Individual Rationality in Constrained Hedonic Games: Additively Separable and Fractional Preferences

Foivos Fioravantes
Czech Technical University in Prague
Prague, Czech Republic
foivos.fioravantes@fit.cvut.cz

Harmender Gahlawat
Ben-Gurion University of the Negev
Beersheba, Israel
harmendergahlawat@gmail.com

Nikolaos Melissinos
Computer Science Institute,
Faculty of Mathematics and Physics
Charles University
Prague, Czech Republic
melissinos@iuuk.mff.cuni.cz

Šimon Schierreich
AGH University of Krakow
Krakow, Poland
Czech Technical University in Prague
Prague, Czech Republic
schiesim@fit.cvut.cz

ABSTRACT

Hedonic games are an archetypal problem in coalition formation, where a set of selfish agents want to partition themselves into *stable* coalitions. In this work, we focus on two natural constraints on the possible outcomes. First, we require that exactly k coalitions are created. Then, loosely following the model of Bilò et al. (AAAI 2022), we assume that each of the k coalitions is additionally associated with a lower and upper bound on its size. The notion of stability that we study is that of *individual rationality* (IR), which requires that no agent strictly prefers to be alone compared to being in his or her coalition.

Although IR is trivially satisfiable even in the most general models of hedonic games, the complexity picture of deciding whether an IR allocation exists, considering the above constraints, is unexpectedly rich. We reveal that tractable fragments of this computational problem require surprisingly nontrivial arguments, even if we restrict ourselves to *additively separable* and *fractional hedonic games*. Our tractability results, achieved by exploiting the structure of the underlying *preference graph*, are also complemented by their intractability counterparts, painting a fairly complete picture of the tractability landscape of this problem.

KEYWORDS

Cooperative Game Theory, Coalition Formation, Individual Rationality, Hedonic Games, Additively Separable Preferences, Fractional Preferences, Fixed-Parameter Tractability, N-fold ILP.

ACM Reference Format:
Foivos Fioravantes, Harmender Gahlawat, Nikolaos Melissinos, and Šimon Schierreich. 2026. Individual Rationality in Constrained Hedonic Games: Additively Separable and Fractional Preferences. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 10 pages. https://doi.org/10.65109/ICRJ7770

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/ICRJ7770

1 INTRODUCTION

Organizing a seating arrangement for a banquet may seem straightforward, but it quickly becomes complex when interpersonal relationships are considered. For example, placing Janine from HR at the same table as Sheldon from the Physics Department could result in a very unpleasant evening—not just for them, but likely for all the attendees. And it is not only about animosities; seating Sheldon at the table only with strangers can be a similarly explosive choice. In addition, there may be certain restrictions on seating, including the number of available tables and their capacity.

The situation introduced above appears naturally in various domains and falls in the area of *coalition formation*. More formally, the goal here is to partition a set of agents into several groups, called *coalitions*, so that no agent is motivated to deviate from its coalition to improve their utility. In more game-theoretical terms, we are interested in *stable* partitions. The prevailing model in this area is that of *hedonic games* [6, 33], where agents care only about other agents in their own coalition and not about how the additional coalitions are formed.

Traditionally, arbitrarily constructed coalition structures are acceptable as long as they meet certain stability criteria. However, in many real-life scenarios, such as the one at the beginning of this section, we have an additional restriction on the outcome. First, as argued by Sless et al. [62], in certain situations, it may be necessary to form an exact number of coalitions; see also [8, 64]. We call this model *k-hedonic games* (k-HG). A different restriction, motivated by the recently proposed model of Bilò et al. [9] (see also [1, 29, 52]), asks for exactly k coalitions, each of size between a given lower and upper bound. We call this model *hedonic games with size-constrained coalitions* (HG-SCC). Several of these works study the computational aspects of finding stable outcomes under these two restrictions. However, they are mostly interested in Nash, individual, or core stability, which are all very demanding, rarely satisfiable, and the associated computational problems are often intractable. In our work, we return to the foundational stability notion of *individual rationality* (IR) [3, 30, 58, 60], which is strictly weaker than the notions mentioned above (see, e.g., [6, Figure 15.1]). Intuitively, a coalition structure is individually rational if each agent weakly prefers to be in its current coalition than to deviate to a

Table 1: Summary of our contributions. We use $\mathrm{vc}(G)$ to denote the vertex cover number of G, and $\mathrm{tw}(G)$ for its treewidth.

Problem	Restrictions	Complexity	Reference
Both k-ASHGs and ASHG-SCC	for every $k \geq 2$, symmetric valuations, G is a split graph with $\mathcal{O}(k^2)$ negative edges and $\mathrm{vc}(G) = 2k$	NP-complete	Theorem 3.2
	for every $k \geq 2$, symmetric valuations, G is a clique, and the valuations use six different unary encoded values	NP-complete	Theorem 3.3
	symmetric, unary encoded valuations, parameterized by the number of coalitions k, the number of negative edges, and $\mathrm{vc}(G)$	W[1]-hard	Theorem 4.2
	unary encoded valuations, parameterized by $\mathrm{vc}(G)$	XP	Theorem 4.1
	for $k = 2$, symmetric and binary valuations, parameterized by the treedepth of G	W[1]-hard	Theorem 5.1
Only k-ASHGs	parameterized by $\mathrm{vc}(G)$ and the maximum weight $\omega_{\max}$	FPT	Theorem 4.3
	binary valuations, parameterized by $\mathrm{tw}(G)$	XP	Theorem 5.2

singleton coalition. It is a well-known fact that every hedonic game admits a trivial IR solution; just place each agent in a coalition on its own. Due to this, individual rationality is completely overlooked in the relevant literature. However, as we show, our restrictions turn IR into an unexpectedly interesting solution concept with a very colorful complexity picture.

1.1 Our Contribution

The main contribution of our work is a detailed algorithmic landscape of the two restrictions of coalition formation introduced above assuming two prominent subclasses of hedonic games. As classical complexity renders the problem very soon intractable, we turn to the finer-grained framework of parameterized complexity.[1]

Intuitively, under this paradigm, we study the complexity of a computational problem not only with respect to the input size, but also assuming some *parameter* p that further restricts the input instance. The ultimate goal of parameterized algorithms is to find (or prove that it is not possible by showing that the problem is W[1]-hard) an algorithm whose running time is exponential only in the parameter, its time dependence on the input size is only polynomial, and the degree of the polynomial is independent of the parameter. Such algorithms are called FPT. Slightly worse, but still positive, are XP algorithms whose running time is also polynomial in input size, but the degree of this polynomial depends on the parameter. One can conditionally rule out the existence of an XP algorithm by showing that a problem is NP-hard already for a constant value of p.

The two classes of hedonic games we study are *additively separable hedonic games* (ASHGs) [14] and *fractional hedonic games* (FHGs) [4]. In these games, each agent i has a specific value $v_i(j)$ for every other agent j. The utility of an agent i in a coalition C, $i \in C$, is then simply the sum of the individual values that agent i has for all other agents in C (divided by the size of C in case of FHGs). We can succinctly encode such preferences using a weighted digraph G over the set of agents containing an arc (i, j) of weight $v_i(j)$ for every pair of distinct agents i, j (and we can remove all zero-weight arcs to reduce the size of the encoding).

First, we observe that if the desired stability notion is individual rationality, these two classes coincide; therefore, it suffices to study only ASHGs. Then, we investigate the complexity of ASHGs with respect to two natural restrictions: the structure of the preference graph G and the maximum value (arc weight) we have in the instance. As we reveal, these two dimensions of the problem interplay in a non-trivial way. Our results are summarized in Table 1.

Specifically, if the weights can be general, both variants of ASHGs are NP-complete, even if the graph G is a split graph of constant vertex cover number. If the weights are polynomially bounded by the number of agents, then there is an XP algorithm with respect to the vertex cover number of G. For FPT tractability, we need to additionally parameterize by the maximum weight, as without this both problems are W[1]-hard already for the combined parameter of the number of coalitions and the vertex cover of G. Next, we show that the XP (or FPT) algorithm for the vertex cover number parameterization cannot be extended to more general graph classes, as we show that both problems are W[1]-hard when parameterized by treedepth, even if $k = 2$ and the valuations are binary. Treedepth is a structural parameter that lies between the vertex cover number and the celebrated treewidth. The complexity picture is completed by an XP algorithm for parameterization by treewidth and binary valuations, and we also show that the problem is highly intractable in *dense* graphs by showing NP-hardness of both variants even if the underlying graph is a clique. Our algorithms use various techniques, notably including the N-fold ILP formulation—a technique which is completely novel in the area of coalition formation. In addition, all our hardness results hold for symmetric preferences, while our algorithms work also for non-symmetric ones. This draws our results even stronger.

[1] A formal introduction to parameterized complexity is provided in Section 2.

1.2 Related Work

Additively separable hedonic games (ASHGs) [14] and *fractional hedonic games* (FHGs) [4] are probably the most studied subclass of hedonic games, as is clear from the number of papers investigating them [5, 15, 22, 35–37, 42, 43, 55–57, 63]. However, all of these works focus on more general notions of stability than IR, as IR is trivially satisfiable without any further restriction of the game, and usually do not consider any restriction on the sizes of the coalitions. An exception is the recent independent paper [21], in which *global* lower bounds are considered; however, it also does not study IR.

