Individual Rationality in Constrained Hedonic Games: Friends, Enemies, and Neutrals

Šimon Schierreich[1,2] and Ildikó Schlotter[3,4]

[1] AGH University of Krakow, Poland
[2] Czech Technical University in Prague, Czechia
[3] ELTE Centre for Economic and Regional Studies, Hungary
[4] Budapest University of Technology and Economics, Hungary
schiesim@fit.cvut.cz, schlotter.ildiko@krtk.elte.hu

Abstract. We study constrained coalition formation in games induced by *friends*, *enemies*, and *neutrals*, under the two standard refinements of additively separable preferences: *friend-oriented* and *enemy-oriented*. We ask for partitions that are *individually rational* (IR), while additionally requiring exactly k non-empty coalitions, each satisfying a prescribed lower and upper bound on its size. Although IR alone is trivial to satisfy for any hedonic game, the size constraints make it computationally intractable to decide whether a feasible partition exists.

The two models tell strikingly different stories. Under enemy-oriented preferences, the problem collapses to size-constrained graph coloring, and its complexity follows accordingly. Under friend-oriented preferences, however, the picture is far more intricate, and is governed by the *enmity structure* rather than the friendships. The complexity is further shaped by two factors: how strict the imposed size requirements are, and whether relationships are *symmetric* or *asymmetric*, with several cases turning out tractable in the symmetric setting but intractable once asymmetry is allowed. Charting this boundary in terms of both classical and parameterized complexity, we provide a complete understanding of which properties of the friend/enemy structure are responsible for hardness.

Keywords: Coalition Formation · Hedonic Games · Friends and Enemies · Individual Rationality · Parameterized Complexity.

1 Introduction

Imagine that you are running a summer camp and must assign the arriving children to the camp's cabins for the week. The children are far from indifferent about who they bunk with: some are close friends who would love to stay together, a few pairs genuinely dislike each other and would make life miserable for their whole cabin if put together, and most are simply indifferent to the rest. The assignment is also bound by hard practical limits. Each cabin has a maximum capacity; no cabin should be opened with only a child or two inside, and the number of cabins available for the week is fixed in advance.

This situation, and many others like it, captures the essence of the broader problem of *coalition formation*. The task is to partition a set of agents into *coalitions* so that the resulting configuration is *stable*[5]—that is, no agent would prefer to deviate from their current coalition. If we further assume that the agents are *self-interested* and care only about the composition of their own coalition, we obtain the well-known model of *hedonic games* [4,23]. In this work, we study hedonic games induced by *friends*, *enemies*, and *neutrals* [16,22,36,39]: each agent partitions the others into those they like, those they dislike, and those they are indifferent to. This class is a special case of a more general model of *additively separable hedonic games* (ASHGs) [5,8]. Following the literature, we consider the two canonical refinements of such preferences, *friend-oriented* and *enemy-oriented*, which differ in whether an agent primarily seeks the company of their friends or primarily seeks to avoid their enemies.

The notion of stability we adopt is that of *individual rationality* (IR) [2,20,38]. Intuitively, IR requires that no agent prefers leaving their current coalition and forming a coalition on their own. Traditionally, any coalition structure is acceptable as long as it meets the chosen stability criterion. In many real-world scenarios, however—from assigning children to camp cabins to forming student project teams or assigning researchers to working groups—we face additional structural constraints on feasible outcomes:

1. We require that the resulting coalition structure consists of exactly k non-empty coalitions. This constraint was introduced by Sless *et al.* [40] and subsequently studied by Waxman *et al.* [41] and Barr *et al.* [6].
2. The size of each of the k coalitions must be within its own lower and upper bound.

It is easy to see that the first constraint is a special case of the second, where the lower bound of each coalition is one, and the upper bound equals the number of agents. The general model combining these constraints with IR was introduced by Fioravantes *et al.* [25] for ASHGs, building on the fixed-size coalitions of Biló *et al.* [7] and a line of related work on fixed- and equal-size coalitions [1,19,35] (see also the recent model of Bullinger *et al.* [13] with global size constraints). The particular structural restrictions we study here were subsequently investigated by Fioravantes *et al.* [26] for anonymous and diversity preferences [10,29].

1.1 Our Contribution

We provide a detailed study of individual rationality under the two structural constraints described above, considering four variants that differ in how the coalition sizes are restricted. In increasing order of generality, these are: the *k-coalition* variant, which only fixes the number of non-empty coalitions; the *balanced-size* (BSC) variant, where any two coalition sizes may differ by at most one; the *fixed-size* (FSC) variant, where each coalition must have a prescribed size; and the *size-constrained* (SCC) variant, where each coalition is assigned its own lower

[5] Other desirable criteria are also studied; e.g., various notions of *efficiency* [11,12,30].

Table 1. An overview of our complexity results for friend-oriented preferences. Here, k is the number of desired coalitions, $\mathrm{ub} = \max_{j\in[k]} \mathrm{ub}(j)$, r denotes the number of recluses (i.e., agents a with $\mathcal{F}_a = \emptyset$ and $\mathcal{E}_a \neq \emptyset$), vc (and $\mathrm{vc}^{\mathcal{E}}$), tw, and cw the vertex cover number of the relationship graph (the enmity graph, resp.), its treewidth, and its clique-width, and ξ the number of conflicting agents (i.e., agents a with $\mathcal{E}_a \neq \emptyset$); in the rows labeled using $|\mathcal{E}_a|$, the stated bound on $|\mathcal{E}_a|$ applies to all agents $a \in A$. Entries of the form "X / Y" display the complexity for symmetric instances (left of the slash) and for asymmetric instances (right); a single entry applies to both settings. NP-h (W[1]-h) means para-NP-hardness (W[1]-hardness) for the given parameter value.

	IR-FO-k	IR-FO-BSC	IR-FO-FSC	IR-FO-SCC
$k = 1$	P [O2]	P [O2]	P [O2]	P [O2]
$k \geq 2$	NP-h [T2]	NP-h [T2]	NP-h [T2]	NP-h [T2]
$\mathrm{ub} \leq 2$	N/A	P [T4]	P [T4]	P [T4]
$\mathrm{ub} \geq 3$	N/A	NP-h [T3]	NP-h [T3]	NP-h [T3]
$r = 0$	P [T8]/NP-h [T9]	NP-h [T3]	NP-h [T3]	NP-h [T3]
$r = 1$	?/NP-h [T9]	NP-h [T3]	NP-h [T3]	NP-h [T3]
$r \geq 2$	NP-h [T2]	NP-h [T3]	NP-h [T3]	NP-h [T3]
vc	FPT [T16]	FPT [T19]	FPT [T19]	?
tw $+k$	FPT [T16]	W[1]-h, XP [T17,T18]	W[1]-h, XP [T17,T18]	W[1]-h, XP [T17,T18]
tw	FPT [T16]	W[1]-h [T18]	W[1]-h [T18]	W[1]-h [T18]
cw	W[1]-h [T18]	NP-h [T18]	NP-h [T18]	NP-h [T18]
$\mathrm{vc}^{\mathcal{E}} = 1$	P [T10]	P [T11]/NP-h [T12]	P [T11]/NP-h [T12]	P [T11]/NP-h [T12]
$\mathrm{vc}^{\mathcal{E}} = 2$	P [T13]/NP-h [T14]	?/NP-h [T14]	?/NP-h [T14]	?/NP-h [T14]
$\mathrm{vc}^{\mathcal{E}} = 3$	P [T13]/NP-h [T14]	NP-h [T14]	NP-h [T14]	NP-h [T14]
$\mathrm{vc}^{\mathcal{E}} \geq 4$	NP-h [T14]	NP-h [T14]	NP-h [T14]	NP-h [T14]
ξ	FPT [T15]	FPT [T15]	FPT [T15]	?
$\lvert\mathcal{E}_a\rvert = 0$	P [O1]	P [O1]	P [O1]	P [O1]
$\lvert\mathcal{E}_a\rvert = 1$	P [T5]/NP-h [T14]	P [T5]/NP-h [T12]	?/NP-h [T12]	?/NP-h [T12]
$\lvert\mathcal{E}_a\rvert = 2$	?/NP-h [T14]	?/NP-h [T12]	?/NP-h [T12]	?/NP-h [T12]
$\lvert\mathcal{E}_a\rvert = 3$	?/NP-h [T14]	?/NP-h [T12]	NP-h [T7]	NP-h [T7]
$\lvert\mathcal{E}_a\rvert \geq 4$	NP-h [T6]	NP-h [T6]	NP-h [T6]	NP-h [T6]

