

Maximizing the Egalitarian Welfare in Friends and Enemies Games

Edith Elkind
Northwestern University
Evanston, United States
edith.elkind@northwestern.edu

Michele Flammini
Gran Sasso Science Institute
L’Aquila, Italy
michele.flammini@gssi.it

Giovanna Varricchio
University of Calabria
Rende, Italy
giovanna.varricchio@unical.it

ABSTRACT

We consider the complexity of maximizing egalitarian welfare in *Friends and Enemies Games*—a subclass of hedonic games in which every agent partitions other agents into *friends* and *enemies*. We investigate two classic scenarios proposed in the literature, namely, *Friends Appreciation* (FA) and *Enemies Aversion* (EA): in the former, each agent primarily cares about the number of friends in her coalition, breaking ties based on the number of enemies, while in the latter, the opposite is true. For EA, we show that our objective is hard to approximate within $O(n^{1-\epsilon})$, for any fixed $\epsilon > 0$, and provide a polynomial-time $(n - 1)$ -approximation. For FA, we obtain an NP-hardness result and a polynomial-time approximation algorithm. Our algorithm achieves a ratio of $2 - \Theta(\frac{1}{n})$ when every agent has at least two friends; however, if some agent has at most one friend, its approximation ratio deteriorates to $n/2$. We recover the $2 - \Theta(\frac{1}{n})$ approximation ratio for two important variants: when randomization is allowed and when the friendship relationship is symmetric. Additionally, for both EA and FA we identify special cases where the optimal egalitarian partition can be computed in polynomial time.

KEYWORDS

Hedonic Games, Friends and Enemies Games, Egalitarian Welfare

ACM Reference Format:

Edith Elkind, Michele Flammini, and Giovanna Varricchio. 2026. Maximizing the Egalitarian Welfare in Friends and Enemies Games. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 9 pages. <https://doi.org/10.65109/YNPD9144>

1 INTRODUCTION

Hedonic Games (HG) [16] provide a game-theoretic framework for analyzing coalition formation among selfish agents and have been extensively studied in the literature (see, e.g., [2, 3, 6, 8, 17, 19, 22]). In these games, the objective is to partition a set of agents into disjoint coalitions, with each agent’s satisfaction determined solely by the members of her coalition.

Different preference models give rise to different subclasses of HGs. For instance, in *additively separable* HGs (ASHGs) [8], agents assign values to each other and evaluate coalitions by summing the values they assign to every other member. A particularly natural

scenario is when agents classify others as either *friends* or *enemies*. Two canonical preference models have been studied in this setting [14]: under *Friends Appreciation* (FA) preferences, agents prioritize coalitions with more friends and, in case of ties, prefer those with fewer enemies; conversely, under *Enemies Aversion* (EA) preferences, agents prefer coalitions with the lowest number of enemies and, if the number of enemies is the same, favor those with more friends.

While most of the existing literature on HGs has focused on stability, i.e., the resilience of an outcome to individual or group deviations, another important concept in coalition formation is the *social welfare*, which captures the quality of a partition. Specifically, the *utilitarian* welfare, which is given by the sum of the agents’ utilities, measures the overall happiness of the agents in a given partition. However, maximizing utilitarian welfare may come at a cost of treating some agents badly if doing so increases the total sum. In contrast, the *egalitarian* welfare is determined by the minimum utility among all agents. Thus, a partition that maximizes the egalitarian welfare is fairer, in the sense that the lowest utility is as high as possible. Despite its appeal, egalitarian welfare has received limited attention in the context of hedonic games.

1.1 Our Contribution

We study the computation of the maximum egalitarian welfare in Friends and Enemies Games.

For EA games, we prove that the problem is hard to approximate within $O(n^{1-\epsilon})$ for any fixed $\epsilon > 0$, where n is the number of agents, and complement this hardness result with a polynomial-time algorithm achieving an $(n - 1)$ -approximation. For FA games, we show that maximizing egalitarian welfare is computationally hard. We then provide a polynomial-time approximation algorithm whose performance depends on the structure of the instance: it achieves an approximation of $2 - \Theta(\frac{1}{n})$, when every agent has at least two friends, and an $n/2$ -approximation in the worst case. We further improve this result in two important settings, namely, (1) when randomization is allowed, and (2) when the preferences are symmetric, that is, the friendship relations are mutual. For both cases, we design algorithms attaining an approximation ratio of $2 - \Theta(\frac{1}{n})$. Finally, we complement these results by identifying special cases of both EA games and FA games where the optimal egalitarian partition can be computed exactly in polynomial time.

1.2 Related Work

There is a substantial body of literature on hedonic games; we refer the interested reader to [5] for a comprehensive overview. In what follows, we focus on work that is directly related to ours.



This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). <https://doi.org/10.65109/YNPD9144>

Aziz et al. [3] consider ASHG, and show that computing an optimal egalitarian partition is strongly NP-hard; moreover, verifying whether a given partition maximizes the egalitarian welfare is coNP-complete. The complexity of this problem is further investigated by Hanaka et al. [25], who identify several tractable cases. Peters [31] obtained a related result, showing that an optimal egalitarian partition is polynomial-time computable in graphical HGs as long as the underlying graph has bounded treewidth. In addition, egalitarian welfare has been studied in online ASHG [11] and in ASHG where at most k coalitions can be formed [33]. For *simple fractional hedonic games* (sFHGs) with symmetric preferences, Aziz et al. [4] show that maximizing the egalitarian welfare is NP-hard, and provide a polynomial-time 3-approximation algorithm. Aziz et al. also state that in instances where the underlying graph is a tree, an optimal egalitarian partition can be computed in polynomial time via dynamic programming. This idea was later formalized by Hanaka et al. [24], who present an efficient algorithm for bounded treewidth graphs. The egalitarian welfare has also been employed as a quality measure for evaluating stable outcomes in modified fractional HGs [29].

A closely related problem is the existence of a wonderful (a.k.a. perfect) partition in HGs where agents have *dichotomous preferences*, that is, each agent values each coalition as either 0 or 1. A wonderful partition is one where each agent assigns value 1 to her coalition. Clearly, a wonderful partition exists if and only if there exists a partition whose egalitarian welfare is 1. Establishing whether such a partition exists is, in general, NP-hard [2, 30], though some tractable cases have been identified [12, 30]. Notably, the positive result of Constantinescu et al. [12] implies that the problem of maximizing the egalitarian welfare in anonymous HGs with single-peaked preferences is polynomial-time solvable.

Friends and Enemies Games constitute a widely studied subclass of hedonic games and have been further extended to capture more complex social contexts [7, 10, 13, 14, 27, 32]. In this line of research, most of the focus has been on finding desirable outcomes according to various stability notions. Beyond stability, aspects related to strategyproofness have also been investigated [15, 20, 21, 28]. From a welfare perspective, however, only the utilitarian welfare has been studied. In particular, for EA preferences, maximum utilitarian welfare is not approximable within a factor of $O(n^{1-\epsilon})$, where n is the number of agents; however, there is a polynomial-time $O(n)$ -approximation algorithm [21]. This bound can be improved, with high probability, in scenarios where friend and enemy relations are generated according to specific probabilistic models [9]. Even under FA preferences, the problem remains NP-hard, though a polynomial-time $(4 + o(1))$ -approximation has been obtained [20].

To the best of our knowledge, egalitarian welfare has not previously been considered in Friends and Enemies Games.

2 MODEL AND PRELIMINARIES

Given a positive integer k , we denote by $[k]$ the set $\{1, \dots, k\}$.

In hedonic games, a set of n agents $\mathcal{N} = \{1, \dots, n\}$ needs to be partitioned into disjoint subsets. We refer to subsets of \mathcal{N} as *coalitions*; the set of all coalitions containing an agent $i \in \mathcal{N}$ is denoted by \mathcal{N}^i . An outcome of a hedonic game with the set of agents \mathcal{N} is a *coalition structure*, i.e., a partition $\pi = \{C_1, \dots, C_m\}$ of \mathcal{N}

such that $\cup_{k=1}^m C_k = \mathcal{N}$ and $C_k \cap C_\ell = \emptyset$ for all $k, \ell \in [m]$ with $k \neq \ell$. We denote by Π the set of all partitions of \mathcal{N} . Given an outcome $\pi \in \Pi$, we denote by $\pi(i)$ the coalition in π that contains agent i . The coalition \mathcal{N} is called the *grand coalition*; abusing terminology, we also refer to the partition $\{\mathcal{N}\}$ as the grand coalition. A coalition of size 1 is called a *singleton coalition*.