There are also other subclasses of hedonic games, such as anonymous [7, 14], diversity [13, 17, 38], social distance [16, 38], friends and enemies games [27, 50, 54, 61], $\mathcal{B}$ games [24, 26], or $\mathcal{W}$ games [25, 26]. However, in these classes of games, either no coalition is worse than the singleton, and therefore, IR is always achievable even under our constraints, or individual rationality was not explored. The only exception is the recent paper of Deligkas et al. [30], who studied individual rationality in *topological distance games* (TDGs) [23]. However, the model of TDGs is much more general compared to ours, and therefore none of their (mostly negative) results carry over to our setting.

Next, k-hedonic games and hedonic games with fixed size coalitions are not the only restrictions studied in the literature. For example, Wright and Vorobeychik [65] introduced hedonic games with restricted maximum coalition size, a variant studied also in several subsequent works [36, 38, 39, 51]. However, even under this restriction, individual rationality is easy as we can form an arbitrary number of coalitions of size one. It is also worth mentioning the model of *constrained coalition formation* [59], where the coalition game comes with certain constraints on the outcome encoded using propositional logic formulas.

Finally, N-fold integer programming [44, 53] is an important technique in the design of parameterized algorithms. It has been used successfully in multiple relevant areas such as scheduling [46, 49], voting theory [10, 47], and fair division [18–20], to name at least a few related to computational social choice, algorithmic game theory, and collective decision making. Nevertheless, we are not aware of any work using this technique in hedonic games.

2 PRELIMINARIES

Let $i \le j \in \mathbb{N}$. We use $[i, j]$ to denote the closed interval $\{i, \ldots, j\}$, and we set $[i] = \{1, \ldots, i\}$ and $[i]_0 = [i] \cup \{0\}$. Additionally, for a set S, we use 2^S to denote the power set of S.

Let N be a finite set of n agents. A nonempty subset $C \subseteq N$ is called a *coalition*. When $|C| = n$, we say that C is the *grand coalition*, while if $|C| = 1$, we call C a *singleton*. A *coalition structure* π is a partition of agents into coalitions, and by $\pi(i)$ we denote the coalition agent $i \in N$ is assigned.

Each agent $i \in N$ provides its subjective *valuation function* $v_i \colon 2^N \to \mathbb{Z}$, which assigns to each coalition C a numerical value expressing i's satisfaction when being part of C. As is usual in hedonic games, we assume that v_i is defined only for coalitions that contain the agent i. Slightly abusing the notation, we extend the valuations from coalitions to coalition structures by setting $v_i(\pi) = v_i(\pi(i))$ and we say that $v_i(\pi)$ is the *utility* the agent i gets in the coalition structure π.

In this work, we investigate two different restrictions of hedonic games. In the first variant, we are interested only in coalition structures consisting of exactly $k \in \mathbb{N}$ non-empty coalitions. Formally, a *k-hedonic game* Γ is a triple $(N, (v_i)_{i \in N}, k)$, where N is a set of agents, v_i is a valuation function for every agent $i \in N$, and k is the desired number of coalitions. In the second variant, not only the number of coalitions, but also their sizes are prescribed. A *hedonic game with size-constrained coalitions* (HG-SCC) Γ is a 5-tuple $(N, (v_i)_{i \in N}, k, \text{lb}, \text{ub})$, where $\text{lb} \colon [k] \to \mathbb{N}$, $\text{ub} \colon [k] \to \mathbb{N}$, and $\text{lb}(j) \le \text{ub}(j)$ for every $j \in [k]$. The goal in HG-SCC is to decide whether a coalition structure $\pi = (\pi_1, \ldots, \pi_k)$ exists such that $\text{lb}(j) \le |\pi_j| \le \text{ub}(j)$ for every $j \in [k]$. Without loss of generality, we assume that function ub satisfies $\text{ub}(1) \ge \text{ub}(2) \ge \cdots \ge \text{ub}(k)$.

Throughout the paper, we are interested in finding coalition structures where no agent can improve its utility by *deviating* from the given coalition structure π. If no beneficial deviation is possible, we say that a coalition structure π is *stable*. Specifically, we adopt the notion of *individual rationality*, which, as argued by Aziz and Savani [6], is a minimal requirement for a solution to be considered stable and is defined as follows.

Definition 2.1. A coalition structure π is said to be *individually rational (IR)* if $v_i(\pi) \ge v_i(\{i\})$ for every agent $i \in N$.

Intuitively, individual rationality requires that no agent prefers to deviate from the current partition and form a coalition alone. Note that the deviation of agent i to a singleton coalition can formally violate our condition on having exactly k coalitions. Therefore, we interpret such deviations so that the deviating agent completely leaves the game instead of actually forming a singleton coalition.

In general, the encoding of the valuation functions may require space exponential in the number of agents—we need to store a value for every agent and every possible coalition. Therefore, we study various restrictions of hedonic games that come with a succinct representation of valuations.

Additively Separable Hedonic Games. In *additively separable hedonic games* (ASHGs), an agent $i \in N$ has a value $v_i(j)$ for each agent $j \in N$. We require $v_i(i) = 0$. Then, the utility of the agent i in a coalition structure π is simply $\sum_{j \in \pi(i)} v_i(j)$. If $v_i(j) = v_j(i)$ for all $i, j \in N$, we say that the valuations are *symmetric*, and if $v_i(j) \in \{-1, 0, 1\}$ for every $i, j \in N$, we call the valuations *binary*. The valuations in ASHGs can be represented using a weighted digraph $G = (N, E, \omega)$ that contains an edge (i, j) of weight $\omega((i, j)) = v_i(j)$ if and only if $v_i(j) \ne 0$. When the valuations are symmetric, the graph G is undirected.

Fractional Hedonic Games. In *fractional hedonic games* (FHGs), similarly to ASHGs, every agent $i \in N$ has a valuation function v_i assigning some value to every agent $j \in N$. These valuations are then extended to coalitions as follows. Let $C \subseteq N$ be a coalition such that $i \in C$. Then, we have $v_i(C) = (\sum_{j \in C} v_i(j))/|C|$. Again, we can capture the valuations using a similar weighted digraph as in the case of ASHGs.

Parameterized Complexity. This is a domain of algorithm design in which a finer analysis of computational complexity is proposed, compared to the classic approach. This is done by considering additional measures of complexity, referred to as *parameters*. In short, this field proposes a multidimensional time-complexity analysis, where each parameter defines its own dimension. Formally, a parameterized problem is a set of instances $(x, k) \in \Sigma^* \times \mathbb{N}$; k is the *parameter*. In this paradigm, an algorithm is considered efficient if it runs in time $f(k)|x|^{O(1)}$ time for any arbitrary computable function $f : \mathbb{N} \to \mathbb{N}$. Such algorithms are known as *fixed-parameter tractable* (FPT). If such an algorithm exists for a problem, we say that this problem *is in FPT*. On the opposite side of the FPT class, we have the class of W[1]-hard problems, and it is widely accepted that if a problem is W[1]-hard then there exists no FPT algorithm to solve it. Finally, the middle ground between these notions is captured by the class XP, which contains all parameterized problems that can be solved in time $|x|^{f(k)}$ for some computable function f. We refer the interested reader to classical monographs in the field [28, 32].

Structural Parameters. Let $G = (V, E)$ be a graph. A set $U \subseteq V$ is a *vertex cover* of G if for every edge $e \in E$ it holds that $U \cap e \neq \emptyset$. The *vertex cover number* of G, denoted $\mathrm{vc}(G)$, is the minimum size of a vertex cover of G.

Definition 2.2. A *tree-decomposition* of G is a pair $(T, \mathcal{B})$, where T is a tree, $\mathcal{B}$ is a family of sets assigning to each node t of T its *bag* $B_t \subseteq V$, and the following conditions hold:

- for every edge $uv \in E(G)$, there is a node $t \in V(T)$ such that $u, v \in B_t$ and
- for every vertex $v \in V$, the set of nodes t with $v \in B_t$ induces a connected subtree of T.

The *width* of a tree-decomposition $(T, \mathcal{B})$ is defined as $\max_{t \in V(T)} |B_t| - 1$, and the treewidth $\mathrm{tw}(G)$ of a graph G is the minimum width of a tree-decomposition of G.

It is well known that computing a tree-decomposition of minimum width is in FPT w.r.t. the treewidth [11, 45]. Recently there are even more efficient algorithms that have been proposed for obtaining near-optimal tree-decompositions [48].

For algorithmic purposes, we use a slight variation of the above definition of the tree-decomposition which is more suitable for a dynamic programming approach.

Definition 2.3. A tree-decomposition $(T, \mathcal{B})$ is *nice* if every node $t \in V(T)$ is exactly of one of the following four types:

Leaf Node: t is a leaf of T and $|B_t| = 0$.

Introduce Node: t has a unique child c and there exists $v \in V$ such that $B_t = B_c \cup \{v\}$.

Forget Node: t has a unique child c and there exists $v \in V$ such that $B_c = B_t \cup \{v\}$.

