

The Efficiency of Best-Response Dynamics

Michal Feldman¹, ✉ Yuval Snappir¹, and Tami Tamir²

¹ School of Computer Science, Tel-Aviv University, Israel.
yuval.snappir@gmail.com

² School of Computer Science, The Interdisciplinary Center, Herzliya, Israel.

Abstract. Best response (BR) dynamics is a natural method by which players proceed toward a pure Nash equilibrium via a local search method. The quality of the equilibrium reached may depend heavily on the order by which players are chosen to perform their best response moves. A *deviator rule* S is a method for selecting the next deviating player. We provide a measure for quantifying the performance of different deviator rules. The *inefficiency* of a deviator rule S with respect to an initial strategy profile p is the ratio between the social cost of the worst equilibrium reachable by S from p and the social cost of the best equilibrium reachable from p . The inefficiency of S is the maximum such ratio over all possible initial profiles. This inefficiency always lies between 1 and the *price of anarchy*.

We study the inefficiency of various deviator rules in network formation games and job scheduling games (both are congestion games, where BR dynamics always converges to a pure NE). For some classes of games, we compute optimal deviator rules. Furthermore, we define and study a new class of deviator rules, called *local* deviator rules. Such rules choose the next deviator as a function of a restricted set of parameters, and satisfy a natural independence condition called *independence of irrelevant players*. We present upper bounds on the inefficiency of some local deviator rules, and also show that for some classes of games, no local deviator rule can guarantee inefficiency lower than the price of anarchy.

Keywords: Congestion games; Best-response dynamics; Deviator rules; Price of anarchy

1 Introduction

Nash equilibrium (NE) is perhaps the most popular solution concept in games. It is a strategy profile from which no individual player can benefit by a unilateral deviation. However, a Nash equilibrium is a declarative notion, not an algorithmic one. To justify equilibrium analysis, we have to come up with a natural behavior model that leads the players of a game to a Nash equilibrium. Otherwise, the prediction that players play an equilibrium is highly questionable. Best response (BR) dynamics is a simple and natural method by which players proceed toward a NE via the following local search method: as long as the strategy profile is not a NE, an arbitrary player is chosen to improve her utility by deviating to her best strategy given the profile of others.

Work on BR dynamics advanced in two main avenues: The first studies whether BR dynamics converge to a NE, if one exists [21, 17]. The second explores how fast it takes until BR dynamics converges to a NE [11, 13, 25, 18]. It is well known that BR dynamics does not always converge to a NE, even if one exists. However, for the class of finite *potential games* [24, 22], a pure NE (PNE) always exists, and BR dynamics is guaranteed to converge to one of the equilibria of the game. A potential game is one that admits a *potential function* — a function that assigns a real value to every strategy profile, and has the miraculous property that for any unilateral deviation, the change in the utility of the deviating player is mirrored accurately in the potential function. This mirroring, combined with the fact that the game is finite, guarantees that any BR sequence must terminate and this happens at some (local) minimum of the potential function, which is a NE by definition. While BRD is guaranteed to converge, convergence may take an exponential number of iterations, even in a potential game [3].

Our focus in this work is different than the directions mentioned above. The description of BR dynamics leaves the choice of the deviating player unspecified. Thus, BR dynamics is essentially a large family of dynamics, differing from one another in the choice of who would be the next player to perform her best response move. In this paper, we study how the choice of the deviating player (henceforth a *deviator rule*) affects the efficiency of the equilibrium reached via BR dynamics. Our contribution is the following: (i) We introduce a new measure for quantifying the inefficiency of deviator rules, (ii) we introduce a natural class of simple and local deviator rules, and (iii) we analyze the inefficiency of deviator rules in network formation games and job scheduling games. Our results distinguish between games where local deviator rules can lead to good outcomes and games for which any local deviator rule performs poorly.

1.1 Model and Problem Statement

A game G has a set N of n players. Each player i has a strategy space P_i , and the player chooses a strategy $p_i \in P_i$. A strategy profile is a vector of strategies for each player, $p = (p_1, \dots, p_n)$. The strategy profile of all players except player i is denoted by p_{-i} , and it is convenient to denote a strategy profile p as $p = (p_i, p_{-i})$. Similarly, for a set of players I , we denote by p_I and p_{-I} the strategy profile of players in I and in $N \setminus I$, respectively, and we write $p = (p_I, p_{-I})$. Each player has a cost function $c_i : P \rightarrow \mathbb{R}^{\geq 0}$, where $c_i(p)$ denotes player i 's cost in the strategy profile p . Every player wishes to minimize her cost. There is also a social objective function, mapping each strategy profile to a social cost.