and upper bounds. For each of these, we chart the complexity of deciding whether an individually rational outcome exists, as summarized in Table 1. We approach the problem from several dimensions: the number of desired coalitions and the permitted coalition sizes; whether relationships are symmetric or asymmetric; the maximum number of enemies per agent; the number of conflicting agents (those having enemies); the number of recluses (agents having enemies but no friends); and structural parameters of the underlying relationship and enmity graphs, namely their vertex cover number, treewidth, and clique-width.

Our first finding is that the two preference orientations could hardly behave more differently. Under enemy-oriented preferences, friendships provably play no role: all four variants collapse to size-constrained versions of the classic k-Coloring problem (Section 3), and their complexity is thus inherited from the graph-coloring literature. Under friend-oriented preferences, no such collapse occurs—the complexity is governed by the enmity structure in a far subtler way, and charting this landscape constitutes the technical core of the paper.

Since friend-oriented (and enemy-oriented) preferences form a special case of additively separable hedonic games, our results directly extend those of Fioravantes *et al.* [25] to this prominent preference domain. Our techniques, however, are quite different. Whereas Fioravantes *et al.* build their algorithms largely on N-fold integer formulations, we draw on a broad range of tools from the parameterized complexity toolbox—including color coding, bounded search trees, and dynamic programming over nice tree decompositions. On the negative side, our hardness reductions start from a variety of source problems and, owing to the nuanced picture we paint, require careful, technically involved constructions.

1.2 Related Work

Friends-and-enemies games were introduced by Dimitrov *et al.* [22]. Ohta *et al.* [36] later extended the model with neutral agents. The existence and verification of stable outcomes in this domain have since been studied for a range of stability notions [16, 33, 39]. Importantly, Chen *et al.* [16] studied parameterized complexity of deciding the existence of stable solutions for the friends–enemies–neutrals model. None of these works, however, considers individual rationality and/or structural constraints on the outcome.

Friends-and-enemies games form a subclass of *additively separable hedonic games* (ASHGs) [8] which, together with *fractional hedonic games*, (FHGs) [3] are among the most studied classes of hedonic games; see, e.g., [9, 24, 37]. As IR is trivially satisfiable in any such game without further restrictions, these works focus on more demanding notions of stability. An exception is the work of Fioravantes *et al.* [25], who studied the same model as we do, albeit for the more general class of ASHGs. Their results, however, do not pertain to our setting: their hardness results do not transfer to a restricted preference domain, and since expressing friend- or enemy-oriented preferences by additive valuations requires weights linear in the number of agents, their algorithmic results—which achieve fixed-parameter tractability only when the maximum weight is taken as an additional parameter—yield at best XP algorithms in our setting. A similar model was also studied by Fioravantes *et al.* [26]; however, for a different preference model incomparable with ours.

More broadly, across the many subclasses of hedonic games, individual rationality is rarely studied as a solution concept: either no coalition is worse than the singleton, so IR is trivially achievable even under our constraints, or it was simply not explored. The one exception we are aware of is the work of Deligkas *et al.* [20], who study IR in *topological distance games* [14]; that model is incomparable to ours, and their (mostly negative) results do not carry over.

Restrictions on the admissible coalition structures other than ours have also been considered. Beyond the k-coalition and fixed-size models discussed above, Wright and Vorobeychik [42] introduced a bounded maximum coalition size, studied further in several follow-ups [24, 29, 34]. Even under this restriction, however, IR remains trivial, as one may always form arbitrarily many singletons.

2 Preliminaries

We assume the reader is familiar with the basics of graph theory [21] and both classical and parameterized complexity [17].

We consider a set A of n *agents*. Each agent $a \in A$ is associated with a set of *friends* $\mathcal{F}_a \subseteq A \setminus \{a\}$ and a set of *enemies* $\mathcal{E}_a \subseteq A \setminus \{a\}$ so that $\mathcal{F}_a \cap \mathcal{E}_a = \emptyset$. Note that $(\mathcal{F}_a, \mathcal{E}_a)$ is not necessarily a partition of $A \setminus \{a\}$: agents in $A \setminus (\mathcal{F}_a \cup \mathcal{E}_a \cup \{a\})$ are *neutrals* for a; we denote their set as $\mathcal{N}_a$. We can encode the friends–enemies–neutral relations using a directed edge-colored *relationship graph* $G = (A, E, \mathrm{c})$, where for each a and every $b \in \mathcal{F}_a \cup \mathcal{E}_a$, E contains an arc (a, b). This arc is colored *blue* if b is a friend of a and *red* if b is a's enemy. We use $G^{\mathcal{F}}$ to denote the *friendship graph*, i.e., the graph G reduced to the blue arcs, and we use $G^{\mathcal{E}}$ to denote the *enmity graph* constructed from G by leaving only the red arcs. Agents' relations are *symmetric* if $b \in \mathcal{F}_a \iff a \in \mathcal{F}_b$ and $b \in \mathcal{E}_a \iff a \in \mathcal{E}_b$ hold for every $a, b \in A$. In this case, the graph G is undirected.

Any subset $C \subseteq A$ of agents A is called a *coalition*. We say that $C \subseteq A$ is a *singleton* if $|C| = 1$, and C is the *grand coalition* if $C = A$. By $\mathcal{C}_a$ we denote the set of all coalitions that contain agent $a \in A$. The preferences of an agent $a \in A$ is a weak ordering $\succeq_a$ over $\mathcal{C}_a$. For two coalitions $C, D \in \mathcal{C}_a$, we use $C \succ_a D$ to represent that agent a strictly prefers coalition C over D, i.e., $C \succeq_a D$ and $D \not\succeq_a C$. Similarly, we use $C \sim D$ if agent a is indifferent between C and D, that is, $C \succeq_a D$ and $D \succeq_a C$. A list $\mathcal{P} = (\succeq_a)_{a \in A}$ is called the *preference profile*. Throughout the paper, we consider the following two restrictions of agents' preferences based on the sets of friends and enemies:

- A preference profile $\mathcal{P}$ is *friend-oriented* if for every $a \in A$ and each pair of $C, D \in \mathcal{C}_a$ it holds that $C \succeq_a D$ if and only if (1) $|C \cap \mathcal{F}_a| > |D \cap \mathcal{F}_a|$ or (2) $|C \cap \mathcal{F}_a| = |D \cap \mathcal{F}_a|$ and $|C \cap \mathcal{E}_a| \leq |D \cap \mathcal{E}_a|$.
- A preference profile $\mathcal{P}$ is *enemy-oriented* if for every $a \in A$ and each pair of $C, D \in \mathcal{C}_a$ it holds that $C \succeq_a D$ if and only if (1) $|C \cap \mathcal{E}_a| < |D \cap \mathcal{E}_a|$ or (2) $|C \cap \mathcal{E}_a| = |D \cap \mathcal{E}_a|$ and $|C \cap \mathcal{F}_a| \geq |D \cap \mathcal{F}_a|$.

