

LP-based Covering Games with Low Price of Anarchy

Georgios Piliouras*

Tomáš Valla[†]

László A. Végh[‡]

March 2, 2012

Abstract

We present a new class of vertex cover and set cover games. The price of anarchy bounds match the best known constant factor approximation guarantees for the centralized optimization problems for linear and also for submodular costs – in contrast to all previously studied covering games, where the price of anarchy cannot be bounded by a constant (e.g. [6, 7, 11, 5, 2]). In particular, we describe a vertex cover game with a price of anarchy of 2. The rules of the games capture the structure of the linear programming relaxations of the underlying optimization problems, and our bounds are established by analyzing these relaxations. Furthermore, for linear costs we exhibit linear time best response dynamics that converge to these almost optimal Nash equilibria. These dynamics mimic the classical greedy approximation algorithm of Bar-Yehuda and Even [3].

1 Introduction

Combinatorial optimization has for several decades dictated the landscape of algorithm design. The extent of its impact can be appreciated by the fact that almost by default the main judging criterion of a polytime algorithmic solution is the approximation guarantee it offers, regardless of other parameters that may affect the applicability of the solution in practice (simplicity of implementation, robustness to input errors, etc.)

One such limiting assumption is the existence of an omnipotent centralized authority that has access to all the relevant information and has the power to enforce any solution of its choice. Over the last decade, the soundness of such assumptions has increasingly come into question following a number of paradigm-shifting socioeconomic events such as the rapid rise of the Internet, the painful realization of the extent of inter-connectivity of the global economy as well as the emergence of global sustainability concerns.

Algorithmic game theory strives for global optimization in such decentralized settings that consist of self-interested individuals. In these more challenging scenarios, tractability can be compromised along two largely independent axes: due to individual incentive issues or due to computability issues.

Price of Anarchy. The competition between individual incentives and social optimality is of fundamental concern in distributed systems as it can lead to highly inefficient outcomes. The price of anarchy literature [26] examines exactly what are the worst case repercussions of such a policy. Formally, price of anarchy is defined as the maximal ratio between the social cost of a Nash equilibrium and that of the global optimal configuration. Intuitively, a low price of anarchy implies that upon converging to a socially stable outcome, the quality of the acquired solution is almost optimal from a central optimization perspective.

Unfortunately, in many cases of interesting games the price of anarchy is prohibitively high. Vertex cover, due to its prominent position within combinatorial optimization, has been studied in the context of game theory from different approaches, all of which so far have shared this limiting characteristic.

Specifically, Cardinal and Hoefer in [6] define a vertex cover game where the edges of a network are owned by k agents. An agent's goal is to have each of his edges supplied by a service point at least one of its endpoints. There is a cost $c(v) \geq 0$ associated to building a service point at vertex v . The strategy of

*School of Electrical & Computer Engineering, Georgia Institute of Technology, Atlanta, GA. Supported by AFOSR projects FA9550-09-1-0538. georgios.piliouras@ece.gatech.edu

[†]Institute for Theoretical Computer Science, Faculty of Mathematics & Physics, Charles University, Prague, Czech Republic. Supported by the GAUK Project 66010 and by ITI, under grant P202/12/G061. valla@kam.mff.cuni.cz

[‡]College of Computing, Georgia Institute of Technology. Supported by NSF Grant CCF-0914732. lvegh@cc.gatech.edu

an agent is a vector consisting of offers to the vertices. Service points will be installed at vertices where the total offer exceeds the cost of the vertex. Similar games are defined by Buchbinder et al. [5] and by Escoffier et al. [11] for the more general set cover problem.

A different approach was followed by Balcan et al. [2]. Here the agents are the vertices of the graph, and their strategies are deciding whether they open a service point. If opening a service point, vertex v incurs a cost $c(v)$. If he decides not to open, he has to pay a penalty for all edges incident to v whose other endpoints are uncovered.

The price of anarchy is $\Theta(k)$ in [6] and $\Theta(n)$ in [2]. Indeed, if the underlying network is a star, and each edge is owned by a different agent in the first case, we get Nash equilibria with all leaves being service points. These guarantees are significantly worse than the ones available in the centralized setting, where simple factor 2-approximation algorithms exist.

In contrast, in our paper, we shall present a simple vertex cover game with a price of anarchy 2. As in [2], the agents are the vertices, and the regulations delegate the responsibility of covering every edge of the network to its two endpoints: both incur a high penalty if the edge is left uncovered. The difference from the setting of [2] is that those who open a service point can demand compensation from their neighbors. This is justified since if u opens a service point, every neighbor v benefits from this as the common responsibility of covering uv is taken over by u .

In the description, we use intuitive terminology of a Mafia (service points) which provides “security” (covers edges). The vertices may choose to join Mafia or to remain civilians. Each edge of the graph has to be “secured”, that is, at least one endpoint must be in Mafia. For agent v , there is an initial cost $c(v)$ to join Mafia. Mafiosi can collect ransoms as the price of security of the incident edges: if a vertex v chooses to be a mafioso, his strategy also includes a ransom vector, so that the total ransom he demands from his neighbors is $c(v)$. It is a one-shot game and mafiosi can ransom both their civilian and mafioso neighbors.

If v is a civilian, he has to pay to his neighbors in the Mafia all ransom they demand. Furthermore, if there is an incident uncovered edge uv , that is, u is also a civilian, both of them have to pay a huge penalty. In contrast, if v is a mafioso, he has to pay $c(v)$ for joining, and he receives whatever he can collect from ransoms. However, mafiosi ransomed excessively obtain a protected status: if the total demand from v is more than $c(v)$, he satisfies only a proportional fraction of the demands. It is important to note that the payoff function is defined locally: besides his own strategy, the payoff of an agent depends only on the strategies of agents at distance at most 2 from him (i.e. immediate neighbors and neighbors of neighbors). Also note that if M is a vertex cover, then the total utility of the agents is $-c(M)$. Consequently, an optimal solution to the optimization problem gives a social optimum of the game.

Our approach avoids bad Nash equilibria that are possible in [6] and [2]. As an example, consider the vertex cover game on a star with all vertices having cost 1. In the models of [6] and [2] there exists a Nash equilibrium where the leaves form the vertex cover. In our model, if all leaves are mafiosi, then all of them would demand ransom from the central agent, who would then have a strong incentive to join the mafia and obtain the protected status. It can be verified that the only Nash equilibria correspond to outcomes where the central vertex and at most one leaf are in the Mafia.

As a different interpretation of the game above, consider a road network with the vertices representing cities. The maintenance of a road must be provided by a facility at one of the endpoints. The cost of opening the facility dominates the operating cost: if city v decides to open one at cost $c(v)$, it is able to maintain all incident roads. As a compensation, the cities can try to recollect the opening cost by asking contributions from the neighboring cities. A city without a facility has to pay all contributions he is asked to pay. However, if a city opens a facility, its liability is limited and has to satisfy demands only up to his opening cost, $c(v)$.

Our approach can be extended to the hitting set problem, which is equivalent to the set cover problem. We are given a hypergraph $G = (V, \mathcal{E})$, and a cost function $c : V \rightarrow \mathbb{R}_+$ on the vertices. Our aim is to find a minimum cost subset M of V intersecting every hyperedge in \mathcal{E} . This problem is known to be approximable within a factor of d , the maximum size of a hyperedge. In the corresponding Mafia game, the hyperedges shall be considered as clubs in need of security. A mafioso can assign ransoms to the clubs he is a member of, that will be distributed equally to all other members of the club.

We shall prove that for the vertex cover and hitting set games, the price of anarchy is 2 and d , respectively. Bar Yehuda and Even gave a simple primal-dual algorithm with this guarantee in 1981 [3]. No better constant factor approximation has been given eversince. Furthermore, assuming the Unique Games Conjecture, Khot and Regev [19] proved that the hitting set problem cannot be approximated by any constant factor smaller

than d .

