Compensate to Not Deviate: On Subsidized Equilibria

Vittorio Bilò[1]*, Gianpiero Monaco[2]*, Luca Moscardelli[2]*

[1]University of Salento, Italy
[2]University of Chieti-Pescara, Italy
vittorio.bilo@unisalento.it, gianpiero.monaco@unich.it, luca.moscardelli@unich.it

Abstract

We introduce a new notion of deterministic stable solution for non-cooperative games, termed subsidized equilibrium. It assumes that an amount of money can be used as a pool of subsidies to stabilize a strategy profile that otherwise would not be accepted by (some of) the players. Roughly speaking, for a given amount of money, a strategy profile is a subsidized equilibrium if the total payoff loss incurred by players not playing best-responses does not exceed that amount, i.e., there is enough money to refund all players experiencing a regret. With respect to many other solution concepts in the literature, the notion of subsidized equilibrium has important advantages. Specifically, for a sufficiently high amount of money, a subsidized equilibrium always exists and can even be computed in polynomial time; also, existence of an efficient subsidized equilibrium can be guaranteed. Thus, determining for which amounts existence, polynomial time computability and efficiency can or cannot be achieved becomes an intriguing question. We provide initial results towards this direction for some widely studied classes of games.

1 Introduction

In this work, we propose a new, simple and deterministic solution concept for non-cooperative games, termed ***subsidized equilibrium***, which can be intuitively illustrated as follows. In a solution not being a pure Nash equilibrium, there is a non-empty set of players who are *unhappy*: it is the set of players who are not playing a best-response and could improve their payoff by deviating. Anyway, it is reasonable to assume that this solution could be accepted by the unhappy players if somebody refunds them what they are losing for not adopting their best-response, i.e., their *regret*.

As is the case in some other closely related works discussed in the next subsection, the notion of subsidized equilibrium assumes that there is a third party (e.g., the game designer, a public entity or even the players themselves), interested in the realization of (efficient) stable outcomes, which invests an amount of money in the game. An interesting application of subsidized equilibria to *project financing* scenarios is also discussed in Section 5.

*These authors contributed equally.

Clearly, for any game with finite payoffs, there always exist sufficiently high subsides for which any outcome becomes acceptable for all players. This immediately implies that existence of a subsidized equilibrium is always guaranteed, computation can be done in polynomial time, and the most efficient solution, for any chosen definition of efficiency, can be enforced. Thus, determining the minimum total amount of subsides for which these properties can or cannot be realized becomes an intriguing question.

1.1 Closely Related Work

Subsidized equilibria share similarities with *approximate pure Nash equilibria*. In both cases, players may experience a regret, but, while in a subsidized equilibrium this regret is cancelled by a monetary refund, in an approximate pure Nash equilibrium it is assumed that players are willing to adopt a suboptimal behaviour as long as their regret is within a bounded distance, either additively or multiplicatively, from the optimal one. This difference already highlights a wider applicability of subsidized equilibria, as they can be accepted by purely selfish players who are not willing to compromise on their well-being for the sake of stability. Moreover, subsidized equilibria also surpass further limitations of approximate pure Nash equilibria. Additive approximation, in fact, makes sense only when the players' payoffs are roughly of the same magnitude. Consider, for instance, a two-player game such that, in *every strategy profile*, player 1 pays a cost of at least 1000, while player 2 pays a cost of at most 10. In this game, an additive approximation of 10 would be happily accepted by player 1, as she would be almost optimal (up to approximately a 1% loss), but could hypothetically be a disaster for player 2 (she could be paying 10, the worst possible cost, while having a deviation yielding zero cost). This is the reason why the literature has mainly focused of multiplicative approximate pure Nash equilibria. These equilibria, however, may not even exist in some important classes of games, while subsidized equilibria are always guaranteed to exist for a sufficiently high subsidy.

Partially related to subsidized equilibria is the concept of k-*implementation* (Monderer and Tennenholtz 2004), which also assumes a third party investing money to enforce desired outcomes. Formally, given a set of strategy profiles T, the third party commits to payments to the players whenever a certain element of T is realized. The sum of these payments must not exceed k, and they must be defined in such a way that only strategy profiles belonging to T can be realized under the assumption that no player chooses a dominated strategy. For normal-form games, Monderer and Tennenholtz (2004) prove that deciding whether a k-implementation of T exists can be solved in polynomial time if T is a singleton, and it is NP-hard otherwise. Subsidized equilibria and k-implementation both rely on the use of money to enforce certain outcomes. However, they approach the same problem under opposite directions. In k-implementation, given a set of outcomes, one wants to bound the budget needed to enforce them; in subsidized equilibria, vice-versa, given a budget, one wants to characterize the set of outcomes that could possibly be enforced. Moreover, while subsidized equilibria are stable outcomes for fully rational players, k-implementation is based on the weaker assumption that players choose non-dominated strategies. So the outcomes of k-implementation might not be stable when players accept best-response strategies only. Finally, while the money needed to stabilize a subsidized equilibrium is implicitly quantified by definition (every player is refunded her regret), in k-implementation the payments for each strategy profile is arbitrary and has to be suitably computed. This creates a complexity barrier, leading to the NP-hardness result in (Monderer and Tennenholtz 2004), which is even amplified when considering multi-player games represented in succinct form, as we do in this work. We conclude that the two notions do not share strong similarities and, in fact, we are not aware of any result for k-implementation that could be comparable with the ones that we achieve for subsidized equilibria. With this respect, it must be stressed that, from Theorem 1 in (Monderer and Tennenholtz 2004), one could argue that, for an amount of money depending on k, a strategy profile is a subsidized equilibrium if and only if it is k-implementable. However, the characterization given in Theorem 1 of (Monderer and Tennenholtz 2004) is not correct (all details will be provided in an extended version).

Another form of subsidy has been considered by Augustine et al. (2015) in the realm of network cost-sharing games. Here, a third party wants to subsidize edges, so that players can buy them at a lower price. Augustine et al. (2015) prove that, in the subclass of undirected broadcast games, a social optimum, i.e. a solution of minimum cost, can be made a pure Nash equilibrium using a budget not exceeding a $1/e$-fraction of its cost. In their model, however, resources (in this case, edges of the given network) rather than players are subsidized; moreover, while Augustine et al. (2015) only allow for enforcing a social optimum, in subsidized equilibria any strategy profile can be possibly subsidized.

1.2 Definitions and Notation

To better discuss our contribution, we need to recall some basic notions from (Algorithmic) Game Theory.

A *game* is a triple $\mathcal{G} = (N, (\Sigma_i)_{i\in N}, (u_i)_{i\in N})$ such that N is a finite set of n players and, for each $i \in N$, Σ_i is the strategy space of player i and $u_i : \times_{i\in N}\Sigma_i \mapsto \mathbb{R}_{\geq 0}$ is her payoff function. Games can be classified into two possible *types*: *profit-maximization games*, where payoffs are profits that every player wants to maximize, and *cost-minimization games*, where payoffs are costs to be minimized.

In the sequel, we introduce a series of definitions which, for the sake of conciseness, will be tailored for profit-maximization games only. Their adaptation to cost-minimization ones is left to the reader.

We denote by $\Sigma = \times_{i\in N}\Sigma_i$ the set of *strategy profiles* (or *outcomes*) of $\mathcal{G}$ and by $\boldsymbol{\sigma} = (\sigma_1, \ldots, \sigma_n)$ the strategy profile in which every player $i \in N$ chooses strategy $\sigma_i \in \Sigma_i$.