A *coalition structure* $\pi = (\pi_1, \ldots, \pi_k)$ is a partition of agents into coalitions, that is, $\bigcup_{i=1}^{k} \pi_i = A$ and $\pi_i \cap \pi_j = \emptyset$ for each pair of distinct $i, j \in [k]$. For an agent $a \in A$ and a coalition structure π, we use $\pi(a)$ to denote the coalition that agent a is part of. We extend the preferences from coalitions to coalition structures by setting $\pi \succeq_a \pi'$ whenever $\pi(a) \succeq_a \pi'(a)$.

The goal of a *hedonic game* $\Gamma = (A, \mathcal{P})$ is to find a *desirable* coalition structure π, where *desirable* is some well-defined requirement on the solution coalition structure π. In this paper, the desirable coalition structures are those that meet certain *stability* criteria. The notion of stability that we adopt in this work is that of *individual rationality* (IR), which requires that no agent strictly prefers to be in a singleton coalition over staying in their current coalition.

Definition 1. *Let Γ be a hedonic game, and π a coalition structure for Γ. We say that π is* individually rational *(IR) if $\pi(a) \succeq_a \{a\}$ for every agent $a \in A$.*

For our two preference restrictions, individual rationality admits a simple combinatorial characterization, which we use throughout the paper in place of the preference relations themselves.

Observation 1. *Let $\Gamma = (A, \mathcal{P})$ be a hedonic game and π a coalition structure.*

(a) If $\mathcal{P}$ is friend-oriented, then π is IR if and only if for each agent $a \in A$ we have that $\pi(a)$ contains a friend of a or no enemy of a.
(b) If $\mathcal{P}$ is enemy-oriented, then π is IR if and only if for each agent $a \in A$ we have that $\pi(a)$ contains no enemy of a.

Nevertheless, it is trivial to find an individually rational coalition structure for any hedonic game—simply place each agent into a coalition by themselves. Therefore, we consider the following four natural *constraints*:

1. A *hedonic game with size-constrained coalition* (HG-SCC) is a tuple $(A, \mathcal{P}, k, \mathrm{lb}, \mathrm{ub})$ where $k \in [n]$ is the desired number of coalitions, and lb and ub are functions from $[k]$ to $\{0, 1, \ldots, n\}$ setting a lower and upper bound on the size of each coalition. The solution structure π of such a game must satisfy $\pi = (\pi_1, \ldots, \pi_k)$ where $\mathrm{lb}(j) \leq |\pi_j| \leq \mathrm{ub}(j)$ for each $j \in [k]$.
2. A *hedonic game with k coalitions* (k-HG) is HG-SCC where $\mathrm{lb}(j) = 1$ and $\mathrm{ub}(j) = n$ for every $j \in [k]$.
3. A *hedonic game with fixed-size coalitions* (HG-FSC) is HG-SCC with $\mathrm{lb}(j) = \mathrm{ub}(j)$ for every $j \in [k]$.
4. A *hedonic game with balanced-size coalitions* (HG-BSC) is HG-FSC such that $\mathrm{ub}(j) \in \{\lfloor n/k \rfloor, \lceil n/k \rceil\}$ for every $j \in [k]$.

The relationships between the constraints are as follows: HG-BSC is a special case of HG-FSC, which in turn is a special case of HG-SCC. k-HG is a special case of HG-SCC and is incomparable with any other restriction. Naturally, a hardness result for a special case implies hardness for a more general model, and algorithmic upper bounds are implied in the opposite direction.

Given an HG-SCC Γ with friend-oriented preferences, the problem of deciding whether Γ admits an IR coalition structure is denoted IR-FO-SCC; the problems IR-FO-FSC, IR-FO-BSC, and IR-FO-k, as well as their enemy-oriented variants, are defined analogously. For fixed-size (or balanced-size) coalitions, we will assume that the given sizes sum up to n. Since we can check efficiently whether a coalition structure is IR, all studied problems are in NP.

3 Enemy-Oriented Preferences

Our first result settles the enemy-oriented case completely, by showing that it collapses to a classic coloring problem. Indeed, if preferences are enemy-oriented, then friendships play no role in whether a given coalition structure is IR: being alone is *not* preferred by an agent a to being in a coalition C if and only if a has no enemies in C. Therefore, individual rationality of a coalition structure π is equivalent to each coalition forming an independent set in the enmity graph $G^{\mathcal{E}}$; notice that enmities can be assumed to be symmetric. That is, our problems are equivalent to the size-constrained variants of the classic k-Coloring problem. Since Equitable 3-Coloring is NP-complete already for 4-regular graphs, this immediately implies the following:

Observation 2. *IR-EO-k and IR-EO-BSC are NP-complete even if $k = 3$ and $|\mathcal{E}_a| = 4$ for every $a \in A$.*

By contrast, it is not hard to see that deciding whether a graph H can be 2-colored so that the color classes respect some size constraint can be done in polynomial time, using a simple dynamic programming approach based on the observation that each connected component of H has a unique 2-coloring up to the reversal of the colors. This yields the following result.

Theorem 1. *IR-EO-SCC for $k = 2$ is polynomial-time solvable.*

The polynomial equivalence with size-constrained coloring settles the enemy-oriented case: any complexity result for these coloring problems, classical or parameterized, translates directly to our setting, and vice versa. No such collapse is possible under friend-oriented preferences, where friendships can compensate for the presence of enemies; the resulting—much richer—complexity landscape is the subject of the rest of the paper.

4 Friend-Oriented Preferences

Directing our attention to friend-oriented preferences, we investigate how different natural parameters influence the computational complexity of our problems.

4.1 Few or Small Coalitions

We show that our problems are already hard if we aim only for two coalitions, yielding a strict dichotomy in terms of the number of desired coalitions.

Theorem 2. *If $k = 1$, IR-FO-SCC can be decided in linear time. IR-FO-k and IR-FO-BSC are NP-complete for $k = 2$, even if preferences are symmetric, each agent has at most 8 friends, and there are only two agents with no friends.*

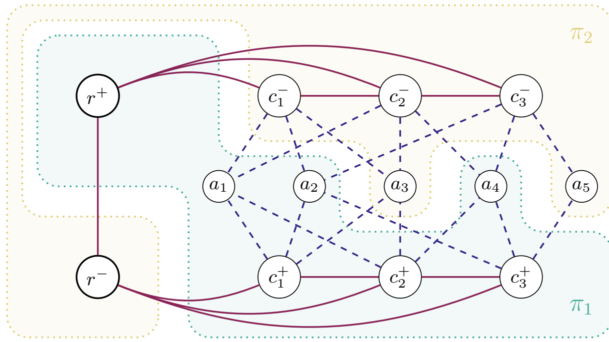

Fig. 1. An illustration of the hardness construction used to prove Theorem 2 for a formula $\varphi = (x_1 \vee x_2 \vee x_3) \wedge (x_1 \vee x_3 \vee x_4) \wedge (x_2 \vee x_4 \vee x_5)$. Red (solid) edges represent that two agents are enemies, while blue (dashed) edges represent friendship. Using colored backgrounds, we highlight one possible IR coalition structure.

