

The Impact of Worst-Case Deviations in Non-Atomic Network Routing Games

Pieter Kleer¹ · Guido Schäfer^{1,2} 

Published online: 26 December 2017

© Springer Science+Business Media, LLC, part of Springer Nature 2017

Abstract We introduce a unifying model to study the impact of worst-case latency deviations in non-atomic selfish routing games. In our model, latencies are subject to (bounded) deviations which are taken into account by the players. The quality deterioration caused by such deviations is assessed by the *Deviation Ratio*, i.e., the worst case ratio of the cost of a Nash flow with respect to deviated latencies and the cost of a Nash flow with respect to the unaltered latencies. This notion is inspired by the *Price of Risk Aversion* recently studied by Nikolova and Stier-Moses (Nikolova and Stier-Moses 2015). Here we generalize their model and results. In particular, we derive tight bounds on the Deviation Ratio for multi-commodity instances with a common source and arbitrary non-negative and non-decreasing latency functions. These bounds exhibit a linear dependency on the size of the network (besides other parameters). In contrast, we show that for general multi-commodity networks an

This article is part of the Topical Collection on *Special Issue on Algorithmic Game Theory (SGT 2016)*

Pieter Kleer is supported by the NWO Gravitation Project NETWORKS, Grant Number 024.002.003.

✉ Guido Schäfer
g.schaefer@cwi.nl

Pieter Kleer
kleer@cwi.nl

¹ Centrum Wiskunde & Informatica (CWI), Networks and Optimization Group, Science Park 123, 1098 XG Amsterdam, The Netherlands

² School of Business and Economics, Vrije Universiteit Amsterdam, De Boelelaan 1105, 1081 HV Amsterdam, The Netherlands

exponential dependency is inevitable. We also improve recent smoothness results to bound the Price of Risk Aversion.

Keywords Selfish routing · Perturbations · Deviation ratio · Price of risk aversion · Biased price of anarchy

1 Introduction

In the classical selfish routing game introduced by Wardrop [20], there is an (infinitely) large population of (non-atomic) players who selfishly choose minimum latency paths in a network with flow-dependent latency functions. An assumption that is made in this model is that the latency functions are “exact” in the sense that they represent the actual delays perceived by the players. Although being a meaningful abstraction, which also facilitates the analysis of such games, this assumption is overly simplistic in situations where latencies are subject to deviations (or perturbations). For example, such deviations might be due to fluctuations in travel times, latency uncertainties of players, approximate estimates of latencies, etc.

In this paper, we study how much the quality of a Nash flow deteriorates in the worst case under bounded deviations of the latency functions. More precisely, given an instance of the selfish routing game with latency functions $(l_a)_{a \in A}$ on the arcs, we define the *Deviation Ratio (DR)* as the worst case ratio $C(f^\delta)/C(f^0)$ of the social cost of a Nash flow f^δ with respect to deviated latency functions $(l_a + \delta_a)_{a \in A}$, where $(\delta_a)_{a \in A}$ are arbitrary deviation functions from a feasible set, and the social cost of a Nash flow f^0 with respect to the unaltered latency functions $(l_a)_{a \in A}$.

Here the social cost function C refers to the total average latency (without the deviations). Our motivation for studying this social cost function is that a central designer usually cares about the long-term performance of the system accounting for the average latency (or pollution). On the other hand, the players typically do not know the exact latencies and use estimates or include “safety margins” in their planning. Similar viewpoints are adopted in [13, 16].

In order to model bounded deviations, we extend an idea previously put forward by Bonifaci, Salek and Schäfer [2] in the context of the *restricted network toll problem*: We assume that for every arc $a \in A$ we are given (flow-dependent) lower and upper bound restrictions θ_a^{\min} and θ_a^{\max} , respectively, and call a deviation δ_a *feasible* if $\theta_a^{\min}(x) \leq \delta_a(x) \leq \theta_a^{\max}(x)$ for all $x \geq 0$.

Our notion of the Deviation Ratio is inspired by and builds upon the *Price of Risk Aversion (PRA)* recently introduced by Nikolova and Stier-Moses [16]. The authors investigate selfish routing games with uncertain latencies by considering deviations of the form $\delta_a = \gamma v_a$, where $\gamma \geq 0$ is the risk-aversion of the players and v_a is the variance of some random variable with mean zero. They derive upper bounds on the Price of Risk Aversion for single-commodity networks with arbitrary non-negative and non-decreasing latency functions if the *variance-to-mean-ratio* v_a/l_a of every arc $a \in A$ is bounded by some constant $\kappa \geq 0$. It is not hard to see that their model is a special case of our model if we choose $\theta_a^{\min} = 0$ and $\theta_a^{\max} = \gamma \kappa l_a$ (see Section 2 for more details).

1.1 Our Contributions

Our contributions presented in this paper are as follows:

1. *Upper bounds on the Deviation Ratio:* We derive a general upper bound on the Deviation Ratio for multi-commodity networks with a common source and arbitrary non-negative and non-decreasing latency functions (Theorem 1). Basically, we show that the social cost of a Nash flow f^δ with respect to feasible deviations $(\delta_a)_{a \in A}$ is at most the social cost of a Nash flow f^0 plus a term that depends on the lower and upper bound restrictions θ^{\min} and θ^{\max} .

In order to prove this bound, we first generalize a result by Bonifaci et al. [2] characterizing the inducibility of a fixed flow by δ -deviations to multi-commodity networks with a common source (Theorem 2). This characterization naturally gives rise to the concept of an *alternating path*, which also plays a crucial role in the work by Nikolova and Stier-Moses [16].

We then study a specific class of latency deviations which we term (α, β) -deviations. Here the latency restrictions are of the form $\theta_a^{\min} = \alpha l_a$ and $\theta_a^{\max} = \beta l_a$ with $-1 < \alpha \leq 0 \leq \beta$. We show that for (α, β) -deviations the Deviation Ratio is at most

$$1 + \frac{\beta - \alpha}{1 + \alpha} \left\lceil \frac{n - 1}{2} \right\rceil r, \quad (1)$$

where n is the number of nodes of the network and r is the sum of the demands of the commodities (Theorem 1). In particular, this reveals that the Deviation Ratio depends linearly on the size of the underlying network (among other parameters).

By using this result, we obtain a bound on the Price of Risk Aversion (Theorem 6) which generalizes the one in [16]. Nikolova and Stier-Moses [16] show that the Price of Risk Aversion for single-commodity networks and non-negative risk-aversion parameter γ is at most $1 + \gamma \kappa \lceil (n - 1)/2 \rceil$. (Here the demand is normalized to one.) We obtain the same bound from (1) with $\alpha = 0$ and $\beta = \gamma \kappa$. Our bound generalizes their result in two ways: (i) it holds for multi-commodity networks with a common source, and (ii) it also holds for negative risk-aversion parameters (capturing risk-taking players). Further, we show that our result can be used to bound the relative error in social cost of Nash flows incurred by small latency perturbations (Theorem 7), which is of independent interest. To the best of our knowledge, this notion has not been studied before in the literature.

2. *Lower bounds on the Deviation Ratio:* We prove that our bound on the Deviation Ratio for (α, β) -deviations is best possible for multi-commodity networks with a common source. We also show that it does not extend to general multi-commodity networks.