As a further extension, we also investigate the submodular hitting set (or set cover) problem, that has received significant attention recently. The goal is to find a hitting set M of a hypergraph minimizing $C(M)$ for a submodular set function C on the ground set. Independently, Koufogiannakis and Young [25] and Iwata and Nagano [18] gave d -approximation algorithms. Our game approach extends even to this setting, with the same price of anarchy d . This involves a new agent, the Godfather, who’s strategy consists of setting a budget vector in the submodular base polyhedron of C . Otherwise, the game is essentially the same as the (linear) hitting set game.

The main results of the paper can be summarized as follows.

Theorem 1.1. *The Mafia games for vertex cover, hitting set and submodular hitting set always have pure Nash equilibria, and the price of anarchy is 2 for vertex cover and d for (submodular) hitting set.*

Recent work of Roughgarden et al. [30, 4, 31] has shown that the majority of positive results in price of anarchy literature can be reduced to a specific common set of structural assumptions. In contrast, in our work, we use a novel approach by exploring connections to the LP relaxations of the underlying centralized optimization problems. This connection raises interesting questions about the limits of its applicability.

Convergence and Complexity of Dynamics. The world of decentralized competition is not immune to the results of computational complexity. Hence, a low price of anarchy although promising does not necessarily yield a usable outcome in the means of the game dynamics, when agents sequentially have the possibility to change their strategies for a better one. The reasons for these inconsistencies fall in one of two possible categories: either non-convergence of the dynamics to a Nash equilibrium or too slow convergence.

Even in a very simple game settings, with a constant number of agents and strategies and no computational complexity issues, it could be the case that games exhibit only highly unstable Nash equilibria. In such settings, numerous learning dynamics, even if they start off from a state close to a Nash equilibrium, they diverge away from it fast [8, 21].

On the other hand, as the games grow in size, even when there exist simple decentralized dynamics which provably converge to Nash equilibria, it is not necessarily the case that this convergence is achieved within polynomial time. For example, in the case of general congestion games, although best response dynamics always converge to a Nash equilibrium, finding any sample equilibrium (even via a centralized algorithm) has been shown to be PLS-hard [12], implying that any decentralized dynamic is bound to fail as well in worst case instances.

In our covering games, we first show that even in simple instances, round robin best response dynamics¹ may end in a loop. However, this can be simply fixed by a slight modification of the payoff. We introduce a secondary utility, that does not affect the price of anarchy results, but merely instigates the mafiosi to use more fair (symmetric) ransoms: $r(u, v) = r(v, u)$. With this secondary objective, we show that actually a single round of best response dynamics under a simple selection rule of the next agent results in a Nash-equilibrium. This dynamics in fact simulates the Bar-Yehuda–Even algorithm. An analogous dynamics is shown in the case of hitting set. Moreover, these dynamics can be interpreted in a distributed manner, enabling several agents to change their strategies at the same time.

In our games, the set of strategies is infinite as ransoms can be arbitrary real numbers. However, if the vertex weights are integers, we can restrict possible ransoms to be integers as well. All results of the paper straightforwardly extend to this finite game.

1.1 Related work

The basic set cover games in [5], [11] and [2] fall into the class of congestion games [29]. In the models of [5], [11], in the hitting set terminology, the agents are the hyperedges that choose a vertex to cover them, and the cost of the vertex is divided among them according to some rule. [5] investigates the influence of a central authority that can influence choices by taxes and subsidies in a best response dynamics; [11] studies different cost sharing rules of the vertices (“local taxes”). However, none of these methods achieve a constant price of anarchy. The model of [2] can achieve a good equilibrium by assuming a central authority that propagates

¹These are the dynamics where each agent takes turn playing his best response in a cyclic ordering according to some fixed permutation.

information on an optimal solution to a fraction of the agents. In contrast to [5] and [2], our model is defined locally, without assuming a central authority.

Cardinal and Hoefer [7] define a general class of covering games, including the vertex cover game [6], and also the selfish network design game by Anshelevich et al. [1]. The game is based on a covering problem given by a linear integer program. Variables represent resources, and the agents correspond to certain sets of constraints they have to satisfy. An agent can offer money for resources needed to satisfy her constraints. From each variable, the number of units covered by the total offers of the agents will be purchased and can be used by all agents simultaneously to satisfy their constraints, regardless to their actual contributions to the resource.

In the vertex cover or hitting set game, the resources are the service points and the set of constraints belonging to the agents express that every (hyper)edge owned by them has to be covered. In the model of [1], agent i wants to connect a set of terminals S_i in a graph $G = (V, E)$ with edge costs c . Hence the variables represent the edges of the graph and the constraints belonging to agent i enforce the connectivity of S_i .

Our games can be seen as the *duals* of these coverings games. That is, the agents correspond to the variables, and are responsible for the satisfaction of the constraints containing them. If a constraint is left unsatisfied, the participating variables get punished. Also, a variable may require compensation (ransoms) from other variables participating in the same constraints. These compensations will correspond to a dual solution in a Nash equilibrium. We hope that our approach of studying dual covering games might be extended to a broader class of problems, with the price of anarchy matching the integrality gap.

Our result and the above papers are focused on noncooperative covering games. A different line of game theoretic study is focused on cost sharing mechanism, e.g. [9, 10, 17, 13, 27, 28].

The performance of behavioral dynamics in games and specifically establishing fast convergence to equilibria of good quality has been the subject of intensive recent research [22, 23, 32]. The importance of such results that go beyond the analysis of performance of Nash equilibria has been stressed in [21] where it has been shown that even in very simple games with constant number of agents and strategies, the performance of simple learning dynamics can be arbitrarily different than (any convex combination of) the payoffs of Nash equilibria.

The rest of the paper is organized as follows. Section 2 defines the Mafia games for vertex cover, hitting set, and submodular hitting set, and proves the existence of Nash equilibria and gives price of anarchy bounds. Section 3 shows that certain simple dynamics rapidly converge to Nash equilibrium for vertex cover and for hitting set. Section 4 discusses possible further research directions.

2 The Mafia games and Price of Anarchy bounds

2.1 Vertex cover

Given a graph $G = (V, E)$, let $c : V \rightarrow \mathbb{R}^+$ be a cost function on the vertices. In the *vertex cover problem*, the task is to find a minimum cost set $M \subseteq V$ containing at least one endpoint of every edge in E . For a vertex $v \in V$, let $N(v) = \{u : uv \in E\}$ denote the set of its neighbors.

Game definition. The *Mafia Vertex Cover Game* is a one-shot game on the agent set V . The basic strategy of an agent is to decide being a civilian or a mafioso. The set of civilians shall be denoted by C , the set of mafiosi (Mafia) by M . For civilians, no further decision has to be made, while for mafiosi, their strategy also contains a ransom vector. Each mafioso $m \in M$ can demand ransoms from his neighbors totaling $c(m)$. The ransom demanded from a neighbor $u \in N(m)$ is $r(m, u) \geq 0$, with $\sum_{u \in N(m)} r(m, u) = c(m)$. The strategy profile $\mathcal{S} = (M, C, r)$ thus consists of the sets of mafiosi and civilians, and the ransom vectors.

Let us call $c(v)$ the *budget* of an agent $v \in V$, and let $T > \sum_{v \in V} c(v)$ be a huge constant. Let $D(v) = \sum_{m \in M} r(m, v)$ be the demand asked from the agent $v \in V$.