Proof sketch. To show NP-hardness of IR-FO-k for $k = 2$, we give a reduction from NP-hard Monotone 4-Regular NAE 3-SAT [18]. The input of this problem is 3-CNF formula $\varphi = C_1 \wedge \cdots \wedge C_m$ over a set $X = \{x_1, \ldots, x_n\}$ of variables such that φ contains no negative literals and each variable appears in exactly four clauses. The question is whether φ admits a *valid* truth assignment, i.e., one in which each clause is satisfied by some but not all of its variables.

Given φ, we construct an equivalent instance $\mathcal{J}$ of FO-k-HG as follows (see Figure 1 for an illustration of the construction). We have one *variable agent* a_i for every variable $x_i \in X$, and two *clause agents* c_j^+ and c_j^- for every clause C_j. In addition, we create two *guard agents* r^+ and r^-, whom we let be enemies. We also let r^+ be enemies with all clause agents c_j^-, and similarly, r^- be enemies with all clause agents c_j^+. The guard agents have no friends. Furthermore, for each $j \in [m-1]$ the pairs (c_j^-, c_{j+1}^-) and (c_j^+, c_{j+1}^+) of clause agents are also enemies. Each clause agent c_j^+ or c_j^- has three variable agents as their friends, namely the ones corresponding to literals present in the clause C_j. All remaining relations are neutral. To finalize the construction, we set $k = 2$.

For correctness, assume first that α is a valid assignment for φ. Then the coalition structure π defined by $\pi_1 = \{r^+\} \cup \{c_j^+ \colon j \in [m]\} \cup \{a_i \colon \alpha(x_i) = \texttt{true}\}$ and $\pi_2 = A \setminus \pi_2$ is IR for $\mathcal{J}$, as can be checked in a straightforward way.

In the opposite direction, let $\pi = (\pi_1, \pi_2)$ be an IR coalition structure for $\mathcal{J}$. We first observe that r^+ and r^- cannot be members of the same coalition, since they do not have friends and are mutual enemies. By the same argument, no clause agent c_j^- can be in the same coalition with r^+, and no clause agent c_j^+ can be in the same coalition as r^-. Hence, without loss of generality, we can assume that $\{r^+, c_1^+, \ldots, c_m^+\} \subseteq \pi_1$ and $\{r^-, c_1^-, \ldots, c_m^-\} \subseteq \pi_2$. We construct an assignment α by setting $\alpha(x_i) = \texttt{true}$ if and only if $a_i \in \pi_1$ for every $i \in [n]$. For contradiction, assume that α is not valid for φ, i.e., there is a clause C_j with either all or none of its literals being true. In the former case, all friends of c_j^- are in π_1. However, c_j^- has at least one enemy in π_2, which contradicts the individual rationality of π. In the latter case, we apply the same argument on c_j^+. Consequently, α is a valid truth assignment, finishing the proof for IR-FO-k. For IR-FO-BSC, we simply add n dummy agents $d_1, \dots, d_n$ that are neutrals from the perspective of all agents. $\triangle$

Next, we show that aiming for small coalitions does not yield tractability: our problems are hard already if we require each coalition to be of size at most 3. The reduction is from the NP-hard $\{P_3, K_3, K_{1,3}\}$-Decomposition problem, which is known to be NP-hard even on graphs of maximum degree 3 [15]. Given a graph $H = (W, F)$ with $|F| = 3m$, this problem asks if F can be partitioned into sets $F_1, \dots, F_m$ of size exactly 3 so that each F_i forms a 3-edge path (P_3), a triangle (K_3), or a star with three edges ($K_{1,3}$). The main idea of our reduction is to introduce an agent corresponding to each edge of the input graph H, with two agents being friends if the corresponding edges share an end-vertex and being enemies otherwise. It is not hard to see that in such a game, a coalition of size 3 is IR iff the corresponding edges in $H \simeq P_3, K_3, K_{1,3}$.

Theorem 3. *IR-FO-BSC is NP-complete even if* $\mathrm{ub}(j) = 3$ *for* $\forall j \in [k]$*, preferences are symmetric,* $|\mathcal{F}_a| \in \{3, 4\}$ *for* $\forall a \in A$*, and there are no neutrals.*

Contrasting Theorem 3, we show that if each coalition must be either a pair of agents or a singleton, then our problems are solvable efficiently. Our algorithm relies on the observation that a coalition formed by two agents is IR if and only if there is no enmity between these two agents (in any direction). By searching for a matching of an appropriate size in the graph where two agents are connected by an edge if and only if neither considers the other an enemy, we can solve our instance in the same time necessary for finding a maximum matching.

Theorem 4. *If* $\mathrm{ub}(j) \leq 2$ *for every* $j \in [k]$*, then IR-FO-SCC* $\in$ P.

4.2 Small Number of Enemies

First, observe that if no agent has enemies, then any coalition structure is IR. Hence, we only need to check whether the size constraints can be respected. In fact, for this, we only need to check whether $n \geq \sum_{i \in [k]} \mathrm{lb}(i)$ and $n \leq \sum_{i \in [k]} \mathrm{ub}(i)$, yielding a linear-time algorithm.

Proposition 1. *If* $|\mathcal{E}_a| = 0$ *for every* $a \in A$*, then IR-FO-SCC* $\in$ P.

If each agent has only one enemy, then we can create a balanced coalition structure (i.e., where all coalition sizes differ by at most 1) using a simple and fast algorithm: It suffices to place each agent that has some enemy into a coalition not shared by their unique mutual enemy. To achieve this, we fix an ordering of the agents so that agents with at least one enemy precede those without enemies, and each agent is next to their enemy (if they have one). Applying the round robin method based on this ordering to place the agents into k coalitions, we obtain an IR coalition structure.

Theorem 5. *If $|\mathcal{E}_a| \leq 1$ for every $a \in A$ and enmities are symmetric, then IR-FO-BSC and IR-FO-k can be solved in polynomial time.*

Theorem 5 is contrasted by our next result, which shows that both IR-FO-BSC and IR-FO-k become intractable once each agent has four enemies, even if the number of coalitions is only $k = 3$. This result follows from the NP-hardness of the classic 3-coloring problem and its equitable variant, where the sizes of the color classes have to differ by at most 1; both of these problems are known to be NP-complete even on 4-regular graphs [28]. Our reduction is very simple: each vertex is represented by an agent with no friends, and two agents are enemies if and only if the corresponding vertices are adjacent in the input graph.

Theorem 6. *IR-FO-BSC and IR-FO-k are NP-complete even if preferences are symmetric, $k = 3$, and $|\mathcal{E}_a| = 4$ for every $a \in A$.*

For the variant where the coalition structure is not required to be balanced but the coalition sizes are still fixed, the problem remains hard even if each agent has exactly three enemies; to show this, we use a reduction from the NP-complete Semi-Equitable 3-Coloring problem [27] where the color classes of an n-vertex input graph need to be of size $4n/10$, $3n/10$, and $3n/10$.

Theorem 7. *IR-FO-FSC is NP-complete even if preferences are symmetric, $k = 3$, and $|\mathcal{E}_a| = 3$ for every $a \in A$.*

4.3 Small Number of Recluses

Notice that under friend-oriented preferences, agents without friends are usually the source of computational intractability, since they must be placed in a coalition that contains none of their enemies, and thus we may need to seek coalitions that form independent sets in the enemy graph—an idea utilized in our reduction from 3-Coloring and its variants in the proofs of Theorems 6 and 7. Hence, it makes sense to study how the number of such agents influences the complexity of our problems, motivating the following definition.

Definition 2. *An agent $a \in A$ is called a* recluse *if $\mathcal{F}_a = \emptyset$ and $|\mathcal{E}_a| \geq 1$.*

Notice that with upper and lower bounds on the number of coalition sizes, Theorem 3 implies that even symmetric instances with no recluses are intractable. For IR-FO-k, i.e., when only the number of coalitions is given, Theorem 2 shows that even two recluses suffice to cause NP-hardness. By contrast, IR-FO-k can be solved if preferences are symmetric and there are no recluses.