More specifically, for *single-commodity networks* we show that our bound is tight in *all* its parameters. Our lower bound construction holds for arbitrary $n \in \mathbb{N}$ and is based on the *generalized Braess graph* [18] (Example 1). In particular, this also complements a recent result by Lianas, Nikolova and Stier-Moses [11] who show that the upper bound on the Price of Risk Aversion in [16] is tight for single-commodity networks with $n = 2^j$ nodes for all $j \in \mathbb{N}$.

Further, for *multi-commodity networks with a common source* we show that our bound is tight in all parameters if n is odd, while a small gap remains if n is even (Theorem 3).

Finally, for *general multi-commodity networks* we establish a lower bound showing that the Deviation Ratio can be exponential in n (Theorem 4). In particular, this shows that there is an exponential gap between the cases of multi-commodity networks with and without a common source. In our proof, we adapt a graph structure used by Lin et al. [12] in their lower bound construction for the network design problem on multi-commodity networks.

3. *Smoothness bounds on the Biased Price of Anarchy:* We improve (and slightly generalize) recent smoothness bounds on the Price of Risk Aversion given by Meir and Parkes [13] and independently by Lianes et al. [11]. In particular, we derive tight bounds for the *Biased Price of Anarchy (BPoA)* [13], i.e., the ratio between the cost of a deviated Nash flow and the cost of a social optimum, for *arbitrary* $(0, \beta)$ -deviations (Theorem 5).¹ Note that the Biased Price of Anarchy yields an upper bound on the Deviation Ratio/Price of Risk Aversion.

It is interesting to note that the smoothness bounds on the Biased Price of Anarchy [13] and the Price of Risk Aversion [11] are independent of the network structure, but dependent on the class of latency functions. In contrast, our bound on the Deviation Ratio holds for arbitrary non-negative and non-decreasing latency functions, but depends on certain parameters of the network.²

4. *Generalizations of our model.* We also consider two natural generalizations of our model for which we derive additional results. In our first generalization, we consider *general path deviations* (which are not representable by arc deviations). We give a smoothness bound on the Biased Price of Anarchy for this setting. As a consequence, we obtain bounds on the Price of Risk Aversion under the non-linear *mean-std model* [11, 16] (Theorem 10). In our second generalization, we consider single-commodity instances with *heterogeneous players*, i.e., where players have different attitudes towards general path deviations. We show that our upper bound on the deviation ratio extends to this setting for certain graph structures. In particular, for series-parallel graphs we obtain a natural generalization of the bound in (1) for the heterogeneous player case.

Our results also answer a question posed by Nikolova and Stier-Moses in [16] regarding possible relations between their Price of Risk Aversion model, the restricted network toll problem [2], and the network design problem [18]. In particular, our results show that the analysis in [16] is not inherent to the used variance function, but rather depends on the restrictions imposed on the feasible deviations.

¹We remark that for certain types of $(0, \beta)$ -deviations, e.g., scaled marginal tolls, better bounds can be obtained; see the section “Relations to network toll problems” in Appendix B for relevant literature.

²For example, there are parallel-arc networks for which the Biased Price of Anarchy is unbounded, whereas the Deviation Ratio is a constant.

1.2 Related Work

The modelling and study of uncertainties in routing games has received a lot of attention in recent years. An extensive survey on this topic is given by Cominetti [5].

As mentioned above, our investigations are inspired by the study of the Price of Risk Aversion by Nikolova and Stier-Moses [16]. They prove that for single-commodity instances with non-negative and non-decreasing latency functions the Price of Risk Aversion is at most $1 + \gamma\kappa\lceil(n-1)/2\rceil$. We generalize this result to multi-commodity networks with a common source and to negative risk-aversion parameters. We elaborate in more detail on the connections to their work in Section 2.

Meir and Parkes [13] and independently Lineas et al. [11] show that for non-atomic network routing games with $(1, \mu)$ -smooth³ latency functions it holds that $\text{PRA} \leq \text{BPoA} \leq (1 + \gamma\kappa)/(1 - \mu)$. An advantage of such bounds is that they hold for general multi-commodity instances (but depend on the class of latency functions). These bounds stand in contrast to the *topological* bounds obtained here and by Nikolova and Stier-Moses [16] which hold for arbitrary non-negative and non-decreasing latency functions (but depend on the size of the network).

Conceptually, our model is related to the *restricted network toll problem* by Bonifaci et al. [2]. The authors study the problem of computing non-negative tolls that have to obey some upper bound restrictions $(\theta_a)_{a \in A}$ such that the cost of the resulting Nash flow is minimized. This is tantamount to computing best-case deviations in our model with $\theta_a^{\min} = 0$ and $\theta_a^{\max} = \theta_a$. In contrast, our focus here is on worst-case deviations. As a side result, we prove that computing such worst-case deviations is NP-hard, even for single-commodity instances with linear latencies (Theorem 12). There are several papers that study the problem of imposing tolls on the arcs of a network to reduce the cost of the resulting Nash flow. Such tolls can naturally be interpreted as latency deviations. We elaborate in more detail on these connections in Appendix B.

Other works also study the relative impact of instance alterations on the resulting Nash flows. For example, Roughgarden [18] studies the *network design problem* of finding a subnetwork that minimizes the latency of all flow-carrying paths of the resulting Nash flow. He introduces the *Braess ratio* which relates the common latency of a Nash flow in the original graph to the common latency of a Nash flow in an (optimal) subgraph. He shows that the *trivial algorithm* (which simply returns the original network) gives an $\lfloor n/2 \rfloor$ -approximation algorithm for single-commodity networks and that this is best possible (unless $\text{P} = \text{NP}$). Later, Lin et al. [12] show that this algorithm can be exponentially bad for multi-commodity networks. The instances that we use in our lower bound constructions are based on the ones used in [12, 18].

Englert, Franke and Olbrich [7] study the sensitivity of Nash flows in non-atomic network routing games. They investigate the relative change in social cost with respect to two alterations: (i) when the demand is perturbed by an additive constant

³Meir and Parkes [13] define a function l to be $(1, \mu)$ -smooth if $xl(y) \leq \mu yl(y) + xl(x)$ for all $x, y \geq 0$ (which is slightly different from Roughgarden's original smoothness definition [19]). Lineas et al. [11] only require *local smoothness* where y is taken fixed.

$\epsilon > 0$, and (ii) when an edge with only an ϵ -fraction of flow is removed. For single-commodity instances with polynomial latency functions of degree at most p , they show that the ratio of the social cost of a Nash flow in the original instance and the social cost of a Nash flow in the instance with demand increased by $\epsilon > 0$, is at most $(1 + \epsilon)^p$. They also show that this bound is tight.

2 Preliminaries

In this section, we introduce our bounded deviation model for non-atomic network routing games, define the Deviation Ratio and elaborate on some related notions. We also derive some preliminary results that are used later.

2.1 Non-Atomic Network Routing Games

An instance of a non-atomic network routing game is given by a tuple $\mathcal{I} = (G = (V, A), (l_a)_{a \in A}, (s_i, t_i)_{i \in [k]}, (r_i)_{i \in [k]})$. Here, $G = (V, A)$ is a directed graph with node set V and arc set $A \subseteq V \times V$, where each arc $a \in A$ has a non-negative, non-decreasing and continuous latency function $l_a : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$. Each commodity $i \in [k]$ is associated with a source-destination pair (s_i, t_i) and has a demand of $r_i \in \mathbb{R}_{> 0}$. We assume without loss of generality that $t_i \neq t_j$ if $i \neq j$ for $i, j \in [k]$. If all commodities share a common source node, i.e., $s_i = s_j = s$ for all $i, j \in [k]$, we call \mathcal{I} a *common source multi-commodity instance (with source s)*. We assume without loss of generality that $1 = r_1 \leq r_2 \leq \dots \leq r_k$ and define $r = \sum_{i \in [k]} r_i$.