In hedonic games, each agent $i \in \mathcal{N}$ is endowed with a preference relation \succeq_i over \mathcal{N}^i . Given two coalitions $X, Y \in \mathcal{N}^i$, we say that i *weakly prefers* X to Y if $X \succeq_i Y$. These preferences are lifted to partitions: an agent $i \in \mathcal{N}$ weakly prefers π to π' if $\pi(i) \succeq_i \pi'(i)$.

2.1 Friends and Enemies Games

In hedonic games with friends and enemies, every agent $i \in \mathcal{N}$ partitions the other agents into a set of friends F_i and a set of enemies E_i , with $F_i \cup E_i = \mathcal{N} \setminus \{i\}$ and $F_i \cap E_i = \emptyset$. We write $f_i := |F_i|$, $e_i := |E_i| = n - f_i - 1$.

The agents' preferences are based on *Friends Appreciation* (FA) when for each $i \in \mathcal{N}$ it holds that $X \succeq_i Y$ iff

$$\begin{aligned} &|X \cap F_i| > |Y \cap F_i| \quad \text{or} \\ &|X \cap F_i| = |Y \cap F_i| \quad \text{and} \quad |X \cap E_i| \leq |Y \cap E_i|, \end{aligned}$$

and they are based on *Enemies Aversion* (EA) when $X \succeq_i Y$ iff

$$\begin{aligned} &|X \cap E_i| < |Y \cap E_i| \quad \text{or} \\ &|X \cap E_i| = |Y \cap E_i| \quad \text{and} \quad |X \cap F_i| \geq |Y \cap F_i|. \end{aligned}$$

In other words, under FA, a coalition is preferred over another one if it contains a higher number of friends; if the number of friends is the same, the coalition with fewer enemies is preferred. On the other hand, under EA, a coalition is preferred if it contains fewer enemies; if the number of enemies is the same, the coalition with more friends is preferred.

Example 1. Let us describe a simple instance for both FA and EA. Note that once the friendship relationships are given, the enemy relationships are uniquely determined. Therefore, we can define an instance by specifying the friendship relations only. Let $\mathcal{N} = \{1, 2, 3\}$ be the set of agents, and let $F_1 = \{2\}$, $F_2 = \{3\}$, $F_3 = \{2\}$; see Figure 1a, where a directed edge from agent i to agent j represents i 's opinion of j , and solid and dashed edges represent friend and enemy relations, respectively.

A well-studied class of hedonic games is *additively separable hedonic games* (ASHGs), where agents assign values $v_i(j)$ to each other, the utility $u_i(C)$ an agent i derives from a coalition $C \in \mathcal{N}^i$ is the sum of her values for the coalition members other than herself, i.e., $u_i(C) = \sum_{j \in C \setminus \{i\}} v_i(j)$, and $C \succeq_i D$ iff $u_i(C) \geq u_i(D)$. Given a partition π , for each $i \in \mathcal{N}$ we write $u_i(\pi) := u_i(\pi(i))$. Both FA games and EA games can be encoded as ASHG, with value functions determined by friendship relationships. Indeed, in the FA case, for every agent $i \in \mathcal{N}$, we can set

$$v_i(j) = \begin{cases} 1, & \text{if } j \in F_i, \\ -\frac{1}{n}, & \text{if } j \in E_i. \end{cases}$$

so that the positive effect of one friend is greater than the overall negative effect of all enemies. Similarly, in the EA case, we can set,




Figure 1: A game instance \mathcal{I} of friends and enemies games and the corresponding graphs G^f and G^{sf} .

for every agent $i \in \mathcal{N}$,

$$v_i(j) = \begin{cases} 1, & \text{if } j \in F_i, \\ -n, & \text{if } j \in E_i. \end{cases}$$

Example 2. Consider the instance described in Example 1 and the partition $\pi = \{\{1, 2, 3\}\}$. Each agent $i \in \mathcal{N}$ has one friend and one enemy in \mathcal{N} , so under FA we have $u_1(\pi) = v_1(2) + v_1(3) = 1 - \frac{1}{3} = \frac{2}{3}$ and, similarly, $u_2(\pi) = u_3(\pi) = \frac{2}{3}$, whereas under EA we obtain $u_1(\pi) = v_1(2) + v_1(3) = 1 - 3 = -2$, and also $u_2(\pi) = u_3(\pi) = -2$.

Observe that, under FA, the presence of an enemy in a coalition may be tolerated as long as there is at least one friend, while, under EA, the presence of an enemy is never acceptable, in the sense that an agent would prefer to be a singleton coalition.

An FA or EA instance \mathcal{I} is specified by a set of agents \mathcal{N} and the set of friendship relations $F = \{(i, j) \mid i \in \mathcal{N}, j \in F_i\}$. We denote by F^s the set of mutual friendships: $F^s = \{(i, j) \mid i \in F_j \wedge j \in F_i\}$.

Graph Representation. In our discussion, we will make use of two different types of graph representation:

- The *friendship graph* $G^f = (\mathcal{N}, F)$, representing all directed friendship relationships, and
- the *strong friendship graph* $G^{sf} = (\mathcal{N}, F^s)$, representing only mutual friendship relationships.

Figures 1b and 1c show the friendship graph and the strong friendship graph, respectively, for the instance described in Example 1. When an instance is *symmetric*, that is, for all $i, j \in \mathcal{N}$ it holds that $i \in F_j$ if and only if $j \in F_i$, we have $G^f \equiv G^{sf}$.

Throughout the paper, we may define a game instance either as $\mathcal{I} = (\mathcal{N}, F)$ or by directly referring to the respective graph G^f or G^{sf} , depending on the context. For simplicity, we will often omit \mathcal{I} from the notation, implicitly assuming that \mathcal{N} and F are given.

2.2 Egalitarian Welfare, Approximation, and Other Guarantees

For an instance \mathcal{I} of an ASHG, we define the *egalitarian welfare* of an outcome π as

$$\text{ESW}^{\mathcal{I}}(\pi) = \min_{i \in \mathcal{N}} u_i(\pi).$$

We are interested in finding outcomes that (approximately) maximize ESW. Accordingly, we define MAXESW as the problem of maximizing the egalitarian welfare. Given an instance \mathcal{I} , we denote by $\text{opt}(\mathcal{I})$ the value of an optimal solution of MAXESW for the instance \mathcal{I} , and by $\text{OPT}(\mathcal{I})$ any partition of the agents such that $\text{ESW}^{\mathcal{I}}(\text{OPT}(\mathcal{I})) = \text{opt}(\mathcal{I})$. When \mathcal{I} is clear from the context, we simply write ESW, opt, and OPT.

Remark 1. In any instance of EA or FA, if $F_i = \emptyset$ for some $i \in \mathcal{N}$, then the maximum attainable egalitarian welfare is necessarily 0, since i cannot achieve a positive utility in any coalition. In this case, partitioning the agents into singletons is optimal. Henceforth, we will always assume that $F_i \neq \emptyset$ for every $i \in \mathcal{N}$.

Approximation Ratio. In this work, we aim to find partitions that maximize the egalitarian welfare. Unfortunately, it turns out that MAXESW is NP-hard, both for EA and for FA. For this reason, we will turn our attention to approximate solutions. Given an algorithm \mathcal{A} , we define its *approximation ratio* as follows:

$$r_n(\mathcal{A}) = \max_{\substack{\mathcal{I}=(\mathcal{N},F): \\ |\mathcal{N}|=n}} \frac{\text{opt}(\mathcal{I})}{\text{ESW}^{\mathcal{I}}(\mathcal{A}(\mathcal{I}))} \quad \text{and} \quad r(\mathcal{A}) = \sup_{n \in \mathbb{N}} r_n(\mathcal{A}),$$

where $\mathcal{A}(\mathcal{I})$ is the partition returned by \mathcal{A} on input \mathcal{I} .

Guarantees for the Agents. In the pursuit of finding efficient algorithms that yield good approximation for MAXESW , we will simultaneously provide some guarantees to (almost) all agents. In particular, given an agent i , we say that a coalition C satisfies for i :

- ALLF, if all friends of i are included in C ;
- ONLYF, if C consists solely of friends of i , though not necessarily all of them;
- ONEF, if C contains at least one friend of i .