Let us now define the payoffs for a given strategy profile \mathcal{S} . For a civilian $v \in C$, let $\text{Pen}(v) = T$ if v is incident to an uncovered edge, that is $C \cap N(v) \neq \emptyset$, and $\text{Pen}(v) = 0$ otherwise. The utility of $v \in C$ is

$$U_{\mathcal{S}}(v) = -D(v) - \text{Pen}(v).$$

If $v \in M$ and the total demand from v is $D(v) > c(v)$ (i.e. v is asked too much), we call v *protected* and denote the set of protected mafiosi by $P \subseteq M$. The real amount of money that the protected mafioso $p \in P$ pays to his neighbors is scaled down to $\frac{c(p)}{D(p)}r(u, p)$. Let $F^-(v) = \min\{D(v), c(v)\}$ be the total amount the mafioso v pays for ransom. Let

$$F^+(v) = \sum_{u \in N(v) \setminus P} r(v, u) + \sum_{u \in N(v) \cap P} \frac{c(u)}{D(u)} r(v, u)$$

denote the income of $v \in M$ from the ransoms. Then the utility of a mafioso $v \in M$ is defined as

$$U_S(v) = -c(v) + F^+(v) - F^-(v).$$

This means v has his initial cost $c(v)$ for entering the Mafia, receives full payment from civilians and unprotected mafiosi, receives reduced payment from protected mafiosi, and pays the full payment to his neighboring mafiosi if v is unprotected, or reduced payment if v is protected.

The existence of pure Nash equilibria. Pure Nash equilibria are (deterministic) strategy outcomes such that no agent can improve her payoff by unilaterally changing her strategy. We will start by establishing that our game always exhibits such states. The following is the standard linear programming relaxation of vertex cover along with its dual.

$$\begin{array}{ll} \min \sum_{v \in V} c(v)x(v) & \text{(P-VC)} \\ x(u) + x(v) \geq 1 \quad \forall uv \in E \\ x \geq 0 \end{array} \qquad \begin{array}{ll} \max \sum_{uv \in E} y(uv) & \text{(D-VC)} \\ \sum_{uv \in E} y(uv) \leq c(u) \quad \forall u \in V \\ y \geq 0 \end{array}$$

For a feasible dual solution y we say that the vertex $v \in V$ is *tight* if $\sum_{uv \in E} y(uv) = c(v)$. We call the pair (M, y) a *complementary pair* if M is a vertex cover, y is a feasible dual solution, and each $v \in M$ is tight with respect to y .

Lemma 2.1. *If (M, y) is a complementary pair, then M is a 2-approximate solution to the vertex cover problem.*

Proof. The primal objective is at most twice the dual objective, as

$$\sum_{v \in M} c(v) = \sum_{v \in M} \sum_{u \in N(v)} y(uv) \leq 2 \sum_{uv \in E} y(uv).$$

The inequality follows as each edge uv is counted at most twice. □

We shall show that the simple approximation algorithm by Bar-Yehuda and Even [3] returns a complementary pair, and therefore has approximation factor 2. Our next lemma proves that a complementary pair provides a Nash equilibrium.

Lemma 2.2. *Let (M, y) be a complementary pair, and consider the strategy profile where the agents in M form the Mafia and $C = V \setminus M$ are the civilians. For $u \in M$, define $r(u, v) = y(uv)$ for every $v \in N(u)$. Then the strategy profile $\mathcal{S} = (M, C, r)$ is a Nash equilibrium.*

Proof. Since $D(v) \leq c(v)$ for all players, there are no protected mafiosi. If v is a civilian, his payoff is $-D(v)$. He would not get a protected status if he entered the Mafia as $D(v) \leq c(v)$, and thus his payoff would be $-c(v) + F^+(v) - D(v) \leq -D(v)$ by arbitrary choice of ransoms. If v is a mafioso, he has $F^+(v) = c(v)$ as none of his neighbors is protected. Thus his utility is $-D(v)$, the maximum he can obtain for any strategy. □

The existence of a complementary pair is provided by the algorithm of Bar-Yehuda and Even [3]. In each step of the algorithm we maintain a feasible dual solution, and M will be the set of tight vertices.

(0) Set $y(uv) := 0$ for each $uv \in E$ and $M = \{v \in V : c(v) = 0\}$.

- (1) While M is not a vertex cover do
 - (1-1) Choose an arbitrary edge $uv \in E$ with $u, v \in V - M$.
 - (1-2) Raise $y(uv)$ until u or v becomes tight.
 - (1-3) Include the new tight endpoint(s) into M .
- (2) Return M .

It is straightforward that the algorithm returns a complementary pair (M, y) . Using Lemma 2.2, we obtain the following.

Theorem 2.3. *The Mafia Vertex Cover Game always has a pure Nash equilibrium.* □

The Price of Anarchy. For a strategy profile \mathcal{S} with α uncovered edges, the sum of the utilities is $-c(M) - 2\alpha T$. The Price of Anarchy compares this sum in a Nash equilibrium at the worst case to the maximum value over all strategy profiles, that corresponds to a minimum cost vertex cover.

Consider a strategy profile \mathcal{S} that encodes a Nash equilibrium. First, observe that Mafia M is a vertex cover. Indeed, if there were an uncovered edge $uv \in E$, both u and v would receive the high penalty T , and therefore they would have incentive to join Mafia. We shall prove that the cost $c(M)$ is at most twice the cost of an optimal vertex cover, consequently, the price of anarchy is at most 2.

Lemma 2.4. *Let the strategy profile $\mathcal{S} = (M, C, r)$ be a Nash equilibrium. Then there are no protected mafiosi.*

Proof. For a contradiction, suppose P is nonempty. First we show there exists an edge $mp \in E$ such that $m \in M \setminus P$, $p \in P$ and $r(m, p) > 0$. Indeed, if there were no such edges, then $\sum_{p \in P} D(p) \leq \sum_{p \in P} c(p)$ as the ransoms demanded from protected mafiosi are all demanded by others P . However, by definition $D(p) > c(p)$ for all $p \in P$, giving $\sum_{p \in P} D(p) > \sum_{p \in P} c(p)$, a contradiction.

Consider the edge $mp \in E$ as above. We claim that m could choose a better strategy, and therefore \mathcal{S} cannot be a Nash equilibrium. If m does not have any civilian neighbors, that is, $N(m) \subseteq M$, then his utility would strictly increase if decides to become a civilian. Indeed, his income now is $F^+(m) < c(m)$ and he has to pay $F^-(m) = D(m) \leq c(m)$. As a civilian, his utility were $-D(m)$.

Next, assume there exists a $v \in C$, $mv \in E$. Then m may decrease $r(m, p)$ to 0 and increase $r(m, v)$ by the same amount. Again, this would be a better strategy for m , as v pays the full amount whereas p payed only a reduced amount. □

Lemma 2.5. *Suppose the strategy profile $\mathcal{S} = (M, C, r)$ is a Nash equilibrium and let $v \in C$. Then $D(v) \leq 2c(v)$.*

Proof. Suppose the contrary: let $D(v) > 2c(v)$ and thus $U_{\mathcal{S}}(v) < -2c(v)$. If joining Mafia, v receives the protected status and thus gains utility at least $-2c(v)$ as $F^-(v) = c(v)$. □

Theorem 2.6. *The price of anarchy in the Mafia game is 2.*

Proof. Let $\mathcal{S} = (M, C, r)$ be a strategy profile in a Nash equilibrium. Using the convention $r(u, v) = 0$ if $u \in C$, let us define $y(uv) = r(u, v) + r(v, u)$ for every edge $uv \in E$. We show that $\sum_{u \in V} y(uv) \leq 2c(v)$ for every $v \in V$. Indeed, if $v \in C$, then $\sum_{u \in V} y(uv) = \sum_{u \in M} r(u, v) = D(v) \leq 2c(v)$ by Lemma 2.5. If $v \in M$, then $\sum_{u \in V} y(uv) = \sum_{u \in N(v)} r(v, u) + D(v) \leq 2c(v)$ by Lemma 2.4. Therefore $\frac{1}{2}y$ is a feasible solution to (D-VC) and

$$\sum_{uv \in E} \frac{1}{2}y(uv) = \frac{1}{2} \sum_{m \in M} \sum_{v \in V} r(m, v) = \frac{1}{2} \sum_{m \in M} c(m).$$

This verifies that the objective value for $\frac{1}{2}y$ is the half of the cost of the primal feasible vertex cover M , proving that M is a 2-approximate vertex cover. □

2.2 Set cover and hitting set

In this section, we generalize our approach to the hitting set problem. Given a hypergraph $\mathcal{G} = (V, \mathcal{E})$ and a cost function $c : V \rightarrow \mathbb{R}_+$, we want to find a minimum cost $M \subseteq V$ intersecting every hyperedge. Let $d = \max\{|S| : S \in \mathcal{E}\}$.

