

Strategyproof Mechanisms for Friends and Enemies Games

Michele Flammini, Bojana Kodric, Giovanna Varricchio

Gran Sasso Science Institute, L'Aquila, Italy

{michele.flammini, bojana.kodric, giovanna.varricchio}@gssi.it

Abstract

We investigate strategyproof mechanisms for *Friends and Enemies Games*, a subclass of Hedonic Games in which every agent classifies any other one as a *friend* or as an *enemy*. In this setting, we consider the two classical scenarios proposed in the literature, called *Friends Appreciation* (FA) and *Enemies Aversion* (EA). Roughly speaking, in the former each agent gives priority to the number of friends in her coalition, while in the latter to the number of enemies.

We provide strategyproof mechanisms for both settings. More precisely, for FA we first present a deterministic n -approximation mechanism, and then show that a much better result can be accomplished by resorting to randomization. Namely, we provide a randomized mechanism whose expected approximation ratio is 4, and arbitrarily close to 4 with high probability. For EA, we give a simple $(1 + \sqrt{2})n$ -approximation mechanism, and show that its performance is asymptotically tight by proving that it is NP-hard to approximate the optimal solution within $O(n^{1-\varepsilon})$ for any fixed $\varepsilon > 0$.

Finally, we show how to extend our results in the presence of *neutrals*, i.e., when agents can also be indifferent about other agents, and we discuss anonymity.

Introduction

Hedonic Games (HG), introduced by (Dreze and Greenberg 1980), are Coalitions Formation Games (CFG) where agents have hedonic preferences, i.e., the preference of each agent depends only on the coalition she belongs to and not on how the other agents aggregate.

HGs have been widely studied in the literature (see for instance (Aziz, Brandt, and Harrenstein 2013; Aziz, Brandt, and Seedig 2013; Banerjee, Konishi, and Sönmez 2001; Bogomolnaia and Jackson 2002; Elkind and Wooldridge 2009; Elkind, Fanelli, and Flammini 2016; Gairing and Savani 2019)). According to the assumptions made on the preference profiles, we distinguish different subclasses, among others Additively Separable Hedonic Games (ASHGs) (Banerjee, Konishi, and Sönmez 2001; Aziz, Brandt, and Seedig 2011) and Fractional Hedonic Games

(FHGs) (Aziz et al. 2019; Bilò et al. 2018). In ASHG the utility achieved by an agent in a coalition is the sum of the values she gives to the single participants, and in FHGs the same summand is divided by the size of her coalition. For a clear picture of HGs, their subclasses, and the main investigates stability concepts, see (Aziz and Savani 2016).

Traditionally, in HGs and CFGs the focus has been put onto the existence and efficiency of several solution concepts, based either on individual (Bloch and Diamantoudi 2011; Feldman, Lewin-Eytan, and Naor 2015; Gairing and Savani 2019) or group deviations (Bogomolnaia and Jackson 2002; Banerjee, Konishi, and Sönmez 2001; Elkind and Wooldridge 2009; Gairing and Savani 2019; Igarashi and Elkind 2016), Nash stability and core stability being the main non-cooperative and cooperative notions, respectively. In this setting, given agents' preferences, one can compute an outcome, i.e., a partition of the agents, which is stable, while trying to maximize the global happiness of the participants, also called social welfare. However, while most of the work in the literature implicitly assumes that preferences are known in advance, agents may act strategically by misreporting them in order to achieve a better individual outcome. To avoid this situation, the mechanism design framework allows to split agents into coalitions in a way such that they are incentivized to truthfully report their preferences (Nisan et al. 2007).

In such a setting, we focus on a subclass of HGs in which agents have only either a positive or a negative opinion of the others, considering them either friends or enemies. This type of games, called *Friends and Enemies Games*, have been introduced by (Dimitrov et al. 2006), where two different types of preference profiles are studied: *Friends Appreciation* (FA) and *Enemies Aversion* (EA). Under FA, agents prefer coalitions with a higher number of friends. When the number of friends is equal, they prefer a coalition with a smaller number of enemies. Conversely, under EA, agents always prefer coalitions with a smaller number of enemies and, in case of a tie, they prefer coalitions with a bigger number of friends.

Our aim is to design strategyproof mechanisms for both FA and EA preference profiles, improving upon the previously proposed solutions to this problem (Dimitrov and Sung 2004), both in terms of social welfare of the computed

outcome and/or time complexity.

Our Contribution We investigate Friends and Enemies Games under FA and EA preference profiles. For FA we provide both a deterministic and a randomized truthful mechanism. While the deterministic mechanism has an approximation ratio of n , where n is the number of agents, the randomized one has an expected approximation ratio of 4 and $4(1 + \varepsilon)$ with high probability, for any fixed $\varepsilon > 0$. For EA profiles, we first show that no polynomial time algorithm can have an approximation ratio that is in $O(n^{1-\varepsilon})$, since this problem is as hard to approximate as the MAX CLIQUE problem. Then, we give a simple polynomial time mechanism reaching an approximation ratio of $(1 + \sqrt{2})n$, which we show is asymptotically tight. Finally, we show how to extend our results to the case of neutrals, and we discuss anonymity.

Related Work Friends and Enemies Games with FA and EA preference profiles are an example of ASHG. This model has been introduced in (Dimitrov et al. 2006), where the authors focus on weak and strong core stability notions. While for FA it is always possible to compute a strict core stable coalition structure in polynomial time, for EA, even if it always exists, it is NP-hard to find a core stable outcome. Moreover, in (Sung and Dimitrov 2007) it is shown that, for EA, determining whether a coalition structure is core stable is co-NP complete. In (Dimitrov and Sung 2004) individual deviations are considered: the authors study Nash stable, individually and contractually individually stable outcomes. While Nash existence is not always guaranteed, individually stable, and thus also contractually individually stable, outcomes always exist.

Subsequent works on Friends and Enemies Games still focus on core and strict core stability notions but allow the presence, beside friends and enemies, also of neutrals that have no impact on agents' preferences (Ohta et al. 2017). In (Barrot et al. 2019) neutrals can have a lower order positive or negative impact on the preferences. More specifically, the authors focus on FA preference profiles and distinguish Friends Appreciation with extroverted and with introverted agents, where for the same number of friends and enemies it is more preferable having a higher or a lower number of neutrals, respectively. They consider core stable and individually stable outcomes, studying also the hardness of deciding their existence.