We denote by \mathcal{P}_i the set of all simple (s_i, t_i) -paths of commodity $i \in [k]$ in G , and we define $\mathcal{P} = \cup_{i \in [k]} \mathcal{P}_i$. An outcome of the game is a feasible flow $f : \mathcal{P} \rightarrow \mathbb{R}_{\geq 0}$ such that $\sum_{P \in \mathcal{P}_i} f_P = r_i$ for every $i \in [k]$. Given a flow $f = (f^i)_{i \in [k]}$, we use f_a^i to denote the total flow on arc $a \in A$ of commodity $i \in [k]$, i.e., $f_a^i = \sum_{P \in \mathcal{P}_i: a \in P} f_P$. The total flow on arc $a \in A$ is defined as $f_a = \sum_{i \in [k]} f_a^i$. The latency of a path $P \in \mathcal{P}$ with respect to f is defined as $l_P(f) := \sum_{a \in P} l_a(f_a)$. The *social cost* $C(f)$ of a flow f is given by its total latency, i.e.,

$$C(f) = \sum_{P \in \mathcal{P}} f_P l_P(f) = \sum_{a \in A} f_a l_a(f_a).$$

A flow that minimizes $C(\cdot)$ is called (*socially*) *optimal*. We use $A_i^+ = \{a \in A : f_a^i > 0\}$ to refer to the support of f^i for commodity $i \in [k]$ and define $A^+ = \cup_{i \in [k]} A_i^+$ as the support of f .

2.2 Bounded Deviation Model

We introduce our bounded deviation model for non-atomic network routing games. For every arc $a \in A$, we have a continuous function $\delta_a : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ modeling the *deviation* on arc a , and we write $\delta = (\delta_a)_{a \in A}$. Note that the deviation δ_a on arc a can be positive or negative. We define the deviation of a path $P \in \mathcal{P}$ as $\delta_P(f) =$

$\sum_{a \in P} \delta_a(f_a)$. The *deviated latency* on arc $a \in A$ is given by $l_a(f_a) + \delta_a(f_a)$; similarly, the deviated latency on path $P \in \mathcal{P}$ is given by $l_P(f) + \delta_P(f)$. We say that f is δ -inducible if and only if it is a *Wardrop flow* (or *Nash flow*) with respect to $l + \delta$, i.e.,

$$\forall i \in [k], \forall P \in \mathcal{P}_i, f_P > 0 : l_P(f) + \delta_P(f) \leq l_{P'}(f) + \delta_{P'}(f) \quad \forall P' \in \mathcal{P}_i. \quad (2)$$

If f is δ -inducible, we also write $f = f^\delta$. Note that a Nash flow f for the unaltered latencies $(l_a)_{a \in A}$ is 0-inducible, i.e., $f = f^0$.

Let $\theta^{\min} = (\theta_a^{\min})_{a \in A}$ and $\theta^{\max} = (\theta_a^{\max})_{a \in A}$ be given threshold functions, where for every $a \in A$, $\theta_a^{\min} : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ is a continuous, non-increasing function and $\theta_a^{\max} : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ is a continuous, non-decreasing function. Further, we assume that $\theta_a^{\min}(x) \leq 0 \leq \theta_a^{\max}(x)$ for all $x \geq 0$ and $a \in A$, and let $\theta = (\theta^{\min}, \theta^{\max})$. We define

$$\Delta(\theta) = \{(\delta_a)_{a \in A} \mid \forall a \in A : \theta_a^{\min}(x) \leq \delta_a(x) \leq \theta_a^{\max}(x), \forall x \geq 0\}$$

as the set of *feasible deviations*. Note that $0 \in \Delta(\theta)$ for all threshold functions θ^{\min} and θ^{\max} . We say that $\delta \in \Delta(\theta)$ is a θ -deviation. Furthermore, f is θ -inducible if there exists a $\delta \in \Delta(\theta)$ such that f is δ -inducible. For $-1 < \alpha \leq 0 \leq \beta$, we call $\delta \in \Delta(\theta)$ an (α, β) -deviation if $\theta^{\min} = \alpha l$ and $\theta^{\max} = \beta l$, and also write $\theta = (\alpha, \beta)$.

We make the following assumption:

Assumption 1 We assume that the function $l_a + \theta_a^{\min}$ is non-negative and non-decreasing for every arc $a \in A$.

Intuitively, the non-negativity property ensures that the deviated latencies $l + \delta$ remain non-negative for all feasible deviations $\delta \in \Delta(\theta)$. The non-decreasingness property is exploited in our upper bound proof on the Deviation Ratio. Note that (α, β) -deviations naturally satisfy this assumption.

Throughout the paper, we (implicitly) only consider deviations δ for which a Nash flow exists. The existence of such flows is always guaranteed under some mild conditions on the threshold functions. As an example, we elaborate on the existence when $\theta^{\min} = 0$ and θ_a^{\max} is non-negative, non-decreasing and continuous for all $a \in A$. It is not hard to see that for a deviated Nash flow f^δ with $\delta \in \Delta(\theta)$ there exists some $0 \leq \lambda_a \leq 1$ for every arc $a \in A$ such that $\delta_a(f_a^\delta) = \lambda_a \theta_a^{\max}(f_a^\delta)$. In particular, this means that the deviations δ' defined as $\delta'_a = \lambda_a \theta_a^{\max}$ satisfies $\delta' \in \Delta(\theta)$ and also induces f^δ . Therefore it is sufficient to consider deviations of the form $\delta_a = \lambda_a \theta_a^{\max}$, where $0 \leq \lambda_a \leq 1$ for all $a \in A$. For such deviations, the deviated latency function $l_a + \delta_a$ is non-negative, non-decreasing and continuous for every $a \in A$. It is well-known that for these types of functions, the existence of a Nash flow is guaranteed (see, e.g., Nisan et al. [17]).

The following lemma shows an equivalence between (α, β) -deviations with $-1 < \alpha \leq 0 \leq \beta$ and $(0, \frac{\beta - \alpha}{1 + \alpha})$ -deviations. In particular, it allows us to assume without loss of generality that $\alpha = 0$. The proof is given in Appendix A.

Lemma 1 *Let $-1 < \alpha \leq 0 \leq \beta$ be fixed. Then f is inducible with an (α, β) -deviation if and only if it is inducible with a $(0, \frac{\beta - \alpha}{1 + \alpha})$ -deviation.*

Our bounded deviation model naturally gives rise to optimization problems where one wants to compute feasible deviations that minimize or maximize the social cost of the resulting Nash flow. We elaborate in more detail on these problems in Appendix B.

2.3 Deviation Ratio

Given an instance \mathcal{I} and threshold functions $\theta = (\theta^{\min}, \theta^{\max})$, we define the *Deviation Ratio* as the worst-case ratio of the cost of a θ -inducible flow and the cost of a 0-inducible flow; more formally,

$$DR(\mathcal{I}, \theta) = \sup_{\delta \in \Delta(\theta)} \left\{ \frac{C(f^\delta)}{C(f^0)} \mid f^\delta \text{ is } \delta\text{-inducible} \right\}.$$

Intuitively, $DR(\mathcal{I}, \theta)$ measures the worst-case deterioration of the social cost of a Nash flow due to (feasible) latency deviations. Note that for a fixed deviation $\delta \in \Delta(\theta)$, there might be multiple Nash flows that are δ -inducible. Unless stated otherwise, we adopt the convention that $C(f^\delta)$ refers to the social cost of the worst Nash flow that is δ -inducible.

