalition, then we would obtain a contradiction regardless of whether $u_i(X(i)) = 0$ or $u_i(X(i)) \geqslant 1$. Since the next part of the proof, which establishes that all $i \in S$ must satisfy the condition of Lemma 1, does not assume $n \geqslant k^2 + k$, we have that Algorithm 1 always finds a solution in the $(k, k-1)$-core. In particular, for $k = 2$, we can efficiently guarantee $(2, 1)$-core. Recall that Theorem 2 provides a slightly better guarantee of $(2, 0)$-core, but not in polynomial time.

Corollary 1. *For* $k = 2$*, Algorithm 1 returns a balanced* 2*-partition in the* $(2, 1)$*-core in polynomial time.*

Next, we show that the guarantee in Theorem 3 is almost tight, at least for balanced min k-cuts, when $k \geqslant 3$.

Theorem 4. For $k \geqslant 3$, there exists an instance with $n \geqslant k^2 + k$ in which some balanced min k-cut is not in the $(\alpha, 0)$-core for $\alpha < 2k - 2$.

Finally, we turn to the case of $n < k^2 + k$. Here, we first notice that one cannot obtain a purely multiplicative guarantee of the form $(\alpha, 0)$-core for any $\alpha \geqslant 1$. This again follows from Theorem 1. In these examples, we showed that one can always find a deviating coalition whose members go from receiving utility 0 to utility 1, preventing us from guaranteeing $(\alpha, 0)$-core for any $\alpha \geqslant 1$. Thus, we turn to additive approximations. We show that any balanced k-partition is in the $(1, k)$-core, while $(1, \beta)$-core may be empty for $\beta < k/2 - 2$.

Theorem 5. Assume $n \leqslant k^2 + k$. Every balanced k-partition is in the $(1, k)$-core, and there exists an instance in which no balanced k-partition is in the $(1, \beta)$-core for $\beta < k/2 - 2$. Further, if $k \geqslant 3$, there exists an instance in which no balanced k-partition is in the $(\alpha, 0)$-core for any $\alpha \geqslant 1$.

Proof. To see the positive result, note that any deviating coalition has size at most $\lceil n/k \rceil \leqslant k+1$. Hence, no agent can improve her utility by an additive factor of more than k when deviating. Hence, every balanced k-partition is trivially in the $(1, k)$-core.

Next, consider the graph G formed by $k+1$ disjoint cliques of size $k-1$ each, denoted $C_0, \ldots, C_k$. Hence, $n = k^2 - 1$. Let X be any balanced k-partition of G. First, we claim that there exists $\ell^* \in [k+1]$ such that $|C_{\ell^*} \cap X_j| \leqslant (k+1)/2$ for all $j \in [k]$. If this is not true, then for every $\ell \in [k+1]$, there exists at least one $j_\ell \in [k]$ with $|C_\ell \cap X_{j_\ell}| > (k+1)/2$. Note that such j_ℓ must be unique. Further, because $|X_j| \leqslant \lceil n/k \rceil \leqslant k+1$ for all $j \in [k]$, we must have $j_\ell \neq j_{\ell'}$ for distinct $\ell, \ell' \in [k+1]$. However, this is not possible as there are $k+1$ cliques but only k parts.

Now, for every agent $i \in C_{\ell^*}$, we have $u_i(X(i)) \leqslant (k-1)/2$. On the other hand, C_{ℓ^*} is a feasible deviating coalition as $|C_{\ell^*}| = k - 1 = \lfloor n/k \rfloor$. Further, for every $i \in C_{\ell^*}$, we have $u_i(C_{\ell^*}) = k - 2 \geqslant u_i(X(i)) + (k-3)/2$. Hence, C_{ℓ^*} is a $(1, k/2 - 2)$-blocking coalition, as desired. We have already argued the impossibility of deriving a multiplicative approximation. □

3.1 Trees

Next, we turn our attention to the special case when the network is a tree. Let us begin with $k = 2$. Recall that for general graphs, we left non-emptiness of the core as an open question, and proved that every balanced min 2-cut is in the $(2, 0)$-core. For trees, we can in fact show that every balanced min 2-cut is in the core. Moreover, the NP-hard problem of finding a balanced min 2-cut in general graphs is known to be polynomial-time solvable for trees [Jansen *et al.*, 2005].

Theorem 6. When $k = 2$ and the network is a tree, every balanced min 2-cut is in the core, and one such solution can be computed in polynomial time.

For $k \geqslant 4$, we show that the core can be empty. In fact, we cannot hope for a multiplicative approximation guarantee of the form $(\alpha, 0)$-core for any $\alpha \geqslant 1$. On the other hand, if we turn to additive approximations, we show that any balanced k-partition of a tree is naturally in the $(1, 1)$-core, which is the best we can hope for. We leave the case of $k = 3$ as an open question.

Theorem 7. Every balanced k-partition of a tree is in the $(1, 1)$-core. For $k \geqslant 4$, there exists a tree for which no balanced k-partition is in the $(\alpha, 0)$-core for any $\alpha \geqslant 1$.

4 Envy-Freeness

We now turn our attention to finding balanced k-partitions that are (approximate) envy-free. We start by showing that EF-1 cannot always be guaranteed.

Theorem 8. Even when $k = 2$, a balanced 2-partition that is EF-1 does not always exist.

To obtain non-trivial approximations to envy-freeness for higher values of k, that too via balanced partitions, we turn to the literature on discrepancy theory. Intuitively, we want to color the elements of a set using k colors such that each pre-determined subset has an approximately equal number of elements of each color. Formally, we are given a universe $\Omega = [n]$ and a set system $\mathcal{S} = \{S_0, \ldots, S_{m-1}\}$, where $S_i \subseteq [n]$ for each $i \in [m]$. The k-color discrepancy of a coloring $\chi : \Omega \to [k]$ on the set system $\mathcal{S}$ is defined as

$$\mathrm{disc}_k(\mathcal{S}, \chi) = \max_{j \in [k], i \in [m]} \left| |\chi^{-1}(j) \cap S_i| - |S_i|/k \right|.$$

The k-discrepancy of $\mathcal{S}$ is then the minimum k-color discrepancy over all χ: $\mathrm{disc}_k(\mathcal{S}) = \min_{\chi:\Omega \to [k]} \mathrm{disc}_k(\mathcal{S}, \chi)$. A celebrated result from this literature is that $\mathrm{disc}_k(\mathcal{S}) = O(\sqrt{\frac{n}{k}} \ln(km/n))$ for any set system $\mathcal{S}$ and a k-coloring achieving this bound can be computed in polynomial time [Chen *et al.*, 2014, Corollary 44].