In the *set cover problem*, we have a ground set U and a collection of subsets \mathcal{S} of U . For a cost function $c : \mathcal{S} \rightarrow \mathbb{R}_+$ we want to find minimum cost collection of subsets whose union is U . This is equivalent to the hitting set problem, where the ground set is \mathcal{S} , and to each $u \in U$, there is a corresponding hyperedge that is the collection of subsets containing u .

For simplicity, we define the hitting set game on a d -uniform hypergraph. This can be done without loss of generality. To verify this, take an arbitrary instance $\mathcal{G} = (V, \mathcal{E})$, and let $T > d \sum_{v \in V} c(v)$. Extend V by $d - 1$ new vertices of cost T , and for every $S \in \mathcal{E}$, extend S by any $d - |S|$ new elements. If there is a d -approximate solution to the modified instance, it cannot contain any of the new elements. Hence finding a d -approximate solution is equivalent in the original and in the modified instance.

Game definition. We define the *Mafia Hitting Set Game* on a d -uniform hypergraph $\mathcal{G} = (V, \mathcal{E})$. The set of agents is V , with $v \in V$ having a *budget* $c(v)$. We shall call the hyperedges *clubs*. For an agent $v \in V$, let $\mathcal{N}(v) \subseteq \mathcal{E}$ denote the set of clubs containing v . The agents again choose from the strategy of being a civilian or being a mafioso, denoting their sets by C and M , respectively. The strategies of the mafioso m incorporates the ransoms $r(m, S)$ for the clubs S containing m , with $\sum_{S \in \mathcal{N}(v)} r(m, S) = c(m)$.

We define the payoffs for the strategy profile $\mathcal{S} = (M, C, r)$ similarly to the vertex cover case. For a civilian $v \in C$, $\text{Pen}(v) = T$ for a large constant T if v participates in a club containing no mafiosi, and 0 otherwise.

In each club S , the ransom $r(m, S)$ of a mafioso $m \in S \cap M$ has to be paid by all other members at equal rate, that is, everyone pays $\frac{r(m, S)}{d-1}$ to m . The demand from an agent is the total amount he has to pay in all clubs he is a member of, that is,

$$D(v) = \frac{1}{d-1} \sum_{S \in \mathcal{N}(v)} \sum_{m \in M \cap S} r(m, S).$$

The utility of a civilian $v \in C$ is defined as $U_{\mathcal{S}}(v) = -D(v) - \text{Pen}(v)$.

A mafioso v receives the protected status if $D(v) > c(v)$. The set of protected mafiosi is denoted by P , and they pay proportionally reduced ransoms. Let $F^-(v) = \min\{D(v), c(v)\}$ be the total amount v pays. The income is defined by

$$F^+(v) = \sum_{S \in \mathcal{N}(v)} \frac{r(v, S)}{d-1} \left(|S \setminus (P \cup \{v\})| + \sum_{u \in (S \cap P) \setminus \{v\}} \frac{c(u)}{D(u)} \right).$$

The utility of a mafioso $v \in M$ is then $U_{\mathcal{S}}(v) = -c(v) + F^+(v) - F^-(v)$.

The existence of pure Nash equilibria. The standard LP-relaxation extends the formulations (P-VC) and (D-VC).

$$\begin{array}{ll} \min \sum_{v \in V} c(v)x(v) & \text{(P-HS)} \\ \sum_{u \in S} x(u) \geq 1 \quad \forall S \in \mathcal{E} & \\ x \geq 0 & \end{array} \qquad \begin{array}{ll} \max \sum_{S \in \mathcal{E}} y(S) & \text{(D-HS)} \\ \sum_{S \in \mathcal{N}(u)} y(S) \leq c(u) \quad \forall u \in V & \\ y \geq 0 & \end{array}$$

Again, for a feasible dual solution y , $v \in V$ is called tight if the corresponding inequality in (D-HS) holds with equality. A pair (M, y) of a hitting set M and a feasible dual y is called a *complementary* pair if the dual inequality corresponding to any $v \in M$ is tight. The following simple claim generalizes Lemma 2.1.

Lemma 2.7. *If (M, y) is a complementary pair, then M is a d -approximate solution to the hitting set problem. \square*

The algorithm of Bar-Yehuda and Even [3], outlined in Section 2.1 naturally extends to the hitting set problem, and delivers a complementary pair.

Lemma 2.8. *Let us define strategies in the Mafia Hitting Set Game based on a complementary pair (M, y) as follows. Let agents in M be the Mafia and $V \setminus M$ be the civilians. For each $v \in M$, define $r(v, S) = y(S)$ for every $S \in \mathcal{E}$ containing v . Then the strategy profile $\mathcal{S} = (M, C, r)$ is a Nash equilibrium.*

Proof. For each $v \in V$, $D(v) \leq \frac{1}{d-1} \sum_{S \in \mathcal{N}(v)} y(S) |(S \cap M) \setminus \{v\}| \leq c(v)$ and therefore there are no protected mafiosi. The proof that nobody has an incentive to change his strategy is the same as for Lemma 2.2. \square

As the algorithm of Bar-Yehuda and Even [3] provides a complementary pair, this immediately yields the following.

Theorem 2.9. *The Mafia Hitting Set Game always has a pure Nash equilibrium.*

The Price of Anarchy.

Lemma 2.10. *Let the strategy profile $\mathcal{S} = (M, C, r)$ be a Nash equilibrium. Then there are no protected mafiosi.*

Proof. The proof follows the same lines as for Lemma 2.4. For a contradiction, assume $P \neq \emptyset$. First, it is easy to show that there exists an unprotected $m \in M \setminus P$ and $S \in \mathcal{E}$, $S \cap P \neq \emptyset$, such that $r(m, S) > 0$ by comparing the total in-demand and out-demand of protected mafiosi. For such m , if there exists no set $S' \in \mathcal{E}$ with $S' \cap M = \{m\}$, then he could increase his utility by leaving the Mafia. Otherwise, he could increase his utility by decreasing $r(m, S)$ and increasing $r(m, S')$. \square

The following lemma is the analogue of Lemma 2.5, yet the proof is more complicated.

Lemma 2.11. *Let the strategy profile $\mathcal{S} = (M, C, r)$ be a Nash equilibrium and let $v \in C$. Then $D(v) \leq \frac{d}{d-1}c(v)$.*

Proof. Suppose the contrary, let there be a $v \in C$ such that $D(v) > \frac{d}{d-1}c(v)$. His current utility is $U_{\mathcal{S}}(v) = -D(v)$.