Join Node: t has exactly two children c_1, c_2 and $B_t = B_{c_1} = B_{c_2}$.

It is known that every graph $G = (V, E)$ admits a nice tree-decomposition that has width equal to $\mathrm{tw}(G)$ and such a decomposition can be found efficiently [12].

The *tree-depth* of G can be defined recursively: if $|V| = 1$ then G has tree-depth 1. Then, G has tree-depth k if there exists a vertex $v \in V$ such that every connected component of $G[V \setminus \{v\}]$ has tree-depth at most $k - 1$.

N-fold ILP. The goal here is to minimize a linear objective f over a set of structured constraints. Formally, let $r, D \in \mathbb{N}$ and $s_i, t_i \in \mathbb{N}$ for every $i \in [D]$. An N-fold ILP contains $d = \sum_{i \in [D]} t_i$ variables partitioned into D *bricks*. Let $x^{(i)}$ denote the i-th brick. Then, the constraints of an N-fold ILP have the following form

$$E_1 x^{(1)} + \cdots + E_D x^{(D)} = \vec{b}_0 \quad (1)$$

$$\forall i \in [D] \qquad A_i x^{(i)} = \vec{b}_i \quad (2)$$

$$\forall i \in [D] \qquad \vec{\ell}_i \leq x^{(i)} \leq \vec{u}_i \quad (3)$$

where $E_i \in \mathbb{Z}^{r \times t_i}$, $A_i \in \mathbb{Z}^{s_i \times t_i}$, $\vec{b}_0 \in \mathbb{Z}^r$, $\vec{b}_i \in \mathbb{Z}^{s_i}$, and $\vec{\ell}_i, \vec{u}_i \in \mathbb{Z}^{t_i}$ for every $i \in [D]$. We call constraints of type (1) *global*, constraints of type (2) *local*, and constraints of type (3) *box*. Observe that r is the total number of global constraints and we use $s = \max_{i \in [D]} s_i$ to denote the maximum number of occurrences of a single variable in local constraints. In our algorithms, we use the following result of Eisenbrand et al. [34].

Theorem 2.4 (Eisenbrand et al. [34]). *An instance of N-fold ILP can be solved in* $a^{O(r^2 s + r s^2)} \cdot d \cdot \log(d) \cdot L$, *where* L *is the maximum feasible value of the objective and* $a = \max_{i \in [D]} \{2, ||E_i||_\infty, ||A_i||_\infty\}$.

3 INDIVIDUAL RATIONALITY CAN BE HARD

Before we start our algorithmic journey, we observe that both classes with graph-restricted preference we assume in this work coincide, if we are interested in IR solutions. Indeed, the only difference between ASHGs and FHGs is that in the latter class, the additive utility of an agent is normalized with respect to the coalition size. This, however, has no effect on the non-negativity of the utility, and therefore, we obtain the following.

Lemma 3.1. *A fractional hedonic game* $\Gamma = (N, (v_i)_{i \in N}, k)$ *admits an IR coalition structure if and only if an additively separable hedonic game* $\Gamma' = (N, (v_i)_{i \in N}, k)$ *admits an IR coalition structure.*

Proof. Let π be an individually rational coalition structure for Γ. Then, since no coalition is of negative size, for every $i \in N$, we have $\sum_{j \in \pi_i} v_i(j) \geq 0$. Therefore, π is also individually rational for the game Γ'. By analogous argumentation, any solution for Γ' is also a solution for Γ. □

Therefore, the complexity picture for FHGs is the same as that for ASHGs. For the rest of this paper, we assume only the case of additively separable hedonic games, but the results directly carry over to the setting of FHGs.

First, we show that if we do not have any restriction on valuation functions, deciding whether an individually rational coalition structure exists is computationally highly intractable.

Theorem 3.2. *For every* $k \geq 2$, *it is NP-complete to decide whether a given* k*-ASHG or ASHG-SCC* Γ *admits an individually rational coalition structure, even if the valuations are symmetric and* G *is a split graph with* $O(k^2)$ *negative edges and of vertex cover number* $2k$.

Proof. We show NP-hardness by a reduction from the Equitable Partition problem. In this problem, we are given a multiset $A = \{a_1, \ldots, a_{2\ell}\}$ of integers such that $\sum_{a \in A} a = 2t$, and the goal is to decide whether $S \subseteq [2\ell]$, $|S| = \ell$, exists so that $\sum_{i \in S} a_i = \sum_{i \in [2\ell] \setminus S} a_i = t$. The problem is known to be (weakly) NP-hard even if each subset $X \subseteq A$ of size at most $\ell - 1$ sums up to at most $t - 1$ [31].

Let $\mathcal{I}$ be an instance of the Equitable Partition problem. We construct an equivalent instance $\mathcal{J} = (N, (v_i)_{i \in N}, k)$ as follows. We describe the construction using the underlying graph. First, we create a clique with four *set agents* x_S, y_S, $x_{\neg S}$, and $y_{\neg S}$. The edges $x_S y_S$ and $x_{\neg S} y_{\neg S}$ are of weight $-t$, and all other edges are of weight $-3t$. Next, for every $a_i \in A$, we introduce an *item agent* v_i and connect it with an edge of weight a_i to *all* set agents. Finally, we set $k = 2$. It is easy to see that the graph G is a split graph with 4 negative edges and the vertex cover number 4.

For correctness, let $\mathcal{I}$ be a Yes-instance and S be a sought solution. We set $\pi_1 = \{x_S, y_S\} \cup \{v_i \mid i \in S\}$ and $\pi_2 = N \setminus \pi_1$. Since the item agents have only positive edges, any coalition structure is individually rational for them. For x_S, we have $v_{x_S}(\pi) = \omega(x_S y_S) + \sum_{i \in S} \omega(x_S, v_i) = -t + \sum_{i \in S} a_i = -t + t = 0$. The same argument holds for all the remaining set agents. That is, $\pi = (\pi_1, \pi_2)$ is an individually rational coalition structure.

In the opposite direction, let $\mathcal{J}$ be a Yes-instance and π be an individually rational coalition structure. First, we observe that x_S is not in the same coalition as $x_{\neg S}$ and $y_{\neg S}$ and, symmetrically, $x_{\neg S}$ is not in the same coalition as x_S and y_S. If this were the case, then the utility of any set agent in its coalition would be at most $-3t + 2t = -t$, which contradicts that π is individually rational. Hence, without loss of generality, $\{x_S, y_S\} \subseteq \pi_1$ and $\{x_{\neg S}, y_{\neg S}\} \subseteq \pi_2$. First, we show that $|\pi_1| = |\pi_2| = \ell + 2$. For the sake of contradiction, assume that it is not the case. Then, without loss of generality, $|\pi_1| \leq \ell + 1$ and consequently, π_1 contains at most $\ell - 1$ item agents. Thus, the utility of x_S is $-t + \sum_{v_i \in \pi_1} a_i = -t + t - \epsilon < 0$, where the $-\epsilon$ part follows by our assumption that the sum of elements of every subset of A of size at most $\ell - 1$ is less than t. This contradicts that π is individually rational, so it must hold that $|\pi_1| = |\pi_2| = \ell + 2$. We set $S = \{i \in [2\ell] \mid v_i \in \pi_1\}$ and claim that S is a solution for $\mathcal{I}$. Indeed, the size of S is exactly ℓ. It remains to show that $\sum_{i \in S} a_i = t$. If $\sum_{i \in S} a_i < t$, then $\sum_{v_i \in \pi_1} \omega(x_S v_i) < t$, meaning that π is not individually rational for x_S. Thus, $\sum_{i \in S} a_i \geq t$. If $\sum_{i \in S} a_i > t$, then $\sum_{v_i \in \pi_2} \omega(w_{\neg S} v_i) = 2t - \sum_{v_i \in \pi_1} \omega(x_S v_i) \leq t - 1$, which implies that π is not IR for $x_{\neg S}$. Therefore, it has to be the case that $\sum_{i \in S} a_i = t$, which shows that S is indeed a solution for $\mathcal{I}$.

For ASHGs-SCC, we use the same reduction and set $\text{lb}(j) = \text{ub}(j) = \ell + 2$ for every $j \in [k]$. It is easy to see that the arguments remain identical, as coalition structures with these sizes are the only possible solution for $\mathcal{J}$. This finishes the proof. □

Our next result again shows intractability of deciding whether an instance admits an individually rational outcome. However, this time, we focus on instances where agents have complete preferences, i.e., each agent has non-zero utility from any other agent.

Theorem 3.3. *For every $k \geq 2$, it is NP-complete to decide whether a k-ASHG or ASHG-SCC Γ admits an individually rational coalition structure, even if the valuations are symmetric, the underlying graph is a clique, and the valuations use six different unary encoded values.*

Proof Sketch. We will provide a reduction from the t-Clique problem, where given an instance (H, t), where H is a graph and $t \in \mathbb{N}$, the goal is to decide if H contains a clique on t vertices. This problem is well known to be NP-complete [40].