In our setting, with $\Omega = V = [n]$, a k-coloring $\chi : \Omega \to [k]$ induces a k-partition X given by $X_j = \chi^{-1}(j)$ for all $j \in [k]$.[4] Further, if we consider the set system $\mathcal{S}$ where $S_i = N_G(i)$ for each $i \in [n]$ (i.e., with $m = n$), then we are guaranteed that agent i can have at most $2\mathrm{disc}_k(\mathcal{S}, \chi)$ more neighbors in any other part than in her own part, implying EF-$2\mathrm{disc}_k(\mathcal{S}, \chi)$. The above discrepancy bound then immediately yields the existence of a k-partition that is EF-$O(\sqrt{\frac{n}{k}} \ln k)$. However, this may not be balanced.

To fix this, we add another set $S_n = V$ to our set system; we now have $m = n + 1$, which does not asymptotically change the discrepancy bound. Now, we obtain a k-partition X that is also approximately balanced: $||X_j| - |X_{j'}|| = O(\sqrt{\frac{n}{k}} \ln k)$ for all $j, j' \in [k]$. By arbitrarily moving $O(\sqrt{\frac{n}{k}} \ln k)$ agents between parts, we can make it perfectly balanced, while only increasing the EF approximation by $O(\sqrt{\frac{n}{k}} \ln k)$. Thus, we get the following.

Theorem 9. For any $k \geqslant 2$, a balanced k-partition that is EF-$O(\sqrt{\frac{n}{k}} \ln k)$ is guaranteed to exist and can be computed in polynomial time.

For discrepancy, the aforementioned upper bound is known to be almost tight: there is a lower bound of $\Omega(\sqrt{n/k})$ [Chen *et al.*, 2014, Theorem 61]. However, for our "one-sided" envy-freeness guarantee, achieving a constant approximation remains an open question.

Open Question 4. *Does every graph admit a balanced k-partition that is EF-2, for all $k \geqslant 2$?*

4.1 Trees

While we proved that EF-1 cannot be achieved for general graphs (Theorem 8), for trees we show that for any instance there exists a balanced EF-1 partition. Let d denote the depth of the tree. Without loss of generality, suppose the tree is labelled as following. Agent 0 is at level 1, agent 1 is the left most node of level 2, agent 2 is the second leftmost node of level 1, and so on, while agent $n - 1$ is the rightmost node of level d. Algorithm 2 first colors the nodes of the tree in a

[4] Technically, we also need to ensure $X_j \neq \emptyset$, but this is guaranteed due to the discrepancy bound.

Algorithm 2 EF1 Trees

1: $\forall j \in [k], X_j \leftarrow \emptyset$;
Phase 1:
2: **for** $i \in N$ **do**
3: $X_{i \mod k} = X_{i \mod k} \cup i$
4: **end for**
Phase 2:
5: **for** $\ell = 2$ to d **do**
6: **for** every $i \in N$ with $level(i) = \ell$ that is envious for more than one agents **do**
7: $i' \leftarrow$ an arbitrary child of i such that $X(i') = X(p(i))$
8: $X(i') = X(i') \cup \{i\} \setminus \{i'\}$
9: $X(i) = X(i) \cup \{i'\} \setminus \{i\}$
10: **end for**
11: **end for**
12: return $X = (X_0, ..., X_{k-1})$

simple round-robin fashion to obtain EF-2 (in fact, it achieves a discrepancy bound of 2, whereby there are at most 2 more nodes of any color than of any other color), and then makes small edits to improve its guarantee to EF-1.

Theorem 10. For any $k \geqslant 2$, Algorithm 2 returns a balanced EF-1 k-partition for every tree, in polynomial time.

Proof. Suppose that at Line 6 of the algorithm, when $\ell = level(i)$, i is not envious for more than one agents. Then, when $\ell = level(i)+1$, a child of i may be moved to the same part with i, but no child of i that is assigned to the the same part with i is removed from it, while afterwards no neighbour of i is never moved to a different part. Hence, clearly, the partition remains EF-1 with respect to i.

Now, suppose i is envious for more than one agents. This means that before Line 5, $|X(i) \cap c(i)| = \lfloor |c(i)|/k \rfloor < |c(i,T)|/k$, and for some $i' \notin N(i)$, $|X(i') \cap c(i,T)| = \lceil |c(i,T)|/k \rceil$ and $X(i') = X(p(i))$. Then, i and one of her children that is assigned to $X(i')$ are swapped. Hence, i is currently assigned to the same group with at least $\lfloor |c(i,T)|/k \rfloor + 1$ of her neighbours while any other part still contains at most $\lceil |c(i,T)|/k \rceil$ neighbours of i. Thus, at Line 6 of the algorithm, when $\ell = level(i) + 1$, i is not envious for more than one agents, and by the same reasoning as above, we have that partition remains EF-1 with respect to i until the end of the algorithm. □

While the above algorithm efficiently computes a balanced EF-1 k-partition, this partition is not too desirable because it unnecessarily divides the friends of each agent between the different parts during the round-robin coloring; note that this coloring actually achieves a discrepancy bound of 2 . More concretely, in the appendix, we provide an instance in which a different balanced k-partition can provide strictly more utility to every agent. A more desirable partition with the same EF-1 guarantee is achieved via balanced min k-cut. While this is NP-hard to compute in general graphs even for $k = 2$, for trees, it is efficiently computable when $k = 2$, but NP-hard when k is part of the input [Fernandes *et al.*, 2015]. Recall

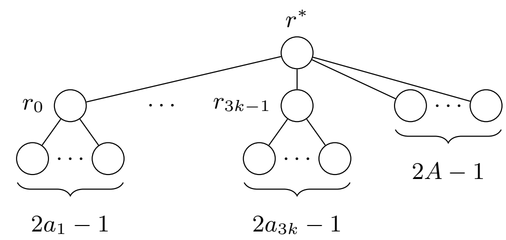

Figure 1: An example of $T = (r^*, G_I)$ given an instance I of 3-Partition problem

that this partition minimizes the cut size and, hence, maximizes the social welfare.

Theorem 11. For any $k \geqslant 2$ and when the network is a tree, every balanced min k-cut is EF-1.

Finally, we consider the complexity of checking if a balanced EF k-partition exists in a given tree. We show that this is NP-hard when k is part of the input[5].

Theorem 12. Checking if a given tree admits a balanced EF k-partition is NP-complete when k is part of the input.