For a strategy profile $\boldsymbol{\sigma} \in \Sigma$, a player $i \in N$ and a strategy $s \in \Sigma_i$, $(\boldsymbol{\sigma}_{-i}, s)$ denotes the strategy profile obtained from $\boldsymbol{\sigma}$ when player i changes her strategy from σ_i to s. A *best-response* for i in $\boldsymbol{\sigma}$ is any strategy in $\text{argmax}_{s\in\Sigma_i} u_i(\boldsymbol{\sigma}_{-i}, s)$. For any $\epsilon \geq 0$, an *ϵ-multiplicative improving deviation* for i in $\boldsymbol{\sigma}$ is any strategy $s \in \Sigma_i$ such that $u_i(\boldsymbol{\sigma}_{-i}, s) > (1+\epsilon)u_i(\boldsymbol{\sigma})$, that is, i improves her payoff by more than a factor $1+\epsilon$ after changing her strategic choice to s. Similarly, an *ϵ-additive improving deviation* for i in $\boldsymbol{\sigma}$ is any strategy $s \in \Sigma_i$ such that $u_i(\boldsymbol{\sigma}_{-i}, s) > u_i(\boldsymbol{\sigma}) + \epsilon$, that is, i improves her payoff by more than ϵ after changing her strategic choice to s. An *ϵ-multiplicative pure Nash equilibrium* (ϵ-PNE*, for short) is a strategy profile in which no player possesses an ϵ-multiplicative improving deviation; similarly, an *ϵ-additive pure Nash equilibrium* (ϵ-PNE$^+$, for short) is a strategy profile in which no player possesses an ϵ-additive improving deviation. We denote by $\mathsf{PNE}^*_\epsilon(\mathcal{G})$ and $\mathsf{PNE}^+_\epsilon(\mathcal{G})$ the sets of ϵ-PNE* and ϵ-PNE$^+$ of $\mathcal{G}$, respectively. For $\epsilon = 0$, the notions of ϵ-PNE* and ϵ-PNE$^+$ collapse to that of *pure Nash equilibrium* (PNE for short) and we denote by $\mathsf{PNE}(\mathcal{G})$ the set of PNE of $\mathcal{G}$.

A game $\mathcal{G}$ is an *exact potential game* if there exists a function $\Phi : \Sigma \mapsto \mathbb{R}_{\geq 0}$, called *exact potential* for $\mathcal{G}$, such that, for any $i \in N$, $\boldsymbol{\sigma} \in \Sigma$ and $s \in \Sigma_i$, $\Phi(\boldsymbol{\sigma}) - \Phi(\boldsymbol{\sigma}_{-i}, s) = u_i(\boldsymbol{\sigma}) - u_i(\boldsymbol{\sigma}_{-i}, s)$. Every game possessing an exact potential admits a PNE (Monderer and Shapley 1996).

The *social value* of a strategy profile $\boldsymbol{\sigma}$ is defined as $\mathsf{SV}(\boldsymbol{\sigma}) = \sum_{i\in N} u_i(\boldsymbol{\sigma})$. We denote by $\boldsymbol{\sigma}^*$ a *social optimum*, that is, a strategy profile maximizing the social value.

Given a non-empty set of equilibria $\mathsf{EQ}(\mathcal{G})$ for game $\mathcal{G}$, the price of anarchy of $\mathcal{G}$ with respect to EQ is defined as $\mathsf{PoA}_{\mathsf{EQ}}(\mathcal{G}) = \max_{\boldsymbol{\sigma}\in\mathsf{EQ}(\mathcal{G})} \frac{\mathsf{SV}(\boldsymbol{\sigma}^*)}{\mathsf{SV}(\boldsymbol{\sigma})}$; the price of stability of $\mathcal{G}$ with respect to EQ is defined as $\mathsf{PoS}_{\mathsf{EQ}}(\mathcal{G}) = \min_{\boldsymbol{\sigma}\in\mathsf{EQ}(\mathcal{G})} \frac{\mathsf{SV}(\boldsymbol{\sigma}^*)}{\mathsf{SV}(\boldsymbol{\sigma})}$. All ratios are interpreted as ∞ when their denominator is zero. Choosing $\mathsf{EQ}(\mathcal{G})$ among $\mathsf{PNE}^*_\epsilon(\mathcal{G})$, $\mathsf{PNE}^+_\epsilon(\mathcal{G})$ and $\mathsf{PNE}(\mathcal{G})$, one obtains the notions of *ϵ-approximate multiplicative price of anarchy* ($\mathsf{PoA}^*_\epsilon(\mathcal{G})$), *$\epsilon$-approximate additive price of anarchy* ($\mathsf{PoA}^+_\epsilon(\mathcal{G})$), *price of anarchy* ($\mathsf{PoA}(\mathcal{G})$) (Koutsoupias and Papadimitriou 1999), *ϵ-approximate multiplicative price of stability* ($\mathsf{PoS}^*_\epsilon(\mathcal{G})$), *$\epsilon$-approximate additive price of stability* ($\mathsf{PoS}^+_\epsilon(\mathcal{G})$), and *price of stability* ($\mathsf{PoS}(\mathcal{G})$) (Anshelevich et al. 2008) of game $\mathcal{G}$. All these notions can be extended to a class of games $\mathcal{C}$, by taking the supremum over all games belonging to $\mathcal{C}$. For instance, the price of anarchy of $\mathcal{C}$ is defined as $\mathsf{PoA}(\mathcal{C}) = \sup_{\mathcal{G}\in\mathcal{C}}\mathsf{PoA}(\mathcal{G})$.

Define the *regret* of a player in a strategy profile as her payoff loss with respect to a best-response.

Definition 1.1. *For a strategy profile $\boldsymbol{\sigma}$ and a player $i \in N$,* *the regret of i in $\boldsymbol{\sigma}$ is $r_i(\boldsymbol{\sigma}) = \max_{s \in \Sigma_i} u_i(\boldsymbol{\sigma}_{-i}, s) - u_i(\boldsymbol{\sigma})$.*

Denote by $R(\boldsymbol{\sigma}) = \{i \in N : r_i(\boldsymbol{\sigma}) > 0\}$ the set of players experiencing a regret in $\boldsymbol{\sigma}$, that is, the set of players requiring a subsidy in $\boldsymbol{\sigma}$. The amount of subsidy needed to enforce a strategy profile cannot be expressed *apriori* in absolute terms, but rather needs to be related to the numbers defining the players' payoffs. Among the many possibilities, we select the reasonable choice of expressing the subsidy in terms of the value of the social optimum.

Definition 1.2. *Given a strategy profile $\boldsymbol{\sigma}$, the subsidy required to stabilize $\boldsymbol{\sigma}$ is $sub(\boldsymbol{\sigma}) = \sum_{i \in N} r_i(\boldsymbol{\sigma}) = \sum_{i \in R(\boldsymbol{\sigma})} r_i(\boldsymbol{\sigma})$. Given $\mu \geq 0$, a strategy profile $\boldsymbol{\sigma}$ is a μ-subsidized equilibrium (μ-SubEq) if $sub(\boldsymbol{\sigma}) \leq \mu \cdot \mathsf{SV}(\boldsymbol{\sigma}^*)$.*

So, any PNE is a 0-SubEq and the set of μ-SubEq, for any $\mu \geq 0$, generalizes that of PNE.

To define the price of stability of SubEq, we elaborate on the following considerations. When comparing the social value of a best-possible SubEq against a social optimum, a price of stability of $\theta \geq 1$ means that there is enough money to subsidize a θ-approximation of a social optimum and make it an equilibrium, without arguing on the amount of subsidy needed to enforce it. Nevertheless, there might exist less efficient strategy profiles requiring a significantly smaller subsidy. Hence, when defining the price of stability of SubEq, that we term *subsidized price of stability*, we also account for the effort required to achieve stability. So, we define the *extended social value* of a strategy profile $\boldsymbol{\sigma}$ as $\mathsf{ESV}(\boldsymbol{\sigma}) = \mathsf{SV}(\boldsymbol{\sigma}) - sub(\boldsymbol{\sigma})$. Note that $\mathsf{ESV}(\boldsymbol{\sigma})$ can be potentially negative. The *subsidized price of stability* of $\mathcal{G}$ is $\mathsf{SubPoS}(\mathcal{G}) = \min_{\boldsymbol{\sigma} \in \Sigma} \frac{\mathsf{SV}(\boldsymbol{\sigma}^*)}{\mathsf{ESV}(\boldsymbol{\sigma})}$[1]. For a class of games $\mathcal{C}$, the subsidized price of stability of $\mathcal{C}$ is defined as $\mathsf{SubPoS}(\mathcal{C}) = \sup_{\mathcal{G} \in \mathcal{C}} \mathsf{SubPoS}(\mathcal{G})$. We stress that our definition allows for fair comparisons between the SubPoS and the PoS of a given class of games, as *(i)* both of them compare against the social optimum and *(ii)* the subsidy required to stabilize a SubEq is deprived from its social value.