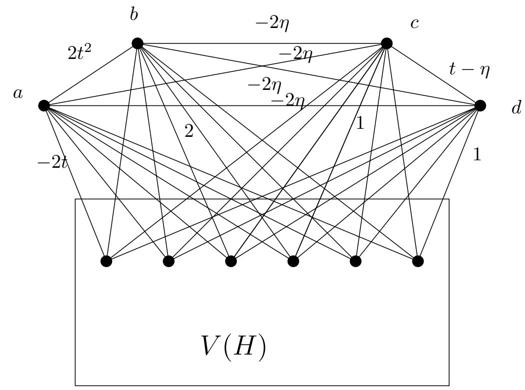

Figure 1: An illustration of the construction used to prove Theorem 3.3.

Construction. See Figure 1 for an illustration. Let $\eta = |V(H)|$. We construct a complete graph G with $V(G) = V(H) \cup \{a, b, c, d\}$. Moreover, for $u, v \in V(H)$, if $uv \in E(H)$ then $\omega(uv) = 2$, else $\omega(uv) = 1$. For each $u \in V(H)$, $\omega(av) = -2t$, $\omega(bv) = 2$, $\omega(cv) = \omega(dv) = 1$. Next, we set $\omega(ab) = 2t^2, \omega(cd) = t - \eta$, and $\omega(ac) = \omega(ad) = \omega(bc) = \omega(bd) = -2\eta$. We remark here that we can safely assume $t - \eta < 0$, as otherwise we can solve our instance of t-Clique in polynomial time. Finally, we have $k = 2$. This completes our reduction.

Correctness. In one direction, we prove that if H has a clique on t vertices, then G admits an IR coalition structure $\pi = (\pi_1, \pi_2)$. Let $U \subset V(H)$ induce a clique on t vertices. Then, $\pi_1 = U \cup \{a, b\}$ and $\pi_2 = V(G) \setminus \pi_1$. First, we show that for each $u \in \pi_1, v_u(\pi_1) \geq 0$. In particular, $v_a(\pi_1) = 0$, $v_b(\pi_1) = 2t^2 + 2t$, and for each $u \in \pi_i \setminus \{a, b\}, v_u(\pi_1) = 0$. Second, to see that for each $u \in \pi_2, v_u(\pi_2) \geq 0$ observe that for each vertex $w \in \pi_2 \setminus \{c, d\}$, only positive edges are incident to w in $G[\pi_2]$ and $v_c(\pi_2) = v_d(\pi_2) = 0$. Thus, $\pi = (\pi_1, \pi_2)$ is individually rational.

In the opposite direction, let $\pi = (\pi_1, \pi_2)$ be an IR solution for Γ. Without loss of generality let us assume that $a \in \pi_1$. We begin by observing some necessary conditions on π in the following claims.

Claim 1. *$b \in \pi_1$ and $c, d \in \pi_2$. Also, $|\pi_2| \geq \eta - t + 2$.*

Now consider π_1. Due to Claim 1, $|\pi_1| \leq t + 2$. Let $U = \pi_1 \setminus \{a, b\}$. We establish in the following that $|U| = k$ and U induces a clique in H to complete our proof. Clearly $|U| \leq t$. Now, for any vertex $u \in U$, we have that $v_u(\pi_1) = v_u(a) + v_u(b) + v_u(U \setminus \{u\}) = -2t + 2 + v_u(U \setminus \{u\})$. Since π is individual rational, we have that $v_u(U \setminus \{u\}) \geq 2t - 2$. Since each vertex in $U \setminus \{u\}$ is coming from H, for any edge uw, where $w \in U \setminus \{u\}$, $\omega(uw) \leq 2$, we have that $|U \setminus \{u\}| \geq t - 1$ and each w is adjacent to u in H (as only then $\omega(uw) \leq 2$). Thus $|U| = t$ and U induces a clique in H. This completes the proof for the hardness of ASHG for cliques, even if the valuations are symmetric, $k = 2$, and the valuations use six different unary encoded values.

Finally, observe that in any individually rational outcome $\pi = (\pi_1, \pi_2)$ we have that $|\pi_1| = t + 2$ and $\pi_2 = \eta - t + 2$. Hence, the same reduction imply the hardness for ASHGs-SCC, since we can set $\text{lb}(1) = \text{ub}(1) = t + 2$ and $\text{lb}(2) = \text{ub}(2) = \eta - t + 2$. □

4 SMALL VERTEX COVER

The proof of Theorem 3.2 relies on valuations that are exponential in the number of agents. If we restrict the valuations to be polynomially bounded by the number of agents, we obtain a polynomial-time algorithm for every graph with constant size vertex cover. The algorithm is based on dynamic programming over the agents outside of the vertex cover.

Theorem 4.1. *If the weights are encoded in unary, there is an algorithm running in $n^{O(\mathrm{vc}(G))}$ time, where $\mathrm{vc}(G)$ is the vertex cover number of G, that decides whether a k-ASHG or ASHG-SCC Γ admits an IR coalition structure.*

Proof. Let C be a vertex cover of G, $I = V(G) \setminus C$, and $\vartheta = |C|$. First, assume that $k \le \vartheta$. In this case, we first guess (by guessing we mean exhaustively trying all possible solutions) a partitioning π_C of the vertices of C in a hypothetical solution. Then, we verify that our guess is correct. To do so, we use a dynamic-programming subprocedure. Specifically, we fix an arbitrary ordering $v_1, \ldots, v_{n-\vartheta}$ of vertices of I and create a dynamic-programming table $\mathrm{DP}[i, \vec{s}, \vec{u}]$, where $i \in [n-\vartheta]_0$, $\vec{s} = (s_1, \ldots, s_k)$ is the number of agents in each coalition, and $\vec{u} = (u_1, \ldots, u_\vartheta)$ is a vector representing the current utility of each agent in the vertex cover. The table stores `true` if it is possible to extend the partition π_C with vertices $v_1, \ldots, v_i$ such that the resulting partition $\pi^{i,\vec{s},\vec{u}}$ is (a) IR for every agent $v_1, \ldots, v_i$, (b) the size of every coalition $\pi_j^{i,\vec{s},\vec{u}}$ in $\pi^{i,\vec{s},\vec{u}}$ is exactly s_j, and (c) $v_{w_j}(\pi^{i,\vec{s},\vec{u}}) = u_j$ for every $w_j \in C$. Otherwise, DP stores `false`.

The computation is then defined as follows. For $i = 0$, which represents an auxiliary basic step when no vertex of I is used, we have

$$\mathrm{DP}[0, \vec{s}, \vec{u}] = \texttt{true} \iff \vec{s} = (|\pi_j^C|)_{j\in[k]} \wedge \vec{u} = (v_{w_i}(\pi^C))_{i\in[\vartheta]}.$$

It is easy to see that for $i = 0$, only one cell is set to `true`; namely, the one corresponding to the guessed partition of vertex cover agents π^C.

For every $i \ge 1$, the computation is then defined as follows. By $(x, \vec{s}_{-i})$ we mean a vector created from $\vec{s}$ by replacing its ith element with the value x.

$$\mathrm{DP}[i, \vec{s}, \vec{u}] = \bigvee_{j\in[k]} \sum_{w\in\pi_j^C} v_{v_i}(w) \ge 0 \wedge$$
$$\mathrm{DP}\left[i-1, (s_j - 1, \vec{s}_{-j}), (u_\ell - v_{w_\ell}(v_i) \cdot [|w_\ell \in \pi_j^C|])_{\ell\in[\vartheta]}\right],$$

where $[|w_\ell \in \pi_j^C|]$ evaluates to 1 if the condition is satisfied and to 0 otherwise.

We prove the correctness using the following two claims. First, we show that whenever $\mathrm{DP}[i, \vec{s}, \vec{u}]$ is set to `true`, there exists a corresponding partial partition $\pi^{i,\vec{s},\vec{u}}$.

Claim 2. *Let $(i, \vec{s}, \vec{u})$ be a triple such that $\mathrm{DP}[i, \vec{s}, \vec{u}] = \texttt{true}$. Then, there exists a partial partition $\pi^{i,\vec{s},\vec{u}}$ of $C \cup \{v_1, \ldots, v_i\}$ such that $\pi_j^C \subseteq \pi_j^{i,\vec{s},\vec{u}}$ for every $j \in [k]$ and satisfying properties (a)-(c).*

In the opposite direction, we show that whenever for some triple $(i, \vec{s}, \vec{u})$ an extension of π^C satisfying properties (a)-(c) exists, our dynamic programming table stores in $\mathrm{DP}[i, \vec{s}, \vec{u}]$ value `true`.

Claim 3. *Let $(i, \vec{s}, \vec{u})$ be a triple such that there exists an extension $\pi^{i,\vec{s},\vec{u}}$ of π^C such that it satisfies all properties (a)-(c). Then, we have* $\mathrm{DP}[i, \vec{s}, \vec{u}] = \texttt{true}$.