Given a strategy profile p , the best response of player i is the set of strategies that minimize player i 's cost, fixing the strategies of all other players, formally $BR_i(p) = \arg \min_{p'_i \in P_i} c_i(p'_i, p_{-i})$. Player i is said to be *suboptimal* in p if the player can reduce her cost by a unilateral deviation, i.e., if $p_i \notin BR_i(p)$. If no player is suboptimal in p , then p is a *Nash equilibrium* (NE) (in this paper we restrict attention to pure NE; i.e., an equilibrium in pure strategies).

Given an initial strategy profile p^0 , a best response (BR) sequence from p^0 is a sequence $\langle p^0, p^1, \dots \rangle$ in which for every $T = 0, 1, \dots$ there exists a player

$i \in N$ s.t. $p^{T+1} = (BR_i(p_{-i}^T), p_{-i}^T)$. In this paper we restrict attention to games in which every BR sequence is guaranteed to converge to a NE.

Deviator rules and their inefficiency. A *deviator rule* is a function $S : P \rightarrow N$ that given a profile p , chooses a deviator among all suboptimal players in p . The chosen player then performs a best response move (breaking ties arbitrarily). Given an initial strategy profile p^0 and a deviator rule S we denote by $NE_S(p^0)$ the set of NE that can be obtained as the final profile of a BR sequence $\langle p^0, p^1, \dots \rangle$, where for every $T \geq 0$, p^{T+1} is a profile resulting from a deviation of $S(p^T)$ (recall that players break ties arbitrarily, thus this is a set of possible Nash equilibria)

Given an initial profile p^0 , let $NE(p^0)$ be the set of Nash equilibria reachable from p^0 via a BR sequence, and let $p^*(p^0)$ be the best NE reachable from p^0 via a BR sequence, that is, $p^*(p^0) = \arg \min_{p \in NE(p^0)} SC(p)$, where $SC : P \rightarrow \mathbb{R}$ is some social cost function.

The *inefficiency* of a deviator rule S in a game G , denoted α_S^G , is defined as the worst ratio, among all initial profiles p^0 , and all NE in $NE_S(p^0)$, between the social cost of the worst NE reachable by S (from p^0) and the social cost of the best NE reachable from p^0 . I.e., $\alpha_S^G = \sup_{p^0} \max_{p \in NE_S(p^0)} \frac{SC(p)}{SC(p^*(p^0))}$.

For a class of games \mathcal{G} , the inefficiency of a deviator rule S with respect to \mathcal{G} is defined as the worst case inefficiency over all games in \mathcal{G} : $\alpha_S^{\mathcal{G}} = \sup_{G \in \mathcal{G}} \{\alpha_S^G\}$. A deviator rule with inefficiency 1 is said to be *optimal*, i.e., an optimal deviator rule is one that for every initial profile reaches a best equilibrium reachable from that initial profile.

The following observation shows that the inefficiency of every deviator rule is bounded from above by the *price of anarchy* (PoA) [20, 23]. Recall that the PoA is the ratio between the cost of the worst NE and the cost of the social optimum, and is used to quantify the loss incurred due to selfish behavior.

Observation 1 *For every game G and for every deviator rule S it holds that the inefficiency of S is at least 1 and bounded from above by the PoA.*

Local deviator rules. We define and study a class of simple deviator rules, called *local* deviator rules. Local deviator rules are defined with respect to state vectors, that represent the state of the players in a particular profile. Given a profile p , every player i is associated with a state vector v_i , consisting of several parameters that describe her state in p and in the strategy profile obtained by her best response. The specific parameters may vary from one application to another. A vector profile is a vector $\mathbf{v} = (v_1, \dots, v_n)$, consisting of the state vectors of all players. A deviator rule is said to be local if it satisfies the independence of irrelevant players condition, defined below.

Definition 1. *A deviator rule S satisfies independence of irrelevant players (IIP) if for every two state vectors v_{i_1}, v_{i_2} , and every two vector profiles \mathbf{v}, \mathbf{v}' such that $\mathbf{v} = (v_{i_1}, v_{i_2}, v_{-i_1, i_2})$, and $\mathbf{v}' = (v_{i_1}, v_{i_2}, v'_{-i_1, i_2})$ ³, if $S(\mathbf{v}) = i_1$, then $S(\mathbf{v}') \neq i_2$.*

³ Note that the vectors \mathbf{v} and \mathbf{v}' may correspond to different sets of players.