2.4 Related Notions

Nikolova and Stier-Moses [16] (see also [11, 15]) consider non-atomic network routing games with uncertain latencies. Here the deviations correspond to variances $(v_a)_{a \in A}$ of some random variable ζ_a (with expectation zero). The *perceived latency* of a path $P \in \mathcal{P}$ with respect to a flow f is then defined as

$$q_P^\gamma(f) = l_P(f) + \gamma v_P(f),$$

where $\gamma \geq 0$ is a parameter representing the *risk-aversion* of the players. They consider two different objectives as to how the deviation $v_P(f)$ of a path P is defined:

1. *mean-var objective*: $v_P(f) = \sum_{a \in P} v_a(f_a)$
2. *mean-std objective*: $v_P(f) = (\sum_{a \in P} v_a(f_a))^{\frac{1}{2}}$.

Note that for the mean-var objective there is an equivalent arc-based definition, where the perceived latency of every arc $a \in A$ is defined as $q_a^\gamma(f_a) = l_a(f_a) + \gamma v_a(f_a)$.

They define the *Price of Risk Aversion* [16] as the worst-case ratio $C(x)/C(z)$, where x is a *risk-averse* Nash flow with respect to $q^\gamma = l + \gamma v$ and z is a *risk-neutral* Nash flow with respect to l .⁴

In their analysis, it is assumed that the *variance-to-mean-ratio* of every arc $a \in A$ under the risk-averse flow x is bounded by some constant $\kappa \geq 0$, i.e., $v_a(x_a) \leq \kappa l_a(x_a)$ for all $a \in A$. Under this assumption, they prove that the Price of Risk Aversion $PRA(\mathcal{I}, \gamma, \kappa)$ of single-commodity instances \mathcal{I} with non-negative and non-decreasing latency functions is at most $1 + \gamma\kappa \lceil (n - 1)/2 \rceil$, where n is the number of nodes.

⁴The existence of a risk-averse Nash flow is proven in [15].

We now elaborate on the relation to our Deviation Ratio. The main technical difference is that in [16] the variance-to-mean ratio is only considered for the respective flow values x_a . Note however that if we write for every $a \in A$, $v_a(x_a) = \lambda_a l_a(x_a)$ for some $0 \leq \lambda_a \leq \kappa$, then the deviation function $\delta_a(y) = \gamma \lambda_a l_a(y)$ has the property that $x = f^\delta$ is δ -inducible with $\delta \in \Delta(0, \gamma\kappa)$. It follows that for every instance \mathcal{I} and parameters γ, κ , $\text{PRA}(\mathcal{I}, \gamma, \kappa) \leq \text{DR}(\mathcal{I}, (0, \gamma\kappa))$.

Another related notion is the *Biased Price of Anarchy (BPoA)* introduced by Meir and Parkes [13]. Adapted to our setting, given an instance \mathcal{I} and threshold functions θ , the Biased Price of Anarchy is defined as $\text{BPoA}(\mathcal{I}, \theta) = \sup_{\delta \in \Delta(\theta)} C(f^\delta) / C(f^*)$, where f^* is a socially optimal flow. Note that because $C(f^*) \leq C(f)$ for every feasible flow f , we have $\text{DR}(\mathcal{I}, \theta) \leq \text{BPoA}(\mathcal{I}, \theta)$.

3 Upper Bounds on the Deviation Ratio

We derive an upper bound on the Deviation Ratio. All results in this section hold for multi-commodity instances with a common source.

The following notion of alternating paths turns out to be crucial. It was first introduced by Lin et al. [12] in the context of the network design problem (see [18]) and is also used by Nikolova and Stier-Moses [16].

Definition 1 (Alternating path [12, 16]) Let x and z be feasible flows. We partition $A = X \cup Z$, where $Z = \{a \in A : z_a \geq x_a \text{ and } z_a > 0\}$ and $X = \{a \in A : z_a < x_a \text{ or } z_a = x_a = 0\}$. We say that $\pi_i = (a_1, \dots, a_r)$ is an *alternating* (s, t_i) -*path* if the arcs in $\pi_i \cap Z$ are oriented in the direction of t_i , and the arcs in $\pi_i \cap X$ are oriented in the direction of s .

An *alternating path tree* π is a tree, rooted at the common source s , which contains an alternating (s, t_i) -path π_i for every commodity i . We show below that an alternating path tree always exists for multi-commodity networks with a common source.

The main theorem which we prove in this section is as follows:

Theorem 1 *Let x be θ -inducible and let z be 0-inducible. Further, let $A = X \cup Z$ be a partition of A as in Definition 1. Let π be an alternating path tree, where π_i denotes the alternating (s, t_i) -path in π .*

- (i) *Suppose $\theta = (\theta^{\min}, \theta^{\max})$. Let X_i be a flow-carrying path of commodity $i \in [k]$ maximizing $l_P(x)$ over all $P \in \mathcal{P}_i$.⁵ Then*

$$C(x) \leq C(z) + \sum_{i \in [k]} r_i \left(\sum_{a \in Z \cap \pi_i} \theta_a^{\max}(z_a) - \sum_{a \in X \cap \pi_i} \theta_a^{\min}(z_a) - \sum_{a \in X_i} \theta_a^{\min}(x_a) \right).$$

⁵Note that the values $l_P(x) + \delta_P(x)$ are the same for all flow-carrying paths, but this is not necessarily true for the values $l_P(x)$.

- (ii) Suppose $\theta = (\alpha, \beta)$ with $-1 < \alpha \leq 0 \leq \beta$. Let η_i be the number of disjoint segments of consecutive arcs in Z on the alternating (s, t_i) -path π_i for $i \in [k]$.⁶ Then

$$C(x) \leq \left(1 + \frac{\beta - \alpha}{1 + \alpha} \sum_{i \in [k]} r_i \eta_i \right) C(z) \leq \left(1 + \frac{\beta - \alpha}{1 + \alpha} \cdot \left\lceil \frac{n - 1}{2} \right\rceil r \right) C(z).$$

We give some interpretation: Theorem 1(i) relates the social cost of a θ -inducible Nash flow x to the social cost of an original Nash flow z . More specifically, it shows that $C(x) - C(z)$ is at most

$$\sum_{i \in [k]} r_i \left(\sum_{a \in Z \cap \pi_i} \theta_a^{\max}(z_a) - \sum_{a \in X \cap \pi_i} \theta_a^{\min}(z_a) - \sum_{a \in X_i} \theta_a^{\min}(x_a) \right),$$

where X_i is a flow-carrying (s, t_i) -path with respect to x and π_i is an alternating (s, t_i) -path for commodity i . Intuitively, the contribution of commodity i to the above term can be seen as the total θ -cost of sending r_i units of flow along the directed cycle C_i which we obtain from π_i and X_i by reversing all arcs in $X \cap \pi_i$ and X_i . Here the θ -cost is defined as $\theta_a^{\max}(z_a)$ for a forward arc $a \in Z \cap \pi_i$, $-\theta_a^{\min}(z_a)$ for a reversed arc $a \in X \cap \pi_i$, and $-\theta_a^{\min}(x_a)$ for a reversed arc $a \in X_i$.