In the study of HGs, another interesting challenge is to control the agents' strategic behavior through mechanism design. The goal here is to provide algorithms resistant to input manipulations. Such algorithms, also called mechanisms, are said to be strategyproof. In the classical framework of mechanism design, payments are used to achieve strategyproofness. However, this is not always possible, both for legal and ethical issues (Nisan et al. 2007), or simply because allowing payments is not feasible (Procaccia and Tennenholtz 2013). Thus, in (Wright and Vorobeychik 2015) the authors focus on strategyproof mechanisms without money for ASHG with positive preferences. They provide a mech-

anism which returns the grand coalition, and investigate the same problem with constraints on the coalitions size and with respect to (approximate) envy-freeness. A study of the properties of strategyproof core stable solutions for hedonic games is also provided in (Rodríguez-Álvarez 2009).

Strategyproof deterministic and randomized mechanisms for ASHG and FHGs are provided in (Flammini, Monaco, and Zhang 2017), where different types of valuations are considered. Of particular interest for our work are ASHG with individual *duplex* valuations, that is, having values in $\{-1, 0, 1\}$. They can be seen as an HG with friends and enemies, where friends have the same impact as enemies. Determining a better mechanism for our EA model but with the inclusion of neutrals would, therefore, yield a better mechanism for duplex valuations in (Flammini, Monaco, and Zhang 2017), because in our EA setting enemies have a negative effect that is higher than the positive effect of friends.

Finally, in (Dimitrov and Sung 2004), the authors also propose strategyproof mechanisms for both FA and EA preference profiles. However, they do not consider the efficiency (both in terms of quality of the returned solutions and of time complexity) of their mechanisms. We significantly improve over their results. For friends appreciation, the algorithm from (Dimitrov and Sung 2004) has an unbounded approximation ratio, while we provide a deterministic n -approximation and a randomized $4(1 + \varepsilon)$ -approximation (both in expectation and with high probability). For enemy aversion, (Dimitrov et al. 2006) only gives a non-polynomial algorithm, while we provide a polynomial one that has bounded approximation. More precisely, we prove that it has linear approximation and show that a sublinear approximation cannot be achieved in polynomial time.

Model and Preliminaries

In the classical framework of HGs we are given a set of selfish agents $N = \{1, \dots, n\}$, and the goal of the game is to partition them into disjoint coalitions $\mathcal{C} = \{C_1, \dots, C_m\}$ such that $\cup_{i=1}^m C_i = N$. Such a partition is also called an *outcome* or a *coalition structure*. The *grand coalition* GC is a coalition structure which consists of only one coalition containing all of the agents and a *singleton coalition* is any coalition of size 1. We denote by \mathcal{C} the set of all possible outcomes, and by $\mathcal{C}(i)$ the coalition that agent i belongs to in an outcome $\mathcal{C} \in \mathcal{C}$. We assume that agent i has a preference relation \prec_i over N_i , where N_i is the family of subsets of N containing i . According to \prec_i , for every $X, Y \in N_i$, we say that agent i prefers X to Y whenever $Y \prec_i X$. A *preference profile* is a collection of agents' preferences $P = \{\prec_i\}_{i \in N}$.

In the special case of HGs with friends and enemies, every agent i partitions the other agents into a set of friends F_i and a set of enemies E_i , with $F_i \cup E_i = N \setminus \{i\}$ and $F_i \cap E_i = \emptyset$. Using such a partition, different settings can be defined.

For every $i \in N$ and for every $X, Y \in N_i$, if coalition X is more or equally preferred than coalition Y by the agent i , we write $X \succeq_i Y$. A preference profile P is based on *Friends Appreciation* (FA) when $X \succeq_i Y$ iff

$$|X \cap F_i| > |Y \cap F_i| \text{ or} \\ |X \cap F_i| = |Y \cap F_i| \text{ and } |X \cap E_i| \leq |Y \cap E_i| ,$$

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

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

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

Friends and Enemies Games are a proper subclass of ASHG, because each agent i has valuation $v_i(j)$ for every other agent j (she considers j to be either a “friend” or an “enemy”), and her utility for being in a given coalition C is $u_i(C) = \sum_{j \in C \setminus \{i\}} v_i(j)$. In the FA case, setting the valuation functions in such a way that, for every agent $i \in N$, $v_i(j) = 1$ if $j \in F_i$, i.e., if j is a friend, and $v_i(j) = -1/n$ if $j \in E_i$, i.e., if j is an enemy, correctly encodes the setting. In other words, the positive effect of one friend is greater than the overall possible negative effect due to the enemies. Similarly, in the EA case, the valuations can be set in such a way that, for every agent $i \in N$, $v_i(j) = 1/n$ if $j \in F_i$, and $v_i(j) = -1$ if $j \in E_i$ (see for instance (Dimitrov et al. 2006)). ASHG have generally been widely investigated in the literature, as they are both able to capture several realistic scenarios in coalition formation games, and they allow for a succinct graph representation in which the node set is N , i.e., nodes represent agents, and every arc (i, j) has weight $v_i(j)$. In this work, we will use the just introduced valuation functions and the aforementioned graph representation.

Each agent i will have to communicate her preferences. Without loss of generality we assume that, both in the FA and the EA case, this is accomplished by declaring $d_i = (F_{d_i}, E_{d_i})$ where F_{d_i} is the declared set of friends, and $E_{d_i} = N \setminus (F_{d_i} \cup \{i\})$ the declared set of enemies of agent i . Notice that, as agents are self-interested entities, each declaration may differ from the actual sets of friends and enemies of agent i , $v_i = (F_i, E_i)$, i.e., it is possible that $F_i \neq F_{d_i}$ and $E_i \neq E_{d_i}$.

Let $\mathbf{d} = (d_1, \dots, d_n)$ be the profile of the agents’ declarations, and \mathbf{d}_{-i} be the subprofile of all the reported valuation functions, except d_i . Let (d'_i, \mathbf{d}_{-i}) be the profile obtained by substituting d'_i to d_i in \mathbf{d} .

A deterministic mechanism \mathcal{M} maps every profile \mathbf{d} to a set of disjoint coalitions $\mathcal{M}(\mathbf{d}) \in \mathcal{C}$. We denote by $\mathcal{M}_i(\mathbf{d})$ the coalition that \mathcal{M} assigned to agent i . The utility of agent i is given by $u_i(\mathcal{M}_i(\mathbf{d}))$. A deterministic mechanism \mathcal{M} is *strategyproof* if for every $i \in N$, every subprofile \mathbf{d}_{-i} , every v_i and every d_i , it holds that $u_i(\mathcal{M}_i(v_i, \mathbf{d}_{-i})) \geq u_i(\mathcal{M}_i(d_i, \mathbf{d}_{-i}))$.