The IIP condition means that if the deviator rule chooses a state vector v_i over a state vector v_j in one profile, then, whenever these two state vectors exist, the deviator rule would not choose vector v_j over v_i . Note that this condition should hold even across different game instances and even when the number of players is different. Many natural deviator rules satisfy the IIP condition. For example, suppose that the state vector of a player contains her cost in the current profile and her cost in the profile obtained by her best response; then, both (i) max-cost, which chooses the player with the maximum current cost, and (ii) max-improvement, which chooses the player with the maximum improvement, are local deviator rules.

Congestion games. A congestion game has a set E of m resources, and the strategy space of every player i is a collection of sets of resources; i.e., $P_i \subseteq 2^E$. Every resource $e \in E$ has a cost function $f_e : \mathbb{N} \rightarrow \mathbb{R}$, where $f_e(\ell)$ is the cost of resource e if ℓ players use resource e . The cost of player i in a strategy profile p is $c_i(p) = \sum_{e \in p_i} f_e(\ell_e(p))$, where $\ell_e(p)$ is the number of players that use resource e in the profile p . Every congestion game is a potential game [22], thus admits a pure NE, and moreover, every BR sequence converges to a pure NE. In this paper we study the efficiency of deviator rules in the following congestion games:

Network formation games [3]: There is an underlying graph, and every player is associated with a pair of source and target nodes s_i, t_i . The strategy space of every player i is the set of paths from s_i to t_i . The resources are the edges of the graph, every edge e is associated with some fixed cost c_e , which is evenly distributed by the players using it. That is, the cost of an edge e in a profile p is $f_e(p) = c_e/\ell(p)$. In network formation games the cost of a resource decreases in the number of players using it. We also consider a weighted version of network formation games on parallel edge networks, where players have weights and the cost of an edge is shared proportionally by its users. The social cost function here is the sum of the players' costs; that is $SC(p) = \sum_{i \in N} c_i(p)$.

The state vector of a player in a network formation game, in a profile p , consists of: (1) player i 's cost in p : $c_i(p)$, (2) the cost of player i 's path: $\sum_{e \in p_i} c_e$, (3) player i 's cost in the profile obtained from a best response of i : $c_i(p'(i))$ (where $p'(i) = (p_{-i}, BR_i(p_{-i}))$ is the profile obtained from a best response of i), and (4) the cost of player i 's path in the profile $p'(i)$: $\sum_{e \in BR_i(p_{-i})} c_e$. In weighted instances, the state vector includes player i 's weight as well.

Job scheduling games [26]: The resources are machines, and players are jobs that need to be processed on one of the machines. Each job has some length, and the strategy space of every player is the set of the machines. The load on a machine in a strategy profile p is the total length of the jobs assigned to it. The cost of a job is the load on its chosen machine. We also consider games with conflicting congestion effect ([14, 7]), where jobs have unit length and in addition to the cost associated with the load, every machine has an activation cost B , shared by the jobs assigned to it. The social cost function here is the *makespan*, that is $SC(p) = \max_{i \in N} c_i(p)$. The state vector of a job (player) in a job scheduling game, in a profile p , consists of the job's length, the job's current machine and the loads on the machines.

1.2 Our Results

In Section 2 we present our results for network formation games. We first study symmetric games, where all the players share the same source and target nodes. We observe that the local *Min-Path* deviator rule, which chooses the player with the cheapest best response path, is optimal. In contrast, the local *Max-Cost* deviator rule has the worst possible inefficiency, n (which matches the PoA for this game). We then consider asymmetric network formation games. Unfortunately, the optimality of Min-Path does not carry over to asymmetric network formation games, even when played on series of parallel paths (SPP) networks. In particular, the inefficiency of Min-Path in single-source multi-target instances is $\theta(|V|)$, and for multi-source multi-target instances, it further grows to $\theta(2^{|V|})$. On the positive side, we show $\text{poly}(n, |V|)$ dynamic-programming algorithms for finding an optimal BR sequence for network formation games played on SPP networks, for single-source multi-target instances, and for multi-source multi-target instances with proper intervals (i.e., where no player's strategy is a subset of another player's strategy). The specification of these algorithms is deferred to the full version due to space constraints. For network formation games played on extension-parallel networks we show that every local deviator rule has an inefficiency of $\Omega(n)$.