If we can bound the total θ -cost by $\lambda C(z) - \mu C(x)$ with $\lambda \geq 0$ and $\mu > -1$, then we obtain an upper bound of $(1 + \lambda)/(1 + \mu)$ on the Deviation Ratio. In particular, for (α, β) -deviations the θ -cost can naturally be related to the latencies. In this case, we obtain the bound stated in Theorem 1(ii).

In order to prove Theorem 1 we proceed as follows: We first derive a characterization of when a given flow f is θ -inducible (Theorem 2). As it turns out, this reduces to a non-negative cycle condition in a suitably defined auxiliary graph $\hat{G}(f)$ with cost function c . In particular, this non-negative cycle condition allows us to relate the cost of a flow-carrying path F_i of f to arbitrary (s, t_i) -paths and (t_i, s) -paths in the auxiliary graph $\hat{G}(f)$ (Lemma 2). We then turn to relating the social cost of a θ -inducible flow x to the one of a 0-inducible flow z . We show that an alternating path tree π with respect to x and z always exist (Lemma 3). With the help of this alternating tree we can then relate the costs of (carefully chosen) flow-carrying paths under x and z for each commodity. Basically, for each commodity i we bound the cost of a flow-carrying path X_i of x by the cost of the alternating path π_i (by applying Lemma 2 to X_i and π_i). The latter in turn can then be bounded by the cost of a flow-carrying path Z_i of z (by applying Lemma 2 to Z_i and π_i).

3.1 Characterization of θ -Inducible Flows

We provide a characterization of the inducibility of a given flow. Let f be a feasible flow. We define an auxiliary graph $\hat{G} = \hat{G}(f) = (V, \hat{A})$ with $\hat{A} = A \cup \bar{A}$, where $\bar{A} = \{(v, u) : a = (u, v) \in A^+\}$, i.e., \hat{A} consists of the set of arcs in A , which we call

⁶Note that $\eta_i \leq \lceil (n - 1)/2 \rceil$.




Fig. 1 The dashed arcs are the reversed arcs in \hat{G} . The black bold arcs indicate the cycle B . We have $(h_0, h_1, h_2, h_3) = (1, 4, 6, 1)$. Note that, for example, it could be the case that $P_1 = P_6 \cup (b_6, b_1)$

forward arcs, and the set \bar{A} of arcs (v, u) with $(u, v) \in A^+$, which we call *reversed arcs*. Further, we define a cost function $c : \hat{A} \rightarrow \mathbb{R}$ as follows:

$$c_a = \begin{cases} l_{(u,v)}(f(u,v)) + \theta_{(u,v)}^{\max}(f(u,v)) & \text{if } (u, v) \in A \\ -l_{(u,v)}(f(u,v)) - \theta_{(u,v)}^{\min}(f(u,v)) & \text{if } (v, u) \in \bar{A}. \end{cases} \tag{3}$$

The following theorem generalizes the characterization result for single-commodity networks in [2] to multi-commodity networks with a common source.

Theorem 2 *Let f be a feasible flow. f is θ -inducible if and only if $\hat{G} = \hat{G}(f)$ does not contain a cycle of negative cost with respect to c .*

Proof Suppose that f is an inducible flow and let δ be a vector of deviations that induce f . Throughout the proof all latency, deviation and threshold functions are evaluated with respect to f . For notational convenience, we omit the explicit reference to f .

Let \hat{B} be a directed cycle in \hat{G} . If \hat{B} only consists of forward arcs, then $\sum_{a \in \hat{B}} (l_a + \theta_a^{\max}) \geq \sum_{a \in \hat{B}} (l_a + \theta_a^{\min}) \geq 0$, where the last inequality holds because of Assumption 1. Next, suppose that there is a reversed arc $a = (v, u) \in \hat{B} \cap \bar{A}$. Then $(u, v) \in A_i^+$ for some commodity $i \in [k]$. Let $B = (b_1, \dots, b_q, b_1)$ be the cycle that we obtain from \hat{B} if all arcs $(v, u) \in \hat{B} \cap \bar{A}$ are replaced by $a = (u, v) \in A^+$ (note that B is contained in G and that it is not a directed cycle). For every arc $b = (b_l, b_{l+1}) \in B \cap A^+$, there is a flow-carrying path P_l from s to b_l for some commodity i (here we use the fact that all commodities share the same source).⁷

Intuitively, the proof is as follows. For all nodes $b \in V(B)$ with two incoming arcs of B , we can find two paths Q_1 and Q_2 leading to that node, using the paths P_l and the cycle B (see also Fig. 1). Furthermore, one of those paths is flow-carrying by construction. We then apply the Nash conditions to those flow-carrying paths

⁷Note that the paths P_l can overlap, use parts of B , or even be subpaths of each other.

(exploiting the common source) and add up the resulting inequalities. The contributions of the paths P_l cancel out in the aggregated inequality, leading to the desired result. We now give a formal proof.

Without loss of generality, we may assume that $(b_1, b_2) \in A^+$. Let $h_1 \in \{2, \dots, q + 1\}$ be the smallest index for which $(b_{h_1}, b_{h_1+1}) \in A^+$ (here we take $b_{q+1} := b_1$ and $P_{q+1} := P_1$). Note that the concatenation of P_{h_1} and $(b_{h_1}, b_{h_1-1}, \dots, b_2)$ is a directed path from s to b_2 . Then we have

$$l_{(b_1, b_2)} + \delta_{(b_1, b_2)} + \sum_{a \in P_1} (l_a + \delta_a) \leq \sum_{j=3}^{h_1} (l_{(b_j, b_{j-1})} + \delta_{(b_j, b_{j-1})}) + \sum_{a \in P_{h_1}} (l_a + \delta_a)$$

by using the fact that a subpath (s, \dots, u) of a shortest (s, t_i) -path $(s, \dots, u, \dots, t_i)$ is a shortest (s, u) -path if G does not contain negative cost cycles under the cost function $l + \delta$ (which is true because of Assumption 1).

We can now repeat this procedure by letting $h_2 \in \{h_1 + 1, \dots, q + 1\}$ be the smallest index for which $(b_{h_2}, b_{h_2+1}) \in A^+$. We then have

$$\begin{aligned} l_{(b_{h_1}, b_{h_1+1})} + \delta_{(b_{h_1}, b_{h_1+1})} + \sum_{a \in P_{h_1}} (l_a + \delta_a) \\ \leq \sum_{j=h_1+2}^{h_2} (l_{(b_j, b_{j-1})} + \delta_{(b_j, b_{j-1})}) + \sum_{a \in P_{h_2}} (l_a + \delta_a). \end{aligned}$$

Continuing this procedure, we find a sequence $1 = h_0 < h_1 < \dots < h_p = q + 1$ such that, for every $0 \leq w \leq p - 1$,

$$\begin{aligned} l_{(b_{h_w}, b_{h_w+1})} + \delta_{(b_{h_w}, b_{h_w+1})} + \sum_{a \in P_{h_w}} (l_a + \delta_a) \\ \leq \sum_{j=h_w+2}^{h_{w+1}} (l_{(b_j, b_{j-1})} + \delta_{(b_j, b_{j-1})}) + \sum_{a \in P_{h_{w+1}}} (l_a + \delta_a). \end{aligned} \tag{4}$$

Note that p is the number of reversed arcs on the cycle \hat{B} .