We conclude this section by formally defining four well-studied classes of games that we are going to use to demonstrate the effectiveness of the notion of SubEq. They are: (polynomial) congestion games, additively-separable hedonic games with symmetric preferences, directed cut games and cost-sharing games.

Congestion games (Rosenthal 1973) are cost-minimization games in which there is a set of resources E and, for each player $i \in N$, $\Sigma_i \subseteq 2^E \setminus \emptyset$. Each resource $e \in E$ has a latency function $\ell_e(x) : \{1, \ldots, n\} \mapsto \mathbb{R}_{\geq 0}$. The cost of player i in a strategy profile $\boldsymbol{\sigma}$ is defined as $u_i(\boldsymbol{\sigma}) = \sum_{e \in \sigma_i} \ell_e(n_e(\boldsymbol{\sigma}))$, where $n_e(\boldsymbol{\sigma}) = |\{i \in N : e \in \sigma_i\}|$ is the *congestion* of resource e in $\boldsymbol{\sigma}$. A *polynomial congestion game* of maximum degree $d \geq 1$ is a congestion game such that every latency function is a polynomial function in the congestion of degree at most d; an *affine congestion game* is a polynomial congestion game of maximum degree 1, that is, $\ell_e(x) = \alpha_e x + \beta_e$, with $\alpha_e, \beta_e \geq 0$, for each $e \in E$.

[1]This ratio is interpreted as ∞ whenever its denominator is non-positive.

Additively-separable hedonic games with symmetric preferences (Bogomolnaia and Jackson 2002) are profit-maximization games in which players must form coalitions. For every pair of players $i, j \in N$, w_{ij} denotes the happiness that both players get when belonging to a same coalition, with $w_{ij} = 0$ when $i = j$. The profit that player i gets when belonging to a coalition C is given by $\sum_{j \in C} w_{ij}$.

Directed cut games (Christodoulou, Mirrokni, and Sidiropoulos 2012; Kun, Powers, and Reyzin 2013) are profit-maximization games defined by an edge-weighted directed graph $G = (N, A, w)$ with $n = |N|$ nodes and $m = |A|$ arcs, where $w : A \to \mathbb{R}_{>0}$ is a function assigning a positive weight (denoted as w_a or w_{ij}) to arc $a = (i, j)$. We write $w(B) = \sum_{a \in B} w_a$ for the total weight of the arcs in $B \subseteq A$ and $\delta_i = \sum_{j:(i,j) \in A} w_{ij}$ for the weighted out-degree of node i. When $w_a = 1$ for all $a \in A$, the game is *unweighted*; when G is undirected, the game is *undirected*. There are n players, each corresponding to a distinct node of G. The players aim to collectively but non-cooperatively build a cut of G, i.e., a partition of its nodes into two sets. Player $i \in N$ has two different strategies: placing her node i at the *left* or at the *right* side of the cut. So we can define the *complement* of strategy σ_i as $\bar{\sigma_i}$, with $\bar{\sigma_i} =$ *left* when $\sigma_i =$ *right* and $\bar{\sigma_i} =$ *right* when $\sigma_i =$ *left*. For an outcome $\boldsymbol{\sigma}$, we denote by $\mathrm{CUT}(\boldsymbol{\sigma})$ the set of edges whose endpoints correspond to players selecting different strategies in $\boldsymbol{\sigma}$. For a subset of players $R \subseteq N$, we denote by $\mathrm{CUT}_R(\boldsymbol{\sigma})$ the subset of $\mathrm{CUT}(\boldsymbol{\sigma})$ that consists of edges starting from a node corresponding to a player in R. With a little abuse of notation, we simplify $\mathrm{CUT}_{\{i\}}(\boldsymbol{\sigma})$ to $\mathrm{CUT}_i(\boldsymbol{\sigma})$. The profit of player i is the total weight of the edges in the cut starting from her node, i.e., $u_i(\boldsymbol{\sigma}) = w(\mathrm{CUT}_i(\boldsymbol{\sigma}))$. For unweighted games, this simplifies to $u_i(\boldsymbol{\sigma}) = |\mathrm{CUT}_i(\boldsymbol{\sigma})|$.

Cost-sharing games (Anshelevich et al. 2008) are cost-minimization games in which there is a set of resources E and, for each player i, $\Sigma_i \subseteq 2^E \setminus \emptyset$. Each resource $e \in E$ has a cost $c(e)$ which is equally shared among all players using it. Formally, $u_i(\boldsymbol{\sigma}) = \sum_{e \in \sigma_i} \frac{c(e)}{n_e(\boldsymbol{\sigma})}$, where $n_e(\boldsymbol{\sigma}) = |\{i \in N : e \in \sigma_i\}|$. Observe that $\mathsf{SV}(\boldsymbol{\sigma})$ coincides with the sum of the costs of all resources used by at least one player in $\boldsymbol{\sigma}$. A *network cost-sharing game* is a cost-sharing game in which the set of resources E coincides with the set of edges of an edge-weighted graph $G = (V, E, c)$ and, for each $i \in N$, Σ_i is the set of simple paths connecting a pair of vertices (s_i, t_i), called, respectively, the source and destination vertices of player i. A(n) *(un)directed network cost-sharing game* is a network cost-sharing game defined by a(n) (un)directed graph. A *multicast game* is a network cost-sharing game in which all players share the same source vertex, that is, $s_i = s$ for each $i \in N$, while a *broadcast game* is a multicast game in which each vertex in $V \setminus \{s\}$ is the destination vertex of exactly one player.

1.3 Our Results

We study existence, poly-time computability and best-case efficiency (i.e., SubPoS) of SubEq. A few preliminary results, some of which can be easily achieved by exploiting relations with (approximate) pure Nash equilibria, will be provided in an extended version. Here, we only state the following:

Proposition 1.3. *For any cost-minimization game $\mathcal{G}$, $\mathsf{SubPoS}(\mathcal{G}) \leq 2$.*

In Section 2, we show how to efficiently compute a SubEq requiring arbitrarily small subsidy for exact potential games under mild assumptions. Specifically, we require that both a best-response and a polynomial approximation of the social optimum can be efficiently computed and that the potential of each strategy profile can be upper bounded by its social value suitably scaled by a factor that is polynomial in the game representation. Under these premises, for any $\mu > 0$, a μ-SubEq can be computed in time which is polynomial in the game representation and in $1/\mu$. We stress that many games of interest satisfy our assumptions. Among them are undirected cut games, polynomial congestion games, cost-sharing games and additively-separable hedonic games with symmetric preferences. For all of them computing a PNE is PLS-complete (see Section 1.4).

In Section 3, we deal with existence of SubEq in non-potential games. We consider the class of directed cut games, which are known to not possess ϵ-PNE* even in the unweighted case. While this instance also implies that no μ-SubEq exists for $\mu < 1/2$, we show that a 1-SubEq always exists in general games, and, through an elaborated application of a probabilistic-like method, we prove the existence of a $3/5$-SubEq for unweighted games. This upper bound can be refined to $1/2$, thus matching the lower bound, under the assumption that the input graph is dense enough so that every node has at least two out-going edges.

In Section 4, we study the SubPoS. We investigate this question with respect to cost-sharing games, affine congestion games, and unweighted directed cut games. For cost-sharing games, we design a cleverly constructed instance showing that the 2 upper bound of Proposition 1.3 is tight, even for directed broadcast games. For undirected games, however, better bounds are possible as, by leveraging a result of Augustine et al. (2015), we show that the SubPoS in broadcast games becomes at most $(e + 1)/e \approx 1.3678$; moreover, if $n = 2$, we prove an exact bound of $8/7$ and, for $n \geq 4$, a lower bound of $6/5$. For affine congestion games, by bounding the amount of subsidies needed to stabilize a social optimum, we derive an upper bound of $3/2$. Moreover, under the assumption that each player has two available strategies only, we show that the SubPoS lies in the interval $[135/112, 10/7]$. Finally, for unweighted directed cut games, we show that the SubPoS is in $[2, 9/2]$.