5 Discussion

In this paper, we considered the problem of partitioning n agents into k almost equal-sized groups, when the agents have binary preferences, induced by a social network. We designed algorithms which approximately satisfy two axiomatic fairness guarantees: the core and envy-freeness.

Our work offers a number of exciting open questions. For example, is the core always non-empty when $k = 2$ or when k divides n? Does a balanced EF-2 partition always exist? We pay special attention to the case of trees in this paper; one can also consider other prominent graph families, such as planar graphs or graphs with bounded maximum degree.

There are two natural ways to extend our model. First, in our model, agents have symmetric binary preferences. One can consider preferences which are asymmetric and/or non-binary. Second, in our model, agents only have preferences over other agents; the groups they are assigned to are apriori identical. A complementary model in the fair division literature considers assigning resources to groups of agents [Segal-Halevi and Suksompong, 2019; Kyropoulou *et al.*, 2020; Manurangsi and Suksompong, 2021], where agents have preferences over the resources, but not over the other agents in their group. An extension of both models would require partitioning n agents into k groups and then allocating resources to these groups, when agents have preferences over both the resources being allocated and the other agents in their group. Such a model would better reflect many real-world scenarios, such as teachers partitioning students into groups and then assigning project ideas to the groups.

[5]When k is a constant, the problem can be solved efficiently via dynamic programming.

References

[Aharoni *et al.*, 1990] Ron Aharoni, Eric C Milner, and Karel Prikry. Unfriendly partitions of a graph. *Journal of Combinatorial Theory, Series B*, 50(1):1–10, 1990.

[Arkin *et al.*, 2009] Esther M Arkin, Sang Won Bae, Alon Efrat, Kazuya Okamoto, Joseph SB Mitchell, and Valentin Polishchuk. Geometric stable roommates. *Information Processing Letters*, 109(4):219–224, 2009.

[Aziz and Savani, 2016] H. Aziz and R. Savani. Hedonic games. In F. Brandt, V. Conitzer, U. Endress, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 15. Cambridge University Press, 2016.

[Aziz *et al.*, 2019] Haris Aziz, Ioannis Caragiannis, Ayumi Igarashi, and Toby Walsh. Fair allocation of indivisible goods and chores. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 53–59, 2019.

[Ban and Linial, 2016] Amir Ban and Nati Linial. Internal partitions of regular graphs. *Journal of Graph Theory*, 83(1):5–18, 2016.

[Barrot and Yokoo, 2019] Nathanaël Barrot and Makoto Yokoo. Stable and envy-free partitions in hedonic games. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 67–73, 2019.

[Bazgan *et al.*, 2007] Cristina Bazgan, Zsolt Tuza, and Daniel Vanderpooten. Efficient algorithms for decomposing graphs under degree constraints. *Discrete Applied Mathematics*, 155(8):979–988, 2007.

[Bazgan *et al.*, 2010] C. Bazgan, Z. Tuza, and D. Vanderpooten. Satisfactory graph partition, variants, and generalizations. *European Journal of Operational Research*, 206(2):271–280, 2010.

[Bogomolnaia and Jackson, 2002] A. Bogomolnaia and M.O. Jackson. Stability of hedonic coalition structures. *Games and Economic Behavior*, 38(2):201–230, 2002.

[Caragiannis *et al.*, 2019] I. Caragiannis, D. Kurokawa, H. Moulin, A. D. Procaccia, N. Shah, and J. Wang. The unreasonable fairness of maximum Nash welfare. *ACM Transactions on Economics and Computation (TEAC)*, 7(3):1–32, 2019.

[Chen and Roy, 2021] Jiehua Chen and Sanjukta Roy. Euclidean 3d stable roommates is np-hard. arXiv:2108.03868, 2021.

[Chen *et al.*, 2014] William Chen, Anand Srivastav, Giancarlo Travaglini, et al. *A panorama of discrepancy theory*. Springer, 2014.

[Conitzer *et al.*, 2019] V. Conitzer, R. Freeman, N. Shah, and J. Wortman-Vaughan. Group fairness for the allocation of indivisible goods. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, pages 1853–1860, 2019.

[Fain *et al.*, 2018] B. Fain, K. Munagala, and N. Shah. Fair allocation of indivisible public goods. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*, pages 575–592, 2018.

[Fernandes *et al.*, 2015] Cristina G Fernandes, Tina Janne Schmidt, and Anusch Taraz. Approximating minimum k-section in trees with linear diameter. *Electronic Notes in Discrete Mathematics*, 50:71–76, 2015.

[Gale and Shapley, 1962] D. Gale and L. S. Shapley. College admissions and the stability of marriage. *Americal Mathematical Monthly*, 69(1):9–15, 1962.

[Garey and Johnson, 1979] M. R. Garey and D. S. Johnson. *Computers and Intractability: a Guide to the Theory of NP-Completeness*. W. H. Freeman and Company, 1979.

[George and Marvin, 1958] Gamow George and Stern Marvin. Puzzle-math, 1958.

[Gillies, 1953] Donald Bruce Gillies. *Some theorems on n-person games*. Princeton University, 1953.

[Irving, 1985] Robert W Irving. An efficient algorithm for the "stable roommates" problem. *Journal of Algorithms*, 6(4):577–595, 1985.

[Jansen *et al.*, 2005] Klaus Jansen, Marek Karpinski, Andrzej Lingas, and Eike Seidel. Polynomial time approximation schemes for max-bisection on planar and geometric graphs. *SIAM Journal on Computing*, 35(1):110–119, 2005.

[Kyropoulou *et al.*, 2020] Maria Kyropoulou, Warut Suksompong, and Alexandros A Voudouris. Almost envy-freeness in group resource allocation. *Theoretical Computer Science*, 841:110–123, 2020.

[Lipton *et al.*, 2004] R. J. Lipton, E. Markakis, E. Mossel, and A. Saberi. On approximately fair allocations of indivisible goods. In *Proceedings of the 6th ACM Conference on Economics and Computation (EC)*, pages 125–131, 2004.

[Manurangsi and Suksompong, 2021] Pasin Manurangsi and Warut Suksompong. Almost envy-freeness for groups: Improved bounds via discrepancy theory. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 335–341, 2021.

[McKay and Manlove, 2021] M. McKay and D. Manlove. The three-dimensional stable roommates problem with additively separable preferences. In *Proceedings of the 14th International Symposium on Algorithmic Game Theory (SAGT)*, 2021. Forthcoming.

[Olsen *et al.*, 2012] Martin Olsen, Lars Bækgaard, and Torben Tambo. On non-trivial nash stable partitions in additive hedonic games with symmetric 0/1-utilities. *Information Processing Letters*, 112(23):903–907, 2012.