Let Γ be a guarantee that a coalition may fulfill for an agent. A partition π *guarantees* Γ for $A \subseteq \mathcal{N}$ if each coalition $C \in \pi$ guarantees Γ to all agents in $A \cap C$. Moreover, π is a *minimal-by-refinement partition guaranteeing* Γ for $A \subseteq \mathcal{N}$ if it is not possible to split C into non-empty coalitions C_1, \dots, C_k so that the partition $\pi' = \pi \setminus \{C\} \cup \{C_1, \dots, C_k\}$ guarantees Γ for $A \subseteq \mathcal{N}$. If $A = \mathcal{N}$, we say that π is a *minimal-by-refinement partition guaranteeing* Γ .

We now present an example to illustrate these notions.

Example 3. The grand coalition $\pi = \{\mathcal{N}\}$ satisfies ALLF. However, it may fail ONLYF, as there may exist an agent i with $F_i \neq \mathcal{N} \setminus \{i\}$. Under our hypothesis that each agent has at least one friend, π also satisfies ONEF. In fact, under this hypothesis ALLF implies ONEF. However, $\{\mathcal{N}\}$ need not be the minimal-by-refinement partition that satisfies ALLF. Indeed, if $\mathcal{N} = \{1, 2, 3, 4\}$ and $F_1 = \{2\}$, $F_2 = \{1\}$, $F_3 = \{4\}$, and $F_4 = \{3\}$, then the unique minimal-by-refinement partition satisfying ALLF is $\{\{1, 2\}, \{3, 4\}\}$.

3 ENEMIES AVERSION

When constructing (approximately) optimal partitions under EA preferences, we can restrict our attention to mutual friendship relationships only. Indeed, if some agent is in a coalition with an enemy, her utility—and hence ESW—is negative, whereas for the coalition structure π^{sing} that consists of n singleton coalitions we

have $\text{ESW}(\pi^{\text{sing}}) = 0$. Consequently, for ESW to be non-negative, an outcome π must satisfy ONLYF for all agents, i.e., each coalition C in π must induce a clique in the strong friendship graph G^{sf} . Moreover, if π contains a singleton coalition, we have $\text{ESW}(\pi) \leq 0$, even if π satisfies ONLYF for all agents. Thus, for $\text{ESW}(\pi)$ to be positive, π must satisfy ONLYF and ONEF for all agents, i.e., induce a partition of G^{sf} into cliques of size at least 2.

Combining these insights, we derive the following lemma.

LEMMA 1. *A partition $\pi \in \Pi$ satisfies $\text{ESW}(\pi) \geq 0$ if and only if π induces a clique partition of G^{sf} . Furthermore, $\text{ESW}(\pi) > 0$ if and only this clique partition contains no cliques of size 1.*

Thus, in an optimal coalition structure, the set of agents must be partitioned into cliques. Moreover, to maximize the ESW, the smallest clique in the partition must be as large as possible.

However, we will now see that this structural result does not mean that maximizing ESW is easy.

3.1 Hardness of Approximation

THEOREM 1. *Under EA, for every $\epsilon > 0$, there is no polynomial-time algorithm that approximates MAXESW within a factor $O(n^{1-\epsilon})$, unless $P = NP$. The hardness result holds even for symmetric instances.*

To prove the theorem, we make use of an inapproximability result for the closely related problem: PARTITION INTO CLIQUES.

PARTITION INTO CLIQUES (PC):

Input: An undirected graph $G = (V, E)$ and a positive integer $K \leq |V|$.

Question: Can the vertices of G be partitioned into $k \leq K$ disjoint sets V_1, \dots, V_k so that, for each $h \in [k]$, the subgraph $G[V_h]$ induced by V_h is a clique?

Let MINPC denote the corresponding minimization problem, i.e., finding the minimum number of cliques in a clique partition of G . Note that every graph admits a clique partition, since a single vertex forms a clique of size 1. The problem MINPC cannot be approximated within a factor $O(|V|^{1-\epsilon})$, for every $\epsilon > 0$, unless $P = NP$ [34].

PROOF OF THEOREM 1. Assume there exists an $n^{1-\epsilon}$ -approximation algorithm \mathcal{A} for MAXESW under EA, for some $\epsilon > 0$. We will show that this assumption contradicts the known inapproximability result for the MINPC problem.

Let $G = (V, E)$ be an arbitrary instance of MINPC. We now construct a corresponding EA instance \mathcal{I} of the MAXESW problem. By Lemma 1, it suffices to describe the friendship graph $G^{\text{sf}} = (\mathcal{N}, F^s)$ for \mathcal{I} , since only coalition structures inducing clique partitions of G^{sf} offer non-negative ESW. To construct our instance \mathcal{I} , we set $\mathcal{N} = V \cup V'$, where $|V'| = |V|$ and V' is disjoint from V . For each $\{u, v\} \in E$, we create an edge $\{u, v\} \in F^s$. Additionally, we create an edge between every pair of vertices in V' and connect every vertex of V to every vertex of V' . Note that \mathcal{I} satisfies $n = |\mathcal{N}| = 2 \cdot |V|$.

Let us now relate $\text{opt}(\mathcal{I})$, i.e., the optimal MAXESW value for \mathcal{I} , to the optimal value of MINPC, denoted opt^{PC} , i.e., the minimum number of cliques needed to cover the vertices of $G = (V, E)$. Let V_1, \dots, V_{k^*} be a partition of G into $k^* = \text{opt}^{\text{PC}}$ disjoint cliques. Partition the vertices in V' into coalitions C'_1, \dots, C'_{k^*} as evenly as

possible; this ensures that $|C'_h| \geq \lfloor \frac{|V|}{k^*} \rfloor$ for each $h \in [k^*]$. Now set $C_h = V_h \cup C'_h$ for all $h \in [k^*]$, and note that $\pi = \{C_1, \dots, C_{k^*}\}$ is a clique partition of G^{sf} . Since $V_h \neq \emptyset$ for all $h \in [k^*]$, we obtain $|C_h| \geq \lfloor \frac{|V|}{k^*} \rfloor + 1 \geq \frac{|V|}{k^*}$ for each $h \in [k^*]$. As the utility of each agent in a clique of size t is $t - 1$, we obtain

$$\text{opt}(\mathcal{I}) \geq \frac{|V|}{k^*} - 1 = \frac{|V|}{\text{opt}^{\text{PC}}} - 1. \quad (1)$$

Consider now the partition π computed by our hypothetical approximation algorithm \mathcal{A} for MAXESW, and let $s = \text{ESW}(\pi)$. By Lemma 1, π must be a clique partition; moreover, $\text{ESW}(\pi) = s$ implies that each coalition in π contains at least $s + 1$ vertices. Hence, π consists of at most $\frac{2 \cdot |V|}{s+1}$ disjoint cliques of G^{sf} . We can now convert π into a partition of G into k cliques with $k \leq \frac{2 \cdot |V|}{s+1}$, as the restriction of each coalition $C \in \pi$ to the nodes of V is either empty, if initially made by only agents in V' , or a clique of G . Since \mathcal{A} computes an $n^{1-\epsilon}$ approximation of MAXESW, we have

$$\frac{\text{opt}(\mathcal{I})}{s} \leq n^{1-\epsilon}. \quad (2)$$

On the other hand, $k \leq \frac{2 \cdot |V|}{s+1}$ implies $s \leq \frac{2 \cdot |V|}{k} - 1$; combining this with (1), we get

$$\frac{\text{opt}(\mathcal{I})}{s} \geq \frac{\frac{|V|}{\text{opt}^{\text{PC}}} - 1}{\frac{2 \cdot |V|}{k} - 1} \geq \frac{\frac{|V|}{\text{opt}^{\text{PC}}}}{\frac{2 \cdot |V|}{k} - 1} - 1 > \frac{k}{2 \cdot \text{opt}^{\text{PC}}} - 1, \quad (3)$$

where the last two inequalities hold since $1 \leq \frac{2 \cdot |V|}{k} - 1 < \frac{2 \cdot |V|}{k}$.

As $n = 2 \cdot |V|$, by putting together (3) and (2), we obtain $\frac{k}{\text{opt}^{\text{PC}}} < 2 \cdot (2 \cdot |V|)^{1-\epsilon} + 2$. For sufficiently large V , this quantity is smaller than $|V|^{1-\epsilon/2}$, a contradiction to the inapproximability of MINPC. \square

3.2 Linear Approximation in Polynomial Time

By Lemma 1, to obtain a bounded approximation for a given EA instance, we need to determine whether the associated strong friendship graph G^{sf} admits a *non-trivial clique partition*, i.e., a clique partition in which each agent belongs to a clique of size at least 2.