In Section 2.4 we study network formation games with weighted players. It turns out that weighted players lead to quite negative results. We show that even in the simplest case of parallel-edge networks, it is NP-hard to find an optimal BR-sequence, and no local deviator rule can ensure a constant inefficiency. Moreover, the Min-Path deviator rule has inefficiency $\Omega(n)$, even in symmetric games on series-parallel graphs, and even if the ratio between the maximal and minimal weights approaches 1.

The analysis of job scheduling games is deferred to the full version. A job's (= player's) state vector in job scheduling games includes the job's lengths and the machines' loads. Local deviator rules capture many natural rules, such as Longest-Job, Max-Cost, Max-Improvement, and more. We show that in an instance with m identical machines, no local deviator rule can guarantee inefficiency better than the *PoA*, which is $\frac{2m}{m+1}$. In contrast, for job scheduling games with conflicting congestion effects [14], we present an optimal local deviator rule.

Positive results on local deviator rules imply that a centralized authority that can control the order of deviations can lead the population to a good outcome, by considering merely local information captured in the close neighborhood of the current state. In contrast, negative results for local deviator rules imply that even if a centralized authority can control the order of deviations, in order to converge to a good outcome, it cannot rely only on local information; rather, it must be able to perform complex calculations and to consider a large search space.

1.3 Related work

Congestion games have been widely studied from a game theoretic perspective. The questions that are most commonly analyzed are the existence of a pure NE,

the convergence of BRD to a NE, and the loss incurred due to selfish behavior – commonly quantified by the *price of anarchy* [20, 23] and *price of stability* [3].

Our work addresses congestion games and some variants thereof. It is well known that every congestion game is a potential game [24, 22] and therefore admits a PNE and possesses the finite improvement property (FIP). In particular, every BRD converges to a PNE. However, the convergence time may, in general, be exponentially long. It is shown in [3, 13] that finding a PNE in network formation games is PLS-complete. Examples for exponential convergence of job scheduling games are presented in [10]. It has been shown in [9] that in random potential games with n players in which every player has at most a strategies, the worst case convergence time is $n \cdot a^{n-1}$.

The observation that the convergence of BRD can be exponentially long has led to a large amount of work aiming to identify special classes of congestion games for which BRD converges to a PNE in polynomial time (or even linear time). Examples include [3] for games with positive congestion effects, and [10, 16] for games with negative congestion effects. For resource selection games (i.e., where feasible strategies are composed of singletons), polynomial convergence has been proven in [18].

Variants of congestion games such as weighted network formation games and resource selection games with player-specific cost functions have been also considered in the literature [17, 21]. Some classes of cost functions that always admit PNE or the FIP were identified in [17]. It was also shown that singleton weighted congestion games that always admit a PNE do not always have the FIP. [4, 12] present variants of network formation games in which computing a player’s best response is NP-hard. A cost-sharing game on unrelated machines has been studied in [5], where it was shown that a PNE exists only for instances with unit-cost machines. Moreover, even when a PNE exists and BRD is guaranteed to converge, the implementation of BRD can be computationally hard.

BRD has been studied also in games that do not converge to a PNE. The notion of *dynamic inefficiency* was defined in [6] as the average social cost in a BR infinite sequence (for games that do not possess the finite improvement property), and different deviator rules are analyzed with respect to the dynamic inefficiency measure.

The effect of the deviator rule on the convergence time of job scheduling games was studied in [10]. This paper considered the convergence time under the Max-Weight-Job, Min-Weight-Job, FIFO and random deviator rules. The Max-Cost deviator rule was considered also in [15] for conflicting congestion games [14, 7], and in [19] for swap-games [2]. In both cases Max-Cost significantly improves convergence time to $\mathcal{O}(n)$.

2 Network Formation Games

In this section we study network formation games. We consider two natural local deviator rules, namely *Max-Cost* and *Min-Path*. The Max-Cost deviator rule chooses a suboptimal player that currently incurs the highest cost, i.e., $Max - Cost(p) \in \arg \max_{\{i \in N \mid p_i \notin BR_i(p)\}} c_i(p)$. The Min-Path deviator rule chooses a

suboptimal player whose path in the profile obtained from a best response move is cheapest, i.e., $Min - Path(p) \in \arg \min_{\{i \in N | p_i \notin BR_i(p)\}} \sum_{e \in BR_i(p)} c_e$.