Once the dynamic programming table DP is computed, we just check whether there exists a pair $(\vec{s}, \vec{u})$ with $s_j \ge 1$ for every $j \in [k]$ ($\vec{s} = \vec{n}$ in the case of ASHGs-SCC) and $u_i \ge 0$ for every $i \in [\vartheta]$ such that $\mathrm{DP}[n-\vartheta, \vec{s}, \vec{u}] = \texttt{true}$. In other words, whether we can extend the initial partial partition π^C with agents of I such that the resulting partition π (a) is IR for all agents in I (by the definition of $\pi^{n-\vartheta,\vec{s},\vec{u}}$) and, by our choice of $\vec{u}$, also for all agents in C and (b) by our choice of $\vec{s}$, all coalitions are of the correct size (non-empty in case of ASHGs and $\vec{n}$ in case of ASHGs-SCC). Also, observe that, apart from the unary encoding, we have no assumption on the valuations, so our algorithm works even in the case of non-symmetric valuations.

For the running time, there are $O(n) \cdot n^k \cdot (\Delta \cdot v_{\max})^{O(\vartheta)}$ different cells of DP. Since $k \le \vartheta$, $\Delta \in O(n)$, and $v_{\max} \in n^{O(1)}$, the number of cells can be upper-bounded by $n^{O(\vartheta)}$. As every cell can be computed in $O(n)$ time, the overall running time is $n^{O(\vartheta)}$, which is clearly in XP.

If $k > \vartheta$, then we distinguish two cases. First, in the case of ASHGs, we observe that if $k > \vartheta$, then $\mathcal{I}$ is always a Yes-instance. Specifically, for every $v \in C$, we create a coalition $\{v\}$. By this, we obtain ϑ coalitions that are clearly individually rational. Then, we arbitrarily partition the vertices of I into $k - \vartheta$ coalitions. Since I is an independent set of G, the utility of every $v \in I$ is zero, which is the same as being in the singleton coalition. Thus, all constructed coalitions are individually rational.

It remains to show how to deal with the case of ASHGs-SCC. For this setting, we use the same dynamic programming as in the case of $k \le \vartheta$. Again, we guess a partial partitioning π^C of C. There are $k^{O(\vartheta)} \in n^{O(\vartheta)}$ different such partitions π_C. Then, we adjust the dynamic programming from the case of $k \le \vartheta$ as follows. Instead of having a vector $\vec{s} = (s_1, \ldots, s_k)$, we have a vector $\vec{s} = (s_1, \ldots, s_{k'}, s_{k'+1})$, where $k' \le \vartheta$ is the number of coalitions that are non-empty according to π^C. Let $\mathrm{id}\colon [k] \to [k'+1]$ be a function such that for every π_j^C that is nonempty, it returns an unique number from the interval $[1, k']$, and for every π_j^C that is empty, it returns $k'+1$. The semantics of $\vec{s}$ is that for every $i \in [k']$, s_i is the number of agents in coalition $\pi_{\mathrm{id}^{-1}(i)}$ (i.e., in coalitions containing at least one vertex cover agent), and $s_{k'+1}$ contains the number of agents that are in coalitions only with other agents of I. The crucial observation here is that we can split the agents which are not in a coalition with vertex cover agents arbitrarily. Now, we run the same dynamic programming algorithm as in the case of $k \le \vartheta$. Finally, we ask whether there exist $\vec{s} = (n_{\mathrm{id}^{-1}(1)}, \ldots, n_{\mathrm{id}^{-1}(k')}, n - \sum_{i=1}^{k'} n_{\mathrm{id}^{-1}(i)})$ and $\vec{u} = (u_1, \ldots, u_\vartheta)$ with $u_i \ge 0$ for every $i \in [\vartheta]$ such that $\mathrm{DP}[n-\vartheta, \vec{n}, \vec{u}] = \texttt{true}$. Again, there are $O(n) \cdot n^{O(\vartheta)} \cdot (\Delta \cdot v_{\max})^{O(\vartheta)} \in n^{O(\vartheta)}$ cells and each can be computed in polynomial time, which leads the promised running time. The correctness of this approach is then analogous to the correctness of the case with $k \le \vartheta$. ◻

The algorithm from the previous result shows that our problems are in XP when parameterized by the vertex cover number of G. In the following, we show that the algorithm is tight, in the sense that, under the standard theoretical assumptions, it cannot be turned into an FPT one.

Theorem 4.2. *It is W[1]-hard when parameterized by the number of coalitions k, the number of negative edges, and the vertex cover of G, combined, to decide whether a k-ASHG or ASHG-SCC Γ admits an individually rational coalition structure, even if the valuations are symmetric, encoded in unary, and G is a split graph.*

Proof Sketch. We show the result by a parameterized reduction from the Balanced Bin Packing problem. Here, we are given a multiset $A = \{a_1, \dots, a_\mu\}$ of integers, the number of bins B, and the capacity C of every bin. The goal is to decide whether an assignment $\alpha\colon A \to [B]$ exists so that for every bin $j \in [B]$ we have $\sum_{a\in\alpha^{-1}(j)} a \le C$ and $|\alpha^{-1}(j)| = \mu/B$. It is known that Balanced Bin Packing is W[1]-hard when parameterized by the number of bins B, even if we assume that $\sum_{a\in A} a = B\cdot C$ and all integers are encoded in unary [49]. Without loss of generality, we can assume that μ is divisible by B, as otherwise, the instance of Balanced Bin Packing is trivially a No-instance.

Given an instance $\mathcal{I} = (A, B, C)$ of the Balanced Bin Packing problem, we construct an equivalent instance $\mathcal{J} = (N, (v_i)_{i\in N}, k)$ of the ASHGs as follows. First, we create $2B$ *bag agents* $b_1^1, b_1^2 \dots, b_B^1, b_B^2$. For every $j \in [B]$, we set $v_{b_j^1}(b_j^2) = v_{b_j^2}(b_j^1) = -C$, and for every other pair, we set the value to $-(B\cdot C + 1)$. Observe that the bag agents induce a complete graph. Next, for every item $a_i \in A$, we create a corresponding *item agent* v_i. Every item agent v_i has a nonzero value only for bag agents. Specifically, we set $v_{v_i}(b_j^\ell) = v_{b_j^\ell}(v_i) = a_i$ for every $j \in [B]$ and $\ell \in [2]$. That is, the item agents induce an independent set of G. To finalize the construction, we set $k = B$.

To wrap up, observe that we used $k = B$, there are $2B$ bag agents who form a vertex cover of G, and all negative edges are inside the clique on bag agents—that is, there are $\binom{B}{2}$ negative edges. The presented reduction can be clearly done in polynomial time, so the theorem follows. For ASHGs-SCC, we create an instance with an identical graph, identical valuations, and identical k, and set $\mathrm{lb}(j) = \mathrm{ub}(j) = \mu/B$ for every $j \in [k]$. ◻

It turns out that for fixed-parameter tractability, we have to restrict the valuations even more. In particular, there is an FPT algorithm for the parameterization by the vertex cover number of G and maximum value in the valuation functions. The algorithm utilize N-fold ILP as its sub-procedure.

Theorem 4.3. *It can be decided in FPT time with respect to the vertex cover number of G and $\omega_{\max}$ whether an instance of k-ASHGs admits an individually rational coalition structure.*

Proof Sketch. For the sake of exposition, assume that $k \le \vartheta$. Our algorithm is based on a guess of a partitioning of the vertex cover agents combined with an ILP formulation that verifies whether the guess can be extended with agents outside C in an individually rational way. Specifically, we model the extension subprocedure as an N-fold ILP.

More formally, let π^C be one of $k^{O(\vartheta)} \in 2^{O(\vartheta\log\vartheta)}$ possible partitionings of the vertex cover agents into k coalitions. For every vertex $v \in V(G)\setminus C$, we compute the set $I_v = \{i \in [k] \mid \sum_{u\in\pi_i^C} \omega(v,u) \ge 0\}$, i.e., I_v is the set of coalitions such that if we add v to π_i^C, it is individually rational for v. Observe that this can be decided only on the basis of the vertex cover vertices in π_i^C, as $V(G)\setminus C$ is an independent set.

In our N-fold ILP, we have a binary variable $x_{v,i}$ for every $v \in V(G)\setminus C$ and every coalition $i \in I_v$. The constraints of the ILP are as follows. First, we have a single local constraint for every $v \in V(G)\setminus C$, which ensures that each agent outside the vertex cover is assigned to exactly one coalition. Formally, the local constraint is

$$\forall v \in V(G)\setminus C \qquad \sum_{i\in I_v} x_{v,i} = 1\,. \tag{4}$$

Next, we add a set of global constraints that ensure that (a) the coalition structure is individually rational for every vertex cover agent $u \in C$ and (b) none of the k coalitions is empty. The first condition can be encoded as

$$\forall u \in C \qquad v_u(\pi^C) + \sum_{\substack{v\in V(G)\setminus C\\ \pi^C(u)\in I_v}} x_{v,\pi^C(u)}\cdot\omega(u,v) \ge 0\,, \tag{5}$$

while the second condition can be enforced by introducing the following set of constraints

$$\forall i \in [k] \qquad |\pi_i^C| + \sum_{\substack{v\in V(G)\setminus C\\ i\in I_v}} x_{v,i} > 0\,. \tag{6}$$

For the running time, observe that each variable participates in exactly one local constraint, there are $\vartheta + k \in O(\vartheta)$ global constraints, and the maximum coefficient of a variable is $\omega_{\max}$. Therefore, the N-fold ILP can be solved in time $(\omega_{\max})^{O(\vartheta^2)}\cdot n^{O(1)}$, which is clearly in FPT. As we need to run the ILP for every possible π^C, the overall running time of the algorithm is $2^{O(\vartheta\log\vartheta)}\cdot(\omega_{\max})^{O(\vartheta^2)}\cdot n^{O(1)} \in 2^{\vartheta^2\cdot\log(\omega_{\max})}\cdot n^{O(1)}$. Also, observe that we do not have any requirement on the symmetry of valuations. ◻

5 TREE-LIKE PREFERENCE GRAPH

In this section we focus on graphs that are sparse (from a graph-theoretical perspective). This is a reasonable approach in view of Theorem 3.3: both variants of our problem become intractable already if the underlying graph G is a complete graph—arguably the "easiest" structure between dense graphs.