[Peters, 2016] Dominik Peters. Graphical hedonic games of bounded treewidth. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, pages 586–593, 2016.

[Segal-Halevi and Suksompong, 2019] Erel Segal-Halevi and Warut Suksompong. Democratic fair allocation of indivisible goods. *Artificial Intelligence*, 277:103167, 2019.

[Steinhaus, 1948] H. Steinhaus. The problem of fair division. *Econometrica*, 16:101–104, 1948.

[Stiebitz, 1996] Michael Stiebitz. Decomposing graphs under degree constraints. *Journal of Graph Theory*, 23(3):321–324, 1996.

Appendix

A Proof of Lemma 1

Proof. Suppose there exists $i \in S$ with $u_i(S \cap X_j) \leqslant u_i(X(i)) + 1$ for each $j \in [k]$ but $u_i(X(i)) \geqslant 1$. Then,

$$\begin{aligned} u_i(S) &= \textstyle\sum_{j \in [k]} u_i(S \cap X_j) \\ &\leqslant (k-1)(u_i(X(i)) + 1) + u_i(X(i)) \\ &\leqslant 2(k-1) \cdot u_i(X(i)) + u_i(X(i)) \\ &= (2k-1) \cdot u_i(X(i)), \end{aligned}$$

which contradicts S being a $(2k-1, 0)$-blocking coalition. □

B Proof of Theorem 4

Proof. Consider a graph G which consists of two disjoint cliques of size $2k-3$ each, denoted by A_1 and A_2; $k-1$ further disjoint cliques of size $2k-4$ each, denote by $B_1, \dots, B_{k-1}$; and another disjoint clique of size 2, denoted by C. Note that $n = 2(2k-3)+(k-1)(2k-4)+2 = 2k^2 - 2k \geqslant k^2 + k$ for $k \geqslant 3$. We start with the following lemma.

Lemma 2. *If X is a balanced min k-cut of G such that for some $j^* \in [k]$, there exist $V \subseteq X_{j^*}$ and a clique $V' \subseteq \cup_{j \in [k]} X_j \setminus X_{j^*}$ with $E(V', X_j \setminus V') = 0$ for every $j \in [k] \setminus j^*$, $E(V', V) = 0$, $E(X_{j^*} \setminus V, V') \geqslant E(X_{j^*} \setminus V, V)$, and $|V| = |V'|$, then swapping the nodes between V and V', using an arbitrary bijection, does not increase the cut size.*

Proof. Since V' is a clique, we easily see that the edges with two endpoints in different parts, except for part X_{j^*} are not increased. On the other hand, as $E(V', X_j \setminus V') = 0$ for any $j \in [k] \setminus j^*$, $E(V', V) = 0$ and $E(X_{j^*} \setminus V, V') \geqslant E(X_{j^*} \setminus V, V)$, we see that all the edges with one endpoint to X_{j^*} and the other endpoint to another part are not increased. Hence, cut is not increased. □

Let $X = (X_0, ..., X_{k-1})$ be an arbitrary balanced min k-cut of G. Suppose that the nodes of A_1 are spread among different parts. Then, there exists a part X_{j_1} that contains at least two nodes of A_1, as $2k-3 > k$ when $k > 3$. Let $S_1 = X_{j_1} \cap A_1$ and $V = X_{j_1} \setminus (S_1 \cup \{i_1\})$ where i_1 is an arbitrary node in $X_{j_1} \setminus S_1$ (such a node always exists as $n/k > |A_1| - 1 = 2k-4$). Notice that $\overline{S}_1 = A_1 \setminus S_1$ is a clique such that $E(\overline{S}_1, X_j \setminus \overline{S}_1) = 0$ for any $j \in [k] \setminus j_1$ and $E(\overline{S}_1, V) = 0$. Moreover, notice that $|\overline{S}_1| = |V|$ as $|A_1| = 2k-3$ and $|X_{j_1}| = 2k-2$. We also see that

$$\begin{aligned} &E(S_1 \cup \{i_1\}, X_{j_1} \setminus (S_1 \cup \{i_1\})) = \\ &E(S_1, X_{j_1} \setminus (S_1 \cup \{i_1\})) + E(i_1, X_{j_1} \setminus (S_1 \cup \{i_1\})) \leqslant 0 + |V| \end{aligned}$$

while

$$E(S_1 \cup \{i_1\}, \overline{S}_1) \geqslant 2 \cdot |\overline{S}_1|$$

as $|S_1| \geqslant 2$ and all the agents in S_1 are connected with all the agents in $\overline{S}_1$. From Lemma 2, we get that if we swap the nodes among V and $\overline{S}_1$, using an arbitrary mapping, the cut of the partition is not increased. Hence, there exists a balanced min k-cut in which all the nodes in A_1 are assigned to the same part X_{j_1}. Given this partition, suppose that the nodes of A_2 are spread among different parts. Then, there exists a part X_{j_2} different than X_{j_1} that contains at least two nodes of A_2, and by a similar argument as above we conclude in a balanced min k-cut in which all the nodes in A_1 are assigned to X_{j_1} and all the nodes in A_2 are assigned to X_{j_2}. Now, starting from this partition, suppose that the nodes of B_1 are spread among different parts. Then, as $|B_1| = 2k-4$, there exists a part $X_{j'_1}$ different that X_{j_1} and X_{j_2} which contains at least two nodes from B_1, when $k > 2$. Hence, by the same arguments as above we can conclude in a balanced min k-cut, in which all the nodes in A_1 are assigned to X_{j_1}, all the nodes in A_2 are assigned to X_{j_2}, and all the nodes in B_1 are assigned to $X_{j'_1}$. By continuing this way and as $|B_\ell| = 2k-4$ for any $\ell \in \{1, ..., k-1\}$, we conclude that there exists a balanced min k-cut, in which all the nodes in A_1 are assigned to X_{j_1}, all the nodes in A_2 are assigned to X_{j_2}, and all the nodes in B_ℓ, for $\ell \in \{1, ..., k-3\}$, are assigned to $X_{j'_\ell}$. Now, the last part different that X_{j_1}, X_{j_2} and $X_{j'_\ell}$ for $\ell \in \{1, ..., k-3\}$, denoted by $X_{j'_{k-2}}$, contains at least two nodes from B_{k-2} or B_{k-1}. Without loss of generality, we assume that $|X_{j'_{k-2}} \cap B_{k-2}| \geqslant 2$. By doing the same arguments as before, we find a balanced min k-cut partition $X' = (X'_0, ..., X'_{k-1})$ in which all the nodes of each clique A_1, A_2, B_1,...,B_{k-2} are assigned to the same group. Let $Q = A_1 \cup A_2 \cup B_1 \cup ... \cup B_{k-2}$. Then notice that for each $j \in [k]$, $|X'_j \cap Q| \geqslant 2k-4$. Hence, the nodes of B_{k-1} and C are spread across different parts and each part contains at most two nodes of B_{k-2}.