We show that v could join Mafia and set ransoms that provide him a strictly larger utility. If $F_{\mathcal{S}'}^+(v)$ is the income for such a strategy profile \mathcal{S}' , then $U_{\mathcal{S}'}(v) = F_{\mathcal{S}'}^+(v) - 2c(v)$, as he would obtain the protected status. To get $U_{\mathcal{S}'}(v) > U_{\mathcal{S}}(v)$, we need to ensure $F_{\mathcal{S}'}^+(v) > 2c(v) - D(v)$. As $D(v) > \frac{d}{d-1}c(v)$ is assumed, it suffices to give an \mathcal{S}' with

$$F_{\mathcal{S}'}^+(v) \geq \frac{d-2}{d-1}c(v). \quad (1)$$

We define the ransoms $r'(v, S)$ by “stealing” the strategies of the other mafiosi. That is, for each club $S \in \mathcal{N}(v)$, $\frac{1}{d-1} \sum_{m \in M \cap S} r(m, S)$ is the total ransom v has to pay to the members of this club. We define $r'(v, S)$ proportionally to this amount:

$$r'(v, S) := \frac{c(v)}{(d-1)D(v)} \sum_{m \in M \cap S} r(m, S).$$

By Lemma 2.10, we know that there are no protected mafiosi in the Nash equilibrium \mathcal{S} . We show that after v enters Mafia, even if some of the old mafiosi become protected, they are only slightly overcharged. More precisely, we shall show that

$$D'(t) \leq \frac{d}{d-1}c(t) \quad \forall t \in M. \quad (2)$$

From this bound, (1) immediately follows. Indeed, everybody will pay at least $\frac{d}{d-1}$ fraction of the demands, and therefore $F_{\mathcal{S}'}^+(v) \geq \frac{d-1}{d}c(v) \geq \frac{d-2}{d-1}c(v)$.

It is left to prove (2). The demand of v from some $t \in M$ can be bounded as follows:

$$\begin{aligned} \frac{1}{d-1} \sum_{S \in \mathcal{N}(v) \cap \mathcal{N}(t)} r'(v, S) &= \frac{c(v)}{(d-1)^2 D(v)} \sum_{S \in \mathcal{N}(v) \cap \mathcal{N}(t)} \left(r(t, S) + \sum_{t' \in (S \cap M) \setminus \{t\}} r(t', S) \right) \\ &< \frac{1}{d(d-1)} (c(t) + (d-1)D(t)) \leq \frac{1}{d(d-1)} d \cdot c(t) = \frac{c(t)}{d-1}. \end{aligned}$$

Here we used that $D(t) \leq c(t)$ as t was not protected in \mathcal{S} . Using this fact once more, we get

$$D'(t) \leq D(t) + \frac{c(t)}{d-1} \leq \frac{d}{d-1} c(t).$$

□

Theorem 2.12. *The price of anarchy for the Mafia Hitting Set Game is d .*

Proof. Let $\mathcal{S} = (M, C, r)$ be a strategy profile in a Nash equilibrium. Then M is a hitting set, as if there was an uncovered club, all members would be unhappy due to the term $\text{Pen}(v)$. We show that the cost of M is within a factor d from the optimum. Let us set $y(S) = \sum_{v \in M \cap S} r(v, S)$ for each $S \in \mathcal{E}$. Lemmas 2.10 and 2.11 easily imply $\sum_{S \in \mathcal{S}: v \in S} y(S) \leq d \cdot c(v)$ for every $v \in V$, and thus $\frac{1}{d}y$ is a feasible dual solution to (D-HS). Then

$$\sum_{S \in \mathcal{S}} \frac{1}{d} y(S) = \sum_{s \in \mathcal{S}} \frac{1}{d} \sum_{m \in M \cap S} r(m, S) = \sum_{m \in M} \frac{1}{d} \sum_{S \in \mathcal{N}(m)} r(m, S) = \frac{1}{d} \sum_{m \in M} c(m),$$

showing that M is a d -approximate solution to (P-HS). □

2.3 Submodular hitting set

In the submodular hitting set problem, we are given a hypergraph $G = (V, \mathcal{E})$ with a submodular set function $C : 2^V \rightarrow \mathbb{R}_+$, that is, $C(\emptyset) = 0$, and

$$C(X) + C(Y) \geq C(X \cap Y) + C(X \cup Y) \quad \forall X, Y \subseteq V.$$

We shall assume also that C is monotone, that is, $C(X) \leq C(Y)$ if $X \subseteq Y$. Our aim is to find a hitting set M minimizing $C(M)$.

Koufogiannakis and Young [25], and Iwata and Nagano [18] obtained d -approximation algorithms for this problem, where d is the maximum size of a hyperedge. We shall present the primal-dual algorithm in [18], a natural extension of the Bar-Yehuda–Even algorithm.

For a submodular function C , it is natural to define the following two polyhedra. The *submodular polyhedron* is

$$P(C) = \{z \in \mathbb{R}^V : z \geq 0, z(Z) \leq C(Z) \quad \forall Z \subseteq V\},$$

and the *submodular base polyhedron* is

$$B(C) = \{z \in \mathbb{R}^V : z \geq 0, z(Z) \leq C(Z) \quad \forall Z \subsetneq V, z(V) = C(V)\}.$$

Given a vector $z \in P(C)$, the set Z is *tight* with respect to z if $z(Z) = C(Z)$. An elementary consequence of submodularity is that for every $z \in P(C)$, there exists a unique maximal tight set. Note that $B(C) \subseteq P(C)$ and $z \in P(C)$ is in $B(C)$ if and only if V is tight.

In the LP relaxation, we assign a primal variable $\xi(Z)$ to every subset $Z \subseteq V$. In an integer solution, $\xi(Z) = 1$ if Z is the chosen hitting set and 0 otherwise.

$$\begin{array}{ll}
\min \sum_{Z \subseteq V} C(Z) \xi(Z) & \max \sum_{S \in \mathcal{E}} y(S) \\
\sum_{Z \in \mathcal{N}(u)} \xi(Z) = x(u) & \forall u \in V \\
\sum_{u \in S} x(u) \geq 1 & \forall S \in \mathcal{E} \quad (\text{P-SHS}) \\
\xi \geq 0 &
\end{array}
\qquad
\begin{array}{ll}
\sum_{S \in \mathcal{N}(u)} y(S) = z(u) & \forall u \in V \\
z \in P(C) & (\text{D-SHS}) \\
y \geq 0 &
\end{array}$$

Note that in the dual program, y uniquely defines z . Therefore we will say that y is a feasible dual solution if the corresponding z is in $P(C)$. For the special case of the (linear) hitting set problem, where $C(Z) = \sum_{v \in Z} c(v)$ for some $c : V \rightarrow \mathbb{R}_+$, this is equivalent to y satisfying (D-HS).

Accordingly, we say that a set Z is tight for a feasible dual y if $z(Z) = P(C)$. For a hitting set M and a feasible dual solution y , we say that (M, y) is a *complementary pair* if M is tight for y . The following is the generalization of Lemmas 2.1 and 2.7.

Lemma 2.13. *If (M, y) is a complementary pair, then M is a d -approximate solution of the hitting set problem.*

Proof. The primal objective is at most d times the dual objective, as

$$C(M) = \sum_{v \in M} z(v) = \sum_{v \in M} \sum_{S \in \mathcal{N}(v)} y(S) \leq d \sum_S y(S).$$

The inequality follows as each S is counted $|S| \leq d$ times. □

The algorithm by Iwata and Nagano [18] is as follows.

- (0) Set $y(S) := 0$ for each $S \in \mathcal{E}$, $z(v) := 0$ for $v \in V$, and let M be the unique maximal set with $C(M) = 0$.
- (1) While M is not a hitting set do
 - (1-1) Choose an arbitrary hyperedge $S \in \mathcal{E}$, $S \cap M = \emptyset$.
 - (1-2) Compute $\varepsilon = \max\{\lambda : z + \lambda \chi_S \in P(C)\}$.
 - (1-3) Increase $y(S)$ and every $z(v)$ for $v \in S$ by ε .
 - (1-4) Replace M by the new unique maximal tight set.
- (2) Return M .

In step (1-2), χ_S is the characteristic function of S . This step can be performed in the same running time as a submodular function minimization (see [14]). Note also that M will always intersect S in step (1-4) and therefore will be strictly extended. It is immediate that it returns a complementary pair (M, y) and thus Lemma 2.13 proves d -approximation.