OBSERVATION 1. *If G^{sf} does not admit a non-trivial clique partition, then a partition into singletons achieves optimal ESW. Otherwise, any non-trivial clique partition guarantees an ESW of at least 1.*

Note that a non-trivial clique partition exists if and only if it is possible to partition the graph into triangles and a matching. This observation enables us to design a polynomial-time algorithm that computes an $(n - 1)$ -approximation to MAXESW.

THEOREM 2. *Under EA, there exists a polynomial-time algorithm computing an $(n - 1)$ -approximation to MAXESW. Furthermore, if $\text{ESW} > 0$, the computed partition satisfies ONEF for all the agents.*

PROOF. The problem of finding a perfect covering of the vertices of a graph by \mathbb{K}_2 and \mathbb{K}_3 , i.e. cliques of size 2 or 3, admits a polynomial-time algorithm [26]. We can execute this algorithm on G^{sf} . If the algorithm reports that there is no perfect covering of G^{sf} with \mathbb{K}_2 and \mathbb{K}_3 , every partition of G^{sf} into cliques contains a singleton, and hence the maximum attainable egalitarian welfare is 0 (so, in particular, a partition into singletons maximizes ESW).

On the other hand, if the algorithm returns a perfect covering, it corresponds to a partition of the agents into coalitions of size 2 and 3, which satisfies ONEF. Any such partition guarantees utility of at least 1 to every agent. Since the maximum utility an agent can derive is at most $n - 1$, the ESW of the partition constructed by our algorithm is within a factor of $n - 1$ from optimal. \square

We also note that MAXESW becomes easy if G^{sf} is *triangle-free*, that is, no cycle in G^{sf} has length 3.

THEOREM 3. *Under EA, if G^{sf} is triangle-free, MAXESW is in P.*

PROOF. In a triangle-free graph, no clique has size larger than 2. Thus, we proceed by computing a maximum matching in G^{sf} . If G^{sf} admits a perfect matching, the maximum ESW is 1, and an optimal partition is attained by creating a coalition for each pair of matched agents. Otherwise, the maximum ESW is 0, and splitting the agents into singletons is an optimal partition. \square

4 FRIENDS APPRECIATION

For FA preferences, unless specified otherwise, we will work with the G^f graph. Let $f_{\min} = \min_{i \in \mathcal{N}} f_i$ denote the minimum number of friends of any agent, and set $\mathcal{N}_{\min} = \{i \in \mathcal{N} \mid f_i = f_{\min}\}$. We start by presenting simple bounds on opt for MAXESW; for the upper bound, note that agents in \mathcal{N}_{\min} have at most f_{\min} friends, and for the lower bound, consider the grand coalition.

OBSERVATION 2. *Under FA, $f_{\min} - 1 < \text{opt} \leq f_{\min}$.*

Consequently, in every optimal partition, every agent $i \in \mathcal{N}_{\min}$ must end up in a coalition that contains all of her friends.

COROLLARY 1. *Under FA, any MAXESW partition satisfies ALLF for all the agents in \mathcal{N}_{\min} .*

4.1 NP-Hardness

It turns out that for FA preferences, too, computing a MAXESW partition is NP-hard; the hardness result persists even if agents' preferences are symmetric.

THEOREM 4. *Under FA, MAXESW is NP-hard. The hardness result holds even for symmetric instances.*

Our proof proceeds by a reduction from the PARTITION INTO TRIANGLES problem, which can be formulated as follows:

PARTITION INTO TRIANGLES:

Input: An undirected graph $G = (V, E)$.

Question: Is it possible to partition V as S_1, \dots, S_k so that for each $S_j, j \in [k]$, its induced subgraph is a triangle?

PARTITION INTO TRIANGLES is known to be NP-hard [23]. Note that we can assume that $|V| = 3k$ for some $k \in \mathbb{N}$, as otherwise we trivially have a “no”-instance. Thus, in what follows, we will make this assumption.

PROOF. Consider an instance $G = (V, E)$ of PARTITION INTO TRIANGLES with $|V| = 3k$. We will construct a (symmetric) instance of FA by specifying its associated graph G^{sf} .

For each vertex $v \in V$ we construct four vertices of G^{sf} : a_v, b_v, c_v and z_v . Thus, $\mathcal{N} = \{a_v, b_v, c_v, z_v \mid v \in V\}$ and $n = |\mathcal{N}| = 4|V|$.

For each $v \in V$, the graph G^{sf} contains edges $\{a_v, z_v\}$, $\{b_v, z_v\}$, and $\{c_v, z_v\}$. Also, for each edge $\{u, v\} \in E$ of G , the graph G^{sf} contains edges $\{a_u, a_v\}$, $\{b_u, b_v\}$, and $\{c_u, c_v\}$. This completes the construction of G^{sf} .

We claim that the instance of FA that corresponds to G^{sf} admits a partition π with $\text{ESW}(\pi) \geq 3 - \frac{8}{n}$ if and only if G is a yes-instance of PARTITION INTO TRIANGLES.

Indeed, let S_1, \dots, S_k be a partition of G into triangles. For each $S_j, j \in [k]$, we create a coalition $C_j = \{a_v, b_v, c_v, z_v \mid v \in S_j\}$. This coalition contains 12 agents, and each agent in it has 3 friends and 8 enemies, so the utility of each agent is $3 - \frac{8}{n}$.

Conversely, suppose our FA game admits a partition π with $\text{ESW}(\pi) \geq 3 - \frac{8}{n}$. Since $n = 4|V| \geq 12$, $3 - \frac{8}{n} > 2$, so in this partition, each agent must be in a coalition with at least 3 friends. Thus, for each $v \in V$ it holds that z_v is in the same coalition as a_v, b_v and c_v (as z_v has no other friends). It follows that each coalition in π is of the form $C = \{a_v, b_v, c_v, z_v \mid v \in S\}$, where S is a subset of V . Moreover, z_v 's coalition must contain at most 8 enemies, implying $|S| \leq 3$. Now, if $|S| < 3$ or if its induced graph is not a triangle, there is an agent a_v in C that has fewer than 3 friends in C , and hence her utility is lower than $3 - \frac{8}{n}$, a contradiction. Thus, each coalition in π corresponds to a triangle in G , and these triangles are pairwise disjoint, i.e., G admits a partition into triangles. \square

In the remainder of this section, we present our approximability results. We describe a simple algorithm whose approximation ratio is linear in the number of agents. While the linear approximation bound is tight for our algorithm, this is caused by instances with $f_{\min} = 1$: we show that as soon as each agent has at least two friends, the approximation ratio of our algorithms becomes $2 - \Theta\left(\frac{1}{n}\right)$. We also obtain $\left(2 - \Theta\left(\frac{1}{n}\right)\right)$ -approximation results for two other settings: for a randomized algorithm (in expectation) and for symmetric instances.

4.2 Deterministic Bounded Approximation

Corollary 1 establishes that, in an optimal partition, the agents with the minimum number of friends, i.e., the ones in \mathcal{N}_{\min} , must be in a coalition with all their friends. As the marginal contribution of all enemies is, in absolute terms, less than that of a single friend, a naive approach would be to partition the agents so that each agent—not only the ones in \mathcal{N}_{\min} —is placed in a coalition together with all of their friends, i.e., to create a partition that satisfies ALLF. Among all such partitions, we prefer ones that avoid placing enemies in the same coalition as much as possible. Thus, an attractive strategy is to create a coalition for each weakly connected component of G^f ; we refer to the resulting algorithm as WeaklyConn:

WeaklyConn: Given a graph G^f , let $\pi = (C_1, \dots, C_m)$ be the list of weakly connected components of G^f ; return the partition π .

Unfortunately, we will show that this algorithm does not always provide a constant-factor approximation to MAXESW.

THEOREM 5. *Under FA, the algorithm WeaklyConn guarantees an approximation of $2 - \frac{6}{n+3}$ to MAXESW if $f_{\min} \geq 2$ and $\frac{n}{2}$ otherwise.*

PROOF. Consider an instance \mathcal{I} of FA. By Remark 1, we can assume that $F_i \neq \emptyset$ for all $i \in \mathcal{N}$, as otherwise we can compute the optimum. Let π denote the outcome returned by WeaklyConn on \mathcal{I} .