Now, we can see that if $C = \{c_1, c_2\}$ and $B_{k-1} = \{b_1, \dots, b_{2k-4}\}$, then X'' given by $X''_0 = A_1 \cup \{c_1\}$, $X''_1 = A_2 \cup \{c_2\}$, and $X''_{j+1} = B_j \cup \{b_{2j-1}, b_{2j}\}$ for $j \in \{1, \dots, k-2\}$ is a balanced min k-cut. But then, the coalition $S = C \cup B_{k-1}$ is a $(2k-3, 0)$-blocking coalition, as the utility of each agent in C increases by an infinite multiplicative factor when deviating with S, while that of each agent in B_{k-1} increases by a multiplicative factor of $2k-3$ when deviating with S. □

C Proof of Theorem 6

Proof. Let $X = (X_0, X_1)$ be a balanced min 2-cut. For the sake of contradiction, assume that there exists a blocking coalition S; we do not even need S to be of size $\lceil n/2 \rceil$ or $\lfloor n/2 \rfloor$ to derive a contradiction.

Let $X_0^* = X_0 \cap S$ and $X_1^* = X_1 \cap S$. Notice that for each agent $i \in X_0^*$, we have $u_i(S) \geqslant u_i(X_0) + 1$, which implies that $u_i(X_1^*) \geqslant u_i(X_0 \setminus S) + 1$. Summing over all $i \in X_0^*$, we have that $|E(X_0^*, X_1^*)| \geqslant |E(X_0^*, X_0 \setminus S) + |X_0^*|$.

Similarly, for each agent $i \in X_1^*$, we have $u_i(X_0^*) \geqslant u_i(X_1 \setminus S) + 1$. Summing over all $i \in X_1^*$, we have $|E(X_0^*, X_1^*)| \geqslant |E(X_1^*, X_1 \setminus S)| + |X_1^*|$.

Adding the two equations together, and noting that $|X_0^*| + |X_1^*| = |S|$, we obtain

$$\begin{aligned} 2 \cdot |E(X_0^*, X_1^*)| \geqslant\ & E(X_0^*, X_0 \setminus S) \\ & + E(X_1^*, X_1 \setminus S) + |S|. \end{aligned} \tag{2}$$

Notice that $X' = (S, V \setminus S) = (X_0^* \cup X_1^*, (X_0 \setminus S) \cup (X_1 \setminus S))$ is also a balanced 2-partition. Since $X = (X_0, X_1)$ is a balanced min 2-cut, we have

$$\begin{aligned} 0 &\leqslant \mathrm{cut}(X') - \mathrm{cut}(X) \\ &= E(X_0^*, X_0 \setminus S) + E(X_1^*, X_1 \setminus S) \\ &\quad - E(X_0^*, X_1^*) - E(X_0 \setminus S, X_1 \setminus S) \\ &\leqslant E(X_0^*, X_0 \setminus S) + E(X_1^*, X_1 \setminus S) - E(X_0^*, X_1^*) \\ &\leqslant |E(X_0^*, X_1^*)| - |S|, \end{aligned}$$

where the final step uses Equation (2).

Hence, we have that $|E(X_0 \cap S, X_1 \cap S)| \geqslant |S|$. But since S is a forest, it can have at most $|S| - 1$ edges, which is the desired contradiction. □

D Proof of Theorem 7

Proof. Let X be any k-partition of a tree. Suppose for contradiction that there exists a $(1,1)$-blocking coalition S. Note that S is a subgraph of a tree, so it must be a forest. Hence, there exists a leaf $i \in S$ with $u_i(S) \leqslant 1$, which contradicts S being a $(1,1)$-blocking coalition.

Now, consider $G = (V, E)$ with $V = \{r, a_1, a_2, b_1, b_2, \ldots, b_{k-2}\}$ and $E = \{(r, a_1), (r, a_2), (a_1, b_1), \cup_{\ell \in \{2,\ldots,k-2\}} (a_2, b_\ell)\}$ as shown in Figure 3 for the lower bound. Note that $n = k + 1$. Let X be any k-partition. Note that it must consist of $k - 1$ parts with a single node each and one part with two nodes. Without loss of generality, assume that $|X_0| = 2$. Like in the proof of Theorem 1, we notice that the smallest maximal matching in this graph has two edges. Hence, there must exist agents $i, i' \notin X_0$ that are connected by an edge. Since the coalition $\{i, i'\}$ is allowed to deviate, agents i and i' can go from receiving utility 0 to utility 1, implying that the partition cannot be in (α, β)-core for any $\alpha \geqslant 1$ and $\beta < 1$.

□

E Proof of Theorem 8

Proof. Consider the $K_{3,3,3}$ graph which consists of three set of three nodes each, denoted by $C_1 = \{c_{11}, c_{12}, c_{13}\}$, $C_2 = \{c_{21}, c_{22}, c_{23}\}$ and $C_3 = \{c_{31}, c_{32}, c_{33}\}$, respectively, and every node of each set is adjacent to every node in the other two sets.

For the sake of contradiction, assume that $X = (X_0, X_1)$ is a partition of the graph that is EF-1. Since the graph is 6-regular, we can see that $|X_0| \geqslant 4$ and $|X_1| \geqslant 4$, as if an agent i is assigned to a part with only at most two of its neighbours, then the other four of its neighbours are assigned to the other part along with an agent i' which is not neighbour of i, and then i envies i' for more than one agent. Without loss of generality, we assume that $|X_0| = 4$. If X_0 contains three nodes of the same set, then we can easily see that this partition is not EF-1, as each of them is assigned to the same group with at most one of its neighbours. As there are three sets and X_0 contains four agents, we see that two agents of the same set, say c_{11} and c_{12}, are assigned to X_0. Then these two agents are in the same part along with at most two of its neighbours, while all the remaining nodes are assigned to X_1. Then, c_{11} and c_{12} envy c_{13} for more than one agents, which is a contradiction. □

F Algorithm 2 is not weak PO

There are cases that while Algorithm 2 returns an EF-1 balanced k-partition $X = (X_1, X_2)$, there exists an EF-1 balanced k-partition $X' = (X_1', X_2')$ under which all the agents receive higher utility. In other words, the algorithm fails to provide *weak Pareto Optimality*. For $k \geqslant 3$, simply consider the path graph P_{2k}, while for $k = 2$, consider the instance shown in Figure 2. The numbers in the nodes illustrate the part that each of them is assigned to according to Algorithm 2, and the red and blue nodes illustrate a balanced EF1 partition in which *all* the agents receive higher utility.