1.4 Further Related Work

Polynomial congestion games of maximum degree d is one of the most studied classes of games (Aland et al. 2011; Bilò and Vinci 2017; Bhawalkar, Gairing, and Roughgarden 2014; Rosenthal 1973; Skopalik and Vöcking 2008). They are exact potential games, but computing a PNE is PLS-complete (Fabrikant, Papadimitriou, and Talwar 2004; Ackermann, Röglin, and Vöcking 2008). For any $d \geq 1$, an $O(d)$-PNE* can be computed in polynomial time (Caragiannis et al. 2015); for affine games (i.e., $d = 1$), the best known guarantee is 1.61 due to (Vijayalakshmi and Skopalik 2020). The PoA^*_ϵ and PoS^*_ϵ have been addressed in (Aland et al. 2011; Awerbuch, Azar, and Epstein 2013; Bilò 2018; Christodoulou, Koutsoupias, and Spirakis 2011; Christodoulou and Gairing 2016; Caragiannis et al. 2011; Christodoulou and Koutsoupias 2005b,a).

Additively-separable hedonic games with symmetric preferences are exact potential games (Bogomolnaia and Jackson 2002), but computing a PNE is PLS-complete (Gairing and Savani 2019). Their PoA is unbounded, while their PoS is 1. These games are widely studied in the AI community (Aziz, Brandt, and Seedig 2013; Brandt, Bullinger, and Tappe 2024; Peters and Elkind 2015).

Kun, Powers, and Reyzin (2013) prove that deciding whether an unweighted directed cut game admits a PNE is NP-complete. D'Ascenzo et al. (2024) detect specific classes of unweighted games for which an ϵ-PNE*, for a very large ϵ, can be efficiently computed. Undirected cut games, instead, are exact potential games, but computing a PNE is PLS-complete (Christodoulou, Mirrokni, and Sidiropoulos 2012; Schäffer and Yannakakis 1991). The best known approximation achievable in polynomial time is a 2.7371-PNE* (Caragiannis and Jiang 2023). For any $\epsilon \geq 0$, the PoA^*_ϵ and PoS^*_ϵ have been shown to be equal to $2 + \epsilon$ (Bilò and Paladini 2016) and 1 (Hoefer 2007), respectively.

Cost-sharing games (Anshelevich et al. 2008) are a special case of congestion games with decreasing latency functions, often used as models of spontaneous network formation. Computing a PNE has been shown to be PLS-complete in (Syrgkanis 2010; Bilò et al. 2021), with Bilò et al. (2021) also contributing an algorithm to compute an ϵ-PNE* in some special cases. The PoA of these games equals the number of players n, while their PoS is $H_n = \sum_{i=1}^n 1/i$ (Anshelevich et al. 2008). For undirected games, however, the characterization of the PoS is still open and has attracted lot of research attention (Anshelevich et al. 2008; Bilò et al. 2013; Bilò, Flammini, and Moscardelli 2020; Fiat et al. 2006; Li 2009; Bilò and Bove 2011; Christodoulou et al. 2009; Mamageishvili, Mihalák, and Montemezzani 2018).

2 Computation of Subsidized Equilibria

In this section, we consider the problem of computing a μ-SubEq for exact potential games. These games admit a PNE and so existence of a μ-SubEq is guaranteed for every $\mu \geq 0$. However, there are several classes of exact potential games for which the computation of a PNE is PLS-complete. We shall surpass this limitation, by showing that, under mild assumptions and for each $\mu > 0$, a μ-SubEq can be computed in time polynomial in the game representation and in $1/\mu$.

Fix an exact potential game $\mathcal{G}$ and observe that, if Φ is an exact potential for $\mathcal{G}$, then, for any number M, Φ' such that $\Phi'(\boldsymbol{\sigma}) = \Phi(\boldsymbol{\sigma}) + M$ is also an exact potential for $\mathcal{G}$. Thus, by a suitable choice of M, we can assume without loss of generality that there exists an exact potential Φ for $\mathcal{G}$ such that $\Phi(\boldsymbol{\sigma}) \geq 0$ for each $\boldsymbol{\sigma} \in \Sigma$. Let c_{move} be the complexity of an algorithm that computes an ϵ-additive improving deviation for a given player, whenever it exists (c_{move} is at most the complexity of computing a best-response). An ϵ-additive improving dynamics, starting from a strategy profile $\boldsymbol{\sigma}^{initial}$, is a sequence of strategy profiles such that each profile is obtained from the previous one as a result of an ϵ-additive improving deviation. It is known (Monderer and Shapley 1996) that, for exact potential profit-maximization (*resp.* cost-minimization) games, Φ increases (*resp.* decreases) by at least ϵ after any ϵ-additive improving deviation. Let A be a polynomial-time algorithm computing an r-approximation of the social optimum and let $O(c_{apx})$ be its complexity. So A outputs a strategy profile $\overline{\boldsymbol{\sigma}}$ such that $\mathsf{SV}(\overline{\boldsymbol{\sigma}}) \leq \mathsf{SV}(\boldsymbol{\sigma}^*) \leq r\mathsf{SV}(\overline{\boldsymbol{\sigma}})$ in profit-maximization games and $\mathsf{SV}(\boldsymbol{\sigma}^*) \leq \mathsf{SV}(\overline{\boldsymbol{\sigma}}) \leq r\mathsf{SV}(\boldsymbol{\sigma}^*)$ in cost-minimization ones. Our idea is to exploit an ϵ-additive improving dynamics starting from $\boldsymbol{\sigma}^{initial} = \overline{\boldsymbol{\sigma}}$ and ending at an ϵ-PNE$^+$ for a suitable choice of ϵ. Since any ϵ-PNE$^+$ requires a subsidy of at most $n\epsilon$, the desired μ-SubEq is obtained.

Next two theorems distinguish between profit-maximization games and cost-minimization ones.

Theorem 2.1. *Given an exact potential profit-maximization game $\mathcal{G}$ and a parameter $\beta > 0$, let $\overline{\boldsymbol{\sigma}}$ be an r-approximation of $\boldsymbol{\sigma}^*$ and b be a number such that $\Phi(\boldsymbol{\sigma}) \leq b\mathsf{SV}(\boldsymbol{\sigma})$ for any $\boldsymbol{\sigma} \in \Sigma$. An $\frac{nr}{\beta}$-SubEq can be computed in $O(c_{apx} + nb\beta c_{move})$ time.*

Proof. Consider an ϵ-additive improving dynamics $\mathcal{D}$ starting from $\overline{\boldsymbol{\sigma}}$, with $\epsilon := \frac{\mathsf{SV}(\overline{\boldsymbol{\sigma}})r}{\beta}$. Recall that $\mathsf{SV}(\overline{\boldsymbol{\sigma}}) \leq \mathsf{SV}(\boldsymbol{\sigma}^*) \leq r\mathsf{SV}(\overline{\boldsymbol{\sigma}})$. Moreover, since Φ is always non-negative and increases at each step by a quantity of more than ϵ, we have that after $b\beta$ steps the value of the potential function is at least $b\beta\epsilon = b\beta\frac{\mathsf{SV}(\overline{\boldsymbol{\sigma}})r}{\beta} \geq b\beta\frac{\mathsf{SV}(\boldsymbol{\sigma}^*)}{\beta} = b\mathsf{SV}(\boldsymbol{\sigma}^*)$. By the assumption $\Phi(\boldsymbol{\sigma}) \leq b\mathsf{SV}(\boldsymbol{\sigma})$ for any $\boldsymbol{\sigma} \in \Sigma$, we derive that, for any $\boldsymbol{\sigma} \in \Sigma$, $\Phi(\boldsymbol{\sigma}) \leq b\mathsf{SV}(\boldsymbol{\sigma}^*)$, which implies that $\mathcal{D}$ reaches an ϵ-PNE$^+$ $\boldsymbol{\sigma}$ after at most $b\beta$ steps. Since at each step of the dynamics we have to check, for each of the n players, whether one of them possesses an ϵ-additive improving deviation, we get that we can compute $\boldsymbol{\sigma}$ in time $O(c_{apx} + nb\beta c_{move})$. Since $sub(\boldsymbol{\sigma}) \leq n\epsilon = nr\frac{\mathsf{SV}(\overline{\boldsymbol{\sigma}})}{\beta} \leq nr\frac{\mathsf{SV}(\boldsymbol{\sigma}^*)}{\beta}$, $\boldsymbol{\sigma}$ is an $\frac{nr}{\beta}$-SubEq. □

The result and the proof for cost-minimization games are similar to the previous ones.