Summing up these inequalities for $0 \leq w \leq p - 1$, we obtain

$$\sum_{(v, u) \in \hat{B} \cap \bar{A}} (l_{(u, v)} + \delta_{(u, v)}) \leq \sum_{a \in \tilde{B} \cap A} (l_a + \delta_a),$$

since all the contributions of the path P_l cancel out. Now using the definition of a θ -deviation, we find

$$\begin{aligned} \sum_{a \in \tilde{B} \cap A} (l_a + \theta_a^{\max}) - \sum_{(v, u) \in \hat{B} \cap \bar{A}} (l_{(u, v)} + \theta_{(u, v)}^{\min}) \\ \geq \sum_{a \in \tilde{B} \cap A} (l_a + \delta_a) - \sum_{(v, u) \in \hat{B} \cap \bar{A}} (l_{(u, v)} + \delta_{(u, v)}) \geq 0. \end{aligned}$$

We have shown that \hat{B} has non-negative cost. Note that \hat{B} has zero cost if all the arcs on the cycle are reversed.

For the other direction of the proof, consider the set $\mathcal{F}(\theta)$ of θ -deviations $\delta \in \Delta(\theta)$ that induce $f = (f_a^i)_{i \in [k], a \in A}$ (see also [12, 18]):

$$\begin{aligned} \mathcal{F}(\theta) = \{ (\delta_a)_{a \in A} \mid & \pi_{i,v} - \pi_{i,u} \leq l_a + \delta_a & \forall a = (u, v) \in A, \forall i \in [k] \\ & \pi_{i,v} - \pi_{i,u} = l_a + \delta_a & \forall a = (u, v) \in A_i^+, \forall i \in [k] \\ & \delta_a \geq \theta_a^{\min} & \forall a \in A \\ & \delta_a \leq \theta_a^{\max} & \forall a \in A \}. \end{aligned} \tag{5}$$

That is, f is θ -inducible if and only if the linear system defining $\mathcal{F}(\theta)$ in (5) has a feasible solution. Now suppose that \hat{G} does not contain a cycle of negative cost. Then we can determine the shortest path distance π_u from s to every node $u \in V$. We define $\pi_{i,u} := \pi_u$ for all $u \in V$ and $i \in [k]$. Furthermore, for $a = (u, v) \in A$, we define $\delta_a := \max\{\theta_a^{\min}, \pi_v - \pi_u - l_a\}$. We will now show that δ induces f by showing that we have constructed a feasible solution for (5). First of all, for all $i \in [k]$ and $a \in A \setminus A_i^+$, we have $\delta_a \geq \pi_v - \pi_u - l_a$, which is equivalent to $\pi_{i,v} - \pi_{i,u} \leq l_a + \delta_a$. Secondly, if $a = (u, v) \in A_i^+$, then $\pi_u - \pi_v \leq -l_a - \theta_a^{\min}$ (which we derive using the reversed arc (v, u)). But this is equivalent to $\pi_{i,v} - \pi_{i,u} - l_a \geq \theta_a^{\min}$. We can conclude that $\delta_a = \pi_{i,v} - \pi_{i,u} - l_a$. Furthermore, we clearly have $\delta_a \geq \theta_a^{\min}$. Lastly, for all $a = (u, v) \in A$ we have $\pi_v - \pi_u \leq l_a + \theta_a^{\max}$ which is equivalent to $\pi_v - \pi_u - l_a \leq \theta_a^{\max}$. Combining this with the trivial inequality $\theta_a^{\min} \leq \theta_a^{\max}$ we can conclude that $\delta_a \leq \theta_a^{\max}$. This completes the proof. \square

The characterization of Theorem 2 applies if all commodities share a common source. In fact, we can show that this characterization does not hold if this assumption is dropped (see Appendix C).

By exploiting the non-negative cycle condition of Theorem 2, we can now establish the following bounds on the cost of a flow-carrying path F_i of a θ -inducible flow f .

Lemma 2 *Let f be θ -inducible and let F_i be a flow-carrying (s, t_i) -path for commodity $i \in [k]$ in G . Let χ and ψ be any (s, t_i) -path and (t_i, s) -path in $\hat{G}(f)$, respectively. Then*

$$\begin{aligned} \sum_{a \in F_i} l_a(f_a) + \theta_a^{\min}(f_a) &\leq \sum_{a \in \chi \cap A} l_a(f_a) + \theta_a^{\max}(f_a) - \sum_{a \in \chi \cap \bar{A}} l_a(f_a) + \theta_a^{\min}(f_a) \\ \sum_{a \in F_i} l_a(f_a) + \theta_a^{\max}(f_a) &\geq \sum_{a \in \psi \cap \bar{A}} l_a(f_a) + \theta_a^{\min}(f_a) - \sum_{a \in \psi \cap A} l_a(f_a) + \theta_a^{\max}(f_a). \end{aligned}$$

We need the following proposition to prove Lemma 2.

Proposition 1 *Let $G = (V, A)$ be a non-empty, directed multigraph with the property that $\delta^-(v) = \delta^+(v)$ for all $v \in V$.⁸ Then G is the union of arc-disjoint*

⁸We use the standard notation $\delta^-(v)$ and $\delta^+(v)$ to refer to the set of outgoing and incoming edges of a node v , respectively.

directed (simple) cycles C_1, \dots, C_l for some l such that $\cup_{j=1}^l V(C_j) = V$ and $\cup_{j=1}^l A(C_j) = A$.

Proof If G is non-empty then we can find a (simple) directed cycle C in G . Removing the arcs of this cycle leads to the graph $G \setminus C := (V, A \setminus A(C))$ that also satisfies $\delta^-(v) = \delta^+(v)$ for all $v \in V$ (note that if there are multiple arcs between two nodes, we only remove the copy on the cycle). By repeating this procedure until G becomes empty, we decompose G into a series of arc-disjoint directed (simple) cycles C_1, \dots, C_l as claimed. \square

Proof (Lemma 2) Since F_i is a flow-carrying path, we know that for every $a = (u, v) \in F_i$ we have a reversed arc $(v, u) \in \tilde{A}$ in \tilde{G} . Let \tilde{F}_i denote the reversed path of F_i . Define \hat{H} as the graph consisting of the (t_i, s) -path \tilde{F}_i and the (s, t_i) -path χ , where we add a copy of an arc if it is used in both paths (i.e., \hat{H} can be a multigraph). Note that \hat{H} satisfies the conditions of Proposition 1. Thus, \hat{H} can be decomposed into arc-disjoint directed cycles C_1, \dots, C_l for some l . By Theorem 2, each such cycle C_j has non-negative cost with respect to c (as defined in (3)). Thus, we have

$$c(C_j) = \sum_{a \in A \cap C_j} (l_a(x_a) + \theta_a^{\max}(x_a)) - \sum_{a \in \tilde{A} \cap C_j} (l_a(x_a) + \theta_a^{\min}(x_a)) \geq 0.$$

By adding these inequalities for all $j = 1, \dots, l$ and rearranging terms, we obtain the first inequality.

The second inequality is proven analogously (applying the same arguments to the graph \hat{H} consisting of paths F_i and ψ .) \square

3.2 Existence of Alternating Path Tree

Let x and z be feasible flows. Recall the definition of an alternating (s, t_i) -path π_i (Definition 1). The following lemma establishes the existence of an alternating path tree π , rooted at the common source s , which contains an alternating (s, t_i) -path π_i for every commodity $i \in [k]$. It is a direct generalization of Lemma 4.6 in [12] and Lemma 4.5 in [16].