G Proof of Theorem 11

Proof. Let $X = (X_0, \ldots, X_{k-1})$ be a balanced min-k cut. Suppose for contradiction that there exists an agent i that is envious for more than one agents. This means that there exists $X_j \neq X(i)$ such that $X_j \cap N(i) > X(i) \cap N(i) + 1$. Let $i' = \arg\max_{t \in X_j} \mathrm{level}(t)$, i.e. there is no other agent in X_j that is located in a higher level than i'. Hence, there is no child of i' in X_j. If we swap i and i', the movement of i' increases the number of edges that cross different parts by at most one, while the movement of agent i decreases the number of these edges by at least two. But, then $X = (X_0, \ldots, X_{k-1})$ would not be a balanced min-k cut which is a contradiction. □

H Proof of Theorem 12

Proof. We reduce from the 3-Partition problem defined as follows. Given $3k$ positive integers a_1,...,a_{3k} and a value A such that $A/4 < a_i < A/2$ for each $i \in [3k]$ and $\sum_{i\in[3k]} a_i = k \cdot A$. A 3-Partition instance admits a solution if the numbers can be partitioned into triples such that each triple adds up to A. Notice that as all the integers are positive, $A \geqslant 3$.

Given an instance I of 3-Partition problem, we construct a tree $G_I = (V_I, E_I)$ as follows. For each a_i, we construct a star with root r_i and $2a_i - 1$ leaves. Notice, that as a_i-s are positive integers $2a_i - 1 \geqslant 1$, and thus each r_i has at least one leaf adjacent to it. Moreover, we add a star with root r^* and $2A - 1$ leaves, and each r_i is connected with r^*. Thus, $|V_i| = 2(k+1)A$. If $T = (G_I, r^*)$, Figure 1 shows T given an instance I of 3-Partition problem.

We show that G_I admits an EF $k+1$-partition if and only if I admits a solution. If I admits a solution, then each r_i along all of its children are assigned to the same part with some $r_{i'}$, if a_i and $a_{i'}$ are assigned to the same triple under the solution of I, and $X_0 = \{\{r^*\} \cup (c(r^*, T) \setminus_{\cup i\in[3k]} \{r_i\})\}$. Each X_j for $j \in \{1, ..., 3k-1\}$ contains exactly three r_i-s. We claim that $X = (X_0, ...X_{k-1})$ is an EF k-partition. Indeed, each node that has as parent some r_j or r^* is assigned to the same group with its unique neighbour, each r_j is assigned to the same group with all of its children, and as each of them has at least one child, they cannot envy any node that is assigned to the same group with r^*, and since $A \geqslant 3$, r^* does not envy any node that is assigned to the same group with three r_j-s.

Now, assume that $X = (X_0, ..., X_{k-1})$ is an EF k-partition. We see that there exists $j \in [k]$ such that $X_j = \{\{r^*\} \cup (c(r^*, T) \setminus_{\cup i\in[3k]} \{r_i\})\}$, as otherwise some node in $c(r^*, T) \setminus_{\cup i\in[3k]} r_i$ is not assigned to the same group with r^*, and then the only way for the partition to be EF is if no other agent is assigned to the same part with r^*, which is not possible. Similarly, each r_j should be assigned to the same group with each of its children. Thus, for each r_i and $r_{i'}$ that are assigned to the same part if we assign a_i and a'_i to the same triple, we find a solution for I. ☐

I Imbalanced k-Partitions

In this section, we consider the case where the groups are required to only be non-empty rather than almost equal in size. Hence, we ask for k-partitions that provide good guarantees with respect to our two notions of fairness without the balancedness requirement.

I.1 Core

Recall that core requires that there be no group of agents (coalition) such that every agent in the coalition prefers to be in that coalition than in her own part. In general, there is no direct correlation between the size of a coalition and the ease with which it can be blocking.[6] Hence, in the imbalanced case, we impose the same restriction on the size of a deviating coalition as we have on the size of a part in a k-partition. Note that all parts in a k-partition X are required to be non-empty, which implies $1 \leqslant |X_j| \leqslant n-k+1$ for all $j \in [k]$; hence, we require a deviating coalition S to also satisfy $1 \leqslant |S| \leqslant n-k+1$. This gives rise to the following variant of the core.

Definition 3. Fix $\alpha \geqslant 1$ and $\beta \geqslant 0$. A coalition $S \subseteq V$ is called (α, β)-blocking for a k-partition X if

$$u_i(S) > \alpha \cdot u_i(X(i)) + \beta$$

for every $i \in S$. A k-partition X is said to be in the (α, β)-*imbalanced core* if there is no (α, β)-blocking coalition S with $1 \leqslant |S| \leqslant n-k+1$. When $\alpha = 1$ and $\beta = 0$, we simply use the terms blocking coalition, and imbalanced core.

Note that the differing size restrictions on the deviating coalitions technically makes our results for the core under imbalanced k-partitions incomparable to our results for the core under balanced k-partitions.

In this section, we study k-partitions in the (approximate) imbalanced core. First, we show that when $k = 2$, a 2-partition in the imbalanced core always exists and Algorithm 3 finds one such partition in polynomial time. When $k > 2$, the same algorithm finds a partition in the $(1, k-2)$-imbalanced core.

Theorem 13. When $k \geqslant 2$, Algorithm 3 finds k-partition in the $(1, k-2)$-imbalanced core in polynomial time. In particular, when $k = 2$, it efficiently finds a 2-partition in the imbalanced core.

Proof. For contradiction, assume that there is a blocking coalition S for the k-partition X computed by Algorithm 3 with $1 \leqslant |S| \leqslant n-k+1$ each of whose agents increased their utility by at least an additive factor of $k-1$.

First, we suppose that G is connected. Notice that every connected graph admits a spanning tree, and that the graph stays connected when deleting a leaf from this tree. Hence, Algorithm 3 is well-defined in this case, and we obtain the guarantee that X_{k-1} is a connected subgraph of G.