Theorem 8. *If there are no recluses, every symmetric IR-FO-k Γ admits an individually rational coalition structure which can be found in polynomial time.*

Proof. Recall that $k \leq n$. We first identify all connected components $K_1, \ldots, K_\ell$ of the friendship graph $G^{\mathcal{F}}$. Due to the absence of recluses, each K_j is either of size at least two or contains a single agent that has neither friends nor enemies.

If $k \leq \ell$, then we set $\pi_j = V(K_j)$ for each $j \in [k-1]$ and $\pi_k = \bigcup_{j=k}^{\ell} V(K_j)$. Then π has k non-empty coalitions and is IR due to our previous observation.

If $k > \ell$, then for every component K_j, $j \in [\ell]$, we compute a spanning tree T_j. We start by setting $\pi_j = V(K_j)$ for every $j \in [\ell]$. This is IR as shown above, but coalitions π_j for $j > \ell$ are empty. Hence, we do the following: let π_j be a non-empty coalition of size at least two and $\pi_{j'}$ be an empty coalition. We find agent $a \in \pi_j$ that is a leaf of the spanning tree T_j; we move a from π_j to $\pi_{j'}$. Now, agent a is in a singleton coalition, so $\pi_{j'}$ is IR for a. Since we removed a leaf from the tree of T_j, the coalition $\pi_j \setminus \{a\}$ remains also IR, since each agent of $\pi_j \setminus \{a\}$ still has at least one friend in $\pi_j \setminus \{a\}$—their neighbor in $T_j \setminus \{a\}$. The procedure terminates once all k coalitions are non-empty. The invariant of individual rationality is satisfied after each transfer. $\square$

Unfortunately, the algorithm of Theorem 8 cannot be extended to the asymmetric case, as our next result shows.

Theorem 9. *IR-FO-k is NP-complete even if there are no recluses.*

Proof sketch. We reduce from 3-Coloring. Let H be the input graph. We construct an instance Γ of IR-FO-k as follows. Let us set $K = |V(H)|^2$. For each vertex $u \in V(H)$, we introduce a *vertex-agent* a_u and $K-1$ additional dummy agents; these K agents together will form a directed cycle of length K in the friendship graph $G^{\mathcal{F}}$. In fact, $G^{\mathcal{F}}$ will be the disjoint union of these $|V(H)|$ cycles, which immediately implies that no recluses are contained in our instance, as each agent has exactly one friend. Besides these friendships, we let each dummy agent consider every agent other than their unique friend and themselves an enemy. Each vertex agent considers all dummy agents other than their unique friend as an enemy. Finally, two vertex agents a_u and a_v consider each other mutual enemies if u and v are adjacent in H. Finally, we set $k = (K-1) \cdot |V(H)| + 3$. It is not hard to verify that H is 3-colorable if and only if Γ admits an IR coalition structure with exactly k non-empty coalitions. $\triangleleft$

Theorems 8 and 9 leave open only the symmetric case with exactly one recluse.

4.4 Restricting the Enmity Graph

Next, we examine how the structure of the enmity graph influences the tractability of our problems. We consider the vertex cover number of $G^{\mathcal{E}}$, denoted by $\mathrm{vc}(G^{\mathcal{E}})$. We start with a simple algorithm for the case when $\mathrm{vc}(G^{\mathcal{E}}) \leq 1$: the main idea is to place the unique problematic agent—the one incident to all edges in $G^{\mathcal{E}}$—into a singleton.

Theorem 10. *If $\mathrm{vc}(G^{\mathcal{E}}) \leq 1$, then IR-FO-$k$ can be solved in linear time.*

For symmetric instances, in the case $\mathrm{vc}(G^{\mathcal{E}}) \leq 1$, a polynomial-time algorithm based on a more careful case analysis can solve the IR-FO-SCC problem.

Theorem 11. *If $\mathrm{vc}(G^{\mathcal{E}}) \leq 1$, then IR-FO-SCC for symmetric instances can be solved in polynomial time.*

Proof. If there are no enmities, then we can solve the instance by Proposition 1. So we assume $\mathrm{vc}(G^{\mathcal{E}}) = 1$, and let c be the agent covering all edges in $G^{\mathcal{E}}$. We start by first guessing the coalition j containing c and its size $s_j \in [n]$ in a solution π. Next, we guess whether c has an enemy in $\pi_j = \pi(c)$.

If c has no enemies in π_j, then we can put an arbitrary set of $s_j - 1$ agents into a coalition C together with c from among $A \setminus \mathcal{E}_c$, and set C is the j-th coalition; if $|A \setminus \mathcal{E}_c| < s_j$, we can reject the current set of guesses. We can then distribute the remaining agents across the remaining coalitions arbitrarily in a way that respects the lower and upper bounds, if this is possible, since there are no enmities between the remaining agents (see the method in Proposition 1).

If c has some enemies in π_j, then we also guess a friend c' of c in π_j. Next, we start adding agents to a coalition C, which initially contains c and c', as follows: as long as $|C| < s_j$, we try to add an agent to C that either does not consider c an enemy or considers some agent in C a friend. If this procedure stops with $|C| < s_j$, then consider the set X of those remaining agents—all of them enemies of c—that have at least one friend. We aim to find a subset $X' \subseteq X$ of s'_j agents that is *valid*, that is, each of them has at least one friend in X'. Adding a valid set of size s'_j to C yields an IR coalition of size s_j.

Observe that if $s'_j = s_j - |C| > |X|$, then no IR coalition exists of size s_j containing c, since every enemy of c in π_j must have a friend in π_j too; thus, we can reject. Otherwise, it is possible to find a valid set X' unless (i) $s'_j = 1$, or (ii) s'_j is odd, but the subgraph of $G^{\mathcal{F}}$ induced by X is the union of independent edges (i.e., a matching). In these cases, if $C = \{c, c'\}$, then no IR coalition of size s_j contains c, because each edge in $G^{\mathcal{F}}[X]$ must have either 0 or 2 of its endpoints in such a coalition. If $|C| > 2$, then it suffices to remove the agent added at the last step to C, and add to C a valid set of $s'_j + 1$ agents from X, obtaining an IR coalition of size s_j containing c. Since there are no enmities between the remaining agents, we can distribute them respecting the upper and lower bounds of the coalitions (or reject if this is not possible). There are at most n^3 guesses, and each of them can be verified in linear time. □

Unfortunately, the algorithm in Theorem 11 cannot be generalized to the asymmetric case, as shown by the following strong intractability result.

Theorem 12. *IR-FO-BSC is* NP*-complete even if* $\mathrm{vc}(G^{\mathcal{E}}) = 1$*,* $k = 2$*, and* $|\mathcal{E}_a| \le 1$ *and* $|\mathcal{F}_a| \le 2$ *for each* $a \in A$*.*

If preferences are symmetric, we can generalize the algorithm of Theorem 10 for a vertex cover of size 3, using an extensive case analysis.

Theorem 13. *If* $\mathrm{vc}(G^{\mathcal{E}}) \le 3$ *and preferences are symmetric, IR-FO-*k *can be solved in polynomial time.*

The next three intractability results show that the algorithm for IR-FO-k in Theorem 13 cannot be generalized to (a) deal with asymmetric preferences, even if $\mathrm{vc}(G^{\mathcal{E}}) = 2$ and we ask for $k = 2$ coalitions, (b) solve IR-FO-k for symmetric instances with $\mathrm{vc}(G^{\mathcal{E}}) = 4$, or (c) solve IR-FO-BSC under the same restrictions. These results are obtained by reductions similar to the one of Theorem 2.