A randomized mechanism \mathcal{M} maps every profile \mathbf{d} to a distribution Δ over the set of all possible outcomes \mathcal{C} . The expected utility of agent i is given by $\mathbb{E}[u_i(\mathcal{M}_i(\mathbf{d}))] = \mathbb{E}_{\mathcal{C} \sim \Delta}[u_i(\mathcal{C}(i))]$. A randomized mechanism \mathcal{M} is *strategyproof in expectation* if for every $i \in N$, every subprofile \mathbf{d}_{-i} , every v_i and every d_i , $\mathbb{E}[u_i(\mathcal{M}_i(v_i, \mathbf{d}_{-i}))] \geq \mathbb{E}[u_i(\mathcal{M}_i(d_i, \mathbf{d}_{-i}))]$.

We are interested in strategyproof mechanisms that perform well with respect to the goal of maximizing the classical utilitarian social welfare, that is, the sum of the utilities achieved by all the agents. Namely, the social welfare of a given outcome \mathcal{C} is $\text{SW}(\mathcal{C}) = \sum_{i \in N} u_i(\mathcal{C}(i))$. We denote by $\text{SW}(\mathbf{d}) = \sum_{i \in \mathcal{C}} u_i(\mathcal{C})$ the overall social welfare achieved by the agents belonging to a given coalition \mathcal{C} .

In the following, given an FA or an EA profile \mathbf{d} , $\mathcal{C}^*(\mathbf{d})$ will denote an optimal outcome for the game instance expressed by \mathbf{d} , and $\text{opt}(\mathbf{d})$ its social welfare.

We measure the performance of a mechanism by comparing the social welfare it achieves to the optimal one. More precisely, the approximation ratio of a deterministic mechanism \mathcal{M} is defined as $r^{\mathcal{M}} = \sup_{\mathbf{d}} \frac{\text{opt}(\mathbf{d})}{\text{SW}(\mathcal{M}(\mathbf{d}))}$. For randomized mechanisms, the approximation ratio is computed with respect to the expected social welfare, that is, $r^{\mathcal{M}} = \sup_{\mathbf{d}} \frac{\text{opt}(\mathbf{d})}{\mathbb{E}[\text{SW}(\mathcal{M}(\mathbf{d}))]}$.

We say that a deterministic mechanism \mathcal{M} is *admissible* if it always guarantees a non negative social welfare, i.e., if $\text{SW}(\mathcal{M}(\mathbf{d})) \geq 0$ for any possible list of preferences \mathbf{d} . Similarly, a randomized mechanism \mathcal{M} is *admissible* if $\mathbb{E}[\text{SW}(\mathcal{M}(\mathbf{d}))] \geq 0$ holds for every \mathbf{d} . In the following, we will always implicitly restrict our attention to admissible mechanisms. In fact, a simple admissible strategyproof mechanism can be trivially obtained by putting every agent into a separate singleton coalition, regardless of the declared valuations.

Graph Representation As already mentioned, Friends and Enemies Games, being a proper subclass of ASHG, can be suitably represented by means of graphs. More specifically, the following representations will be useful for our purposes. For a given profile \mathbf{d} ,

- $G_{\mathbf{d}}^+ = (N, F_{\mathbf{d}})$ is a directed graph with the edge set $F_{\mathbf{d}} = \{(i, j) \mid i, j \in N, j \in F_{d_i}\}$, i.e., $G_{\mathbf{d}}^+$ contains only edges corresponding to friendship relations.
- $G_{\mathbf{d}}^* = (N, F_{\mathbf{d}}^+)$ is an undirected graph with the edge set $F_{\mathbf{d}}^+ = \{\{i, j\} \mid i, j \in N, j \in F_{d_i} \wedge i \in F_{d_j}\}$, i.e., each edge corresponds to a mutual friendship relation.

Examples of $G_{\mathbf{d}}^+$ and $G_{\mathbf{d}}^*$ w.r.t. \mathbf{d} are depicted in Figure 1.




Figure 1: An example of graphs $G_{\mathbf{d}}^+$ and $G_{\mathbf{d}}^*$. Solid (resp. dashed) edges represent friendship (resp. enemy) relations.

Given a directed graph G and a pair of nodes x, y in G , x and y are *weakly connected* if they are connected in the undirected version of G . Moreover, a *weakly connected component* in G is a maximal subset of nodes in G that are weakly connected. Similarly, given a directed graph G and a pair of nodes x, y in G , x and y are *strongly connected* if there




Figure 2: The lower bound instances for deterministic strategyproof mechanisms in the FA case.

exists a directed path both from x to y and from y to x in G . Thus, a *strongly connected component* in G is a maximal subset of nodes in G that are strongly connected.

In what follows, for the sake of simplicity, we will often identify a coalition C with the subgraph it induces in $G_{\mathbf{d}}^+$ or $G_{\mathbf{d}}^*$. Furthermore, we will denote by $f = f_{\mathbf{d}} = |F_{\mathbf{d}}|$ the overall number of friendship relations in the profile \mathbf{d} .

Friends Appreciation

FA preference profiles can be suitably analyzed exploiting the representation graph $G_{\mathbf{d}}^+$. The following simple lemma will prove to be useful in the sequel.

Lemma 1. *Let C be any coalition inducing a subgraph of g edges in $G_{\mathbf{d}}^+$, or analogously containing g positive relations, and let $k = |C|$. Then, $SW(C) = g \cdot (1 + \frac{1}{n}) - \frac{k(k-1)}{n}$.*

A useful topology for proving some of our results is given by a *directed star*, consisting of one central node and $n - 1$ remaining leaf nodes having an edge toward it (see Figure 2a). The following lemma concerns optimal solutions in directed stars.

Lemma 2. *Given an FA profile \mathbf{d} , if $G_{\mathbf{d}}^+$ is a directed star of n agents, then $\frac{n^2-1}{4n} \leq \text{opt}(\mathbf{d}) \leq \frac{n}{4}$.*

Proof. We show that an optimal solution $C^*(\mathbf{d})$ can be obtained by an outcome with one coalition containing the center and $\frac{n}{2}$ or $\frac{n\pm 1}{2}$ leaf agents (depending on the parity of n), while all other agents are in singleton coalitions. To this aim, notice first that any coalition without the central agent does not have a strictly positive social welfare.

Given a coalition C_k of size k that forms a star (or, in fact, any tree) in $G_{\mathbf{d}}^+$, by applying Lemma 1 with $g = k - 1$, we obtain

$$u(C_k) = (k - 1) \left(1 - \frac{k - 1}{n} \right) \quad (1)$$

A standard mathematical argument shows that $u(C_k)$ is maximized for $k = \frac{n}{2} + 1$ if n is even and $k = \frac{n\pm 1}{2} + 1$ if n is odd, thus, yielding the lemma. \square

Strategyproof Mechanisms

In (Dimitrov and Sung 2004) the authors show that for FA returning the strongly connected components of $G_{\mathbf{d}}^+$ is a strategyproof mechanism. However, this mechanism has a very bad performance. In fact, if we consider the example of a

directed star depicted in Figure 2a, the achieved social welfare is 0, since agents are split into singleton coalitions. On the other hand, as shown in Lemma 2, the social optimum is linear in the number of agents, and, thus, the mechanism has an unbounded approximation ratio.