Lemma 3 *Let x and z be feasible flows and let $A = X \cup Z$ be a partition of A as in Definition 1. Then there exists an alternating path tree.*

Proof Let $G' = (V', A')$ be the graph defined by $V = V \cup \{t\}$ and $A' = A \cup \{(t_i, t) : i \in [k]\}$. Let x', z' be the flows defined by

$$x'_a = \begin{cases} x_a & \text{for } a = (u, v) \in A \\ r_i & \text{for } a = (t_i, t) \text{ with } i \in [k], \end{cases}$$

$$z'_a = \begin{cases} z_a & \text{for } a = (u, v) \in A \\ r_i & \text{for } a = (t_i, t) \text{ with } i \in [k]. \end{cases}$$

Then x' and z' are feasible (s, t) -flows in G' . We can write $A = Z' \cup X'$ with $Z' = Z \cup \{(t_i, t) : i \in [k]\}$ and X' having the same properties as Z and X in G (which follows from $x'_a = z'_a = r_i > 0$ for all $a = (t_i, t)$).

We can now apply the same argument as in the proof of Lemma 4.5 in [16] of which we will give a short summary (for sake of completeness). For any s - t cut defined by $S \cup V'$ with $s \in S$ we claim that we can cross S with an arc in Z' , or a reversed arc in X' . Suppose that this would not be the case, i.e., all arcs into S are in the set Z' and all the outgoing arcs of S are in X' . Let $x_{Z'}$ and $z_{Z'}$ be the total incoming flows from S , and $x_{X'}$ and $z_{X'}$ the total outgoing flows from S (for flows x and z , respectively). From the definition of Z' it follows that $x_{Z'} \leq z_{Z'}$. From conservation of flow it follows that $x_{X'} - x_{Z'} = z_{X'} - z_{Z'}$. Combining these two observations, we find that $x_{X'} \leq z_{X'}$. However, by definition of X' , we have $x_{X'} > z_{X'}$ (since we removed all arcs a with $z_a = x_a = 0$). We find a contradiction.

Having proved the claim that we can always cross with an arc in Z' or a reversed arc in X' , we can now easily construct a spanning tree π' consisting of alternating paths, by starting with the cut $(S, G \setminus S)$ given by $S = \{s\}$.

Note that t cannot be an interior point of π' , since t is only adjacent to incoming arcs of the set Z' . This means that if we remove (t_j, t) from π' (where j is the index for which (t_j, t) is in the tree π'), we have found an alternating path tree π for the graph G , under the flows x and z . □

3.3 Proof of Theorem 1

We now have all the ingredients to prove Theorem 1.

Throughout this section, let x be a θ -inducible flow and let z be a 0-inducible. Let π be an alternating path tree (which exists by Lemma 3). Without loss of generality we may remove all arcs with $z_a = x_a = 0$ (as they do not contribute to the social cost). Note that if along the alternating (s, t_i) -path π_i we reverse the arcs of Z then the resulting path is a directed (t_i, s) -path in $\hat{G}(z)$ (which we call the s -oriented version of π_i); similarly, if we reverse the arcs of X then the resulting path is an (s, t_i) -path in $\hat{G}(x)$ (which we call the t_i -oriented version of π_i).

We start with the proof of Theorem 1(i).

Proof (Theorem 1(i)) Let X_i be a flow-carrying path of commodity $i \in [k]$ maximizing $l_P(x)$ over all $P \in \mathcal{P}_i$. Note that by our choice of X_i , we have

$$C(x) = \sum_{i \in [k]} \sum_{P \in \mathcal{P}_i} x_P^i l_P(x) \leq \sum_{i \in [k]} r_i \sum_{a \in X_i} l_a(x_a).$$

Let Z_i be an arbitrary flow-carrying path of commodity $i \in [k]$ with respect to z . We have

$$C(z) = \sum_{i \in [k]} r_i \sum_{a \in Z_i} l_a(z_a).$$

By applying the first inequality of Lemma 2 to the flow x in the graph $\hat{G}(x)$, where we choose χ to be the t_i -oriented version of π_i , we obtain

$$\sum_{a \in X_i} l_a(x_a) + \theta_a^{\min}(x_a) \leq \sum_{a \in Z \cap \pi_i} l_a(x_a) + \theta_a^{\max}(x_a) - \sum_{a \in X \cap \pi_i} l_a(x_a) + \theta_a^{\min}(x_a). \tag{6}$$

By applying the second inequality of Lemma 2 to the flow z in the graph $\hat{G}(z)$ with $\theta^{\max} = \theta^{\min} = 0$, where we choose ψ to be the s -oriented version of π_i , we obtain

$$\sum_{a \in Z_i} l_a(z_a) \geq \sum_{a \in Z \cap \pi_i} l_a(z_a) - \sum_{a \in X \cap \pi_i} l_a(z_a). \tag{7}$$

By combining these inequalities, we obtain

$$\begin{aligned} \sum_{a \in X_i} l_a(x_a) + \theta_a^{\min}(x_a) &\leq \sum_{a \in Z \cap \pi_i} l_a(x_a) + \theta_a^{\max}(x_a) - \sum_{a \in X \cap \pi_i} l_a(x_a) + \theta_a^{\min}(x_a) \\ &\leq \sum_{a \in Z \cap \pi_i} l_a(z_a) + \theta_a^{\max}(z_a) - \sum_{a \in X \cap \pi_i} l_a(z_a) + \theta_a^{\min}(z_a) \\ &\leq \sum_{a \in Z_i} l_a(z_a) + \sum_{a \in Z \cap \pi_i} \theta_a^{\max}(z_a) - \sum_{a \in X \cap \pi_i} \theta_a^{\min}(z_a). \end{aligned}$$

Here the first inequality follows from (6). The second inequality holds because of the definition of X and Z and the non-decreasingness of $l_a + \theta_a^{\max}$ and $l_a + \theta_a^{\min}$ (Assumption 1) for every $a \in A$. The last inequality holds because of (7).

The claim now follows by multiplying the above inequality with r_i and summing over all commodities $i \in [k]$. □

We need the following proposition for the proof of Theorem 1(ii).

Proposition 2 *Let $z = f^0$ be a Nash flow for a multi-commodity instance with a common source. Let $v \in V$ and let $i, j \in [k]$ be two commodities for which there exist flow-carrying (s, v) -paths $P_1 \in \mathcal{P}_i$ and $P_2 \in \mathcal{P}_j$, respectively. Then there exists a feasible Nash flow \bar{z} with $\bar{z}_a = z_a$ for all $a \in A$ such that both paths P_1, P_2 are flow-carrying for commodity i , and both paths P_1, P_2 are flow-carrying for commodity j , i.e., we have $\bar{z}_{P_1}^i, \bar{z}_{P_2}^i, \bar{z}_{P_1}^j, \bar{z}_{P_2}^j > 0$.*

Proof Intuitively, we shift an ϵ amount of flow of commodity i to path P_2 and an ϵ amount of flow of commodity j to path P_1 . Formally, choose $\epsilon > 0$ small enough such that $z_{P_1}^i - \epsilon, z_{P_2}^j - \epsilon > 0$. We define

$$\bar{z}_P^l = \begin{cases} z_{P_1}^i - \epsilon & \text{if } P = P_1 \text{ and } l = i \\ z_{P_1}^j + \epsilon & \text{if } P = P_1 \text{ and } l = j \\ z_{P_2}^i + \epsilon & \text{if } P = P_2 \text{ and } l = i \\ z_{P_2}^j - \epsilon & \text{if } P = P_2 \text{ and } l = j \end{cases}$$

and let all the other flow-carrying paths remain unchanged. It then immediately follows that $z_a = \bar{z}_a$ for all $a \in A$, and in the resulting feasible flow \bar{z} , both commodities i and j are flow-carrying for both paths P_1 and P_2 . The feasibility of \bar{z} follows

because both commodities have the same source. Moreover, the common source also implies that if z is a Nash flow, then \bar{z} is also a Nash flow (since commodity i implies that $l_{P_1}(z) \leq l_{P_2}(z)$, and commodity j implies that $l_{P_2}(z) \leq l_{P_1}(z)$). \square

We now give the proof of Theorem 1(ii).