By Observation 2, we have $\text{opt} \leq f_{\min}$. On the other hand, in π , every agent $i \in \mathcal{N}$ is in a coalition with f_i friends and at most $n - f_i - 1$ enemies. Therefore,

$$u_i(\pi) \geq f_i - \frac{n - f_i - 1}{n} = f_i \cdot \frac{n + 1}{n} - \frac{n - 1}{n}.$$

Let j be an agent with minimum utility in π ; note that $j \in \mathcal{N}_{\min}$. The approximation ratio of `WeaklyConn` on \mathcal{I} can be bounded as

$$\frac{\text{opt}(\mathcal{I})}{\min_{i \in \mathcal{N}} u_i(\pi)} \leq \frac{f_{\min}}{f_{\min} \cdot \frac{n+1}{n} - \frac{n-1}{n}} = \frac{f_{\min} \cdot n}{f_{\min}(n+1) - (n-1)}.$$

The ratio $\frac{f_{\min} \cdot n}{f_{\min}(n+1) - (n-1)}$ is a decreasing function of f_{\min} , so it takes its maximum value, which is $\frac{n}{2}$, if $f_{\min} = 1$. If $f_{\min} \geq 2$, we have

$$\frac{f_j \cdot n}{f_j(n+1) - (n-1)} \leq \frac{2n}{2(n+1) - (n-1)} = \frac{2n}{n+3} = 2 - \frac{6}{n+3}.$$

This completes our proof. \square

In the full version of this work [18], we give an example showing that both bounds in the statement of Theorem 5 are tight.

4.3 Assigning Agents with Only One Friend

Theorem 5 establishes that the most challenging instances to approximate are those containing an agent who only has one friend. Indeed, a constant approximation can be achieved as soon as $|F_i| \geq 2$ for all $i \in \mathcal{N}$. In the next sections, we show that constant-factor approximation can be recovered in randomized and symmetric settings. To this end, we need to tackle the problem of how to assign agents who only have one friend. Hereafter, we denote by $\mathcal{N}_1(\mathcal{I}) = \{i \in \mathcal{N} \mid f_i = 1\}$ the set of agents with a unique friend. When \mathcal{I} is clear from the context, we simply write \mathcal{N}_1 .

If $\mathcal{N}_1 \neq \emptyset$, Observation 2 and Corollary 1 establish that, to obtain any approximation, we must satisfy $\text{ONEF} \equiv \text{ALLF}$ for all agents in \mathcal{N}_1 and, moreover, $\text{opt} \in (0, 1]$. Given an FA instance \mathcal{I} , we can guarantee ONEF to all agents in \mathcal{N}_1 by placing all agents in the same coalition. A more sophisticated, yet tractable, approach is to create a minimal-by-refinement partition that enforces ONEF for all agents in \mathcal{N}_1 ; we will refer to this algorithm as `OneWeaklyConn`.

`OneWeaklyConn`: On an input $G^f = (\mathcal{N}, F)$, the algorithm creates the graph $G_1^f = (\mathcal{N}, F_1)$ by only maintaining the friend relationships of the agents in \mathcal{N}_1 , i.e., it sets $F_1 = \{(i, j) \in F \mid i \in \mathcal{N}_1\}$. Then it splits the agents into the weakly-connected components of G_1^f , that is, applies `WeaklyConn` on G_1^f .

By construction, the partition computed by `OneWeaklyConn` is the unique minimal-by-refinement partition that guarantees ONEF to all agents in \mathcal{N}_1 . Indeed, if a partition π satisfies ONEF for all agents in \mathcal{N}_1 , then every coalition $S \in \pi$ must contain a weakly connected component of G_1^f . Moreover, if S contains several weakly connected components of G_1^f , then π is not a minimal-by-refinement partition that satisfies ONEF , as S can be split into smaller coalitions—one for each weakly connected component of G_1^f that it contains.

Next, we establish a useful connection between the utility of agents in \mathcal{N}_1 in the minimal-by-refinement ONEF partition and in an optimal partition.

LEMMA 2. *Let π be the unique minimal-by-refinement partition that provides ONEF to all agents in \mathcal{N}_1 . Then, for all $i \in \mathcal{N}_1$ we have $u_i(\pi) \geq \text{opt}$, and $\pi(i) \subseteq \pi^*(i)$ for every optimal partition π^* .*

PROOF. Assume $\mathcal{N}_1 \neq \emptyset$; otherwise, the statement is immediate.

Let π^* be an optimal partition. By Observation 2 we have $\text{opt} > 0$; hence, π^* satisfies condition ONEF for all agents. As previously observed, for ONEF to hold for all agents in \mathcal{N}_1 , the members of each weakly connected component of G_1^f must belong to the same coalition. Moreover, we have already established that the unique minimal-by-refinement partition π satisfying ONEF for the agents in \mathcal{N}_1 is precisely the partition where each coalition is a weakly connected component of G_1^f . Consequently, for every $i \in \mathcal{N}_1$ we have $\pi(i) \subseteq \pi^*(i)$. Moreover, since $F_i \subseteq \pi(i) \subseteq \pi^*(i)$, we obtain $u_i(\pi) \geq u_i(\pi^*) \geq \text{opt}$. \square

By Lemma 2, `OneWeaklyConn`(\mathcal{I}) guarantees utility of at least opt to the agents in \mathcal{N}_1 ; however, this algorithm may place some agents in $\mathcal{N} \setminus \mathcal{N}_1$ in singleton coalitions, so that their utility is 0.

Clearly, if $\mathcal{N}_1 = \mathcal{N}$, then, by Lemma 2, `OneWeaklyConn` computes a `MAXESW` partition. This leads to the following theorem.

THEOREM 6. *Under FA, if $f_i = 1$ for all $i \in \mathcal{N}$ then `MAXESW` is in \mathcal{P} . Furthermore, `OPT` is the unique minimal-by-refinement partition satisfying ONEF (which in this case coincides with `ALLF`) for all agents.*

4.4 Randomized Constant Approximation

In what follows, we combine `OneWeaklyConn` and `WeaklyConn` to obtain a randomized algorithm whose approximation ratio is at most 2. To start, we need to define the approximation ratio in the randomized setting. Since for a given instance \mathcal{I} the space of all possible outcomes is finite, a randomized algorithm \mathcal{A} can be viewed as a probability distribution over a finite set of deterministic algorithms Σ , where each algorithm in Σ computes a partition of the agents in \mathcal{I} ; we write $\mathcal{A} \sim \Delta(\Sigma)$.

There are two natural approaches to measuring the performance of a randomized algorithm in our setting: (1) the expected minimum utility, over the random choices of the algorithm, as given by $\mathbb{E}_{\mathcal{A} \sim \Delta(\Sigma)} [\min_{i \in \mathcal{N}} u_i(\mathcal{A}(\mathcal{I}))]$, and (2) the minimum expected utility over all agents, as given by $\min_{i \in \mathcal{N}} \mathbb{E}_{\mathcal{A} \sim \Delta(\Sigma)} [u_i(\mathcal{A}(\mathcal{I}))]$.

It is easy to see that the first measure is more stringent, i.e.,

$$\mathbb{E}_{\mathcal{A} \sim \Delta(\Sigma)} \left[\min_{i \in \mathcal{N}} u_i(\mathcal{A}(\mathcal{I})) \right] \leq \min_{i \in \mathcal{N}} \mathbb{E}_{\mathcal{A} \sim \Delta(\Sigma)} [u_i(\mathcal{A}(\mathcal{I}))].$$

Indeed, under the first approach, we cannot expect to get an improved approximation ratio from randomization, since for every instance \mathcal{I} we have

$$\mathbb{E}_{\mathcal{A} \sim \Delta(\Sigma)} \left[\min_{i \in \mathcal{N}} u_i(\mathcal{A}(\mathcal{I})) \right] \leq \max_{\mathcal{A} \in \Sigma} \min_{i \in \mathcal{N}} u_i(\mathcal{A}(\mathcal{I})),$$

i.e., the expected egalitarian welfare of a randomized algorithm $\mathcal{A} \sim \Delta(\Sigma)$ cannot be better than the egalitarian welfare achieved deterministically by the best algorithm in Σ .

Thus, we pursue the second approach and define the approximation ratio as follows:

$$r_n(\mathcal{A} \sim \Delta(\Sigma)) = \sup_{\substack{\mathcal{I}=(\mathcal{N}, F) \\ |\mathcal{N}|=n}} \frac{\text{opt}(\mathcal{I})}{\min_{i \in \mathcal{N}} \mathbb{E}_{\mathcal{A} \sim \Delta(\Sigma)} [u_i(\mathcal{A}(\mathcal{I}))]},$$

and $r(\mathcal{A} \sim \Delta(\Sigma)) = \sup_{n \in \mathbb{N}} r_n(\mathcal{A} \sim \Delta(\Sigma))$.