Game definition. The vector z in (D-SHS) plays an analogous role to the budgets c in the (linear) Mafia Hitting Set Game. We introduce a new agent, the *Godfather* to set the budgets of the agents.

The *Submodular Mafia Hitting Set Game* is defined on a hypergraph $\mathcal{G} = (V, \mathcal{E})$ and a monotone submodular set function $C : 2^V \rightarrow \mathbb{R}_+$. There are $|V| + 1$ agents, one for each vertex and a special agent g , called the *Godfather*.

The strategy of the Godfather is to return a budget vector $\tilde{c} \in B(C)$. The basic strategy of an agent $v \in V$ is to decide being a civilian or being a mafioso. The strategy of a mafioso m further incorporates normalized ransoms $r_0(m, S) \geq 0$ for clubs $S \in \mathcal{N}(m)$ with $\sum_{S \in \mathcal{N}(m)} r_0(m, S) = 1$, that is, $r_0(m, S)$ expresses the fraction of the budget of m he is willing to charge on S .

The sets of civilians and mafiosi will again be denoted by C and M , respectively. Hence a strategy profile is given as $\mathcal{S} = (M, C, \tilde{c}, r_0)$. The actual ransoms will be $r(m, S) = r_0(m, S) \cdot \tilde{c}(m)$.

The utility of the Godfather is the total budget of the Mafia: $U_{\mathcal{S}}(g) = C(M)$. The utility of the vertex agents is defined the same way as for the linear Mafia Hitting Set Game in Section 2.2, with replacing $c(v)$ by $\tilde{c}(v)$ everywhere.

For linear cost functions, we have $C(Z) = \sum_{v \in Z} c(v)$. Then the only vector in $B(C)$ is c , hence the Godfather has only one strategy to choose. Therefore we obtain the same game as described in Section 2.2.

Existence of a Nash equilibrium and bounding the Price of Anarchy. As for vertex cover and hitting set, we show that a complementary solution (M, y) to (P-SHS) and (D-SHS) provides a solution in Nash equilibrium. Let $z(u) = \sum_{S \in \mathcal{N}(u)} y(S)$. Note that $z \in P(C)$ and M is tight for z . Let us raise the $z(v)$ values for $v \in C$ arbitrarily in order to get a vector in the base polyhedron $B(C)$. Let \tilde{c} denote such a vector.

Lemma 2.14. *Let us define strategies in the Mafia Hitting Set Game based on a complementary pair (M, y) with M being the Mafia and $V \setminus M$ the civilians. Let the Godfather assign the budget vector \tilde{c} as defined above. For $u \in M$ and $S \in \mathcal{N}(u)$, define $r_0(u, S) = y(S)/\tilde{c}(u)$. Then the strategy profile $\mathcal{S} = (M, C, \tilde{c}, r_0)$ is a Nash equilibrium.*

Proof. The Godfather has no incentive to change as by $\tilde{c}(M) = C(M)$, he already receives the maximum possible utility for the given M . By the definition, $r(u, S) = y(S)$ for each $v \in V$, hence $D(v) \leq \frac{1}{d-1} \sum_{S \in \mathcal{N}(v)} y(S) |S \cap M| \leq \tilde{c}(v)$ and therefore there are no protected mafiosi. The proof that nobody has an incentive to change his strategy is the same as for Lemma 2.2. \square

Theorem 2.15. *The price of anarchy for the Submodular Mafia Hitting Set Game is d .*

Proof. Consider a strategy profile $\mathcal{S} = (M, C, \tilde{c}, r_0)$ in a Nash equilibrium. We can repeat the entire argument of Section 2.2 to show that there are no protected mafiosi and that every civilian is demanded at most $\frac{d}{d-1} \tilde{c}$. This is since if the Godfather does not change his strategy, the game is identical to the linear game with fixed budgets \tilde{c} from the perspective of the vertex agents.

Finally we define $y(S) = \sum_{m \in M \cap S} r(m, S)$, and then $\frac{1}{d} y$ is a feasible dual solution for (D-SHS). Observe that $\tilde{c}(M) = C(M)$, as otherwise the Godfather would have incentive to decrease the budgets of certain civilians and increase for certain mafiosi. Consequently, $C(M) = \tilde{c}(M) = \sum_{S \in \mathcal{E}} y(S)$, showing that M is a d -approximate solution for (P-SHS). \square

3 Convergence to Nash equilibrium

In this section, we investigate if the Mafia Games defined in the previous section converge under certain best response dynamics. We first show that already in the Mafia Vertex Cover Game, a round robin best response dynamics may run into a loop.

Motivated by this example, we modify the utilities by adding a secondary payoff, that instigates the mafiosi to use symmetric ransoms: $r(u, v) = r(v, u)$. With this secondary objective, we show that a single round of best response dynamics under a simple selection rule results in a Nash-equilibrium. This dynamics simulates the Bar-Yehuda–Even algorithm. An analogous result is then proved for hitting set. Finally, we discuss possible extensions for the submodular case.

3.1 Vertex cover

Let us now show an example where a round robin dynamics does not necessarily converge. Consider a star on the vertices v_1, v_2, v_3, v_4 and the central vertex z . Assume we are playing round robin in the order z, v_1, v_2, v_3, v_4 . Let $c(v_i) = 1$ for $i = 1, 2, 3$ and $c(z) = 2$, and let us start with the strategy profile where $M = \{z\}$, $r(z, v_1) = 2$. Assume that whenever z can change his strategy to get a higher utility, he always chooses to demand his entire budget 2 from one of the civilians among v_1, v_2, v_3, v_4 (this is always a best response).

We claim that this will always be possible as z always stays in the Mafia, and at most 3 vertices among v_1, v_2, v_3, v_4 will be in the Mafia at the same time. Indeed, a civilian will enter only if being ransomed by z . If v_i is in the Mafia then his only option is setting $r(v_i, z) = 1$, and thus if z has at least 3 neighbors in the Mafia, he becomes protected and thus all his neighbors he is not actually ransoming will have an incentive to leave.

The dynamics never reaches a Nash equilibrium, as if z is ransoming a mafioso v_i , he has incentive to change to ransoming a civilian as v_i is protected. On the other hand, if z ransoms a civilian v_i , v_i has an incentive to join the Mafia to obtain the protected status.

If we could incentivize z to change his strategy less drastically and ransom the other players by at most 1, we could rapidly reach a Nash-equilibrium. To enforce such a behavior, we introduce a secondary utility function. For a strategy profile $\mathcal{S} = (M, C, r)$, $U_{\mathcal{S}}(v)$ is the utility as defined in Section 2.1. Let us define $\tilde{U}_{\mathcal{S}}(v) = 0$ if $v \in C$ and

$$\tilde{U}_{\mathcal{S}}(v) = - \sum_{uv \in E, u \in M} |r(u, v) - r(v, u)|$$

if $u \in M$. The total utility is then $(U_{\mathcal{S}}(v), \tilde{U}_{\mathcal{S}}(v))$ in the lexicographic ordering: the agents' main objective is to maximize $U_{\mathcal{S}}(v)$, and if that is the same for two outcomes, they choose the one maximizing $\tilde{U}_{\mathcal{S}}(v)$. In the above example, the dynamics would reach an equilibrium in the second round, with $r(z, v_i) \leq 1$ for all i .

$\tilde{U}_{\mathcal{S}}(v) \leq 0$ and equality holds if $r(u, v) = r(v, u)$ for every $uv \in M$, $u, v \in M$. Therefore all results in Section 2.1 remain valid: in Lemma 2.2 we define a strategy profile where $\tilde{U}_{\mathcal{S}}(v) = 0$ for all agents, hence it also gives a Nash equilibrium for the extended definition of utilities. The secondary utility term \tilde{U} does not affect the proofs in Section 2.1.