2.1 Warmup: Symmetric Network Formation Games

A network-formation game is *symmetric* if all the players have the same source and target nodes. Recall that the inefficiency of any deviator rule is upper bounded by the price of anarchy (PoA) of the game. It is well known that the PoA of network formation games is n .

We first show that the Max-Cost rule may perform as poorly as the PoA, even in symmetric games on parallel-edge networks.

Observation 2 *The inefficiency of Max-Cost in symmetric network formation games on parallel-edge networks is n .*

On the other hand, we show that Min-Path is an optimal deviator rule, i.e., it always reaches the best NE reachable from any initial profile. Our analysis of Min-Path is based on the following Lemma:

Lemma 1. *In symmetric network formation games, the path chosen by the first deviator is the unique path that will be chosen by all subsequent players, regardless of the order in which they deviate.*

Lemma 1 directly implies the optimality of Min-Path:

Theorem 1. *Min-Path is an optimal deviator rule for symmetric network formation games.*

Proof. By Lemma 1 the first deviation dictates the NE to be reached. Thus, the set of BR paths in p^0 are the set of reachable NE. Clearly choosing the cheapest one among them is optimal. \square

2.2 Series of Parallel Paths (SPP) Networks

In this section we study NFGs played on SPP networks. An SPP network consists of m segments, where each segment is a parallel-edge network. Let $\{u_0, \dots, u_m\}$ denote the vertex set, and for every $j \leq m$, let E_j denote the set of edges in segment j (i.e., the parallel edges connecting u_{j-1} and u_j). For a player i , let $E(i) = \cup_{s_i < k \leq t_i} E_k$, denote the set of edges player i may choose.

Note that in an SPP network, a player's choice of an edge in E_j is independent of any other segment in her path. This implies that a NFG on an SPP network consists of a sequence of symmetric games, where the set of players participating in each game varies. Combining this observation with Lemma 1 implies:

Lemma 2. *In every network formation game played on an SPP network with m segments, for every $1 \leq j \leq m$, and every BR sequence, let i be the first player in the sequence such that $E_j \in E(i)$, and let e be the edge in E_j chosen by player i . Then e is the unique edge in E_j players deviate to.*

Based on the above lemma, it is possible to develop polynomial-time algorithms, based on dynamic programming, for finding optimal BR sequences for SPP networks, for both single-source multi-targets games and multi-source multi-target games with proper intervals. Due to space constraints, the algorithms are omitted.

The performance of Min-Path in SPP networks. Recall that Min-Path was shown to be optimal for symmetric NFGs. We now analyze its inefficiency for SPP networks. Given an SPP network and a BR-sequence, we say that a segment is *unresolved* if there are at least two players whose intervals include the segment, and each of them will select a different edge in the segment if chosen to perform a BR next. The other segments are denoted *resolved*. By Lemma 2, after player i performs her best response, all the segments in her interval are resolved. Thus, no player migrates more than once. The reachable NEs have the same edges in the resolved segments and can only differ in unresolved segments. Therefore, the migrations of players who use only resolved segments does not influence the reachable NEs and in the following analyses we ignore them. Thus, all the deviations we consider resolve at least one segment. We denote by R_i the resolved segments after i such deviations. Let OPT denote the minimal cost of a NE reachable from p^0 by some BR sequence. Formally, $OPT = SC(p^*(p_0))$.

Lemma 3. *For any BR sequence of an SPP network instance, as long as there are unresolved segments, there exists a suboptimal player whose interval includes unresolved segments, and if this player is chosen next, then the cost of the unresolved segments she would set is at most OPT .*

Proof. Let p be an intermediate strategy profile in the BR sequence. Consider the players according to the order they deviate in some optimal BR sequence. Let i' be the first player in this order who is suboptimal in p . Since no player prior to i' in the optimal sequence is suboptimal, the segments that i' would resolve by a deviation from p are a subset of the segments she resolves in the optimal sequence. In the optimal sequence she obviously resolves these segments such that the selected edges are of total cost at most OPT , and therefore this is an upper bound on the total cost of unresolved segments she would set by deviating from p . \square

Using the above lemma, we provide tight analysis on the performance of Min-Path for SPPs with multi-targets and single or multiple sources. Note that in a single-source instance, every player resolves the prefix of the network corresponding to her interval.