We are ready to define our randomized algorithm, RandAlgo.

RandAlgo: With probability α , executes algorithm WeaklyConn, and with probability $1 - \alpha$ it executes algorithm OneWeaklyConn, where the parameter $\alpha \in [0, 1]$ will be set later and will be dependent on the input instance.

THEOREM 7. *Under FA, there exists a choice of α such that*

$$r_n(\text{RandAlgo}) \leq 2 - \frac{5}{n+3}.$$

PROOF. Given an instance \mathcal{I} , let us set $\pi = \text{WeaklyConn}(\mathcal{I})$ and $\pi' = \text{OneWeaklyConn}(\mathcal{I})$. Moreover, let $v = \min_{i \in \mathcal{N}_1} u_i(\pi')$.

If $\mathcal{N}_1 = \emptyset$, we set $\alpha = 1$: then RandAlgo coincides with WeaklyConn, so its approximation ratio for the case $\mathcal{N}_1 = \emptyset$ is bounded by $2 - \frac{6}{n+3}$.

Assume now $\mathcal{N}_1 \neq \emptyset$. We will estimate the expected utility of an agent $i \in \mathcal{N}$ depending on whether $i \in \mathcal{N}_1$ or $i \in \mathcal{N} \setminus \mathcal{N}_1$.

If $i \in \mathcal{N}_1$, then $u_i(\pi) \geq \frac{2}{n}$, as shown in the proof of Theorem 5, while $u_i(\pi') \geq v$. Hence, $\mathbb{E}[u_i(\text{RandAlgo})] \geq \alpha \cdot \frac{2}{n} + (1 - \alpha) \cdot v$.

If $i \in \mathcal{N} \setminus \mathcal{N}_1$, then $u_i(\pi) \geq 1 + \frac{3}{n}$, as in the worst case i has only two friends and is put in a coalition with all her enemies, while $u_i(\pi')$ may be 0. Consequently, $\mathbb{E}[u_i(\text{RandAlgo})] \geq \alpha \cdot \left(1 + \frac{3}{n}\right)$.

Putting these two scenarios together, we get

$$\min_{i \in \mathcal{N}} \mathbb{E}[u_i(\text{RandAlgo}(\mathcal{I}))] \geq \min \left\{ v + \alpha \cdot \left(\frac{2}{n} - v \right), \alpha \cdot \left(1 + \frac{3}{n} \right) \right\}.$$

Since $v \geq \frac{2}{n}$, the function $v + \alpha \cdot \left(\frac{2}{n} - v \right)$ is decreasing in α , whereas $\alpha \cdot \left(1 + \frac{3}{n} \right)$ is increasing in α . The minimum of the two is attained when they are equal, which happens for $\alpha = \frac{v}{1 + \frac{1}{n} + v}$. Hence,

$$\min_{i \in \mathcal{N}} \mathbb{E}[u_i(\text{RandAlgo}(\mathcal{I}))] \geq \left(1 + \frac{3}{n} \right) \cdot \frac{v}{1 + \frac{1}{n} + v}. \quad (4)$$

The right-hand side of (4) is an increasing function of v and by Lemma 2 we have $v \geq \text{opt}(\mathcal{I})$. Hence,

$$\min_{i \in \mathcal{N}} \mathbb{E}[u_i(\text{RandAlgo}(\mathcal{I}))] \geq \left(1 + \frac{3}{n} \right) \cdot \frac{\text{opt}(\mathcal{I})}{1 + \frac{1}{n} + \text{opt}(\mathcal{I})}.$$

Now, $\mathcal{N}_1 \neq \emptyset$ implies $\text{opt}(\mathcal{I}) \leq 1$, so we obtain

$$\begin{aligned} r_n(\text{RandAlgo}) &= \sup_{\substack{\mathcal{I}=(\mathcal{N},F) \\ |\mathcal{N}|=n}} \frac{\text{opt}(\mathcal{I})}{\min_{i \in \mathcal{N}} \mathbb{E}[u_i(\text{RandAlgo}(\mathcal{I}))]} \\ &\leq \frac{1 + \frac{1}{n} + \text{opt}(\mathcal{I})}{1 + \frac{3}{n}} \leq \frac{2n+1}{n+3} = 2 - \frac{5}{n+3} \end{aligned}$$

which is what we wanted to prove. \square

In the full version of this work [18], we give an example showing the tightness of the analysis.

4.5 Deterministic Constant Approximation in the Symmetric Case

In this section, we will be working with symmetric instances; therefore, we will only consider the graph G^{sf} .

THEOREM 8. *Under FA, there exists a polynomial-time algorithm that computes a $(2 - \frac{4}{n+2})$ -approximation for symmetric instances.*

PROOF. If $\mathcal{N}_1 = \emptyset$, then the approximation ratio of WeaklyConn (which in this case simply computes the connected components of G^{sf}) is $2 - \frac{6}{n+3} < 2 - \frac{4}{n+2}$. Thus, hereafter, we assume $\mathcal{N}_1 \neq \emptyset$. Our algorithm proceeds in three stages.

Stage 1: Let $\mathcal{R} = \{j \mid j \in F_i \text{ for some } i \in \mathcal{N}_1\}$ and $\overline{\mathcal{N}}_1 = \mathcal{N}_1 \cup \mathcal{R}$; that is, $\overline{\mathcal{N}}_1$ contains the agents having only one friend as well as their unique friends. In the first stage, we construct the minimal partition of the agents in $\overline{\mathcal{N}}_1$ that places each agent in \mathcal{N}_1 together with her unique friend. To this end, for each $r \in \mathcal{R}$ let $S(r) = \{i \in \mathcal{N}_1 \mid r \in F_i\}$, and note that the sets $\{S(r)\}_{r \in \mathcal{R}}$ form a partition of $\overline{\mathcal{N}}_1$, which we will denote by π_1 (if two agents $i_1, i_2 \in \mathcal{N}_1$ are each other's unique friends, they both appear in \mathcal{R} ; the sets $S(i_1)$ and $S(i_2)$ then coincide, and π_1 includes a single copy of this set). By construction, each $S(r)$, $r \in \mathcal{R}$, induces a star in G^{sf} . Furthermore, the utility of the root of a star is lower-bounded by the utility of its leaves, all of whom share the same utility. Note that, by Lemma 2, for each partition π^* of \mathcal{N} that maximizes ESW and for each $i \in \mathcal{N}_1$ we have $\pi_1(i) \subseteq \pi^*(i)$ and hence $u_i(\pi_1) \geq \text{opt}$.

Stage 2: Next, let $\mathcal{L} = \{i \in \mathcal{N} \setminus \overline{\mathcal{N}}_1 \mid F_i \subseteq \overline{\mathcal{N}}_1\}$ be the set of *lonely* agents in $\mathcal{N} \setminus \overline{\mathcal{N}}_1$, i.e., agents with no friends in $\mathcal{N} \setminus \overline{\mathcal{N}}_1$. In the second stage, we deal with agents in \mathcal{L} . Note that, for each $i \in \mathcal{L}$ we have $F_i \subseteq \mathcal{R} \setminus \mathcal{N}_1$, so, to obtain positive utility, agents in \mathcal{L} need to be added to coalitions in π_1 . We add agents in \mathcal{L} to coalitions in π_1 so as to minimize the size of the largest coalition in the resulting partition of $\overline{\mathcal{N}}_1 \cup \mathcal{L}$, subject to the condition that each agent in \mathcal{L} is placed in a coalition with one of her friends.

To this end, we construct a bipartite graph with parts \mathcal{L} and \mathcal{R} (where there is an edge between $i \in \mathcal{L}$ and $r \in \mathcal{R}$ if and only if $\{i, r\} \in G^{\text{sf}}$), and solve a sequence of b-matching problems on this graph. Specifically, we iterate through $\kappa = \max_{r \in \mathcal{R}} |S_r|, \dots, n$ and check if this graph admits a b-matching of size $|\mathcal{L}|$ where the capacity of each $i \in \mathcal{L}$ is 1 and the capacity of $r \in \mathcal{R}$ is $\kappa - |S_r|$; we identify the smallest κ for which such a b-matching exists and place each $i \in \mathcal{L}$ into the coalition $S(r)$ such that i is matched to r . This step can be executed in polynomial time, as the problem of finding a b-matching in a bipartite graph is polynomial-time solvable.

Let π_2 be the resulting partition of $\overline{\mathcal{N}}_1 \cup \mathcal{L}$; by construction, for each agent $i \in \overline{\mathcal{N}}_1 \cup \mathcal{L}$ we have $u_i(\pi_2) \geq \text{opt}$.