Theorem 14. *The followings hold:*
(a) IR-FO-k is NP*-complete even if* $\mathrm{vc}(G^{\mathcal{E}}) = 2$*,* $k = 2$ *and* $\max_{a \in A} |\mathcal{E}_a| \leq 1$*.*
(b) IR-FO-k is NP*-complete even if* $\mathrm{vc}(G^{\mathcal{E}}) = 4$ *and the instance is symmetric.*
(c) IR-FO-BSC is NP*-complete even on symmetric instances with* $\mathrm{vc}(G^{\mathcal{E}}) = 3$*.*

We close this section with a result that shows that taking the number ξ of agents that have enemies as a parameter, IR-FO-FSC is fixed-parameter tractable. The result is based on the observation that friendship relations between agents who have no enemies can be removed; hence, the underlying relationship graph is guaranteed to have a vertex cover of size ξ. As we will see in Theorem 19, such instances can be solved in FPT time with respect to ξ.

Theorem 15. *IR-FO-FSC is in* FPT *when parameterized by the number of conflicting agents* $\xi = |\{a \in A : \mathcal{E}_a \neq \emptyset\}|$*.*

4.5 Restricting the Relationship Graph

We now turn our attention to the relationship graph G—including both friendships and enmities—to study how classic structural graph parameters such as treewidth and clique-width influence the computational complexity of our problems. To start with, we develop a fixed-parameter tractable algorithm for IR-FO-k with parameter $\mathrm{tw}(G)$, the treewidth of G. Although our algorithm uses the standard method of dynamic programming along a tree-decomposition of G, the details of the computation and the proof of its correctness are far from trivial.

Theorem 16. *IR-FO-k admits an algorithm with* $2^{\mathcal{O}(\mathrm{tw}(G) \cdot \log \mathrm{tw}(G))} \cdot n$ *running time. Thus, IR-FO-k is in* FPT *parameterized by the treewidth of* G*.*

Proof sketch. For the sake of exposition, we assume that $k \leq \mathrm{tw}(G)$ in this proof sketch. We formulate an explicit bottom-up dynamic programming algorithm over a *nice* tree decomposition $\mathcal{T}$ of G—a rooted tree decomposition in a standard normal form whose nodes are of four types (leaf, introduce, forget, and join nodes), simplifying the description of the computation; see the textbook of Cygan *et al.* [17] for a formal definition.

For every node x of $\mathcal{T}$, let $\beta(x)$ denote the agents in the bag of x, and let A^x denote the set of agents in the union of all bags in the subtree of $\mathcal{T}$ rooted at x. We compute a dynamic programming table $\mathrm{DP}_x[p, f, s]$, where $p\colon \beta(x) \to [k]$ is a partition of bag agents between coalitions, $f\colon \beta(x) \to \{-1, 0, 1\}$ assigns to each bag agent a *flag* based on their neighborhood, and $s\colon [k] \to \{0, 1\}$ represents whether the coalition $j \in [k]$ is non-empty. We call the triple (p, f, s) a *signature*. Observe that for each node x, there are at most $k^{\mathcal{O}(\mathrm{tw})} \cdot 3^{\mathcal{O}(\mathrm{tw})} \cdot 2^{\mathcal{O}(k)} = 2^{\mathcal{O}(\mathrm{tw} \cdot \log \mathrm{tw})}$ different signatures. For a signature (p, f, s), the table DP_x stores `true` if there exists a corresponding *partial partition* π' of A^x such that

- (i) for every agent $a \in \beta(x)$ we have $\pi'(a) = \pi'_{p(a)}$, that is, each bag agent is allocated to the correct coalition,
- (ii) for every agent $a \in \beta(x)$ such that $f(a) = -1$ holds that $|\pi'(a) \cap \mathcal{E}_a| \geq 1$ and $|\pi'(a) \cap \mathcal{F}_a| = 0$,

(iii) for every agent $a \in \beta(x)$ such that $f(a) = 0$ holds that $|\pi'(a) \cap \mathcal{E}_a| = |\pi'(a) \cap \mathcal{F}_a| = 0$,
(iv) for every agent $a \in \beta(x)$ such that $f(a) = 1$ holds that $|\pi'(a) \cap \mathcal{F}_a| \geq 1$,
(v) for every agent $a \in A^x \setminus \beta(x)$ the partition π' is IR, and
(vi) for every $j \in [k]$ it holds that $|\pi_j| \geq 1$ if $s(j) = 1$ and $|\pi_j| = 0$ otherwise.

If no such partial partition exists for (p, f, s), the dynamic programming table stores `false` for this signature.

Assume now that the dynamic programming table is computed correctly for the root node r. By definition, the bag of the root node is empty. Therefore, we ask if $\mathrm{DP}_r[\emptyset, \emptyset, \mathbf{1}]$, where $\emptyset$ represents the function with an empty domain and $\mathbf{1}$ is a constant-1 function, is set to `true`. If it is the case, then there exists a partition of $A^r = A$ which is individually rational for every agent $a \in A$ by property (v) and all coalitions are non-empty by property (vi). So, the algorithm returns *yes*. Otherwise, no such partition exists, and, therefore, the algorithm correctly returns *no*. $\lhd$

We can extend our dynamic programming algorithm on the tree-decomposition of the relationship graph G so that for each coalition that has an agent in the current bag, we store the number of agents contained in a partial solution; this approach yields an XP algorithm for our size-constrained problem variants with parameter $k + \mathrm{tw}(G)$.

Theorem 17. *IR-FO-SCC is in XP parameterized by $k + \mathrm{tw}(G)$.*

Based on our reductions from the q-Coloring problem—the generalization of 3-Coloring in which the vertices must be properly colored using q colors—we can show that our algorithm presented in Theorem 16 has optimal running time, assuming the Exponential Time Hypothesis (ETH) [32]. In addition, the same reduction allows us to observe that IR-FO-k is W[1]-hard when parameterized by the clique-width of G. These intractability results also extend to the IR-FO-BSC problem, due to known results on the Equitable q-Coloring problem.

Theorem 18. *The followings hold:*
(a) Unless ETH fails, no algorithm solves FO-k-HG in $2^{o(\mathrm{tw}(G)\cdot\log\mathrm{tw}(G))}\cdot n$ time.
(b) IR-FO-k is W[1]-hard when parameterized by the clique-width of G.
(c) IR-FO-BSC is W[1]-hard when parameterized by the treewidth of G and the number k of coalitions, combined, and para-NP-hard when parameterized by the clique-width of G.

Motivated by the intractability of IR-FO-BSC with respect to treewidth, as shown in Theorem 18, we consider the vertex cover number of the relationship graph as a stronger parameter, leading to an FPT algorithm for IR-FO-FSC.

Theorem 19. *When parameterized by $\mathrm{vc}(G)$, IR-FO-FSC is in FPT.*

Proof. Let M denote a vertex cover of G of minimum size; we can compute M in $1.25284^{|M|} \cdot (n + |E|)$ time [31]. We will assume that there is an IR coalition structure π for our instance. We call an agent $a \in M$ *lonely* in π if $\mathcal{F}_a \cap \pi(a) = \emptyset$.

For each agent $a \in M$ that is not lonely, we fix some agent $a_c \in \mathcal{F}_a \cap \pi(a)$. Let $A' = \{a_c : c \in M, c \text{ is not lonely in } \pi\}$.