Theorem 2. *The inefficiency of Min-Path in SPP NFGs with single-source and multi-targets is $\theta(m)$.*

Proof. We show that the total cost determined for the segments resolved in every iteration is at most OPT . Since at least one segment is resolved in each iteration, the whole network's cost is bounded by $m \cdot OPT$. Let i be the i -th player chosen to deviate by Min-Path and assume i has unresolved segments. Let i' be the player guaranteed by Lemma 3. It may be that $i = i'$. Both players have the same BR path in the resolved segments and therefore differ only in their unresolved segments. Since Min-Path chose i , the cost of her unresolved segments is at most the cost of i' 's unresolved segments, which is at most OPT by Lemma 3.




Fig. 1. A network on which Min-Path has inefficiency $\Omega(m)$

We show that the analysis is tight: Consider the network depicted in Figure 1. There are $n = m$ players, where t_i is the target of player i . In the initial strategy profile for every $1 \leq i \leq n$, $p_i^0 = \langle e_1, e_2, \dots, e_{i-1}, e'_i \rangle$. Note that in every segment, E_i , connecting t_{i-1} and t_i , the upper edge costs $n - i$ and is used by the $n - i$ players $i + 1, \dots, n$ and the lower edge costs $1 + \epsilon$ and is used only by player i .

In every segment, the players using the upper edge will benefit from deviating to the lower one and the player using the lower edge will benefit from deviating to the upper one. By Lemma 2, the first deviation will determine the edge that will be used in the NE reached. Therefore, a first deviation of player n to $\langle e'_1, \dots, e'_n \rangle$ will result in the NE corresponding to that path and has a social cost $n \cdot (1 + \epsilon)$.

Player i 's BR path's cost is $\sum_{1 \leq t \leq i-1} (1 + \epsilon) + (n - i) = (i - 1) \cdot (1 + \epsilon) + (n - i) = (n - 1) + (i - 1) \cdot \epsilon$. Therefore, Min-Path chooses Player 1 to deviate first. After her deviation all the rest of the players would use e_1 in their BR and treating t_1 as source shows that the next Player to deviate will be Player 2 and then Player 3 etc. Min-Path's BR sequence's NE will consist of $\langle e_1, \dots, e_{n-1}, e'_n \rangle$ and therefore its inefficiency for this strategy profile is $\frac{(n-1)\frac{n}{2} + 1 + \epsilon}{n(1 + \epsilon)} \xrightarrow{\epsilon \rightarrow 0} \frac{(n-1)\frac{n}{2} + 1}{n} \approx \frac{n}{2}$. Since $n = m$, we conclude that the inefficiency of Min-Path in SPP networks with single-source and multi-targets is $\theta(m)$. \square

We next show that Min-Path performs poorly on more general instances.

Theorem 3. *The inefficiency of Min-Path in SPP network formation games with multi-sources and multi-targets is $\theta(2^m)$.*

Proof. Let $c(R_i)$ denote the total cost of resolved segments after i deviations of players whose deviation resolved at least one segment. Since the initially resolved segments has to be included in the NE reached, it holds that $c(R_0) \leq OPT$. We prove that $c(R_i) - c(R_{i-1}) \leq c(R_{i-1}) + OPT$ for every i ; i.e., $c(R_i) \leq 2c(R_{i-1}) + OPT$. This implies that the total network's cost is $c(R_m) \leq 2^{m+1} \cdot OPT$.

Let i be the i^{th} player chosen to deviate by Min-Path that has some unresolved segments. The total cost of the unresolved segments that i resolves is $c(R_i) - c(R_{i-1})$. Let i' be a player guaranteed by Lemma 3. Since i was chosen by Min-Path, the cost of i 's BR path is lower than the cost of i' 's BR path. But the cost of i 's BR path is at least the cost of the unresolved segments in

i 's BR path. On the other hand, the cost of i 's BR path equals the sum of the cost of her resolved segments, which is bounded by $c(R_{i-1})$ and the cost she would set to her unresolved segments, bounded by OPT (by Lemma 3). Putting it all together, we get $c(R_i) - c(R_{i-1}) \leq c(R_{i-1}) + OPT$, as required. In the full version, we present a matching lower bound. \square

2.3 Local Rules for Extension Parallel (EP) Graphs

We now show that the inefficiency of any local deviator rule is $\Omega(n)$, even in the restricted class of EP networks. Recall that the state vector of a player consists of the player's cost in her current profile and in the profile obtained by a deviation of the player, and the total cost of the path used by the player in the two profiles.