We claim that $S \cap X_{k-1} \neq \emptyset$ and $X_{k-1} \setminus S \neq \emptyset$. To see the former claim, note that if $|S| \leqslant k-1$, then $u_i(S) \leqslant k-2 \leqslant u_i(X(i)) + k-2$ for any $i \in S$, which would be a contradiction. Hence, we must have $|S| \geqslant k$, which implies $S \cap X_{k-1} \neq \emptyset$. To see the latter claim, note that $|S| \leqslant n-k+1$. Also, $|X_{k-1}| = n-k+1$. Thus, if $S \supseteq X_{k-1}$, then we would have $S = X_{k-1}$. This would imply $u_i(S) = u_i(X(i))$ for all $i \in S$, which would again be a contradiction. Hence, we must have $X_{k-1} \setminus S \neq \emptyset$.

Fix $i_1^* \in S \cap X_{k-1}$ and $i_2^* \in X_{k-1} \setminus S$. Because X_{k-1} is a connected subgraph of G, there exists a path from i_1^* to i_2^* using only the nodes in X_{k-1}. Consider the first edge of this path to travel out of S; say this edge is (i', i'') with $i' \in S \cap X_{k-1}$ and $i'' \in X_{k-1} \setminus S$. When deviating from X_{k-1} to S, agent i' loses at least one neighbor (namely i'') from X_{k-1} and may gain up to $k-1$ neighbors (the nodes in $\cup_{r\in[k-1]} X_r$). This implies $u_{i'}(S) \leqslant u_{i'}(X_{k-1}) + k - 2$, which is a contradiction.

Next, suppose G is not connected. Since no connected component can contain all nodes of G, the algorithm must have moved at most $k-2$ nodes from X_0 to $\cup_{t\in\{1,...,k-1\}} X_t$. Hence, none of the agents who are in X_0 in the final solution can join coalition S as their utility cannot improve by more than an additive factor of $k-2$ when doing so. Further, if there exists $i \in S \cap X_{k-1}$, then $u_i(S) \leqslant$

[6] Smaller coalitions have the advantage of only having to improve the utility of fewer agents, whereas larger coalitions can include more friends of their members.

Algorithm 3 (Approximate) Imbalanced Core

if G is connected **then**
  for $r = 0, \ldots, k-2$ **do**
    $i_r \leftarrow$ a leaf node in a spanning tree of $G \setminus \cup_{t \in [r]} X_t$
    $X_r \leftarrow \{i_r\}$
  end for
else
  $X_0 \leftarrow$ any connected component of G
  for $r = 1, \ldots, k-2$ **do**
    if $V \setminus \cup_{t \in [r]} X_t \neq \emptyset$ **then**
      $X_r \leftarrow \{i\}$, for an arbitrary $i \in V \setminus \cup_{t \in [r]} X_t$
    else
      $X_r \leftarrow \{i\}$, for an arbitrary $i \in X_0$
      $X_0 \leftarrow X_0 \setminus \{i\}$
    end if
  end for
end if
if $V \setminus \cup_{t \in [k-1]} X_t \neq \emptyset$ **then**
  $X_{k-1} = V \setminus \cup_{t \in [k-2]} X_t$
else
  $X_{k-1} = \{i\}$, for an arbitrary $i \in X_0$
  $X_0 \leftarrow X_0 \setminus \{i\}$
end if
return $X = (X_0, \ldots, X_{k-1})$

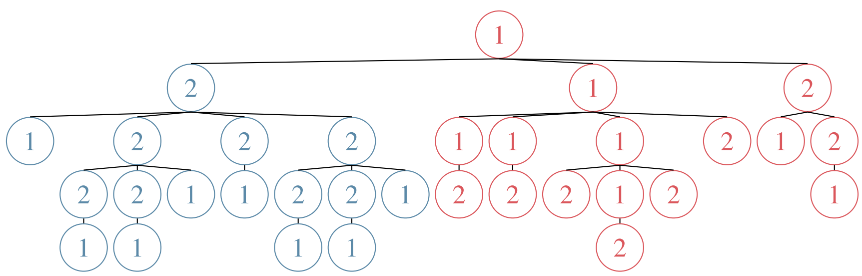

Figure 2: Instance that Algorithm 2 fails to provide weak Pareto Optimality, when k=2. Numbers illustrate the partition of Algorithm 2 and colors illustrate an EF-1 2-partition under which all agents improve their utility.

$u_i(X_{k-1}) + \sum_{r=1}^{k-2} u_i(X_r) \leqslant u_i(X_{k-1}) + k - 2$ (as we have already established $S \cap X_0 = \emptyset$), which is again a contradiction. Hence, we must have $S \subseteq \cup_{r \in \{1,\ldots,k-2\}} X_r$, implying that $|S| \leqslant k-2$. But then, $u_i(S) \leqslant k - 3 < u_i(X(i)) + k - 2$ for all $i \in S$, which is again a contradiction. □

For $k \geqslant 3$, notice that in the examples used in the proof of Theorem 1 to show the possible emptiness of the core, we had $n = k+1$. In this case, all k-partitions are balanced, and the imbalanced core becomes equivalent to the core. Hence, the imbalanced core can be empty when $k \geqslant 3$. This means that the approximation guarantee provided in Theorem 13 is tight in two ways. First, one cannot hope to achieve a multiplicative approximation guarantee of the form $(\alpha, 0)$-imbalanced core, for any $\alpha \geqslant 1$. Second, one also cannot hope to achieve a better additive approximation guarantee of the form $(1, \beta)$-imbalanced core for $\beta < k - 2$.

Trees

In this subsection, we consider the special case where the social network G is a tree. We show that this substantially improves how well we can approximate the imbalanced core. In particular, now a k-partition in the imbalanced core exists for $k \in \{2, 3\}$ (as opposed to just for $k = 2$ in the general case), and for $k > 3$, the best possible guarantee is the $(1,1)$-imbalanced core (as opposed to the $(1, k-2)$-imbalanced core in the general case).

We remark that all our techniques extend to the case of forests; we consider trees for ease of exposition.

Theorem 14. When $k \leqslant 3$ and the network is a tree, a k-partition in the imbalanced core always exists and can be found in polynomial time.

Proof. For $k = 2$, this follows from Theorem 13.

Let $k = 3$. Let (i_0, i_1) be a pair of nodes that are the farthest apart; note that both must be leaves. Let p_0 and p_1 be the unique neighbors of i_0 and i_1, respectively. Let $X_0 = \{i_0\}$, $X_1 = \{i_1\}$, and $X_2 = V \setminus \{i_0, i_1\}$. Clearly, this can be computed in polynomial time. Suppose for contradiction that this is not in the imbalanced core and S is a blocking coalition. Note that if $i \in S \setminus \{i_0, i_1\}$, then $i \in \{p_0, p_1\}$, otherwise $u_i(X(i)) = |N(i)|$, preventing i from gaining by deviating with S.