One might wonder whether returning the social optimum is strategyproof. Unfortunately, this is not the case. It is easy to check that a mechanism always returning an optimal outcome is not strategyproof.

Let us start with one of the simplest examples of a deterministic strategyproof mechanism.

Mechanism \mathcal{M}_1 . *Given an FA preference profile \mathbf{d} , \mathcal{M}_1 returns the grand coalition GC .*

Mechanism \mathcal{M}_1 is strategyproof, as the outcome does not depend on the agents' declared profile \mathbf{d} . However, it does not always return an admissible outcome, e.g., for any profile with the total number of friendship relations $f < n - 1$. Indeed, by applying Lemma 1 with $g = f$ and $k = n$, we see that $u(GC) = f \cdot (1 + \frac{1}{n}) - (n - 1) < 0$.

Our goal is not only to find admissible, but also, and more importantly, to find good approximation mechanisms. To this aim, we will first identify a broad class of strategyproof mechanisms, to which our aforementioned deterministic and randomized mechanism for FA profiles will belong to.

Definition 3. *We define \mathcal{M} as the class of mechanisms \mathcal{M} that, given an FA preference profile \mathbf{d} , work as follows:*

1. \mathcal{M} selects, independently from \mathbf{d} , deterministically or at random, a partition \mathcal{P} of N ;
2. \mathcal{M} computes a coalition structure \mathcal{C} s.t. for each agent $i \in N$, $\mathcal{C}(i)$ is the (maximal) weakly connected component containing i in the subgraph of $G_{\mathbf{d}}^+$ induced by $\mathcal{P}(i)$.

In other words, once \mathcal{P} is computed, an agent will be assigned to a subset of $\mathcal{P}(i)$ containing all friends she has in $\mathcal{P}(i)$, and such a set is the minimal one guaranteeing the same property for all the other agents in $\mathcal{C}(i)$.

Theorem 4. *\mathcal{M} is an admissible and strategyproof class for FA profiles.*

Proof. First, note that for any $\mathcal{M} \in \mathcal{M}$, if \mathcal{M} returns an agent i in a coalition that is larger than a singleton, this means that her contribution to the social welfare of the coalition is at least 1 (as there exists at least one agent who considers her to be a friend) and $u_i(\mathcal{M}_i(\mathbf{d})) \geq -\frac{n-1}{n}$. Therefore, \mathcal{M} is admissible.

Next, note that by the definition of the class \mathcal{M} , the declaration d_i of i cannot influence the partition \mathcal{P} selected by \mathcal{M} . Given any subprofile \mathbf{d}_{-i} declared by the remaining agents, consider now the weakly connected components in the subgraph of $G_{\mathbf{d}}^+$ induced by the agents in $\mathcal{P}(i)$ and \mathbf{d}_{-i} , that is, by deleting the outgoing edges of i in $G_{\mathbf{d}}^+$. By the definition of \mathcal{M} , all the agents belonging to each of these components will never be split into different coalitions. Since one friend contributes more to the utility of i than all her enemies subtract, the best utility that agent i can hope to achieve is obtained when she is put together with all the agents in the weakly connected component determined by $\mathcal{P}(i)$, \mathbf{d}_{-i} and F_i , which will contain all her friends in $\mathcal{P}(i)$. But this

exactly is the outcome selected by \mathcal{M} if i simply declares her true valuation v_i , and therefore reporting any different declaration d_i cannot achieve a better utility. \square

While every mechanism in the class \mathcal{M} is always strategyproof, it is not true that every strategyproof mechanism is in \mathcal{M} . A simple counter-example is given by the mechanism of (Dimitrov and Sung 2004), which computes the strongly connected components.

Deterministic Mechanisms We present a mechanism from the just described class \mathcal{M} and give a bound on its approximation ratio.

Mechanism \mathcal{M}_2 . Given an FA preference profile \mathbf{d} , \mathcal{M}_2 outputs as coalitions the weakly connected components of $G_{\mathbf{d}}^+$.

Theorem 5. Mechanism \mathcal{M}_2 is admissible, strategyproof and has an approximation ratio $r^{\mathcal{M}_2} \leq n$.

Proof. Mechanism \mathcal{M}_2 belongs to the class \mathcal{M} , and thus, by Theorem 4, it is admissible and strategyproof.

For what concerns its approximation ratio, let us first assume that $G_{\mathbf{d}}^+$ is weakly connected, so then $f \geq n - 1$. Since, by Lemma 1, $\text{SW}(\mathcal{M}_2(\mathbf{d})) = f \cdot (1 + \frac{1}{n}) - (n - 1)$ and $\text{opt}(\mathbf{d}) \leq f$, then $r^{\mathcal{M}_2} \leq \frac{f}{f \cdot (1 + \frac{1}{n}) - (n - 1)}$. Such a ratio is decreasing in f , and at most equal to n for $f = n - 1$, thus, implying the claim.

If $G_{\mathbf{d}}^+$ is not weakly connected, the same argument can be repeated on each weakly connected component. In fact, the worst case is reached when $G_{\mathbf{d}}^+$ is weakly connected. \square

Notice that the poor performance of mechanism \mathcal{M}_2 is due to instances for which the number of friendship relations f is very close to n . However, as soon as $G_{\mathbf{d}}^+$ becomes denser, and in particular when f is at least $c \cdot n$ for any fixed constant $c > 1$, the ratio becomes constant. That is, for $f \geq 2(n - 2)$, the ratio is $\frac{f}{f \cdot (1 + \frac{1}{n}) - (n - 1)} < 2$.

It is possible to show the following lower bound on the approximation ratio achievable by deterministic mechanisms. The lower bound example uses the two instances shown in Figure 2a and 2b, which differ only by the reported preferences of agent i .

Theorem 6. No deterministic strategyproof mechanism for FA profiles can have an approximation ratio of less than 2.

Randomized Mechanisms So far, we have provided a deterministic strategyproof mechanism achieving a linear approximation ratio. We now present a randomized mechanism which improves upon this ratio, reaching a constant approximation factor in expectation and with high probability.

Mechanism \mathcal{M}_3 . Given an FA preference profile \mathbf{d} , \mathcal{M}_3 first randomly generates a partition $\mathcal{P} = \{P_1, P_2\}$, by placing every agent $i \in N$ in P_1 or P_2 uniformly and independently at random (i.e., i is put in P_1 with probability $\frac{1}{2}$, and in P_2 otherwise), and then outputs the weakly connected components of P_1 and P_2 in $G_{\mathbf{d}}^+$.