Theorem 2.2. *Given an exact potential cost-minimization game $\mathcal{G}$ and a parameter $\beta > 0$, let $\overline{\boldsymbol{\sigma}}$ be an r-approximation of $\boldsymbol{\sigma}^*$ and b be a number such that $\Phi(\boldsymbol{\sigma}) \leq b\mathsf{SV}(\boldsymbol{\sigma})$ for any $\boldsymbol{\sigma} \in \Sigma$. An $\frac{nr}{\beta}$-SubEq can be computed in $O(c_{apx} + nbr\beta c_{move})$ time.*

Using Theorems 2.1 and 2.2, we can efficiently compute arbitrarily good SubEq for several classes of exact potential games for which computing a PNE is PLS-complete.

Corollary 2.3. *For any $\mu > 0$, a μ-SubEq can be computed in time polynomial in the game representation and in $1/\mu$ for undirected cut games, polynomial congestion games, cost-sharing games, and additively-separable hedonic games with symmetric preferences.*

3 Existence of Subsidized Equilibria

In this section, to show the power of SubEq, we focus on directed cut games, which, we recall, do not even admit ϵ-PNE* for any finite ϵ. We first show that a μ-SubEq cannot exist for $\mu < 1/2$, even in the unweighted case.

Theorem 3.1. *There exists an unweighted directed cut game admitting no μ-SubEq with $\mu < 1/2$.*

On the positive side, a 1-SubEq always exists.

Theorem 3.2. *Every directed cut game admits a 1-SubEq.*

Proof sketch. Consider a directed cut game $\mathcal{G}$ defined by graph $G = (N, A, w)$ and let $\boldsymbol{\sigma}^*$ be a social optimum for $\mathcal{G}$. We show that it is possible to subsidize $\boldsymbol{\sigma}^*$ by using a subsidy at most equal to $\mathsf{SV}(\boldsymbol{\sigma}^*)$, i.e., that $sub(\boldsymbol{\sigma}^*) \leq \mathsf{SV}(\boldsymbol{\sigma}^*)$, thus proving the claim. The key property is that, for any player i, $r_i(\boldsymbol{\sigma}^*)$ cannot be larger than the overall weight of all arcs starting from i and not belonging to $\mathrm{CUT}(\boldsymbol{\sigma}^*)$. Thus, summing $r_i(\boldsymbol{\sigma}^*)$ over all players gives at most $w(A) - \mathrm{CUT}(\boldsymbol{\sigma}^*)$, which, given that $\mathrm{CUT}(\boldsymbol{\sigma}^*) = \mathsf{SV}(\boldsymbol{\sigma}^*)$ and $w(A) - \mathrm{CUT}(\boldsymbol{\sigma}^*) \leq \mathsf{SV}(\boldsymbol{\sigma}^*)$ otherwise $\boldsymbol{\sigma}^*$ would not be the social optimum, yields the claim. □

For unweighted directed cut games, by exploiting a probabilistic-like method, we obtain a much stronger result.

In order to prove it, we need the following technical lemma.

Lemma 3.3. *For any integer $d \geq 2$, it holds that $\sum_{\ell=\lceil \frac{d}{2} \rceil}^{d} \left(\binom{d}{\ell} \cdot (2\ell - d)\right) \leq \frac{2^d \cdot d}{4}$.*

We are now ready to prove the following theorem.

Theorem 3.4. *Any unweighted directed cut game admits a $3/5$-SubEq.*

Proof sketch. In this sketch, we assume that, for any $i \in N$, $\delta_i \geq 2$. In the full proof, we carefully manage the case $\delta_i = 1$, by exploiting a suitable decomposition of the graph into components called gadgets. Actually, when $\delta_i \geq 2$ for any $i \in N$, it is possible to prove an even stronger and tight result, showing the existence of a $1/2$-SubEq.

Let $G = (N, A)$ be the directed unweighted graph defining $\mathcal{G}$. We exploit a lower bound to the social optimum that can be derived from a result in (Poljak and Turzík 1986) and is very similar to the Edwards-Erdös bound (Edwards 1975) holding for undirected graphs: there must exist an outcome $\boldsymbol{\sigma}^*$ for $\mathcal{G}$ such that $\mathsf{SV}(\boldsymbol{\sigma}^*) = |\mathrm{CUT}(\boldsymbol{\sigma}^*)| \geq \frac{m}{2} + \frac{n-1}{4}$.

To prove the existence of a $1/2$-SubEq $\boldsymbol{\sigma'}$, we consider the nodes one-by-one: since we have that $\mathsf{SV}(\boldsymbol{\sigma}^*) \geq \frac{m}{2} + \frac{n-1}{4} \geq \sum_{i=1}^{n} \frac{\delta_i}{2}$, in the following we will provide a vector of values $\alpha_1, \ldots, \alpha_n$ such that $\sum_{i=1}^{n} \alpha_i = sub(\boldsymbol{\sigma'})$ and, for any $i \in N$, it holds that $\alpha_i \leq \frac{\delta_i}{4}$, thus proving the theorem because $sub(\boldsymbol{\sigma'}) = \sum_{i=1}^{n} \alpha_i \leq \sum_{i=1}^{n} \frac{\delta_i}{4} \leq \frac{1}{2}\mathsf{SV}(\boldsymbol{\sigma}^*)$.

Our goal is to provide an upper bound to the average subsidy α required to stabilize an outcome, calculated over the set Σ containing all possible 2^n outcomes that can be obtained by the strategic choices of the players. For any $i \in N$, let α_i be the average subsidy of node i. Clearly, $\alpha = \sum_{i \in N} \alpha_i$ and there must exist an outcome $\boldsymbol{\sigma'}$ requiring a subsidy at most α. In order to calculate this average subsidy, we proceed by considering one node at a time: let i be any node. Consider any outcome $\boldsymbol{\sigma} \in \Sigma$ and let $\delta'_i(\boldsymbol{\sigma})$ be the number of nodes being adjacent to i and selecting the same strategy of i, i.e., $\delta'_i(\boldsymbol{\sigma}) = |\{j : (i,j) \in A \wedge \sigma_j = \sigma_i\}|$. Clearly, $r_i(\boldsymbol{\sigma}) = \max\{0, \delta'_i(\boldsymbol{\sigma}) - (\delta_i - \delta'_i(\boldsymbol{\sigma}))\}$, or equivalently $r_i(\boldsymbol{\sigma}) = 2\delta'_i(\boldsymbol{\sigma}) - \delta_i$ if $\delta'_i(\boldsymbol{\sigma}) > \frac{\delta_i}{2}$ and $r_i(\boldsymbol{\sigma}) = 0$ otherwise. For any integer $\ell = 0, 1, \ldots, \delta_i$, there are $\binom{\delta_i}{\ell} \cdot 2^{n-\delta_i}$ outcomes in Σ with $\delta'_i(\cdot) = \ell$, i.e., having ℓ nodes being adjacent to i and selecting the same strategy of i. We therefore obtain

$$\begin{aligned} \alpha_i &= \frac{\sum_{\ell=\lceil \frac{\delta_i}{2} \rceil}^{\delta_i} \left(\binom{\delta_i}{\ell} \cdot 2^{n-\delta_i} \cdot (2\ell - \delta_i) \right)}{2^n} \\ &= \frac{\sum_{\ell=\lceil \frac{\delta_i}{2} \rceil}^{\delta_i} \left(\binom{\delta_i}{\ell} \cdot (2\ell - \delta_i) \right)}{2^{\delta_i}} \leq \frac{\delta_i}{4}, \end{aligned}$$

where the last inequality holds by Lemma 3.3. □

The proof of Theorem 3.4 naturally yields a greedy algorithm for computing a $3/5$-SubEq in polynomial time. Roughly speaking, and always assuming that $\delta_i \geq 2$ for any $i \in N$, this algorithm proceeds in n steps, and at step i it decides the strategy for the i-th player by selecting the one for which the average subsidy required to stabilize an outcome (among all outcomes in which the strategies of the first $i-1$ players are already fixed) is minimized. This last calculation can be performed by adapting the calculations for $\alpha_1, \ldots, \alpha_n$ presented in the proof of Theorem 3.4 in order to deal with the fact that the strategies of some players have been fixed. When considering also nodes with out-degree equal to one, a similar approach can be exploited and the following theorem follows.