Stage 3: Finally, let $\mathcal{U} = \mathcal{N} \setminus (\overline{\mathcal{N}}_1 \cup \mathcal{L})$ be the set of agents who remained unassigned by the end of Stage 2. In the third stage, we deal with agents in \mathcal{U} . Note that, by construction of \mathcal{L} (and by symmetry), no agent in \mathcal{U} has a friend in \mathcal{L} , and hence each agent in \mathcal{U} has at least one friend in \mathcal{U} .

If every agent in \mathcal{U} has at least 2 friends in \mathcal{U} , or if $|\mathcal{U}| \leq \frac{n}{2} + 1$, we put all agents in \mathcal{U} in the same coalition $\mathcal{U}_0 := \mathcal{U}$, add this coalition to π_2 , and return the resulting partition $\pi_3 = \pi_2 \cup \{\mathcal{U}_0\}$.

Otherwise, we repeat the step (*) described below until either $|\mathcal{U}| \leq \frac{n}{2} + 1$ or all agents in \mathcal{U} have at least two friends in \mathcal{U} :

- (*) Pick an agent $i \in \mathcal{U}$ with at most one friend in \mathcal{U} . If the unique friend of i in \mathcal{U} is j and $F_j \cap \mathcal{U} = \{i\}$, add $\{i, j\}$ to the partition and remove i, j from \mathcal{U} . Otherwise, identify a coalition S in the current partition that contains a friend of i

(this is always possible since by construction of \mathcal{U} we have $|F_i| \geq 2$ for each $i \in \mathcal{U}$), add i to S , and remove i from \mathcal{U} .

Observe that after step (*) our partition may contain coalitions that are not stars: e.g., the first time we execute this step, we may add an agent i to a star $S(r)$ such that r is a friend of i , but at a later iteration of step (*) we may add a friend of i to this coalition.

Let \mathcal{U}_0 be what remains of \mathcal{U} at the end of this process. Place all agents in \mathcal{U}_0 in the same coalition, and add it to the partition; let π_3 be the resulting coalition structure. This completes the description of our algorithm. We now claim that its approximation ratio is at most $2 - \frac{4}{n+2}$.

Note that in π_3 each agent is in a coalition with at least one friend, and consider a coalition $S \in \pi_3$.

If $|S| \leq \frac{n}{2} + 1$, then for each $i \in S$ we have

$$u_i(S) \geq 1 - \frac{|S| - 2}{n} \geq 1 - \frac{\frac{n}{2} - 1}{n} = \frac{1}{2} + \frac{1}{n} = \frac{n+2}{2n}. \quad (5)$$

On the other hand, suppose that $|S| > \frac{n}{2} + 1$. This can happen if (i) S was constructed during stages 1 and 2, or (ii) $S = \mathcal{U}_0$ and each agent in \mathcal{U}_0 has at least two friends in \mathcal{U}_0 . Note, in particular, that we cannot create coalitions of that size in step (*), since this step is only executed when there are at least $\frac{n}{2} + 1$ agents in \mathcal{U} . In case (i), we have argued that $u_i(S) \geq \text{opt}$ for each $i \in S$, and in case (ii), the utility of each agent in S is at least $2 - \frac{|S|-3}{n} > 1 \geq \text{opt}$.

Thus, if for some agent $i \in \mathcal{N}$ we have $u_i(\pi_3) < \text{opt}$, then $u_i(\pi_3) \geq \frac{n+2}{2n}$. On the other hand, $\mathcal{N}_1 \neq \emptyset$ implies $\text{opt} \leq 1$. Hence, the approximation ratio of our algorithm does not exceed $\frac{2n}{n+2} = 2 - \frac{4}{n+2}$, as desired. \square

In the full version of this work [18], the pseudocode of the algorithm described in the proof of Theorem 8 is presented together with a discussion of its polynomial-time complexity. Moreover, we also demonstrate that our analysis is tight.

Finally, we show that, for symmetric instances where G^{sf} is a forest, the problem of finding a MAXESW partition is polynomial-time solvable. To this end, we establish a connection to a different class of hedonic games, namely, *fractional hedonic games (FHGs)* [1]. In these games, too, agents assign values to each other, with agent i assigning value $v_i(j) \in \mathbb{R}$ to agent j , and the utility an agent $i \in \mathcal{N}$ assigns to a coalition $C \in \mathcal{N}^i$ is computed as $\frac{1}{|C|} \sum_{j \in C \setminus \{i\}} v_i(j)$. An FHG with a set of agents \mathcal{N} is *simple* if $v_i(j) \in \{0, 1\}$ for all $i, j \in \mathcal{N}$; it is *symmetric* if $v_i(j) = v_j(i)$ for all $i, j \in \mathcal{N}$. A simple symmetric FHG \mathcal{G} can be represented by an undirected graph $G(\mathcal{G}) = (\mathcal{N}, \mathcal{E})$ where $\{i, j\} \in \mathcal{E}$ iff $v_i(j) = v_j(i) = 1$. Thus, an undirected graph G with vertex set \mathcal{N} induces two different games on the set of agents \mathcal{N} : a symmetric FA game and a simple symmetric FHG. We will now establish a tight connection between the structure of optimal solutions of these games when G is acyclic.

LEMMA 3. *Consider a simple symmetric FHG \mathcal{G}_1 and a symmetric FA \mathcal{G}_2 with the same set of agents \mathcal{N} such that $G(\mathcal{G}_1)$ is acyclic and coincides with the strong friendship graph of \mathcal{G}_2 . Then any partition π^* that maximizes the ESW in \mathcal{G}_1 can be transformed in polynomial time into a partition π' that maximizes the ESW in \mathcal{G}_2 .*

PROOF. Assume without loss of generality that $G(\mathcal{G}_1)$ is a tree; otherwise, we apply the subsequent argument to each connected component of $G(\mathcal{G}_1)$.

Let π be a partition of \mathcal{N} , and consider a coalition $C \in \pi$. Suppose first that C is not connected. Then splitting C into its connected components is a Pareto improvement for all agents in both \mathcal{G}_1 and \mathcal{G}_2 and hence does not decrease the ESW in either game. Now, suppose C is not a star. Then C has two non-leaf agents connected by an edge; let e be some such edge, and let C' and C'' be the coalitions obtained by splitting C along e . In each of C , C' and C'' , the lowest-utility vertices are leaves (this holds both in FA games and in FHGs), and the utility of a leaf in C' and C'' is higher than the utility of a leaf in C . Thus, this split, too, does not decrease the ESW in either game. It follows that we can split each coalition in π into stars without lowering the ESW in either game.

Now, let π^* be a partition that maximizes ESW in \mathcal{G}_1 . Let π' be a partition obtained from π^* by splitting each non-star coalition in π^* into stars of size at least 2; note that π' can be obtained from π^* in polynomial time. Moreover, the argument above shows that π' maximizes ESW in \mathcal{G}_1 . Observe that, in both games, the agents with the lowest utility in π' are the leaves of the largest stars in π' .

We claim that π' is optimal for \mathcal{G}_2 . Indeed, suppose not, and let s be the size of the largest star in π' . Let π'' be an optimal partition for \mathcal{G}_2 . By the argument above, we can assume that π'' is a collection of stars, so for it to have a higher ESW in \mathcal{G}_2 than π' , the largest star in π'' must be of size $s'' < s'$. But then the ESW in \mathcal{G}_1 is higher than that of π' , a contradiction. \square

If the graph $G(\mathcal{G})$ associated with an FHG \mathcal{G} is acyclic, a MAXESW partition can be computed in polynomial time [4, 25]. Together with Lemma 3, this implies the following result.

THEOREM 9. *Under FA, if the game instance is symmetric and G^{sf} is a forest, then MAXESW is in P.*

While we have just shown that, when the underlying graph is a tree, solving MAXESW is equivalent in FA and in sFHGs, this equivalence no longer holds for general graph structures.

Example 4. Consider a symmetric instance where $n \geq 4$ is an even number and G^{sf} is an n -vertex cycle. Under FA, the unique optimum is the grand coalition, whereas in sFHGs, every optimal partition corresponds to a perfect matching in G^{sf} .