Our algorithm makes four guesses in total: the partitioning of M induced by π, the set of lonely agents, an auxiliary graph H, and, later on, an ordering of the parts of the guessed partitioning; the number of possible combinations is bounded in the running-time analysis at the end of the proof. We start by guessing the partitioning $\widetilde{\pi}$ of M induced by π. Next, for each agent a in M, we guess whether a is lonely in π; let M' denote the agents in M that are *not* lonely. Then, we guess the following auxiliary graph H whose vertex set is $M' \cup A'$ and whose edge set is $\{(c, a_c) : c \in M'\}$. For each coalition $\widetilde{C} \in \widetilde{\pi}$, we let $S(\widetilde{C}) = \widetilde{C} \cup \{a_c : c \in M' \cap \widetilde{C}\}$. Assuming that our guesses are correct, we have $S(\widetilde{C}) \subseteq \pi(\widetilde{C})$ where $\pi(\widetilde{C})$ denotes the unique coalition in π that contains all agents in $\widetilde{C}$. Although we do not know the identity of the agents in $S(\widetilde{C})$, we know the cardinality $|S(\widetilde{C})|$.

Next, we decide which coalitions we should allocate the agents of $S(\widetilde{C})$ to. We may assume that $\mathrm{ub}_i \leq \mathrm{ub}_j$ whenever $1 \leq i \leq j \leq k$. Let us now guess the ordering $\widetilde{C}_1, \ldots, \widetilde{C}_\ell$ of the sets in $\widetilde{\pi}$ according to the size of the coalitions they are contained in under π, so that $|\pi(\widetilde{C}_i)| \leq |\pi(\widetilde{C}_j)|$ whenever $1 \leq i < j \leq \ell$. Now, we compute a mapping $\mathsf{ind} : [\ell] \to [k]$ as follows. For $i = 1, \ldots, \ell$, we find a coalition index $\mathsf{ind}(i) \in [k]$ such that (i) $\mathsf{ind}(i) > \mathsf{ind}(i-1)$ for every $i > 1$, (ii) $\mathrm{ub}_{\mathsf{ind}(i)} \geq |S(\widetilde{C}_i)|$, and (iii) $\mathrm{ub}_{\mathsf{ind}(i)}$ is as small as possible.

Claim 1. *Suppose that an IR coalition structure π exists, and the algorithm made correct guesses with respect to π. Then there exists an IR coalition structure π' where, for each $i \in [\ell]$, the coalition $\pi(\widetilde{C}_i)$ has size $\mathrm{ub}_{\mathsf{ind}(i)}$.*

Due to Claim 1, we may restrict our search for a solution in which $\widetilde{C}_i$ is allocated to the coalition with index $\mathsf{ind}(i)$. To compute a suitable IR coalition structure, we further need to allocate the agents in $A \setminus M$ to the coalition indices so that (i) each non-lonely agent in M has a friend in their coalition (recall that we do not know the identity of agents in $A' \setminus M$), (ii) no lonely agent in M shares their coalition with an enemy, and (iii) the size constraints are respected.

To this end, we construct a flow network N as follows, with source s and target t. Let $I = \{\mathsf{ind}(i) : i \in [\ell]\}$ be the indices of coalitions that will hold conflicting agents. For all coalition index $i \in [\ell]$, we set the *demand* of $\mathsf{ind}(i) \in I$ as $\mathrm{ub}_{\mathsf{ind}(i)} - |S(\widetilde{C}_i)|$, and we set the demand of indices $j \in [k] \setminus I$ simply as ub_j.

First, we create an *agent vertex* u_a for each agent $a \in A \setminus M$. Second, for each coalition index $j \in [k]$, we create a *coalition vertex* v_j. Third, we also define a *guard vertex* $g_{a'}$ for each $a' \in A' \setminus M$.

There is an arc (s, u_a) with capacity 1 for each agent $a \in A$, an arc (v_j, t) for each coalition index j whose capacity is the demand of j. We also add an arc (g_a, t) for each agent $a \in A' \setminus M$ with capacity 1. Now, for each $a \in A \setminus M$, we add an arc of capacity 1 from agent vertex u_a to each coalition vertex v_j where either $j \notin I$ or $j = \mathsf{ind}(i)$ for some $i \in [\ell]$ such that $a \notin \mathcal{E}_c$ for any lonely agent $c \in \widetilde{C}_i$. We also add an arc of capacity 1 from u_a to $g_{a'}$ for some $a' \in A' \setminus M$

if a is *suitable for* a', meaning that a is considered a friend by all agents c for which $a' = a_c$. Recall that even though we do not know the identity of a', we can decide for each $a \in A \setminus M$ whether a is suitable for a' using the guessed auxiliary graph H—hence, our network is well defined. We finish our algorithm by finding a maximum flow in N; if its value is $|A \setminus M|$, we return 'yes'; otherwise, we proceed with our next set of guesses.

Claim 2. *Suppose there exists an IR coalition structure π that, for each $i \in [\ell]$, allocates all agents in $S(\widetilde{C}_i)$ to coalition index $\mathsf{ind}(i)$, and the algorithm made correct guesses with respect to π. Then there exists a flow of size $|A \setminus M|$ in the network N. Conversely, if N admits a flow of size $|A \setminus M|$, then there exists an IR coalition structure π.*

Claims 1 and 2 show the correctness of our algorithm.

Running time. Computing M is in FPT. Guessing the partitioning of M yields $\mathrm{vc}(G)^{\mathrm{vc}(G)}$ possibilities, while guessing the lonely agents yields $2^{\mathrm{vc}(G)}$ possibilities. Guessing the auxiliary graph H results in at most $\mathrm{vc}(G)^{\mathrm{vc}(G)}$ possibilities, since there are at most $|M|$ agents in A', and each agent in M' is connected to exactly one agent in A' in H. Finally, guessing the ordering of the sets in π' yields at most $\mathrm{vc}(G)! = \mathcal{O}\big(\mathrm{vc}(G)^{\mathrm{vc}(G)}\big)$ possibilities. Hence, the number of possible guesses is $2^{\mathrm{vc}(G)\log(\mathrm{vc}(G))}$. Since for each fixed set of guesses, the bottleneck is to compute a maximum flow with lower bounds, which can be done in polynomial time with standard techniques, the running time of the presented algorithm is indeed fixed-parameter tractable with parameter $\mathrm{vc}(G)$. □

5 Conclusion

We studied individual rationality in friend-oriented and enemy-oriented hedonic games under four variants of size constraints, charting the boundary between tractable and intractable cases across a range of natural parameters.

Two messages emerge from our results. First, under friend-oriented preferences, it is the enmity structure alone that governs the complexity: whenever enmities are scarce or well-structured, tractability can often be recovered, whereas rich friendship structures never cause hardness on their own. Second, symmetry acts as a genuine computational resource: several polynomial-time solvable cases for symmetric relations become NP-hard as soon as asymmetry is allowed.

Beyond resolving the remaining open cases in Table 1—the most intriguing one, in our view, being whether the size-constrained (SCC) variant is fixed-parameter tractable with respect to the vertex cover number of the relationship graph—we see two natural directions for future research. One is to combine the size constraints studied here with more demanding stability notions, such as Nash or core stability, for which even the existence of unconstrained stable outcomes is non-trivial in our preference domain. Another is optimization: among all feasible individually rational coalition structures, find one that maximizes social welfare.

Acknowledgments. This research was supported in part by the National Science Centre, Poland, grant number UMO-2025/58/A/ST6/00371, and co-funded by the European Union under the project Roboprox (reg. no. CZ.02.01.01/00/22_008/0004590). Ildikó Schlotter is supported by the Hungarian Academy of Sciences under its Momentum Programme (LP2021-2) and its János Bolyai Research Scholarship.