Theorem 3.5. *For unweighted directed cut games, a $3/5$-SubEq can be computed in polytime.*

4 Efficiency of Subsidized Equilibria

In this section, we provide bounds on the SubPoS for three widely-studied classes of games, namely, cost-sharing games, affine congestion games and directed cut games.

4.1 Cost-Sharing Games

We start by showing that the basic 2-upper bound on SubPoS coming from Proposition 1.3 is tight for a large fraction of cost-sharing games. With this respect, we stress that designing lower bounding instances for the SubPoS requires proving that the extended social welfare of *each strategy profile* is at least at a certain distance from the social optimum.

Theorem 4.1. *Let $\mathcal{C}$ be the class of cost-sharing games. We have $\mathsf{SubPoS}(\mathcal{C}) = 2$ and the bound is tight even for the subclass of directed broadcast games.*

The lower bound of Theorem 4.1 cannot be extended to undirected network cost-sharing games, for which better results may be possible. Towards this end, we consider the case of undirected broadcast games and prove a much better upper bound of approximately 1.368.

Theorem 4.2. *Let $\mathcal{B}^u$ be the class of undirected broadcast games. We have $\mathsf{SubPoS}(\mathcal{B}^u) \leq \frac{e+1}{e}$.*

Proof. Augustine et al. (2015) show that, by subsidizing edges rather than players, i.e., if money can be used to lower the cost of a set of edges, a social optimum $\boldsymbol{\sigma}^*$ can be made a PNE at a cost of at most $\mathsf{SV}(\boldsymbol{\sigma}^*)/e$. Let $s_{\tilde{e}}$ be the subsidy put on edge $\tilde{e}$ by their algorithm. Since $\boldsymbol{\sigma}^*$ is a PNE in the modified game in which the cost of an edge $\tilde{e}$ is redefined as $c_{\tilde{e}} - s_{\tilde{e}}$, it follows that $\sum_{\tilde{e} \in \sigma_i^*} \frac{c_{\tilde{e}} - s_{\tilde{e}}}{n_{\tilde{e}}(\boldsymbol{\sigma}^*)} \leq \sum_{\tilde{e} \in \tau} \frac{c_{\tilde{e}} - s_{\tilde{e}}}{n_{\tilde{e}}(\boldsymbol{\sigma}^*_{-i}, \tau)}$ for every $i \in N$ and $\tau \in \Sigma_i$. This implies that $\sum_{\tilde{e} \in \sigma_i^*} \frac{c_{\tilde{e}}}{n_{\tilde{e}}(\boldsymbol{\sigma}^*)} - \sum_{\tilde{e} \in \tau} \frac{c_{\tilde{e}}}{n_{\tilde{e}}(\boldsymbol{\sigma}^*_{-i}, \tau)} \leq \sum_{\tilde{e} \in \sigma_i^*} \frac{s_{\tilde{e}}}{n_{\tilde{e}}(\boldsymbol{\sigma}^*)}$ for every $i \in N$ and $\tau \in \Sigma_i$. Thus, $r_i(\boldsymbol{\sigma}^*) \leq \sum_{\tilde{e} \in \sigma_i^*} \frac{s_{\tilde{e}}}{n_{\tilde{e}}(\boldsymbol{\sigma}^*)}$ for each $i \in N$ and so $sub(\boldsymbol{\sigma}^*) \leq \sum_{i \in N} \sum_{\tilde{e} \in \sigma_i^*} \frac{s_{\tilde{e}}}{n_{\tilde{e}}(\boldsymbol{\sigma}^*)} = \sum_{\tilde{e} \in E : n_{\tilde{e}}(\boldsymbol{\sigma}^*) > 0} s_{\tilde{e}} \leq \mathsf{SV}(\boldsymbol{\sigma}^*)/e$. □

Determining a general, and possibly tight, lower bound for undirected broadcast games is an interesting and challenging open problem. By restricting to an instance with four players only, we are able to design a lower bound larger than 1.2. Despite the small number of players, this already requires analyzing a significant number of strategy profiles.

Theorem 4.3. *Let $\mathcal{B}^u$ be the class of undirected broadcast games. We have $\mathsf{SubPoS}(\mathcal{B}^u) > 6/5$.*

Towards a better understanding of the SubPoS of these games, we start by attacking the 2-player case which, despite its simplicity, already requires an involved case-analysis.

Theorem 4.4. *Let $\mathcal{B}_2^u$ be the class of undirected broadcast games with two players. We have $\mathsf{SubPoS}(\mathcal{B}_2^u) = 8/7$.*

Our results show a significantly better efficiency of SubEq with respect to PNE. In fact, $\mathsf{PoS}(\mathcal{C}) = H_n \geq \ln n$, with the lower bound being tight even for directed broadcast games (Anshelevich et al. 2008). Also for the undirected case, SubEq perform better, as $\mathsf{PoS}(\mathcal{B}^u) \geq 20/11$ (Bilò et al. 2013) and $\mathsf{PoS}(\mathcal{B}_2^u) = 4/3$ (Anshelevich et al. 2008).

4.2 Affine Congestion Games

We first show an upper bound of $3/2$ on the SubPoS by bounding the subsidy needed to stabilize a social optimum.

Theorem 4.5. *Let $\mathcal{C}$ be the class of affine congestion games. We have $\mathsf{SubPoS}(\mathcal{C}) \leq 3/2$.*

Proof. We show that $sub(\boldsymbol{\sigma}^*) \leq \mathsf{SV}(\boldsymbol{\sigma}^*)/2$, for a fixed social optimum $\boldsymbol{\sigma}^*$, by making use of a primal-dual approach, as in the spirit of (Bilò 2018; Bilò and Vinci 2023).

Fix an affine congestion game and a social optimum $\boldsymbol{\sigma}^*$. For each $i \in R(\boldsymbol{\sigma}^*)$, let σ_i be a best-response for i in $\boldsymbol{\sigma}^*$. Set $o_e := n_e(\boldsymbol{\sigma}^*)$. For each $i \in R(\boldsymbol{\sigma}^*)$, we have $r_i(\boldsymbol{\sigma}^*) = \sum_{e \in \sigma_i^* \setminus \sigma_i} (\alpha_e o_e + \beta_e) - \sum_{e \in \sigma_i \setminus \sigma_i^*} (\alpha_e(o_e + 1) + \beta_e)$. Set $h_e := |\{i \in R(\boldsymbol{\sigma}^*) : e \in \sigma_i^* \setminus \sigma_i\}|$ and $f_e := |\{i \in R(\boldsymbol{\sigma}^*) : e \in \sigma_i \setminus \sigma_i^*\}|$. Observe that $h_e \leq o_e$ holds for each $e \in E$ by definition. By summing the previous equality for all players $i \in R(\boldsymbol{\sigma}^*)$, we obtain

$$sub(\boldsymbol{\sigma}^*) = \sum_{e \in E} \left(\alpha_e(h_e o_e - f_e(o_e + 1)) + \beta_e(h_e - f_e) \right).$$

As $\boldsymbol{\sigma}^*$ is a social optimum, $\mathsf{SV}(\boldsymbol{\sigma}^*) - \mathsf{SV}(\boldsymbol{\sigma}^*_{-i}, \sigma_i) \leq 0$ for each $i \in R(\boldsymbol{\sigma}^*)$. By recalling that, for each $\boldsymbol{\sigma} \in$