The idea underlying the definition of \mathcal{M}_3 is firstly maintaining the good performance of the deterministic mechanism \mathcal{M}_2 when $f \geq 2(n - 2)$. This is accomplished by retaining, with respect to the coalitions formed by \mathcal{M}_2 , each friendship relation with probability 1/2 and each enemy relation with probability at most 1/2, thus, yielding an expected social welfare which is at least half of the one of \mathcal{M}_2 . At the same time, \mathcal{M}_3 avoids the pathological instances causing \mathcal{M}_2 to have an approximation ratio linear in n . Such instances occur when $f \approx n$ and the graph contains a large weakly connected component of size close to n . In these cases, it is possible to achieve a better social welfare by splitting the large weakly connected component into smaller coalitions. This can be accomplished by creating a partition \mathcal{P} whose expected number of agents on one side is $n/2$. A paradigmatic example of the difference between \mathcal{M}_2 and \mathcal{M}_3 is exhibited by considering directed stars, for which \mathcal{M}_2 achieves an approximation ratio of n , while \mathcal{M}_3 returns an almost optimal solution in expectation and also with high probability. In fact, as shown in Lemma 2, the optimal solution for a star is obtained by putting roughly half of the leaves in the same coalition with the center and the other leaves in singletons. Mechanism \mathcal{M}_3 does the same by exploiting the partition \mathcal{P} .

Theorem 7. \mathcal{M}_3 is admissible and strategyproof, and $r^{\mathcal{M}_3} \leq 4$ in expectation. Moreover, for any fixed $\varepsilon > 0$, $r^{\mathcal{M}_3} \leq 4(1 + \varepsilon)$ with high probability.

Proof (Sketch). The mechanism belongs to the class \mathcal{M} , and therefore it is admissible and strategyproof. We now prove the upper bound on the expected approximation ratio.

Given the FA profile \mathbf{d} as input, let \bar{f} be the expected number of positive relations with both endpoints either in P_1 or in P_2 . Then, $\bar{f} = \frac{f}{2}$. Moreover, the expected number of agents in P_1 is $n/2$. Recalling that $\text{opt}(\mathbf{d}) \leq f$, in our analysis we distinguish 3 cases.

Case 1: $\bar{f} \leq \frac{n}{2} - 1$ (and thus, $f \leq n - 2$). In this scenario the worst case outcome \mathcal{I} occurs when all of the \bar{f} positive relations contribute to just one weakly connected component T that forms a tree on one side of the partition (either in P_1 or P_2), while the other side is completely disconnected. Then, $\text{SW}(\mathcal{I}) = \text{SW}(T)$ and, according to Eq. (1), $\text{SW}(T) = \bar{f} \left(1 - \frac{\bar{f}}{n}\right)$. Thus, $r^{\mathcal{M}_3} \leq \frac{2}{1 - \frac{\bar{f}}{2n}} \leq \frac{2}{1 - \frac{1}{2}} = 4$.

Case 2: $\frac{n}{2} \leq \bar{f} \leq n - 1$. Here, the worst case outcome \mathcal{I} occurs when $\frac{n}{2} - 1$ of the \bar{f} positive relations form a tree T_1 connecting all the $\frac{n}{2}$ nodes on one side of the partition and on the other side the remaining $\bar{f} - \frac{n}{2} + 1$ positive relations form a tree T_2 of $\bar{f} - \frac{n}{2} + 2$ nodes. Thus, applying Eq. (1), the social welfare of the outcome is $\frac{n}{4} - \frac{1}{n} + (\bar{f} - \frac{n}{2} + 1) \left(1 - \frac{\bar{f} - \frac{n}{2} + 1}{n}\right) \geq 2\bar{f} \left(1 + \frac{1}{n}\right) - \frac{n}{2} - \frac{\bar{f}^2}{n}$. Then, $r^{\mathcal{M}_3} \leq \frac{f}{f \cdot (1 + \frac{1}{n}) - \frac{n}{2} - \frac{f^2}{4n}}$. This ratio is a convex function for $n \leq f \leq 2n - 2$ and takes values $4\frac{n}{n+4}$ and $\frac{4(n-1)}{n+4-\frac{6}{n}}$ for $f = n$ and $f = 2n - 2$, respectively. Thus, $r^{\mathcal{M}_3} \leq 4$.

Case 3: $\bar{f} \geq n$. The worst case outcome \mathcal{I} occurs when

both P_1 and P_2 are weakly connected by the \bar{f} retained positive relations. In that case, $\text{SW}(\mathcal{I}) \geq \bar{f} \left(1 + \frac{1}{n}\right) - \frac{n}{2} + 1 \geq \frac{1}{2}(f - n)$ thus, $r^{\mathcal{M}_3} \leq \frac{2f}{f-n} \leq 4 \frac{n-1}{n-2} \leq 4$.

The $4(1 + \varepsilon)$ approximation bound with high probability holds by applying tail distribution bounds to the random variables representing the sizes of P_1 and P_2 and the number of positive relations in P_1 and P_2 . \square

Before concluding the section, let us remark that, even if we did not state it explicitly, all of our mechanisms for FA profiles are efficient, that is, they compute their outcomes in polynomial time.

Enemies Aversion

In this section we consider EA preference profiles. Unlike in the FA case, here we will mostly use the undirected graph representation $G_{\mathbf{d}}^*$. Again, we will often identify a coalition C with the subgraph it induces in $G_{\mathbf{d}}^*$.

Different from the FA case, we show that for EA profiles good approximation mechanisms cannot be found not due to strategyproofness, but because of the inherent hardness of polynomial time approximation of the optimal solution below a ratio linear in n .

To this aim we provide an approximation preserving reduction from the MAXCLIQUE problem, in order to transfer the well known $O(n^{1-\varepsilon})$ inapproximability result to our problem. To this aim, the following definitions and facts will show to be useful.

Definition 8. *Given an undirected graph $G = (V, E)$, a (disjoint) clique partition of G is a collection of cliques $\mathcal{K} = \{K_1, \dots, K_m\}$ such that for each $i, j \in \{1, \dots, m\}$ where $i \neq j$, $K_i \cap K_j = \emptyset$ and $\bigcup_{i=1}^m K_i = V$.*

We are interested in clique partitions that induce good outcomes for EA profiles.

Definition 9. *A best clique partition for an EA profile \mathbf{d} is any clique partition \mathcal{K}^* of $G_{\mathbf{d}}^* = (N, F_{\mathbf{d}}^+)$ such that \mathcal{K}^* achieves the highest possible social welfare in the instance induced by \mathbf{d} .*

The following two lemmas concern the structure of optimal outcomes.