Theorem 4. *For the class of single-source network formation games played on extension-parallel networks, the inefficiency of every local deviator rule is $\Omega(n)$.*

Proof. Consider the network depicted in Figure 2(a). There are n players, all sharing the source s , and the targets are as depicted in the figure. Consider the following profile: (i) Player 1 uses the path $\langle e_1 \rangle$. (ii) Player 2 uses the path $\langle e_1, e_2 \rangle$. (iii) Players 3, 4 use the path $\langle e_3 \rangle$. (iv) Players 5 to n use the path $\langle e_5 \rangle$.




Fig. 2. A local deviator rule fails (a) if v_2 is preferred, and (b) if v_3 is preferred. Every edge is labelled by the edge cost and the number of players using it in the initial strategy profile. E.g., the edge e_1 costs 24 and is used by 2 players in the initial strategy profile.

Recall that the state vector of player i consists of her current cost, the total cost of her current path, her post-deviation cost, and the total cost of her post-deviation path. Consider player 2, who uses the path $\langle e_1, e_2 \rangle$. Her current cost is 22 (she shares the cost of edge e_1 with player 1 and pays fully for edge e_2), the total cost of her path is 34, her post-deviation cost is 10 (obtained by deviating to e_3 , and sharing this cost with players 3, 4), and the total cost of her post-deviation path is 30. Thus, the state vector of player 2 is $v_2 = (22, 34, 10, 30)$. Similarly, one can verify that the state vector of player 3 (or layer 4) is $v_3 = (15, 30, 13, 34)$ (obtained by deviating to the path $\langle e_1, e_2 \rangle$). The suboptimal players in this profile are players 2 and 3 (or 4). If the deviator rule chooses the state vector v_2 over v_3 , then player 2 will deviate to e_3 , reaching a NE whose social cost is $54 + 7.4(n-4)$. On the other hand, if the deviator rule chooses the state vector v_3 over v_2 , then player 3 will deviate to $\langle e_1, e_2 \rangle$, and from this point on all players will deviate to $\langle e_1, e_2 \rangle$, reaching a NE whose social

cost is 34. We conclude that a deviator rule that prefers v_2 over v_3 reaches an inefficiency of $\frac{54+7.4(n-4)}{34} = \Omega(n)$.

Consider next the network depicted in Figure 2(b). There are n players, all sharing the source s , with targets as depicted in the figure. Consider the following profile: (i) Player 1 uses the path $\langle e_1, e_2 \rangle$. (ii) Player 2 uses the path $\langle e_1, e_6 \rangle$. (iii) Players 3, 4 use the path $\langle e_3 \rangle$. (iv) Players 5 to n use the path $\langle e_5 \rangle$.

One can verify in a similar analysis to the one showed for the game in 2(a) that if a deviator rule prefer v_3 over v_2 then it has inefficiency of $\frac{34+6.5(n-4)}{30} = \Omega(n)$. We conclude that any local deviator rule has inefficiency of $\Omega(n)$. \square

2.4 Weighted Symmetric Network Formation Games

In this section we consider network formation games with weighted players [1, 17], where every player is associated with a cost w_i . If an edge of cost c_e is shared by k players with weights w_1, w_2, \dots, w_k , then player i pays $\frac{w_i}{\sum_{j=1}^k w_j} \cdot c_e$.

For weighted NFGs on parallel-edge graphs, we prove the following.

Theorem 5. *In weighted network formation games on parallel edge networks: (a) it is NP-hard to calculate the social cost of an optimal reachable NE from a given profile; (b) finding a reachable NE that approximates the social optimum by factor $\frac{3}{2}$ is NP-hard; (c) any local deviator rule has inefficiency $\Omega(\sqrt{n})$.*

A direct corollary of the last theorem (part (a)) is that the problem of finding an optimal BR sequence is NP-hard.

Weighted symmetric network formation games with strategies consisting of two resources (i.e., a 2-segment SPP) are potential games [3]. Theorem 1 shows that in unweighted symmetric games, the Min-Path deviator rule ensures convergence to the optimal reachable NE. Here we show that in weighted games the efficiency of Min-Path can be as poor as the PoA, even if the weights are arbitrarily close to each other, and the strategies are sets of two resources.