Σ, $\mathsf{SV}(\boldsymbol{\sigma}) = \sum_{e\in E}\left(\alpha_e n_e(\boldsymbol{\sigma})^2 + \beta_e n_e(\boldsymbol{\sigma})\right)$, inequality $\mathsf{SV}(\boldsymbol{\sigma}^*) - \mathsf{SV}(\boldsymbol{\sigma}^*_{-i}, \sigma_i) \leq 0$ can be rewritten as

$$\begin{aligned} &\sum_{e\in\sigma_i^*\setminus\sigma_i}\left(\alpha_e(o_e^2-(o_e-1)^2)+\beta_e(o_e-(o_e-1))\right)\\ &\quad+\sum_{e\in\sigma_i\setminus\sigma_i^*}\left(\alpha_e(o_e^2-(o_e+1)^2)\right)\\ &\quad+\sum_{e\in\sigma_i\setminus\sigma_i^*}\left(\beta_e(o_e-(o_e+1))\right)\\ =&\sum_{e\in\sigma_i^*\setminus\sigma_i}\left(\alpha_e(2o_e-1)+\beta_e\right)\\ &\quad-\sum_{e\in\sigma_i\setminus\sigma_i^*}\left(\alpha_e(2o_e+1)+\beta_e\right)\leq 0. \end{aligned}$$

By summing this inequality for each $i \in R(\boldsymbol{\sigma}^*)$, we obtain

$$\sum_{e\in E}\left(\alpha_e(2h_eo_e-h_e-2f_eo_e-f_e)+\beta_e(h_e-f_e)\right)\leq 0. \tag{1}$$

Now we apply the primal-dual method. We want suitable values for the coefficients α_e and β_e for each $e \in E$, so that $sub(\boldsymbol{\sigma}^*)$ is maximized, constrained to the fact that $\mathsf{SV}(\boldsymbol{\sigma}^*)$ is normalized to 1 and equation (1) holds. Towards this end, we formulate this problem as a linear program whose variables are the coefficients α_e and β_e for each $e \in E$. The program, named LP, is the following:

$$\begin{aligned} \max\quad &\sum_{e\in E}\left(\alpha_e(h_eo_e-f_e(o_e+1))+\beta_e(h_e-f_e)\right)\\ s.t.\quad &\sum_{e\in E}\left(\alpha_e(2h_eo_e-h_e-2f_eo_e-f_e)\right)\\ &\quad+\sum_{e\in E}\left(\beta_e(h_e-f_e)\right)\leq 0\\ &\sum_{e\in E}(\alpha_eo_e^2+\beta_eo_e)=1\\ &\alpha_e,\beta_e\geq 0 \qquad \forall e\in E. \end{aligned}$$

The dual program of LP, obtained by associating variables y and z with the first two constraints of LP respectively, is:

$$\begin{aligned} \min\quad & z\\ s.t.\quad & y(2h_eo_e-h_e-2f_eo_e-f_e)\\ &\quad+zo_e^2\geq h_eo_e-f_e(o_e+1) && \forall e\in E\\ & y(h_e-f_e)+zo_e\geq h_e-f_e && \forall e\in E\\ & y\geq 0. \end{aligned}$$

Any feasible solution for this dual problem yields an upper bound to the optimal solution of LP and so an upper bound on $sub(\boldsymbol{\sigma}^*)$. We claim that the solution obtained by setting $y := 1/2$ and $z := 1/2$ is feasible for the dual program. By substituting in the first and second constraint, respectively, we obtain the inequalities $o_e^2+f_e-h_e \geq 0$ and $o_e+f_e-h_e \geq 0$ which are both true, as $h_e \leq o_e$. The last constraint is trivially satisfied. Thus, we have $sub(\boldsymbol{\sigma}^*) \leq \mathsf{SV}(\boldsymbol{\sigma}^*)/2$. □

We now consider the case in which all players have two available strategies only. Under this restriction, we show that the SubPoS is between $135/112 \approx 1.205$ and $10/7 \approx 1.429$. Clearly, the lower bound also extends to the case without any restriction.

Theorem 4.6. *Let $\mathcal{C}^{(2)}$ be the class of affine congestion games in which all players have at most two strategies each. Then,* $\mathsf{SubPoS}(\mathcal{C}^{(2)}) \in [135/112, 10/7]$.

Again, our results show a better efficiency of SubEq with respect to PNE, as $\mathsf{PoS}(\mathcal{C}) = 1 + 1/\sqrt{3} \approx 1.577$, with the lower bound being tight even for $\mathcal{C}^{(2)}$ (Caragiannis et al. 2011; Christodoulou and Koutsoupias 2005a).

4.3 Cut Games

By exploiting techniques similar to the ones used in the proofs of Theorems 3.1 and 3.4, it is possible to show that the SubPoS for unweighted directed cut games is constant.

Theorem 4.7. *For the class of unweighted directed cut games $\mathcal{C}$,* $\mathsf{SubPoS}(\mathcal{C}) \in [2, 9/2]$.

Recall that no finite upper bound on $\mathsf{PoS}^*_\epsilon(\mathcal{C})$ is possible as ϵ-PNE$^+(\mathcal{C}) = \emptyset$ for any finite ϵ.

5 Conclusions and Future Work

We introduced a new, simple and deterministic stability notion for non-cooperative games which, in our humble opinion, may open interesting research challenges in the field of Algorithmic Game Theory, as partially demonstrated by the variety and the technical depth of (some of) our results.

Besides closing the gaps of our non-tight bounds and providing results for other classes of games not considered in this paper, it would be nice to extend the efficient computation of a SubEq requiring almost negligible subsidy from exact potential games to general potential games and, more ambitiously, to even games not admitting a PNE. Also applying the idea of using subsidies to stabilize mixed strategy profiles is a worth-to-be-explored direction.

Furthermore, it may be interesting to study variants of SubEq in which the regret of a player is quantified with respect to notions of collective deviations. In cooperative game theory and computational social choice, for instance, there are notions of stability (e.g., core stability) based on the non-existence of *blocking coalitions*, which are coalitions of players that can improve their condition by collectively deviating. A SubEq in this setting would be a solution in which there is enough money to neutralize all blocking coalitions through an appropriate refund.

We conclude with a further motivation on how and why money can be invested in a game. A widely adopted practice in the realization of public infrastructures is *project financing*, according to which a public entity (e.g., the state) delegates the realization of a costly work (e.g., a railroad network) to private companies. These companies entirely or partially finance the realization of the work, but the ownership is transferred to the public entity. The companies, in return, retain the rights of either using the work or collecting the profits made from it (*concession deed*) for a certain time period. Thus, a public entity which cannot afford a certain work and wants to delegate it to a pool of strategic non-cooperative companies can obtain the realization of an economically efficient solution at the expenses of a fraction of its cost (for instance, by virtue of Theorem 4.2, in the case of an undirected broadcast game, the public entity obtains an optimal solution by investing only 37% of its cost).

Acknowledgments

Vittorio Bilò acknowledges support by the PNRR MIUR project FAIR - Future AI Research (PE00000013), Spoke 9 - Green-aware AI, WP 9.2 “Interactions of green-aware agents”, cascade sub-project “Green Aware Mechanisms” - GAMES, CUP H23C22000860006, and by the GNCS group of INdAM.

Gianpiero Monaco and Luca Moscardelli acknowledge support by the PNRR MIUR project FAIR - Future AI Research (PE00000013), Spoke 9 - Green-aware AI, WP 9.2 “Interactions of green-aware agents”, cascade sub-project “Existence, Complexity and Efficiency of Stable Solutions in Green-Oriented Games” – ECOGAMES, CUP D23C24000210006, and by the GNCS group of INdAM.

References

Ackermann, H.; Röglin, H.; and Vöcking, B. 2008. On the impact of combinatorial structure on congestion games. *Journal of the ACM*, 55(6): 25:1–25:22.

Aland, S.; Dumrauf, D.; Gairing, M.; Monien, B.; and Schoppmann, F. 2011. Exact Price of Anarchy for Polynomial Congestion Games. *SIAM Journal on Computing*, 40(5): 1211–1233.

Anshelevich, E.; Dasgupta, A.; Kleinberg, J. M.; Tardos, É.; Wexler, T.; and Roughgarden, T. 2008. The Price of Stability for Network Design with Fair Cost Allocation. *SIAM Journal on Computing*, 38(4): 1602–1623.