We explore two different dimensions. First, the following hardness proves that the algorithm of Theorem 4.3 cannot be strengthened to more general graph families, as both variants of ASHGs are W[1]-hard when parameterized by the treedepth of G, even if the valuations are binary. Moreover, the reduction asks for partitioning into two coalitions. This strengthen hardness from Theorem 3.2 and shows that even if the number of coalitions is constant, polynomial-time algorithm cannot exist regardless of how restricted the preferences are. Additionally, the result is tight, as if $k = 1$, the only possible solution is the grand coalition, and we can check whether it is IR in polynomial time. The reduction is from the General Factors problem [41], where we are given a bipartite graph $H = (X, Y)$, along with a list function $L\colon V(H) \to 2^{[\Delta(H)]}$, and the goal is to decide if there exists a subset $S \subseteq E(H)$ such that $d_{H-S}(u) \in L(u)$ for all $u \in V(H)$.

Theorem 5.1. *When parameterized by the treedepth of G, it is W[1]-hard to decide if a k-ASHG or ASHG-SCC Γ admits an individually rational coalition structure, even if the valuations are symmetric, binary, and $k = 2$.*

To finalize the complexity picture, in our next result we show that when the valuations are binary, there is an XP algorithm for the parameterization by the celebrated treewidth. As is common for such result, we employ dynamic programming over a nice tree-decomposition of our graph. However, before we can do so, we settle a relation between a chromatic number of G and our problem, which allow us to significantly restrict the instance.

Theorem 5.2. *When the valuations are binary and the problem is parameterized by the treewidth* tw *of G, it is in XP to decide whether k-ASHG Γ admits an individually rational coalition structure.*

Proof Sketch. If $k \geq \text{tw} + 1$, then we can employ the following.

Claim 4. *Given a graph G and a proper coloring of $G = (V, E)$ with c colors, we can compute, in polynomial time, a coalition structure with $k \geq c$ coalitions that is individually rational.*

Proof. Observe first that $k < |V|$ as otherwise a coalition structure that contains all vertices of V as singletons is trivially IR. We will transform the proper c-coloring of G into a proper k-coloring by employing the following procedure. First, we find a color $i \in [c]$ such that there exist at least two vertices u, v in G colored i. This color exists since $c \leq k$ and $k < |V|$. We then recolor u by using a new color, and update c accordingly. We repeat this procedure until $c = k$. At this stage we have a proper k-coloring of G. We define the coalition structure π so that π_i contains the vertices colored i. As each color induces an independent set, π is IR. ◀

Thus, if $k \geq \text{tw} + 1$, then (G, k) is a Yes-instance of ASHGs. Furthermore, we can compute an individually rational coalition structure with k coalitions in FPT time w.r.t. tw by computing the chromatic number of G [2] and employing the previous claim. To do this it suffices to compute the chromatic number of G. This can be done in FPT time w.r.t. tw using a known algorithm [2].

So we may assume that $k \leq \text{tw}$. We solve this case through a bottom-up dynamic programming algorithm on the nice tree-decomposition $\mathcal{T}$ of the graph G rooted at a node r. For a node t of $\mathcal{T}$, we denote by B_t the bag of this node and by $B_t^{\downarrow}$ the set of vertices of the graph that appears in the bags of the nodes of the subtree with t as a root. Observe that $B_t \subseteq B_t^{\downarrow}$.

On a high level, we build a set of partial solutions for each node t of the tree-decomposition. We then extend these partial solutions so that they are also partial solutions of the parent of t. In particular, we keep coalition structures that may not necessarily be individually rational, but *have the potential* to be in the future. We will also show that it suffices to keep at most $(tw \cdot \max\{|X_p|, |X_n|\})^{tw}$ representatives of these coalition structures, where $X_p = \max\{\sum_{w(e)>0} w(e) \mid e \in E\}$ and $X_n = \max\{-\sum_{w(e)<0} w(e) \mid e \in E\}$.

Let B_t be a bag of T. We denote by $C_t = (C_i)_{i \in [k]}$ the projection of a coalition structure of G on $B_t^{\downarrow}$, and we say that C_t is a *partial coalition structure* of G. Also, for any $u \in B_t$, we denote by $id(u)$ the index $i \in [k]$ such that $u \in C_i$ according to C_t. The information we keep for each node t of the tree-decomposition is as follows. For each bag B_t of the tree-decomposition and *partial coalition structure* C_t we store:

- a table I_t such that $I_t[i] = C_i \cap B_t$, for each $i \in [k]$, and
- and a table W_t such that $W_t[u] = \sum_{v \in N(u) \cap C_{id[u]}} w(uv)$, for each $u \in B_t$.

Then, (G, k) is a Yes-instance of ASHGs if and only if there is a "partial" coalition structure of B_r whose table W_r contains only non-negative values. Also, observe that two partial coalition structures that agree on the vertices of B_t, for any t, will behave in the same way regarding the final coalition structures of G. Therefore it suffices to consider only one of such partial coalition structures. In total, we need to keep at most $(tw \cdot \max\{X_p, X_n\})^{O(tw)}$ partial coalition structures per node t.

As for the running time, it suffices to consider the join nodes, as they are the most computationally demanding types of nodes. For any pair of tables (I_{t_1}, W_{t_1}) we consider all pairs of tables (I_{t_2}, W_{t_2}) such that $I_{t_1} = I_{t_2}$. Thus, for any of the $(tw \cdot \max\{X_p, X_n\})^{O(tw)}$ partial coalition structures of t_1, we consider at most $(2 \cdot \max\{X_p, X_n\})^{tw+1}$ partial coalition structures of t_2. In total, this takes $(tw \cdot \max\{X_p, X_n\})^{O(tw)}$ time. □

It is not hard to see that, when the valuations are binary, if we additionally parameterize by the maximum degree of G, the previous algorithm becomes fixed-parameter tractable.

Corollary 5.3. *When the valuations are binary and the problem is parameterized by the treewidth of G and Δ, it is in FPT to decide whether a k-ASHGs Γ admits an IR coalition structure.*

6 CONCLUSIONS

In our work, we study two natural restrictions of hedonic games, namely k-hedonic games, where we are asked to form exactly k stable coalitions, and hedonic games with size-constrained coalitions, where the goal is to construct k coalitions whose sizes are between the given lower and upper bound. We provide a fairly complete algorithmic landscape from the perspective of both classical and parameterized complexity for individual rationality and two restricted variants of hedonic games where valuations are encoded using an underlying preference graph.

However, we believe that our contribution is much broader. First, we bring to light that individual rationality, a traditionally neglected stability notion in coalition formation, can be surprisingly interesting. For example, in the case of k-ASHGs, the complexity picture of our problem is identical to the core stability verification problem in (unrestricted) ASHGs [43], even though the techniques used are very different. We believe it is worth revisiting how traditional stability notions interplay. Moreover, to the best of our knowledge, we are the first using the N-fold ILP technique in the area of coalition formation, which turned out to be a very powerful tool in the design of a tractable algorithm for this domain.

An immediate open question left after our work is the complexity classification of ASHGs with fixed size coalitions under binary preferences and bounded treewidth preference graph. Specifically, the algorithm provided in Theorem 5.2 relies on a connection between ASHGs and the chromatic number of bounded treewidth graphs, which can no longer be easily exploited if we are in the setting of size-constrained coalitions.

ACKNOWLEDGMENTS

This project was supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 101002854), by the European Research Council (ERC) grant titled PARAPATH (grant agreement number 101039913), and co-funded by the European Union under the project Robotics and advanced industrial production (reg. no. CZ.02.01.01/00/22_008/0004590). Foivos Fioravantes is supported by the International Mobility of Researchers MSCA-F-CZ-III at Czech Technical University in Prague, CZ.02.01.01/00/22_010/0008601 Programme Johannes Amos Comenius. Nikolaos Melissinos is partially supported by Charles University projects UNCE 24/SCI/008 and PRIMUS 24/SCI/012, and by the project 25-17221S of GAČR.