Theorem 6. *In weighted network formation games on a 2-segment SPP, the Min-Path deviator rule has inefficiency $\Omega(n)$. This bound is valid even if the ratio $\max_i w_i / \min_i w_i$ is arbitrarily close to 1.*

References

1. H. Ackermann, H. Röglin, B. Vöcking. Pure Nash equilibria in player-specific and weighted congestion games. *Theor. Comput. Sci.*, 410(17):1552–1563, 2009.
2. N. Alon, E.D. Demaine, M. Hajiaghayi, T. Leighton. Basic Network Creation Games. *Proc. of the 22nd ACM Symposium on Parallelism in Algorithms and Architectures*, pages 106–113, 2010.
3. E. Anshelevich, A. Dasgupta, J. Kleinberg, E. Tardos, T. Wexler, T. Roughgarden. The price of stability for network design with fair cost allocation. *SIAM J. Comput.*, 38(4):1602–1623, 2008.
4. G. Avni, O. Kupferman, and T. Tamir. Network-formation games with regular objectives. *J. of Information and Computation*, 251:165–178, 2016.

5. G. Avni and T. Tamir, Cost-Sharing Scheduling Games on Restricted Unrelated Machines. *Theoretical Computer Science*, 646:26-39, 2016.
6. N. Berger, M. Feldman, O. Neiman, M. Rosenthal Dynamic Inefficiency: Anarchy without Stability In *Proc. of SAGT*, 2011.
7. B. Chen and S. Gürel. Efficiency analysis of load balancing games with and without activation costs. *Journal of Scheduling*, 15(2), pp. 157–164, 2011.
8. H. Chen and T. Roughgarden. Network Design with Weighted Players, *Theory of Computing Systems*, 45(2), 302–324, 2009.
9. S. Durand and B. Gaujal. Complexity and Optimality of the Best Response Algorithm in Random Potential Games. In *Proc. of SAGT*, 2016.
10. E. Even-Dar, A. Kesselman, Y. Mansour. Convergence time to nash equilibria. In *Proc of ICALP*, pp. 502–513, 2003.
11. E. Even-Dar and Y. Mansour. Fast Convergence of Selfish Rerouting. In *Proc. of SODA*, pp. 772–781, 2005.
12. A. Fabrikant, A. Luthra, E. Maneva, C. Papadimitriou, and S. Shenker. On a network creation game. In *Proc. of PODC*, pages 347-351, 2003.
13. A. Fabrikant, C. Papadimitriou, K. Talwar. The complexity of pure Nash equilibria. In *Proc. of STOC*, pp. 604–612, 2004.
14. M. Feldman and T. Tamir. Conflicting Congestion Effects in Resource Allocation Games. *Journal of Operation Research*. 60(3):529–540, 2012.
15. M. Feldman and T. Tamir. Convergence of Best-Response Dynamics in Games with Conflicting Congestion Effects. *Information Processing Letters* 115(2):112–118, 2015.
16. D. Fotakis. Congestion games with linearly independent paths: convergence time and price of anarchy. *Theory Comput. Syst.*, 47(1):113–136, 2010.
17. T. Harks and M.Klimm. On the Existence of Pure Nash Equilibria in Weighted Congestion Games *Mathematics of Operations Research*, 37(3):419–436, 2012.
18. S. Ieong, R.McGrew, E. Nudelman, Y. Shoham, Q. Sun, Fast and Compact: A Simple Class of Congestion Games. In *Proc. of AAAI*, pp. 489–494, 2005.
19. B. Kawald, P. Lenzner. 2013. On Dynamics in Selfish Network Creation. In *Proc. of SPAA*, pp. 83-92, 2013.
20. E. Koutsoupias and C. Papadimitriou. Worst-case Equilibria. *Computer Science Review*, 3(2): 65-69, 1999.
21. I. Milchtaich. Congestion games with player specific payoff functions. *Games and Economic Behavior*, 13:111-124, 1996.
22. D. Monderer and L. S. Shapley. Potential Games. *Games and Economic Behavior*, 14: 124–143, 1996.
23. C. H. Papadimitriou. Algorithms, games, and the internet. In *Proc. 33rd STOC*, pp. 749–753, 2001.
24. R.W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2:65–67, 1973.
25. V. Syrgkanis. The complexity of equilibria in cost sharing games. In *Proc. 6th WINE*, pp. 366–377. 2010.
26. B. Vöcking. Chapter 20: Selfish load balancing. In *Algorithmic Game Theory*, Cambridge University Press, 2007.