References

1. Agarwal, P., Agarwal, H., Raj, V., Nath, S.: Harmonious balanced partitioning of a network of agents. In: AAMAS '25. pp. 41–49 (2025)
2. Ashlagi, I., Roth, A.E.: Individual rationality and participation in large scale, multi-hospital kidney exchange. In: EC '11. pp. 321–322 (2011)
3. Aziz, H., Brandl, F., Brandt, F., Harrenstein, P., Olsen, M., Peters, D.: Fractional hedonic games. ACM Trans. Econ. Comput. **7**(2), 6:1–6:29 (2019)
4. Aziz, H., Savani, R.: Hedonic games. In: Handbook of Computational Social Choice, pp. 356–376. Cambridge University Press (2016)
5. Banerjee, S., Konishi, H., Sönmez, T.: Core in a simple coalition formation game. Soc. Choice Welf. **18**(1), 135–153 (2001)
6. Barr, H., Trabelsi, Y., Kraus, S., Roditty, L., Hazon, N.: The complexity of manipulation of k-coalitional games on graphs. In: ECAI '24. pp. 3388–3396 (2024)
7. Bilò, V., Monaco, G., Moscardelli, L.: Hedonic games with fixed-size coalitions. In: AAAI '22. pp. 9287–9295 (2022)
8. Bogomolnaia, A., Jackson, M.O.: The stability of hedonic coalition structures. Games Econ. Behav. **38**(2), 201–230 (2002)
9. Brandt, F., Bullinger, M., Tappe, L.: Stability based on single-agent deviations in additively separable hedonic games. Artif. Intell. **334**, 104160 (2024)
10. Bredereck, R., Elkind, E., Igarashi, A.: Hedonic diversity games. In: AAMAS '19. pp. 565–573 (2019)
11. Bullinger, M.: Pareto-optimality in cardinal hedonic games. In: AAMAS '20. pp. 213–221 (2020)
12. Bullinger, M., Chatziafratis, V., Shahkar, P.: Welfare approximation in additively separable hedonic games. In: AAMAS '25. pp. 418–426 (2025)
13. Bullinger, M., Dunajski, A., Elkind, E., Gilboa, M.: Single-deviation stability in additively separable hedonic games with constrained coalition sizes. In: AAMAS '26. pp. 3624–3626 (2026)
14. Bullinger, M., Suksompong, W.: Topological distance games. Theor. Comput. Sci. **981**, 114238 (2024)
15. Bulteau, L., Fertin, G., Labarre, A., Rizzi, R., Rusu, I.: Decomposing subcubic graphs into claws, paths or triangles. J. Graph Theory. **98**(4), 557–588 (2021)
16. Chen, J., Csáji, G., Roy, S., Simola, S.: Hedonic games with friends, enemies, and neutrals: Resolving open questions and fine-grained complexity. In: AAMAS '23. pp. 251–259 (2023)
17. Cygan, M., Fomin, F.V., Kowalik, Ł., Lokshtanov, D., Marx, D., Pilipczuk, M., Pilipczuk, M., Saurabh, S.: Parameterized Algorithms. Springer (2015)
18. Darmann, A., Döcker, J.: On a simple hard variant of Not-All-Equal 3-Sat. Theor. Comput. Sci. **815**, 147–152 (2020)
19. Deligkas, A., Eiben, E., Ioannidis, S.D., Knop, D., Schierreich, Š.: Balanced and fair partitioning of friends. In: AAAI '25. pp. 13754–13762 (2025)
20. Deligkas, A., Eiben, E., Knop, D., Schierreich, Š.: Individual rationality in topological distance games is surprisingly hard. In: IJCAI '24. pp. 2782–2790 (2024)

21. Diestel, R.: Graph Theory. Springer, Berlin, Heidelberg, 6th edn. (2025)
22. Dimitrov, D., Borm, P., Hendrickx, R., Sung, S.C.: Simple priorities and core stability in hedonic games. Soc. Choice Welf. **26**(2), 421–433 (2006)
23. Drèze, J.H., Greenberg, J.: Hedonic coalitions: Optimality and stability. Econometrica **47**(4), 987–1003 (1980)
24. Fioravantes, F., Gahlawat, H., Melissinos, N.: Exact algorithms and lower bounds for forming coalitions of constrained maximum size. In: AAAI '25. pp. 13847–13855 (2025)
25. Fioravantes, F., Gahlawat, H., Melissinos, N., Schierreich, Š.: Individual rationality in constrained hedonic games: Additively separable and fractional preferences. In: AAMAS '26. pp. 2848–2857 (2026)
26. Fioravantes, F., Gahlawat, H., Melissinos, N., Schierreich, Š.: Individual rationality in constrained hedonic games: Anonymous and diversity preferences. In: GAIW '26 (2026)
27. Furmanczyk, H., Kubale, M.: Equitable coloring of corona products of cubic graphs is harder than ordinary coloring. Ars Math. Contemp. **10**(2), 333–347 (2016)
28. Furmanczyk, H., Mkrtchyan, V.: Graph theoretic and algorithmic aspect of the equitable coloring problem in block graphs. Discret. Math. Theor. Comput. Sci. **23**(2) (2022)
29. Ganian, R., Hamm, T., Knop, D., Schierreich, Š., Suchý, O.: Hedonic diversity games: A complexity picture with more than two colors. Artif. Intell. **325**, 104017 (2023)
30. Ganian, R., Hamm, T., Knop, D., Roy, S., Šimon Schierreich, Suchý, O.: Maximizing social welfare in score-based social distance games. In: TARK '23. pp. 272–286 (2023)
31. Harris, D.G., Narayanaswamy, N.S.: A faster algorithm for vertex cover parameterized by solution size. In: STACS '24. pp. 40:1–40:18 (2024)
32. Impagliazzo, R., Paturi, R., Zane, F.: Which problems have strongly exponential complexity? J. Comput. Syst. Sci. **63**(4), 512–530 (2001)
33. Lang, J., Rey, A., Rothe, J., Schadrack, H., Schend, L.: Representing and solving hedonic games with ordinal preferences and thresholds. In: AAMAS '15. pp. 1229–1237 (2015)
34. Levinger, C., Azaria, A., Hazon, N.: Social aware coalition formation with bounded coalition size. In: AAMAS '23. pp. 2667–2669 (2023)
35. Li, L., Micha, E., Nikolov, A., Shah, N.: Partitioning friends fairly. In: AAAI '23. pp. 5747–5754 (2023)
36. Ohta, K., Barrot, N., Ismaili, A., Sakurai, Y., Yokoo, M.: Core stability in hedonic games among friends and enemies: Impact of neutrals. In: IJCAI '17. pp. 359–365 (2017)
37. Peters, D., Elkind, E.: Simple causes of complexity in hedonic games. In: IJCAI '15. pp. 617–623 (2015)
38. Roth, A.E.: Individual rationality and Nash's solution to the bargaining problem. Math. Oper. Res. **2**(1), 64–65 (1977)
39. Rothe, J., Schadrack, H., Schend, L.: Borda-induced hedonic games with friends, enemies, and neutral players. Math. Soc. Sci. **96**, 21–36 (2018)
40. Sless, L., Hazon, N., Kraus, S., Wooldridge, M.J.: Forming k coalitions and facilitating relationships in social networks. Artif. Intell. **259**, 217–245 (2018)
41. Waxman, N., Kraus, S., Hazon, N.: Manipulation of k-coalitional games on social networks. In: IJCAI '21. pp. 450–457 (2021)
42. Wright, M., Vorobeychik, Y.: Mechanism design for team formation. In: AAAI '15. pp. 1050–1056 (2015)