Lemma 10. *Given an EA preference profile \mathbf{d} , no agent in an optimal outcome $C^*(\mathbf{d})$ can be the endpoint of more than one negative relation in her coalition.*

Lemma 11. *Given an EA preference profile \mathbf{d} , there exists at most one coalition in $C^*(\mathbf{d})$ that is not a clique in $G_{\mathbf{d}}^*$.*

Proof. Let us assume that $C^* \in C^*(\mathbf{d})$ is not a clique in $G_{\mathbf{d}}^*$, and let $k = |C^*|$ and k_1 be the number of negative relations in C^* . Since by Lemma 10 such relations do not share any endpoint, picking one endpoint per relation and putting it in a new coalition K_1 , we can split C^* into two cliques K_1 and K_2 of sizes k_1 and $k_2 = k - k_1$, respectively, so that all the negative relations now lie in the cut between K_1 and K_2 . Thus, since this deletes k_1 negative relations and $2k_1(k - k_1) - k_1$ positive relations from C^* , $\text{SW}(C^*)$ is equal to $\text{SW}(K_1) + \text{SW}(K_2) + \frac{1}{n} \cdot$

$(2k_1(k - k_1) - k_1) - k_1$. Furthermore, since C^* belongs to $C^*(\mathbf{d})$, $\frac{1}{n}(2k_1(k - k_1) - k_1) - k_1 \geq 0$ must hold, which implies $k \geq \frac{n+1}{2} + k_1$. Therefore, because a coalition in $C^*(\mathbf{d})$ that is not a clique must contain more than half of the agents from N , it must also be unique. \square

By Lemma 11, even if in general the optimal solution for an EA preference profile \mathbf{d} is not a partition into cliques of $G_{\mathbf{d}}^*$, it is ‘‘almost a clique partition’’. So, a natural next step is trying to quantify how far $\text{opt}(\mathbf{d})$ is from the social welfare of a best clique partition for an EA profile \mathbf{d} .

Lemma 12. *Given an EA preference profile \mathbf{d} and a best clique partition \mathcal{K}^* of $G_{\mathbf{d}}^*$, $\text{opt}(\mathbf{d}) \leq \left(\frac{1+\sqrt{2}}{2}\right) \cdot \text{SW}(\mathcal{K}^*)$.*

Proof. Let us assume that $C^*(\mathbf{d}) = \{C_1^*, \dots, C_m^*\}$. By Lemma 11, there exists at most one coalition in $C^*(\mathbf{d})$ that is not a clique. W.l.o.g. let us assume that C_1^* is such a coalition, and let $k = |C_1^*|$. Recalling the proof of Lemma 11, it is possible to split C_1^* into two cliques K_1 and K_2 of respective sizes k_1 and k_2 , where k_1 is the number of negative relations in C_1^* , and $\text{SW}(C_1^*) = \text{SW}(K_1) + \text{SW}(K_2) + \frac{1}{n}(2k_1k_2 - k_1) - k_1$. Therefore, taking into account that $\{K_1, K_2, C_2^*, \dots, C_m^*\}$ is a clique partition of $G_{\mathbf{d}}^*$, we know that $\text{SW}(\mathcal{K}^*) \geq \text{SW}(K_1) + \text{SW}(K_2) + \sum_{\ell=2}^m \text{SW}(C_{\ell}^*)$. Thus, $\frac{\text{opt}(\mathbf{d})}{\text{SW}(\mathcal{K}^*)} \leq \frac{\text{SW}(C_1^*) + \sum_{\ell=2}^m \text{SW}(C_{\ell}^*)}{\text{SW}(K_1) + \text{SW}(K_2) + \sum_{\ell=2}^m \text{SW}(C_{\ell}^*)}$, so $\frac{\text{opt}(\mathbf{d})}{\text{SW}(\mathcal{K}^*)} \leq \frac{\text{SW}(C_1^*)}{\text{SW}(K_1) + \text{SW}(K_2)} = 1 + s$, where $s = \frac{\frac{1}{n} \cdot k_1(2k_2 - 1 - n)}{\frac{1}{n} \cdot (k_1(k_1 - 1) + k_2(k_2 - 1))}$. Let $\alpha < 1$ be the positive number such that $k_1 = \alpha k$ and $k_2 = (1 - \alpha)k$. Then, $s \leq \frac{2\alpha(1-\alpha)n - \alpha - \alpha n}{(\alpha^2 + (1-\alpha)^2)n - 1} \leq \frac{\alpha(1-2\alpha)}{2\alpha(\alpha-1)+1} \leq \frac{1}{\sqrt{2}} - \frac{1}{2}$. \square

It is possible to show that the above bound is tight. However, even if the best clique partition gives a $\frac{1}{\sqrt{2}} + \frac{1}{2}$ approximation of the optimum, no polynomial time algorithm can compute it. In fact, we now prove that the social welfare maximization problem is as hard to approximate as the MAXCLIQUE problem.

Theorem 13. *No polynomial time algorithm can approximate the optimal social welfare for EA preference profiles with an approximation ratio $O(n^{1-\varepsilon})$ for any fixed $\varepsilon > 0$, unless $P = NP$.*

Proof. For a fixed $\varepsilon > 0$, let us assume that there exists a polynomial time algorithm A that, given as input an EA preference profile \mathbf{d} , always returns a partition $\mathcal{C}^A(\mathbf{d})$ such that $\text{SW}(\mathcal{C}^A(\mathbf{d}))$ approximates the optimal social welfare $\text{opt}(\mathbf{d})$ within an approximation ratio $o(n^{1-\varepsilon})$. We now show how algorithm A can be exploited to find a good approximation for MAXCLIQUE.

By assumption, for each $c > 0$, there exists $n_0 \in \mathbb{N}$ s.t. $\frac{\text{opt}(\mathbf{d})}{\text{SW}(\mathcal{C}^A(\mathbf{d}))} < cn^{1-\varepsilon}$, $\forall n > n_0$, where $\mathbf{d} = (d_1, \dots, d_n)$.

Consider then an instance G of MAXCLIQUE and any EA profile \mathbf{d} such that $G_{\mathbf{d}}^* \simeq G$. Let k^* be the size of a maximum clique in $G_{\mathbf{d}}^*$. Since a maximum clique of $G_{\mathbf{d}}^*$ completed with singleton coalitions containing all the other agents is a possible outcome, $k^*(k^* - 1)/n \leq \text{opt}(\mathbf{d})$.