Consider now the following simple dynamics: *Start from the strategy profile where all agents are civilians. In each step, take an agent who is incident to uncovered edge and subject to this, minimizes $c(v) - D(v)$, and give him the opportunity to change his strategy.*

Theorem 3.1. *After each agent changing his strategy at most once, we obtain a strategy profile in Nash equilibrium.*

Proof. By induction, we shall prove that in every step, $c(v) \geq D(v)$ and $\tilde{U}_{\mathcal{S}}(v) = 0$ for all $v \in V$. Consider the next move, when a player v incident to some uncovered edges minimizing $c(v) - D(v)$ moves. He obviously has to enter the Mafia, and can achieve a maximal (primary and secondary) utility if he sets $r(v, u) = r(u, v)$ for any $u \in M \cap N(v)$, and distributes the rest of his ransoms arbitrarily to his civilian neighbors. Note that this can always be done because $c(v) \geq D(v)$. Also, note that the total ransom v will demand from other civilians is $c(v) - D(v)$. By the extremal choice of v , it follows that none of his civilian neighbors z will violate $c(z) \geq D(z)$. This also remains true if $z \in M$, as $D(z)$ is at most the total ransom z demands due to the symmetry of the ransoms.

Hence the induction hypothesis is maintained by an arbitrary best response of v . A mafioso who is not protected and has secondary objective 0 has no incentive to change his strategy. Also, a civilian v with $c(v) \geq D(v)$ has no incentive to join the Mafia if there are no uncovered edges incident to v . Consequently, the game ends after all uncovered edges are gone, and once an agent joins to Mafia, he would not change his strategy anymore. \square

Observe that the dynamics is closely related to the Bar-Yehuda–Even algorithm: if the next agent always ransoms only one of its civilian neighbors, then it corresponds to a possible performance of the algorithm.

The above dynamics can be naturally interpreted in a distributive manner. In the proof of Theorem 3.1, we only use that the vertex v changing his strategy *is a local minimizer* of $c(v) - D(v)$. The simultaneous move of two agents u and v could interfere only if $uv \in E$ or they have a neighbor t in common. In this case, $c(t) < D(t)$ could result if both u and v start ransoming t simultaneously.

We assume that the agents have a hierarchical ordering \prec : $u \prec v$ expresses that v is more powerful than u . We call an agent v a *local minimizer* if $v \in C$, v is incident to some uncovered edges, and $c(v) - D(v) \leq c(u) - D(u)$ whenever $u \in C$, $uv \in E$. A local minimizer v is then called *eligible* if $u \prec v$ for all local minimizers u whose distance from v is at most 2.

We start from $C = V$. In each iteration of the dynamics, we let all eligible agents change their strategy to a best response simultaneously. As in the proof of Theorem 3.1, $c(v) - D(v) \geq 0$ is maintained for all $v \in V$, and thus the dynamics terminates after each agent changes his strategy at most once.

There are multiple distributed algorithms in the literature for vertex cover, e.g. [20, 15, 24]. The distributed algorithm by Koufogiannakis and Young [24] computes in $O(\log n)$ rounds a 2-approximation in expectation with high probability. In contrast, we cannot give good bounds on the number of iterations of our distributed dynamics. For example, if the graph is a path $v_1 \dots v_n$, and the budgets are $c(v_i) = i$, then only agent i will move in step i . Yet we believe that our dynamics could be practically efficient.

3.2 Hitting set

The natural generalization of the secondary objective for hitting set is as follows. For a club $S \in \mathcal{E}$, let $var(S)$ denote the maximum difference between ransoms on this edge. That is, $var(S) = 0$ if $|S \cap M| \leq 1$ and $var(S) = \max_{v \in S \cap M} r(v, S) - \min_{v \in S \cap M} r(v, S)$ otherwise. For a strategy profile $\mathcal{S} = (M, C, r)$, let $\tilde{U}_{\mathcal{S}}(v) = 0$ if $v \in C$ and

$$\tilde{U}_{\mathcal{S}}(v) = - \sum_{S \in \mathcal{N}(v)} var(S)$$

if $v \in M$. The utility of an agent is then $(U_{\mathcal{S}}(v), \tilde{U}_{\mathcal{S}}(v))$, under lexicographic ordering.

A natural expectation would be to prove rapid convergence as for vertex cover, if always the agent minimizing $c(v) - D(v)$ is allowed to play. However, the Bar-Yehuda–Even algorithm does not seem to be modeled by this dynamics. Instead, we define a slightly different extremal choice of the next agent. Let

$$D^*(v) = \sum_{S \in \mathcal{N}(v)} \max_{m \in S \cap M} r(m, S),$$

that is, for each club S we consider the largest ransom demanded in this club. Note that $D(v) \leq D^*(v)$. Let us consider the following dynamics. *We start from the strategy profile where everyone is civilian, and we always let a civilian play next who is contained in an uncovered club. Among them, we let the one play who minimizes $c(v) - D^*(v)$.*

Theorem 3.2. *After each agent changing his strategy at most once, we obtain a strategy profile in Nash equilibrium.*

Proof. We prove by induction, that in every step, $c(v) \geq D^*(v)$ and $\tilde{U}_{\mathcal{S}}(v) = 0$ for all $v \in V$. Note that this implies that there are no protected mafiosi. If M is not a hitting set, we let a v minimizing $c(v) - D^*(v)$ play. $\tilde{U}_{\mathcal{S}}(m) = 0$ for all $m \in M$ means that for every club S , $r(m, S)$ is equal for every $m \in M \cap S$; let r_S denote this common value. As for the vertex cover case, the best responses of v are to set $r(v, S) = r_S$ whenever S was already covered by the Mafia, and to distribute the remaining ransoms arbitrarily on the hyperedges covered only by v .

As $D^*(v) = \sum_{S \in \mathcal{N}(v)} r_S$, the remaining amount v distributes is exactly $c(v) - D^*(v)$. Then by the choice of v , $c(z) \geq D^*(z)$ shall be maintained for every civilian z , and also for other mafiosi (note that if $z \in M$, then $D^*(z)$ does not change). It can be seen analogously as for vertex cover, that we have a Nash equilibrium if there are no more uncovered clubs. \square

Similarly to the vertex cover case, this dynamics essentially simulates the Bar-Yehuda–Even algorithm. Also, an analogous distributed interpretation can be given.

3.3 Submodular hitting set

One would expect that the Submodular Mafia Hitting Set Game also converges under some dynamics that simulates the primal-dual algorithm by Iwata and Nagano [18]. However, if the Godfather does not have a secondary utility, the following example shows that it can run into a loop even in very simple instances.

Let $V = \{a, b\}$, $C(\{a\}) = C(\{b\}) = C(\{a, b\}) = 1$, and let g be the Godfather. Let $\mathcal{E} = \{\{a, b\}\}$; for simplicity, we use the notation of vertex cover, e.g. $r_0(a, b)$ denotes $r_0(a, \{a, b\})$. Let us start from the strategy profile $C = V$, $\tilde{c}(a) = 1$, $\tilde{c}(b) = 0$, and play a round robin in the order a, b, g .

First, a enters Mafia and sets $r_0(a, b) = 1$. Then b also enters to receive the protected status and sets $r_0(b, a) = 1$. g has no incentive to move as $\tilde{c}(M) = 1$ is already maximal. In the next round, a is happier if he leaves Mafia; b has no incentive to change, however g modifies to $\tilde{c}(a) = 0$ and $\tilde{c}(b) = 1$. This will lead to a loop: a enters again in next round, b leaves, \tilde{c} is changed again, etc.