REFERENCES

[1] Pulkit Agarwal, Harshvardhan Agarwal, Vaibhav Raj, and Swaprava Nath. 2025. Harmonious Balanced Partitioning of a Network of Agents. In *Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems, AAMAS '25*. IFAAMAS, Richland, SC, USA, 41–49. https://doi.org/10.5555/3709347.3743515

[2] Stefan Arnborg and Andrzej Proskurowski. 1989. Linear Time Algorithms for NP-hard Problems Restricted to Partial k-Trees. *Discrete Applied Mathematics* 23, 1 (1989), 11–24. https://doi.org/10.1016/0166-218X(89)90031-0

[3] Itai Ashlagi and Alvin E. Roth. 2011. Individual Rationality and Participation in Large Scale, Multi-Hospital Kidney Exchange. In *Proceedings of the 12th ACM Conference on Electronic Commerce, EC '11*. ACM, New York, NY, USA, 321–322. https://doi.org/10.1145/1993574.1993625

[4] Haris Aziz, Florian Brandl, Felix Brandt, Paul Harrenstein, Martin Olsen, and Dominik Peters. 2019. Fractional Hedonic Games. *ACM Transactions on Economics and Computation* 7, 2 (2019), 6:1–6:29. https://doi.org/10.1145/3327970

[5] Haris Aziz, Felix Brandt, and Hans Georg Seedig. 2013. Computing Desirable Partitions in Additively Separable Hedonic Games. *Artificial Intelligence* 195 (2013), 316–334. https://doi.org/10.1016/j.artint.2012.09.006

[6] Haris Aziz and Rahul Savani. 2016. Hedonic Games. In *Handbook of Computational Social Choice*, Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia (Eds.). Cambridge University Press, 356–376. https://doi.org/10.1017/CBO9781107446984.016

[7] Suryapratim Banerjee, Hideo Konishi, and Tayfun Sönmez. 2001. Core in a Simple Coalition Formation Game. *Social Choice and Welfare* 18, 1 (2001), 135–153. https://doi.org/10.1007/s003550000067

[8] Hodaya Barr, Yohai Trabelsi, Sarit Kraus, Liam Roditty, and Noam Hazon. 2024. The Complexity of Manipulation of k-Coalitional Games on Graphs. In *Proceedings of the 27th European Conference on Artificial Intelligence, ECAI '24 (FAIA, Vol. 392)*. IOS Press, Amsterdam, Netherlands, 3388–3396. https://doi.org/10.3233/FAIA240889

[9] Vittorio Bilò, Gianpiero Monaco, and Luca Moscardelli. 2022. Hedonic Games With Fixed-Size Coalitions. In *Proceedings of the 36h AAAI Conference on Artificial Intelligence, AAAI '22*. AAAI Press, Palo Alto, CA, USA, 9287–9295. https://doi.org/10.1609/aaai.v36i9.21156

[10] Václav Blažej, Dušan Knop, Jan Pokorný, and Šimon Schierreich. 2024. Equitable Connected Partition and Structural Parameters Revisited: N-Fold Beats Lenstra. In *Proceedings of 49th International Symposium on Mathematical Foundations of Computer Science, MFCS '24 (LIPIcs, Vol. 306)*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 29:1–29:16. https://doi.org/10.4230/LIPIcs.MFCS.2024.29

[11] Hans L. Bodlaender. 1996. A Linear-Time Algorithm for Finding Tree-Decompositions of Small Treewidth. *SIAM Journal on Computing* 25, 6 (1996), 1305–1317. https://doi.org/10.1137/s0097539793251219

[12] Hans L. Bodlaender. 1998. A Partial k-Arboretum of Graphs With Bounded Treewidth. *Theoretical Computer Science* 209, 1 (1998), 1–45. https://doi.org/10.1016/s0304-3975(97)00228-4

[13] Niclas Boehmer and Edith Elkind. 2020. Individual-Based Stability in Hedonic Diversity Games. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence, AAAI '20*. AAAI Press, Palo Alto, CA, USA, 1822–1829. https://doi.org/10.1609/aaai.v34i02.5549

[14] Anna Bogomolnaia and Matthew O. Jackson. 2002. The Stability of Hedonic Coalition Structures. *Games and Economic Behavior* 38, 2 (2002), 201–230. https://doi.org/10.1006/game.2001.0877

[15] Felix Brandt, Martin Bullinger, and Leo Tappe. 2024. Stability Based on Single-Agent Deviations in Additively Separable Hedonic Games. *Artificial Intelligence* 334 (2024), 104160. https://doi.org/10.1016/j.artint.2024.104160

[16] Simina Brânzei and Kate Larson. 2011. Social Distance Games. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence, IJCAI '11*. AAAI Press, Menlo Park, CA, USA, 91–96. https://doi.org/10.5591/978-1-57735-516-8/IJCAI11-027

[17] Robert Bredereck, Edith Elkind, and Ayumi Igarashi. 2019. Hedonic Diversity Games. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS '19*. IFAAMAS, Richland, SC, USA, 565–573. https://dl.acm.org/doi/10.5555/3306127.3331741

[18] Robert Bredereck, Aleksander Figiel, Andrzej Kaczmarczyk, Dušan Knop, and Rolf Niedermeier. 2021. High-Multiplicity Fair Allocation Made More Practical. In *Proceedings of the 20th International Conference on Autonomous Agents and Multiagent Systems, AAMAS '21*. IFAAMAS, Richland, SC, USA, 260–268. https://www.ifaamas.org/Proceedings/aamas2021/pdfs/p260.pdf

[19] Robert Bredereck, Andrzej Kaczmarczyk, Dušan Knop, and Rolf Niedermeier. 2019. High-Multiplicity Fair Allocation: Lenstra Empowered by N-Fold Integer Programming. In *Proceedings of the 20th ACM Conference on Economics and Computation, EC '19*. ACM, New York, NY, USA, 505–523. https://doi.org/10.1145/3328526.3329649

[20] Robert Bredereck, Andrzej Kaczmarczyk, Dušan Knop, and Rolf Niedermeier. 2023. High-Multiplicity Fair Allocation Using Parametric Integer Linear Programming. In *Proceedings of the 26th European Conference on Artificial Intelligence, ECAI '23 (FAIA, Vol. 372)*. IOS Press, Amsterdam, Netherlands, 303–310. https://doi.org/10.3233/FAIA230284

[21] 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 '25*. IFAAMAS, Richland, SC, USA, 418–426. https://dl.acm.org/doi/10.5555/3709347.3743556

[22] Martin Bullinger, Adam Dunajski, Edith Elkind, and Matan Gilboa. 2025. *Single-Deviation Stability in Additively Separable Hedonic Games With Constrained Coalition Sizes*. Technical Report. arXiv. arXiv:2510.12641 [cs.GT] https://arxiv.org/abs/2510.12641

[23] Martin Bullinger and Warut Suksompong. 2024. Topological Distance Games. *Theoretical Computer Science* 981 (2024), 114238. https://doi.org/10.1016/j.tcs.2023.114238

[24] Katarína Cechlárová and Jana Hajduková. 2003. Computational Complexity of Stable Partitions With B-Preferences. *International Journal of Game Theory* 31, 3 (2003), 353–364. https://doi.org/10.1007/s001820200124

[25] Katarína Cechlárová and Jana Hajduková. 2004. Stable Partitions With W-Preferences. *Discrete Applied Mathematics* 138, 3 (2004), 333–347. https://doi.org/10.1016/s0166-218X(03)00464-5

[26] Katarína Cechlárová and Antonio Romero-Medina. 2001. Stability in Coalition Formation Games. *International Journal of Game Theory* 29, 4 (2001), 487–494. https://doi.org/10.1007/s001820000053

[27] 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 '23*. IFAAMAS, Richland, SC, USA, 251–259. https://dl.acm.org/doi/10.5555/3545946.3598644

[28] Marek Cygan, Fedor V. Fomin, Łukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michał Pilipczuk, and Saket Saurabh. 2015. *Parameterized Algorithms*. Springer, Cham. https://doi.org/10.1007/978-3-319-21275-3

[29] Argyrios Deligkas, Eduard Eiben, Stavros D. Ioannidis, Dušan Knop, and Šimon Schierreich. 2025. Balanced and Fair Partitioning of Friends. In *Proceedings of the 39th AAAI Conference on Artificial Intelligence, AAAI '25*. AAAI Press, Washington, DC, USA, 13754–13762. https://doi.org/10.1609/aaai.v39i13.33503

[30] Argyrios Deligkas, Eduard Eiben, Dušan Knop, and Šimon Schierreich. 2024. Individual Rationality in Topological Distance Games Is Surprisingly Hard. In *Proceedings of the 33rd International Joint Conference on Artificial Intelligence, IJCAI '24*. ijcai.org, 2782–2790. https://doi.org/10.24963/ijcai.2024/308

[31] Argyrios Deligkas, Eduard Eiben, Viktoriia Korchemna, and Šimon Schierreich. 2024. The Complexity of Fair Division of Indivisible Items With Externalities. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence, AAAI '24*. AAAI Press, Washington, DC, USA, 9653–9661. https://doi.org/10.1609/aaai.v38i9.28822

[32] Rodney G. Downey and Michael R. Fellows. 2013. *Fundamentals of Parameterized Complexity*. Springer, London. https://doi.org/10.1007/978-1-4471-5559-1