Consider now the outcome $\mathcal{C}^A(\mathbf{d}) = \{C_1^A, \dots, C_m^A\}$ returned by A . W.l.o.g. assume that no agent in $\mathcal{C}^A(\mathbf{d})$ participates in more than one negative relation in her coalition, otherwise we can put such an agent alone in a singleton coalition, increasing the overall social welfare (as in Lemma 10). Then, similarly as in Lemma 11, we can split every coalition $C_i^A \in \mathcal{C}^A(\mathbf{d})$ that contains negative relations into two cliques K_i^1 and K_i^2 . If C_i^A is already a clique, we define $K_i^1 = C_i^A$ and $K_i^2 = \emptyset$. Let \mathcal{K} be the coalition structure consisting of all the cliques K_i^1 and $K_i^2 \neq \emptyset$. By exploiting the same arguments as in the proof of Lemma 11, it is possible to show that $\text{SW}(\mathcal{C}^A(\mathbf{d})) \leq \left(\frac{1+\sqrt{2}}{2}\right) \text{SW}(\mathcal{K})$. Then,

if $k_{max} = \max_{i,j} \{|K_i^j|\}$ and $q = |\{K_i^j \mid K_i^j \neq \emptyset\}|$, where $i \in \{1, \dots, m\}$ and $j \in \{1, 2\}$, as $\text{SW}(\mathcal{K}) \leq q \cdot k_{max}(k_{max} - 1)/n$, and $k^*(k^* - 1)/n \leq \text{opt}(\mathbf{d})$, we obtain $\frac{k^*(k^* - 1)/n}{q \cdot k_{max}(k_{max} - 1)/n} \leq \frac{\text{opt}(\mathbf{d})}{\text{SW}(\mathcal{K})} \leq \left(\frac{1}{\sqrt{2}} + \frac{1}{2}\right) \frac{\text{opt}(\mathbf{d})}{\text{SW}(\mathcal{C}^A(\mathbf{d}))} < \left(\frac{1}{\sqrt{2}} + \frac{1}{2}\right) cn^{1-\varepsilon}$.

By using $q \leq n$ and that for $k_{max} \leq k^*$ it holds that $\frac{k^*}{k_{max}^2} \leq \frac{k^*(k^* - 1)}{k_{max}(k_{max} - 1)}$, we finally get $k^* < k_{max} \sqrt{c \cdot \frac{1+\sqrt{2}}{2} n^{1-\frac{\varepsilon}{2}}}$.

In other words, using algorithm A , we can extract in polynomial time from $\mathcal{C}^A(\mathbf{d})$ a clique of size k_{max} in G that approximates the optimal solution of MAXCLIQUE with an approximation ratio $o(n^{1-\frac{\varepsilon}{2}})$, which is not possible unless $P = NP$. \square

Strategyproof Mechanisms

We start this subsection by first focusing on efficient mechanisms, that is, mechanism running in polynomial time.

Mechanism \mathcal{M}_4 . Given an EA preference profile \mathbf{d} , \mathcal{M}_4

1. enumerates the agents in N from 1 up to n ;
2. sets $\mathcal{C} = \emptyset$;
3. for $i = 1$ up to n
 - if there exists $j > i$ in the neighborhood of i in $G_{\mathbf{d}}^*$ not matched yet, then $\mathcal{C} = \mathcal{C} \leftarrow \{i, j\}$,
 - otherwise, $\mathcal{C} = \mathcal{C} \leftarrow \{i\}$;
4. returns \mathcal{C} .

According to Theorem 13, the approximation factor of \mathcal{M}_4 that we prove next is asymptotically optimal.

Theorem 14. \mathcal{M}_4 is strategyproof and $r^{\mathcal{M}_4} \leq (1 + \sqrt{2}) \cdot n$.

Proof. For what concerns the strategyproofness, we observe that only unassigned agents are interested in manipulating the mechanism. Let i be an unassigned agent. By the definition of \mathcal{M}_4 , this means that all her neighbors in $G_{\mathbf{d}}^*$ have been assigned in the previous rounds. Agent i can manipulate in the following ways: 1) declare an enemy j as a friend, or, 2) declare a friend j as an enemy. In case 1), either i will still not be assigned by the mechanism or, even worse, she will be assigned to her enemy j . In case 2), the outcome does not change. Thus, \mathcal{M}_4 is strategyproof.

In order to establish the approximation ratio of \mathcal{M}_4 , we first observe that, given any EA profile \mathbf{d} , the returned coalition structure $\mathcal{C} = \{C_1, \dots, C_m\}$ forms a maximal matching in $G_{\mathbf{d}}^*$, and that, as it is well-known, such a matching consists of at least half of the edges of a maximum matching of $G_{\mathbf{d}}^*$. Moreover, given a best clique partition $\mathcal{K}^* = \{K_1^*, \dots, K_m^*\}$ for $G_{\mathbf{d}}^*$, in a maximum matching there are at least $\sum_{\ell=1}^m \lfloor \frac{k_{\ell}^*}{2} \rfloor$ edges, where $k_{\ell}^* = |K_{\ell}^*|$. Therefore, since each coalition $C_{\ell} \in \mathcal{C}$ contains a pair of opposite positive relations and thus contributes $2/n$ to the overall social welfare, $\text{SW}(\mathcal{C}) \geq \frac{2}{n} \sum_{\ell=1}^m \lfloor \frac{k_{\ell}^*}{2} \rfloor$, so that

$$\frac{\text{SW}(\mathcal{K}^*)}{\text{SW}(\mathcal{C})} \leq \frac{\frac{1}{n} \sum_{\ell=1}^m k_{\ell}^*(k_{\ell}^* - 1)}{\frac{1}{n} \sum_{\ell=1}^m \lfloor \frac{k_{\ell}^*}{2} \rfloor} \leq \frac{\frac{1}{n} k_{max}^*(k_{max}^* - 1)}{\frac{1}{n} \lfloor \frac{k_{max}^*}{2} \rfloor} \leq 2n,$$

where k_{max}^* is the size of the biggest clique in \mathcal{K}^* . By Lemma 12, $\text{opt}(\mathbf{d}) \leq \left(\frac{1+\sqrt{2}}{2}\right) \cdot \text{SW}(\mathcal{K}^*)$ and the approximation ratio follows. \square

As already observed, the approximation ratio of mechanism \mathcal{M}_4 is high due to the inherent difficulty of computing a sublinear approximation in polynomial time. However, if efficiency is not a concern, a constant non-polynomial approximation mechanism exists.

It is easy to check that the mechanism that always returns an optimal outcome is not strategyproof. In (Dimitrov and Sung 2004), the authors consider a strategyproof mechanism that iteratively extracts from $G_{\mathbf{d}}^*$ the clique of maximum size which comes first in the lexicographic order with respect to the enumeration of the agents. This mechanism has approximation ratio of at least 2, as it can be easily checked by considering the example in Figure 3. However, it can also be shown that the approximation ratio of the just described mechanism is constant.