Augustine, J.; Caragiannis, I.; Fanelli, A.; and Kalaitzis, C. 2015. Enforcing Efficient Equilibria in Network Design Games via Subsidies. *Algorithmica*, 72(1): 44–82.

Awerbuch, B.; Azar, Y.; and Epstein, A. 2013. The Price of Routing Unsplittable Flow. *SIAM J. Comput.*, 32(1): 160–177.

Aziz, H.; Brandt, F.; and Seedig, H. G. 2013. Computing desirable partitions in additively separable hedonic games. *Artif. Intell.*, 195: 316–334.

Bhawalkar, K.; Gairing, M.; and Roughgarden, T. 2014. The Price of Anarchy, Universal Worst-Case Examples, and Tightness. *ACM Trans. Economics and Comput.*, 2(4): 14:1–14:23.

Bilò, V. 2018. A Unifying Tool for Bounding the Quality of Non-Cooperative Solutions in Weighted Congestion Games. *Theory of Computing Systems*, 62(5): 1288–1317.

Bilò, V.; Caragiannis, I.; Fanelli, A.; and Monaco, G. 2013. Improved Lower Bounds on the Price of Stability of Undirected Network Design Games. *Theory Comput. Syst.*, 52(4): 668–686.

Bilò, V.; Flammini, M.; and Moscardelli, L. 2020. The price of stability for undirected broadcast network design with fair cost allocation is constant. *Games Econ. Behav.*, 123: 359–376.

Bilò, V.; and Vinci, C. 2023. *Coping with Selfishness in Congestion Games - Analysis and Design via LP Duality*. Monographs in Theoretical Computer Science. An EATCS Series. Springer.

Bilò, V.; and Bove, R. 2011. Bounds on the Price of stability of Undirected Network Design Games with Three Players. *J. Interconnect. Networks*, 12(1–2): 1–17.

Bilò, V.; Flammini, M.; Monaco, G.; and Moscardelli, L. 2021. Computing approximate Nash equilibria in network congestion games with polynomially decreasing cost functions. *Distributed Computing*, 34(1): 1–14.

Bilò, V.; and Paladini, M. 2016. On the performance of mildly greedy players in cut games. *Journal of Combinatorial Optimization*, 32(4): 1036–1051.

Bilò, V.; and Vinci, C. 2017. On the Impact of Singleton Strategies in Congestion Games. In *Proceedings of the 25th Annual European Symposium on Algorithms (ESA)*, 17:1–17:14.

Bogomolnaia, A.; and Jackson, M. O. 2002. The Stability of Hedonic Coalition Structures. *Games and Economic Behavior*, 38(2): 201–230.

Brandt, F.; Bullinger, M.; and Tappe, L. 2024. Stability based on single-agent deviations in additively separable hedonic games. *Artif. Intell.*, 334: 104160.

Caragiannis, I.; Fanelli, A.; Gravin, N.; and Skopalik, A. 2015. Approximate Pure Nash Equilibria in Weighted Congestion Games: Existence, Efficient Computation, and Structure. *ACM Trans. Economics and Comput.*, 3(1): 2:1–2:32.

Caragiannis, I.; Flammini, M.; Kaklamanis, C.; Kanellopoulos, P.; and Moscardelli, L. 2011. Tight bounds for selfish and greedy load balancing. *Algorithmica*, 61(3): 606–637.

Caragiannis, I.; and Jiang, Z. 2023. Computing Better Approximate Pure Nash Equilibria in Cut Games via Semidefinite Programming. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing (STOC)*, 710–722. ACM.

Christodoulou, G.; Chung, C.; Ligett, K.; Pyrga, E.; and van Stee, R. 2009. On the Price of Stability for Undirected Network Design. In *Proceedings of the 7th International Workshop on Approximation and Online Algorithms (WAOA)*, 86–97.

Christodoulou, G.; and Gairing, M. 2016. Price of Stability in Polynomial Congestion Games. *ACM Trans. Economics and Comput.*, 4(2): 10:1–10:17.

Christodoulou, G.; and Koutsoupias, E. 2005a. On the price of anarchy and stability of correlated equilibria of linear congestion games. In *Proceedings of the 13th Annual European Symposium on Algorithms (ESA)*, 59–70.

Christodoulou, G.; and Koutsoupias, E. 2005b. The price of anarchy of finite congestion games. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC)*, 67–73.

Christodoulou, G.; Koutsoupias, E.; and Spirakis, P. G. 2011. On the Performance of Approximate Equilibria in Congestion Games. *Algorithmica*, 61(1): 116–140.

Christodoulou, G.; Mirrokni, V. S.; and Sidiropoulos, A. 2012. Convergence and approximation in potential games. *Theoretical Computer Science*, 438: 13–27.

D'Ascenzo, A.; D'Emidio, M.; Flammini, M.; and Monaco, G. 2024. Digraph k-Coloring Games: New Algorithms and Experiments. *J. Artif. Intell. Res.*, 81: 163–202.

Edwards, C. S. 1975. An improved lower bound for the number of edges in a largest bipartite subgraph. *Recent Advances in Graph Theory*, 167–181.

Fabrikant, A.; Papadimitriou, C. H.; and Talwar, K. 2004. The complexity of pure Nash equilibria. In *Proceedings of the 36th ACM Symposium on Theory of Computing (STOC)*, 604–612.

Fiat, A.; Kaplan, H.; Levy, M.; Olonetsky, S.; and Shabo, R. 2006. On the Price of Stability for Designing Undirected Networks with Fair Cost Allocations. In *Proceedings of the 33rd International Colloquium on Automata, Languages and Programming (ICALP)*, volume 4051 of *LNCS*, 608–618. Springer.

Gairing, M.; and Savani, R. 2019. Computing Stable Outcomes in Symmetric Additively Separable Hedonic Games. *Mathematics of Operations Research*, 44(3): 1101–1121.

Hoefer, M. 2007. *Cost sharing and clustering under distributed competition*. Ph.D. thesis, University of Konstanz.

Koutsoupias, E.; and Papadimitriou, C. H. 1999. Worst-case Equilibria. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, 404–413.

Kun, J.; Powers, B.; and Reyzin, L. 2013. Anti-coordination Games and Stable Graph Colorings. In *Proceedings of the 6th International Symposium on Algorithmic Game Theory (SAGT)*, volume 8146 of *LNCS*, 122–133. Springer.

Li, J. 2009. An $O(\log n/ \log\log n)$ upper bound on the price of stability for undirected Shapley network design games. *Information Processing Letters*, 109(15): 876–878.

Mamageishvili, A.; Mihalák, M.; and Montemezzani, S. 2018. Improved bounds on equilibria solutions in the network design game. *Int. J. Game Theory*, 47(4): 1113–1135.

Monderer, D.; and Shapley, L. S. 1996. Potential Games. *Games and Economic Behavior*, 14: 124–143.

Monderer, D.; and Tennenholtz, M. 2004. K-Implementation. *J. Artif. Intell. Res.*, 21: 37–62.

Peters, D.; and Elkind, E. 2015. Simple Causes of Complexity in Hedonic Games. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence, (IJCAI)*, 617–623.

Poljak, S.; and Turzík, D. 1986. A polynomial time heuristic for certain subgraph optimization problems with guaranteed worst case bound. *Discrete Mathematics*, 58(1): 99–104.

Rosenthal, R. W. 1973. A Class of Games Possessing Pure-Strategy Nash Equilibria. *International Journal of Game Theory*, 2: 65–67.

Schäffer, A. A.; and Yannakakis, M. 1991. Simple Local Search Problems That are Hard to Solve. *SIAM J. Comput.*, 20(1): 56–87.

Skopalik, A.; and Vöcking, B. 2008. Inapproximability of pure nash equilibria. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC)*, 355–364. ACM.

Syrgkanis, V. 2010. The complexity of equilibria in cost sharing games. In *Proceedings of the 6th Workshop on Internet and Network Economics (WINE)*, 366–377.

Vijayalakshmi, V. R.; and Skopalik, A. 2020. Improving Approximate Pure Nash Equilibria in Congestion Games. In *Proceedings of the 16th International Conference on Web and Internet Economics (WINE)*, 280–291.