Let L be the path between i_0 and i_1, and $|L|$ denote the number of edges in this path. Note that $n \geqslant 3$ implies $|L| \geqslant 2$. If $|L| = 2$, then the graph is a star. It is easy to check that the center of the star (equal to both p_0 and p_1) cannot gain from joining any coalition S of size at most $n-2$, so there is no blocking coalition.

Next, suppose $|L| \geqslant 3$, so p_0 and p_1 are distinct. In particular, note that for each $t \in \{0, 1\}$, we have $u_{p_t}(X_2) = |N(p_t)| - 1$, so for p_t to deviate with S, we need $N(p_t) \subseteq S$.

When $|L| \geqslant 4$, each p_t has an adjacent node on L other than i_t and p_{1-t}. Since this node is adjacent to neither i_0 nor i_1, it is not in S. Hence, $p_0, p_1 \notin S$, which implies $i_0, i_1 \notin S$, which is a contradiction.

When $|L| = 3$, we have $L = (i_0, p_0, p_1, i_1)$. If p_0 and p_1 have no neighbors other than each other, i_0, or i_1, then the tree consists of only these four nodes; in this case, it is easy to see that the claimed partition is in the imbalanced core. Otherwise, without loss of generality, assume that p_0 has a neighbour $j \notin \{i_0, p_1\}$. From the previous argument, since j is not adjacent to i_0 or i_1, $j \notin S$. Since $u_{p_t}(X_2) = |N(p_t)| - 1$ for each $t \in \{0, 1\}$, this implies that p_0 would not join S, which in turn implies that p_1 would not join S. Then, $i_0, i_1 \notin S$ as well, which is a contradiction. Hence, the partition is in the imbalanced core.

□

Finally, for $k \geqslant 4$, we show that the best approximation we can guarantee is the $(1, 1)$-imbalanced core.

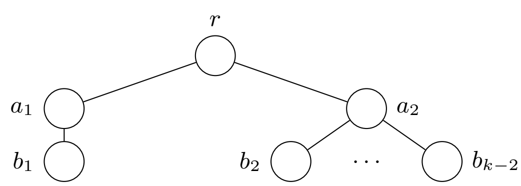

Figure 3: A tree in which no k-partition is in the $(\alpha, 0)$-imbalanced core for any $\alpha \geqslant 1$ and $k \geqslant 4$.

Theorem 15. Every k-partition of a tree is in the $(1, 1)$-imbalanced core. When $k \geqslant 4$, there exists a tree in which no k-partition is in the (α, β)-imbalanced core with any $\alpha \geqslant 1$ and $\beta < 1$.

Proof. The proof for the positive result follows the same reasoning that we used in the proof of Theorem 7 to argue that every balanced k-partition of a tree is in the $(1, 1)$-core. Since any deviating coalition S is a subgraph of the tree, there must be $i \in S$ with $u_i(S) \leqslant 1$. Hence S cannot be a $(1, 1)$-blocking coalition.

Let us turn to the negative result for $k \geqslant 4$. Recall that in the proof of Theorem 7, we provided an example tree in which any balanced k-partition admits a deviating coalition of size 2 whose members go from receiving utility 0 to utility 1. Since this example used $n = k + 1$, a deviating coalition of size 2 is also allowed under the imbalanced core. Hence, this example shows the impossibility of achieving $(\alpha, 0)$-imbalanced core for any $\alpha \geqslant 1$. □

I.2 Envy-Freeness

Finally, we turn our attention to envy-freeness. Luckily, the definition of envy-freeness does not require any modification to make it meaningful for imbalanced k-partitions.

First, we use the following result from the literature on satisfactory partition, stated in our framework, to establish the existence of an EF-2 partition when $k = 2$.

Theorem 16 (Stiebitz 1996, Bazgan *et al.* 2007)**.** Given a graph $G = (V, E)$ and functions $a, b : V \to \mathbb{N}$ such that $d(i) \geqslant a(i) + b(i) + 1$ for every $i \in V$, there exists a 2-partition $X = (X_0, X_1)$ of V such that $u_i(X_0) \geqslant a(i)$ for each $i \in X_0$ and $u_i(X_1) \geqslant b(i)$ for all $i \in X_1$, and it can be computed in polynomial time.

In our case, we use functions $a(i) = b(i) = \lfloor (d(i) - 1)/2 \rfloor$ for all $i \in V$. Note that these satisfy the condition $d(i) \geqslant a(i) + b(i) + 1$. Hence, the above result allows us to efficiently compute a 2-partition X satisfying $u_i(X(i)) \geqslant \lfloor (d(i) - 1)/2 \rfloor$ for all $i \in V$. Since there are only two parts, this also implies that for all $i, i' \in V$,

$$\begin{aligned} u_i(X(i')) - u_i(X(i)) &\leqslant d(i) - 2 \cdot \lfloor (d(i) - 1)/2 \rfloor \\ &\leqslant d(i) - 2 \cdot (d(i) - 2)/2 = 2, \end{aligned}$$

which implies that X is EF-2.

Theorem 17. A 2-partition that is EF-2 always exists and can be computed in polynomial time.

Theorem 16 admits an extension to $k > 2$ parts, but in our case, this only guarantees that $u_i(X(i)) \geqslant \lfloor (d(i) - k + 1)/k \rfloor$ for all $i \in V$ [Bazgan *et al.*, 2007]. This does not meaningfully limit the number of neighbors that agent i has in another part and, therefore, fails to provide a non-trivial approximation to envy-freeness. That said, if one is interested in the slightly weaker guarantee of proportionality [Steinhaus, 1948], which, in our setting, would require $u_i(X(i)) \geqslant d(i)/k$, then this would provide an additive 1-approximation.

For the satisfactory partition problem, where the goal is to indeed minimize $u_i(X(i')) - u_i(X(i))$, as in the equation above, it is easy to see that an additive error of 2 is the best possible. Consider dividing any clique with an odd number of nodes into two parts. An agent i in the smaller part will have at least two more neighbors in the larger part than in her own part. However, this does not hold for envy-freeness: if i envisions *swapping places* with an agent i' from the other part, then $X(i') \cup \{i\} \setminus \{i'\}$ will only contain one more neighbor of i than

$X(i)$ does. Nonetheless, notice that the example that is used in the proof of Theorem 8 can also be used to show that EF-1 cannot always be guaranteed even in the imbalanced case when $k = 2$.

However, if we restrict our attention to trees, we can achieve EF-1 even with balanced k-partitions for all $k \geqslant 2$ (see Theorem 10).