[33] J. H. Drèze and J. Greenberg. 1980. Hedonic Coalitions: Optimality and Stability. *Econometrica* 47, 4 (1980), 987–1003. https://doi.org/10.2307/1912943

[34] Friedrich Eisenbrand, Christoph Hunkenschröder, Kim-Manuel Klein, Martin Koutecký, Asaf Levin, and Shmuel Onn. 2025. Sparse Integer Programming Is Fixed-Parameter Tractable. *Mathematics of Operations Research* 50, 3 (2025), 2141–2156. https://doi.org/10.1287/moor.2023.0162

- [35] Angelo Fanelli, Gianpiero Monaco, and Luca Moscardelli. 2025. Relaxed Core Stability in Hedonic Games. *Artificial Intelligence* 348 (2025), 104394. https://doi.org/10.1016/j.artint.2025.104394
- [36] Foivos Fioravantes, Harmender Gahlawat, and Nikolaos Melissinos. 2025. Exact Algorithms and Lower Bounds for Forming Coalitions of Constrained Maximum Size. In *Proceedings of the 39th AAAI Conference on Artificial Intelligence, AAAI '25*. AAAI Press, Washington, DC, USA, 13847–13855. https://doi.org/10.1609/aaai.v39i13.33514
- [37] Michele Flammini, Gianpiero Monaco, and Qiang Zhang. 2017. Strategyproof Mechanisms for Additively Separable Hedonic Games and Fractional Hedonic Games. In *Proceedings of the 15th International Workshop on Approximation and Online Algorithms, WAOA '17 (LNCS, Vol. 10787)*. Springer, Cham, 301–316. https://doi.org/10.1007/978-3-319-89441-6_22
- [38] Robert Ganian, Thekla Hamm, Dušan Knop, Sanjukta Roy, Šimon Schierreich, and Ondřej Suchý. 2023. Maximizing Social Welfare in Score-Based Social Distance Games. In *Proceedings of the 19th conference on Theoretical Aspects of Rationality and Knowledge, TARK '23 (EPTCS, Vol. 379)*. OPA, 272–286. https://doi.org/10.4204/EPTCS.379.22
- [39] Robert Ganian, Thekla Hamm, Dušan Knop, Šimon Schierreich, and Ondřej Suchý. 2023. Hedonic Diversity Games: A Complexity Picture With More Than Two Colors. *Artificial Intelligence* 325 (2023), 104017. https://doi.org/10.1016/j.artint.2023.104017
- [40] Michael R. Garey and David S. Johnson. 1979. *Computers and Intractability: A Guide to the Theory of NP-Completeness.* W. H. Freeman.
- [41] Gregory Z. Gutin, Eun Jung Kim, Arezou Soleimanfallah, Stefan Szeider, and Anders Yeo. 2012. Parameterized Complexity Results for General Factors in Bipartite Graphs With an Application to Constraint Programming. *Algorithmica* 64, 1 (2012), 112–125. https://doi.org/10.1007/s00453-011-9548-8
- [42] 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. https://doi.org/10.1007/s10878-025-01283-6
- [43] Tesshu Hanaka, Noleen Köhler, and Michael Lampis. 2026. Core Stability in Additively Separable Hedonic Games of Low Treewidth. *Journal of Computer and System Sciences* 157 (2026), 103748. https://doi.org/10.1016/j.jcss.2025.103748
- [44] Raymond Hemmecke, Shmuel Onn, and Lyubov Romanchuk. 2013. N-Fold Integer Programming in Cubic Time. *Mathematical Programming* 137, 1-2 (2013), 325–341. https://doi.org/10.1007/s10107-011-0490-y
- [45] Ton Kloks. 1994. *Treewidth, Computations and Approximations*. Lecture Notes in Computer Science, Vol. 842. Springer, Berlin, Heidelberg. https://doi.org/10.1007/BFB0045375
- [46] Dušan Knop and Martin Koutecký. 2018. Scheduling Meets N-Fold Integer Programming. *Journal of Scheduling* 21, 5 (2018), 493–503. https://doi.org/10.1007/s10951-017-0550-0
- [47] Dušan Knop, Martin Koutecký, and Matthias Mnich. 2020. Voting and Bribing in Single-Exponential Time. *ACM Transactions on Economics and Computation* 8, 3 (2020), 12:1–12:28. https://doi.org/10.1145/3396855
- [48] Tuukka Korhonen and Daniel Lokshtanov. 2023. An Improved Parameterized Algorithm for Treewidth. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing, STOC '23*. ACM, New York, NY, USA, 528–541. https://doi.org/10.1145/3564246.3585245
- [49] Martin Koutecký and Johannes Zink. 2025. Complexity of Scheduling Few Types of Jobs on Related and Unrelated Machines. *Journal of Scheduling* 28, 1 (2025), 139–156. https://doi.org/10.1007/s10951-024-00827-8
- [50] Jérôme Lang, Anja Rey, Jörg Rothe, Hilmar Schadrack, and Lena Schend. 2015. Representing and Solving Hedonic Games With Ordinal Preferences and Thresholds. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems, AAMAS '15*. IFAAMAS, Richland, SC, USA, 1229–1237. https://dl.acm.org/doi/10.5555/2772879.2773308
- [51] Chaya Levinger, Amos Azaria, and Noam Hazon. 2023. Social Aware Coalition Formation With Bounded Coalition Size. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems, AAMAS '23*. IFAAMAS, Richland, SC, USA, 2667–2669. https://dl.acm.org/doi/10.5555/3545946.3599037
- [52] Lily Li, Evi Micha, Aleksandar Nikolov, and Nisarg Shah. 2023. Partitioning Friends Fairly. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence, AAAI '23*. AAAI Press, Washington, DC, USA, 5747–5754. https://doi.org/10.1609/aaai.v37i5.25713
- [53] Jesús A. De Loera, Raymond Hemmecke, Shmuel Onn, and Robert Weismantel. 2008. N-Fold Integer Programming. *Discrete Optimization* 5, 2 (2008), 231–241. https://doi.org/10.1016/j.disopt.2006.06.006
- [54] Kazunori Ohta, Nathanaël Barrot, Anisse Ismaili, Yuko Sakurai, and Makoto Yokoo. 2017. Core Stability in Hedonic Games Among Friends and Enemies: Impact of Neutrals. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence, IJCAI '17*. ijcai.org, 359–365. https://doi.org/10.24963/ijcai.2017/51
- [55] Dominik Peters. 2016. Complexity of Hedonic Games With Dichotomous Preferences. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence, AAAI '16*. AAAI Press, Palo Alto, CA, USA, 579–585. https://doi.org/10.1609/aaai.v30i1.10047
- [56] Dominik Peters. 2016. Graphical Hedonic Games of Bounded Treewidth. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence, AAAI '16*. AAAI Press, Palo Alto, CA, USA, 586–593. https://doi.org/10.1609/aaai.v30i1.10046
- [57] Dominik Peters and Edith Elkind. 2015. Simple Causes of Complexity in Hedonic Games. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence, IJCAI '15*. AAAI Press, Palo Alto, CA, USA, 617–623. http://ijcai.org/Abstract/15/093
- [58] Gleb Polevoy and Marcin Dziubinski. 2022. Fair, Individually Rational and Cheap Adjustment. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence, IJCAI '22*. ijcai.org, 447–453. https://doi.org/10.24963/ijcai.2022/64
- [59] Talal Rahwan, Tomasz P. Michalak, Edith Elkind, Piotr Faliszewski, Jacek Sroka, Michael J. Wooldridge, and Nicholas R. Jennings. 2011. Constrained Coalition Formation. In *Proceedings of the 25th AAAI Conference on Artificial Intelligence, AAAI '11*. AAAI Press, Palo Alto, CA, USA, 719–725. https://doi.org/10.1609/aaai.v25i1.7888
- [60] Alvin E. Roth. 1977. Individual Rationality and Nash's Solution to the Bargaining Problem. *Mathematics of Operations Research* 2, 1 (1977), 64–65. https://doi.org/10.1287/moor.2.1.64
- [61] 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. https://doi.org/10.1016/j.mathsocsci.2018.08.003
- [62] Liat Sless, Noam Hazon, Sarit Kraus, and Michael J. Wooldridge. 2018. Forming k Coalitions and Facilitating Relationships in Social Networks. *Artificial Intelligence* 259 (2018), 217–245. https://doi.org/10.1016/j.artint.2018.03.004
- [63] Shao Chin Sung and Dinko Dimitrov. 2010. Computational Complexity in Additive Hedonic Games. *European Journal of Operational Research* 203, 3 (2010), 635–639. https://doi.org/10.1016/j.ejor.2009.09.004
- [64] Naftali Waxman, Sarit Kraus, and Noam Hazon. 2021. Manipulation of k-Coalitional Games on Social Networks. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence, IJCAI '21*. ijcai.org, 450–457. https://doi.org/10.24963/ijcai.2021/63
- [65] Mason Wright and Yevgeniy Vorobeychik. 2015. Mechanism Design for Team Formation. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence, AAAI '15*. AAAI Press, Palo Alto, CA, USA, 1050–1056. https://doi.org/10.1609/aaai.v29i1.9326