The above behavior can be avoided by introducing a secondary utility for g : let $\tilde{U}_{\mathcal{S}}(g) = \sum_{v \in M} F^+(v)$, that is, the sum of the actual incomes of the mafiosi. Note that $\tilde{U}_{\mathcal{S}}(g) \leq \tilde{c}(M)$ and equality holds if and only if there are no protected mafiosi. With this secondary utility, after both a and b enter Mafia, g will modify to $\tilde{c}(a) = \tilde{c}(b) = 0.5$, giving a Nash equilibrium.

We conjecture that with this secondary utility and the secondary utilities for the vertex agents as for hitting set, rapid convergence can be shown under an appropriate choice of the next agent.

4 Conclusions and further research

We have defined games whose Nash equilibria correspond to certain covering problems, with the price of anarchy matching the best constant factor approximations. The payoffs in these games are locally defined, and the analysis is based on the LP relaxations of the corresponding covering problems. An intriguing question is if a similar game theoretic approach could be applied for further combinatorial optimization problems.

The first natural direction would be to extend our approach to a broader class of covering games. The most general approximation result on covering games is [25], giving a d -approximation algorithm for minimizing a submodular function under monotone constraints, each constraint dependent on at most d variables. As a first step, one could study hitting set with the requirement that each hyperedge S must be covered by at least $h(S) \geq 1$ elements; a simple primal-dual algorithm was given in [16]. However, extending our game even to this setting does not seem straightforward.

One could also try to formulate analogous settings for classical optimization problems such as facility location, Steiner-tree or knapsack. One inherent difficulty is that in our analysis, it seems to be crucial that any greedily chosen maximal feasible dual solution gives a good approximation. Also, we heavily rely on the fact that each constraint contains at most d variables.

In Section 3, we have shown that the best response dynamics rapidly converges for vertex cover and hitting set under certain assumptions. Stronger convergence results might hold: for example, it is open if arbitrary round robin best response dynamics converge to a Nash equilibrium. For the Submodular Mafia Hitting Set Game, we do not even have the weaker convergence result.

Acknowledgements We would like to thank Jarik Nešetřil for inspiring us to work on this problem and for a generous support in all directions.

References

- [1] E. Anshelevich, A. Dasgupta, É. Tardos, and T. Wexler. Near-optimal network design with selfish agents. *Theory of Computing*, 4(1):77–109, 2008.
- [2] M. Balcan, S. Krehbiel, G. Piliouras, and J. Shin. Near optimality in covering and packing games by exposing global information. *Arxiv preprint arXiv:1109.3606*, 2011.
- [3] R. Bar-Yehuda and S. Even. A linear-time approximation algorithm for the weighted vertex cover problem. *Journal of Algorithms*, 2(2):198–203, 1981.
- [4] K. Bhawalkar, M. Gairing, and T. Roughgarden. Weighted congestion games: Price of anarchy, universal worst-case examples, and tightness. *Algorithms–ESA 2010*, pages 17–28, 2010.
- [5] N. Buchbinder, L. Lewin-Eytan, J. Naor, and A. Orda. Non-cooperative cost sharing games via subsidies. *Algorithmic Game Theory*, pages 337–349, 2008.
- [6] J. Cardinal and M. Hoefer. Selfish service installation in networks. *Internet and Network Economics*, pages 174–185, 2006.
- [7] J. Cardinal and M. Hoefer. Non-cooperative facility location and covering games. *Theoretical Computer Science*, 411(16-18):1855–1876, 2010.
- [8] C. Daskalakis, R. Frongillo, C. Papadimitriou, G. Pierrakos, and G. Valiant. On learning algorithms for nash equilibria. *Algorithmic Game Theory*, pages 114–125, 2010.
- [9] X. Deng, T. Ibaraki, and H. Nagamochi. Algorithmic aspects of the core of combinatorial optimization games. *Mathematics of Operations Research*, pages 751–766, 1999.
- [10] N. Devanur, M. Mihail, and V. Vazirani. Strategyproof cost-sharing mechanisms for set cover and facility location games. *Decision Support Systems*, 39(1):11–22, 2005.

- [11] B. Escoffier, L. Gourves, and J. Monnot. On the impact of local taxes in a set cover game. In *Structural Information and Communication Complexity (SIROCCO)*, volume 6058, page 2. Springer-Verlag New York Inc, 2010.
- [12] A. Fabrikant, C. Papadimitriou, and K. Talwar. The complexity of pure nash equilibria. In *ACM Symposium on Theory of Computing (STOC)*, pages 604–612. ACM, 2004.
- [13] Q. Fang and L. Kong. Core stability of vertex cover games. In *International Conference on Internet and Network Economics*, pages 482–490. Springer-Verlag, 2007.
- [14] L. Fleischer and S. Iwata. A push-relabel framework for submodular function minimization and applications to parametric optimization. *Discrete Applied Mathematics*, 131(2):311–322, 2003.
- [15] F. Grandoni, J. Könemann, and A. Panconesi. Distributed weighted vertex cover via maximal matchings. *Computing and Combinatorics*, pages 839–848, 2005.
- [16] N. Hall and D. Hochbaum. A fast approximation algorithm for the multicovering problem. *Discrete Applied Mathematics*, 15(1):35–40, 1986.
- [17] N. Immorlica, M. Mahdian, and V. Mirrokni. Limitations of cross-monotonic cost-sharing schemes. *ACM Transactions on Algorithms (TALG)*, 4(2):24, 2008.
- [18] S. Iwata and K. Nagano. Submodular function minimization under covering constraints. In *Foundations of Computer Science (FOCS)*, pages 671–680. IEEE, 2009.
- [19] S. Khot and O. Regev. Vertex cover might be hard to approximate to within $2-\epsilon$. In *Computational Complexity*, pages 379–386. IEEE, 2003.
- [20] S. Khuller, U. Vishkin, and N. Young. A primal-dual parallel approximation technique applied to weighted set and vertex covers. *J. Algorithms*, 17(2):280–289, 1994.
- [21] R. Kleinberg, K. Ligett, G. Piliouras, and É. Tardos. Beyond the nash equilibrium barrier. In *Symposium on Innovations in Computer Science (ICS)*, 2011.
- [22] R. Kleinberg, G. Piliouras, and E. Tardos. Multiplicative updates outperform generic no-regret learning in congestion games. In *ACM Symposium on Theory of Computing (STOC)*, 2009.
- [23] R. Kleinberg, G. Piliouras, and É. Tardos. Load balancing without regret in the bulletin board model. *Distributed Computing*, 24(1):21–29, 2011.
- [24] C. Koufogiannakis and N. Young. Distributed and parallel algorithms for weighted vertex cover and other covering problems. In *ACM Symposium on Principles of Distributed Computing (PODC)*, pages 171–179. ACM, 2009.
- [25] C. Koufogiannakis and N. Young. Greedy Δ -approximation algorithm for covering with arbitrary constraints and submodular cost. *Automata, Languages and Programming*, pages 634–652, 2009.
- [26] E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *Symposium on Theoretical Aspects of Computer Science (STACS)*, pages 404–413. Springer-Verlag, 1999.
- [27] X. Li, Z. Sun, and W. Wang. Cost sharing and strategyproof mechanisms for set cover games. In *Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 3404, page 218. Springer, 2005.
- [28] X. Li, Z. Sun, W. Wang, X. Chu, S. Tang, and P. Xu. Mechanism design for set cover games with selfish element agents. *Theoretical Computer Science*, 411(1):174–187, 2010.
- [29] R. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2(1):65–67, 1973.

- [30] T. Roughgarden. Intrinsic robustness of the price of anarchy. In *ACM Symposium on Theory of Computing (STOC)*, pages 513–522. ACM, 2009.
- [31] T. Roughgarden and F. Schoppmann. Local smoothness and the price of anarchy in atomic splittable congestion games. *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 255–267, 2011.
- [32] D. Shah and J. Shin. Dynamics in congestion games. In *ACM SIGMETRICS Conference on Measurement and Modeling of Computer Systems*, page 107, 2010.