Figure 3: Lower bound instance for the greedy mechanism that iteratively extract a maximum clique provided in previous work (Dimitrov and Sung 2004).

Conclusions and Future Work

A natural extension of our work is considering the possibility of neutrals, that is, allowing agents to also be indifferent to others. First, we observe that for FA profiles with neutrals, all of our mechanisms maintain the same approximation ratio. In fact, it is easy to see that, once the positive relations are fixed, the worst cases occur when the remaining relations are negative. For EA with neutrals, our mechanisms have an unbounded ratio, but an $O(n^2)$ approximation can be determined mimicking the mechanism of (Flammini, Monaco, and Zhang 2017) for *duplex* valuations $-1, 0$ or 1 . In fact, such a mechanism matches just one pair of agents i and j such that i considers j to be a friend and j consider i either to be a friend or a neutral. Providing a mechanism with a better approximation ratio would also yield a better result

for their setting. A further extension in this direction might be neutrals having a negligible yet non-null positive or negative effect on the utilities.

Anonymity is another interesting feature of strategyproof mechanisms. Roughly speaking, a mechanism is anonymous if it does not rely on the agents' identities. While for FA our mechanisms are anonymous, for EA they are not. In fact, for EA it is not possible to have an anonymous deterministic mechanism, as it can be checked considering an instance with 3 agents where $(1, 3)$, $(3, 1)$ are the negative relations, and all others are positive. In this example, only the outcomes $\{\{1, 2\}, \{3\}\}$ and $\{\{1\}, \{2, 3\}\}$ achieve a positive social welfare. However, both of these outcomes cannot be consistently reached by an anonymous mechanism. Thus, under EA a deterministic anonymous mechanism cannot be admissible. Anonymity for randomized mechanism remains as an interesting open question.

A natural open problem is reducing our approximation gaps, especially in the FA deterministic case. This appears to be a challenging task, and bears similarities with analogous gaps in (Flammini, Monaco, and Zhang 2017).

Acknowledgments

This work was partially supported by the Italian MIUR PRIN 2017 Project ALGADIMAR “Algorithms, Games, and Digital Markets”.

References

- Aziz, H., and Savani, R. 2016. Hedonic games. In *Handbook of Computational Social Choice*. 356–376.
- Aziz, H.; Brandl, F.; Brandt, F.; Harrenstein, P.; Olsen, M.; and Peters, D. 2019. Fractional hedonic games. *ACM Trans. Economics and Comput.* 7(2):6:1–6:29.
- Aziz, H.; Brandt, F.; and Harrenstein, P. 2013. Pareto optimality in coalition formation. *Games and Economic Behavior* 82:562–581.
- Aziz, H.; Brandt, F.; and Seedig, H. G. 2011. Stable partitions in additively separable hedonic games. In *10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2011), Taipei, Taiwan, May 2-6, 2011, Volume 1-3*, 183–190.
- Aziz, H.; Brandt, F.; and Seedig, H. G. 2013. Computing desirable partitions in additively separable hedonic games. *Artif. Intell.* 195:316–334.
- Banerjee, S.; Konishi, H.; and Sönmez, T. 2001. Core in a simple coalition formation game. *Social Choice and Welfare* 18(1):135–153.
- Barrot, N.; Ota, K.; Sakurai, Y.; and Yokoo, M. 2019. Unknown agents in friends oriented hedonic games: Stability and complexity. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*, 1756–1763.
- Bilò, V.; Fanelli, A.; Flammini, M.; Monaco, G.; and Moscardelli, L. 2018. Nash stable outcomes in fractional hedonic games: Existence, efficiency and computation. *J. Artif. Intell. Res.* 62:315–371.
- Bloch, F., and Diamantoudi, E. 2011. Noncooperative formation of coalitions in hedonic games. *Int. J. Game Theory* 40(2):263–280.
- Bogomolnaia, A., and Jackson, M. O. 2002. The stability of hedonic coalition structures. *Games and Economic Behavior* 38(2):201–230.
- Dimitrov, D., and Sung, S. C. 2004. Enemies and friends in hedonic games: individual deviations, stability and manipulation. *CentER Discussion Paper Series*.
- Dimitrov, D.; Borm, P.; Hendrickx, R.; and Sung, S. C. 2006. Simple priorities and core stability in hedonic games. *Social Choice and Welfare* 26(2):421–433.
- Dreze, J. H., and Greenberg, J. 1980. Hedonic coalitions: Optimality and stability. *Econometrica (pre-1986)* 48(4):987.
- Elkind, E., and Wooldridge, M. 2009. Hedonic coalition nets. In *8th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2009), Budapest, Hungary, May 10-15, 2009, Volume 1*, 417–424.
- Elkind, E.; Fanelli, A.; and Flammini, M. 2016. Price of pareto optimality in hedonic games. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA*, 475–481.
- Feldman, M.; Lewin-Eytan, L.; and Naor, J. 2015. Hedonic clustering games. *TOPC* 2(1):4:1–4:48.
- Flammini, M.; Monaco, G.; and Zhang, Q. 2017. Strategyproof mechanisms for additively separable hedonic games and fractional hedonic games. In *Approximation and Online Algorithms - 15th International Workshop, WAOA 2017, Vienna, Austria, September 7-8*, 301–316.
- Gairing, M., and Savani, R. 2019. Computing stable outcomes in symmetric additively separable hedonic games. volume 44, 1101–1121.
- Igarashi, A., and Elkind, E. 2016. Hedonic games with graph-restricted communication. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems, Singapore, May 9-13, 2016*, 242–250.
- Nisan, N.; Roughgarden, T.; Tardos, É.; and Vazirani, V. V., eds. 2007. *Algorithmic Game Theory*. Cambridge University Press.
- Ohta, K.; Barrot, N.; Ismaili, A.; Sakurai, Y.; and Yokoo, M. 2017. Core stability in hedonic games among friends and enemies: Impact of neutrals. In *IJCAI*, 359–365.
- Procaccia, A. D., and Tennenholtz, M. 2013. Approximate mechanism design without money. *ACM Trans. Economics and Comput.* 1(4):18:1–18:26.
- Rodríguez-Álvarez, C. 2009. Strategy-proof coalition formation. *Int. J. Game Theory* 38(3):431–452.
- Sung, S. C., and Dimitrov, D. 2007. On core membership testing for hedonic coalition formation games. *Oper. Res. Lett.* 35(2):155–158.
- Wright, M., and Vorobeychik, Y. 2015. Mechanism design for team formation. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, January 25-30, 2015, Austin, Texas, USA*, 1050–1056.