5 CONCLUSIONS AND FUTURE WORK

We studied the problem of maximizing the egalitarian welfare in Friends and Enemies Games, focusing on two classical preference models: Enemies Aversion (EA) and Friends Appreciation (FA). We provided an almost complete picture of the complexity of this problem, by establishing both hardness results and polynomial-time approximation guarantees. A natural next step is to determine whether a constant-factor approximation is indeed possible for FA or to derive an inapproximability result. Moreover, employing the egalitarian welfare as a measure of the quality of stable outcomes, similarly to the work of Monaco et al. [29], also represents a promising direction for future work. More broadly, it would be interesting to explore egalitarian welfare in other variants of Friends and Enemies Games, such as those that incorporate neutral agents (with small positive or negative effects) or assign a different negative value for enemies instead of the standard $-\frac{1}{n}$ or $-n$. Finally, while the utilitarian and egalitarian objectives have received attention, the Nash welfare in hedonic games remains unexplored.

ACKNOWLEDGMENTS

This work was partially supported by: PNRR MIUR project FAIR - Future AI Research (PE00000013), Spoke 9 - Green-aware AI; PNRR MIUR project VITALITY (ECS00000041), Spoke 2 - Advanced Space Technologies and Research Alliance (ASTRA); the European Union - Next Generation EU under the Italian PNRR, Mission 4, Component 2, Investment 1.3, CUP J33C22002880001, partnership on “Telecommunications of the Future”(PE00000001 - program “RESTART”), project MoVeOver/SCHEDULE (“Smart interseCtions with conneCteD and aUtonomous vehicLEs”, CUP J33C22002880001); and GNCS-INdAM project, CUP_E53C24001950001.

REFERENCES

- [1] Haris Aziz, Florian Brandl, Felix Brandt, Paul Harrenstein, Martin Olsen, and Dominik Peters. 2019. Fractional hedonic games. *ACM Transactions on Economics and Computation (TEAC)* 7, 2 (2019), 1–29.
- [2] Haris Aziz, Felix Brandt, and Paul Harrenstein. 2013. Pareto optimality in coalition formation. *Games and Economic Behavior* 82 (2013), 562–581.
- [3] Haris Aziz, Felix Brandt, and Hans Georg Seedig. 2013. Computing desirable partitions in additively separable hedonic games. *Artificial Intelligence* 195 (2013), 316–334.
- [4] Haris Aziz, Serge Gaspers, Joachim Gudmundsson, Julián Mestre, and Hanjo Täubig. 2015. Welfare maximization in fractional hedonic games. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence, IJCAI 2015*. 461–467.
- [5] Haris Aziz and Rahul Savani. 2016. Hedonic games. In *Handbook of Computational Social Choice*. 356–376.
- [6] Suryapratim Banerjee, Hideo Konishi, and Tayfun Sönmez. 2001. Core in a simple coalition formation game. *Social Choice and Welfare* 18, 1 (2001), 135–153.
- [7] Nathanaël Barrot, Kazunori Ota, Yuko Sakurai, and Makoto Yokoo. 2019. Unknown agents in friends oriented hedonic games: Stability and complexity. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence, AAAI 2019*. 1756–1763.
- [8] Anna Bogomolnaia and Matthew O Jackson. 2002. The stability of hedonic coalition structures. *Games and Economic Behavior* 38, 2 (2002), 201–230.
- [9] Martin Bullinger, Vaggos Chatziafratis, and Parnian Shahkar. 2025. Welfare approximation in additively separable hedonic games. In *Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2025*. 418–426.
- [10] Jiehua Chen, Gergely Csáji, Sanjukta Roy, and Sofia Simola. 2023. Hedonic games With friends, enemies, and neutrals: Resolving open questions and fine-grained complexity. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2023*. 251–259.
- [11] Saar Cohen and Noa Agmon. 2025. Egalitarianism in online coalition formation. In *Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2025*. 2475–2477.
- [12] Andrei Constantinescu, Pascal Lenzen, Rebecca Reiffenhäuser, Daniel Schmand, and Giovanna Varricchio. 2024. Solving Woeginger’s hiking problem: Wonderful partitions in anonymous hedonic games. In *Proceedings of the 51st International Colloquium on Automata, Languages, and Programming, ICALP 2024*. 48:1–48:18.
- [13] Argyrios Deligkas, Eduard Eiben, Stavros D Ioannidis, Dušan Knop, and Simon Schierreich. 2025. Balanced and fair partitioning of friends. In *Proceedings of the 39th AAAI Conference on Artificial Intelligence, AAAI 2025*. 13754–13762.
- [14] Dinko Dimitrov, Peter Borm, Ruud Hendrickx, and Shao Chin Sung. 2006. Simple priorities and core stability in hedonic games. *Social Choice and Welfare* 26, 2 (2006), 421–433.
- [15] Dinko Dimitrov and Shao Chin Sung. 2004. Enemies and friends in hedonic games: Individual deviations, stability and manipulation. *SSRN Electronic Journal* (02 2004). <https://doi.org/10.2139/ssrn.639483>
- [16] Jacques H Dreze and Joseph Greenberg. 1980. Hedonic coalitions: Optimality and stability. *Econometrica: Journal of the Econometric Society* (1980), 987–1003.
- [17] Edith Elkind, Angelo Fanelli, and Michele Flammini. 2020. Price of Pareto optimality in hedonic games. *Artificial Intelligence* 288 (2020), 103357.
- [18] Edith Elkind, Michele Flammini, and Giovanna Varricchio. 2025. Maximizing the Egalitarian Welfare in Friends and Enemies Games. *arXiv preprint arXiv:2512.20261* (2025).
- [19] Edith Elkind and Michael Wooldridge. 2009. Hedonic coalition nets. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2009*. 417–424.
- [20] Michele Flammini, Maria Fomenko, and Giovanna Varricchio. 2025. Non-obvious manipulability in hedonic games with friends appreciation preferences. In *Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2025*. 767–775.
- [21] Michele Flammini, Bojana Kodric, and Giovanna Varricchio. 2022. Strategyproof mechanisms for friends and enemies games. *Artificial Intelligence* 302 (2022), 103610.
- [22] Martin Gairing and Rahul Savani. 2010. Computing stable outcomes in hedonic games. In *Proceedings of the 3rd International Symposium on Algorithmic Game Theory, SAGT 2010*. 174–185.
- [23] M. R. Garey and David S. Johnson. 1979. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman.
- [24] Tesshu Hanaka, Airi Ikeyama, and Hirotaka Ono. 2025. Maximizing utilitarian and egalitarian welfare of fractional hedonic games on tree-like graphs. *Journal of Combinatorial Optimization* 49, 3 (2025), 53.
- [25] Tesshu Hanaka, Hironori Kiya, Yasuhide Maei, and Hirotaka Ono. 2019. Computational complexity of hedonic games on sparse graphs. In *Proceedings of the 22nd International Conference on Principles and Practice of Multi-Agent Systems, PRIMA 2019*. 576–584.
- [26] Pavol Hell and David G Kirkpatrick. 1984. Packings by cliques and by finite families of graphs. *Discrete Mathematics* 49, 1 (1984), 45–59.
- [27] Anna Maria Kerkmann, Nhan-Tam Nguyen, Anja Rey, Lisa Rey, Jörg Rothe, Lena Schend, and Alessandra Wiechers. 2022. Altruistic hedonic games. *Journal of Artificial Intelligence Research* 75 (2022), 129–169.
- [28] Bettina Klaus, Flip Klijn, and Seçkin Özbilen. 2025. Core stability and strategy-proofness in hedonic coalition formation problems with friend-oriented Preferences. *Games and Economic Behavior* 154 (2025), 16–52.
- [29] Gianpiero Monaco, Luca Moscardelli, and Yllka Velaj. 2019. On the performance of stable outcomes in modified fractional hedonic games with egalitarian social welfare. In *Proceedings of the 18th International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2019*. 873–881.
- [30] Dominik Peters. 2016. Complexity of hedonic games with dichotomous preferences. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence, AAAI 2016*. 579–585.
- [31] Dominik Peters. 2016. Graphical hedonic games of bounded treewidth. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence, AAAI 2016*. 586–593.
- [32] Jörg Rothe, Hilmar Schadrack, and Lena Schend. 2018. Borda-induced hedonic games with friends, enemies, and neutral players. *Mathematical Social Sciences* 96 (2018), 21–36.
- [33] Naftali Waxman, Sarit Kraus, and Noam Hazon. 2020. On maximizing egalitarian value in k-coalitional hedonic games. *arXiv preprint arXiv:2001.10772* (2020).
- [34] David Zuckerman. 2007. Linear degree extractors and the inapproximability of max clique and chromatic number. *Theory of Computing* 3, 1 (2007), 103–128.