Proof (Theorem 1(ii)) By Lemma 1 we can assume without loss of generality that for every arc $a \in A$:

$$\theta_a^{\min} = 0 \quad \text{and} \quad \theta_a^{\max} = \frac{\beta - \alpha}{1 + \alpha} l_a.$$

Fix a commodity i and consider the alternating (s, t_i) -path π_i . Let a *segment* of π be a maximal sequence of consecutive arcs on π_i which belong to Z . Suppose π consists of η_i segments. Let A_{ij} denote the j -th segment of π_i .

Using Theorem 1(i) and the definition of A_{ij} , we obtain

$$\begin{aligned} C(x) &\leq C(z) + \frac{\beta - \alpha}{1 + \alpha} \sum_{i \in [k]} r_i \sum_{a \in Z \cap \pi_i} l_a(z_a) \\ &\leq C(z) + \frac{\beta - \alpha}{1 + \alpha} \sum_{i \in [k]} r_i \left(\eta_i \cdot \max_{j=1, \dots, \eta_i} \sum_{a \in A_{ij}} l_a(z_a) \right). \end{aligned}$$

Note that the claim follows if we can prove that $\sum_{a \in A_{ij}} l_a(z_a) \leq C(z)$ for all $j = 1, \dots, \eta_i$ and $i \in [k]$.

Fix a segment A_{ij} . Below we argue that there always exists a commodity $w \in [k]$ (possibly $w \neq i$) such that every $a \in A_{ij}$ is flow-carrying for commodity w , i.e., $z_a^w > 0$ for every $a \in A_{ij}$. By choosing a suitable path decomposition of z for commodity w , we can thus assume that A_{ij} is contained in some flow-carrying path $P \in \mathcal{P}_w$ and thus $\sum_{a \in A_{ij}} l_a(z_a) \leq l_P(z)$. Recall that $C(z) = \sum_{i \in [k]} r_i l_{Z_i}(z)$, where $Z_i \in \mathcal{P}_i$ is an arbitrary flow-carrying path for commodity $i \in [k]$. By exploiting that $r_i \geq 1$ for every $i \in [k]$, we obtain

$$\sum_{a \in A_{ij}} l_a(z_a) \leq l_P(z) \leq \sum_{i \in [k]} r_i l_{Z_i}(z) = C(z).$$

We now prove that there always exists a commodity w as claimed above. Suppose there are two consecutive edges $a_1 = (u, v)$ and $a_2 = (v, w)$ in A_{ij} that are flow-carrying for commodities w_1 and w_2 in z , respectively. Then there are two (s, v) -paths W_1 and W_2 which are flow-carrying with respect to commodities w_1 and w_2 , respectively. The existence of W_1 is clear. The existence of W_2 follows from flow-conservation applied to commodity w_2 (because some positive amount of flow leaves node v). But then, by Proposition 2, we may assume that a_1 is also flow-carrying for commodity w_2 . By applying this argument repeatedly, starting with the last two arcs on A_{ij} and proceeding towards the front, we can show that there is a commodity for which the whole segment A_{ij} is flow-carrying. \square

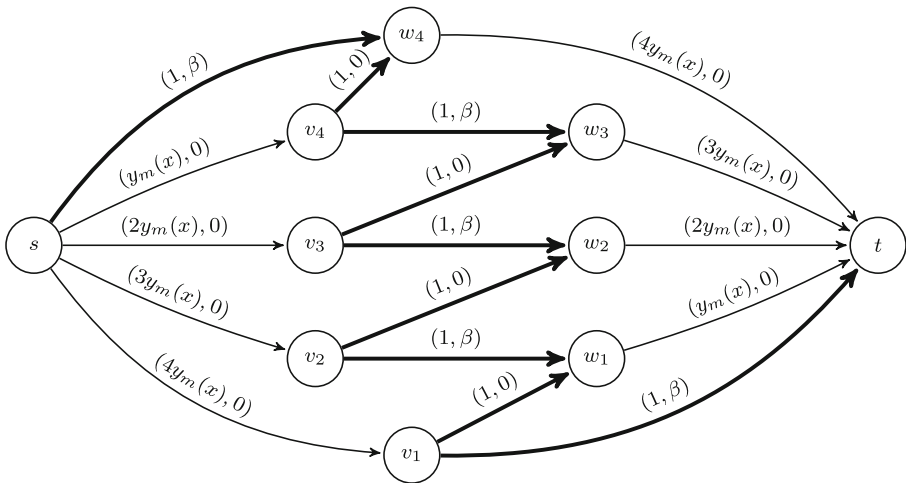


Fig. 2 The fifth Braess graph with (l_a^5, δ_a^5) on the arcs as defined in Example 1. The bold arcs indicate the alternating path π_1

4 Lower Bounds on the Deviation Ratio for (α, β) -Deviations

In this section, we provide lower bounds on the Deviation Ratio for (α, β) -deviations. We first consider single-commodity instances and prove that the bound given in Theorem 1 is tight in all its parameters. We then extend this result to instances with a common source. In contrast, for general multi-commodity instances the situation is much worse. In particular, we establish an exponential lower bound on the Deviation Ratio.

4.1 Single-Commodity Instances

Our instance is based on the generalized Braess graph [18]. The m -th Braess graph $G^m = (V^m, A^m)$ is defined by $V^m = \{s, v_1, \dots, v_{m-1}, w_1, \dots, w_{m-1}, t\}$ and A^m as the union of three sets: $E_1^m = \{(s, v_j), (v_j, w_j), (w_j, t) : 1 \leq j \leq m - 1\}$, $E_2^m = \{(v_j, w_{j-1}) : 2 \leq j \leq m\}$ and $E_3^m = \{(v_1, t) \cup \{(s, w_{m-1})\}\}$. See Fig. 2 for an example. The rough idea behind the lower bound construction is that in the unaltered Nash flow all players spread out evenly over the m paths *not* involving the arcs of the form (v_i, w_i) . However, as a result of introducing deviations on the arcs of the form (v_i, w_{i-1}) the players switch to the paths involving the arcs (v_i, w_i) , but this increases the latencies on all arcs adjacent to s and t .

Example 1 By Lemma 1, we can assume without loss of generality that $\alpha = 0$. Let $\beta \geq 0$ be a fixed constant and let $n = 2m \geq 4 \in \mathbb{N}$.⁹ Let G^m be the m -th Braess

⁹Note that the value $\lceil (n - 1)/2 \rceil$ is the same for $n \in \{2m, 2m + 1\}$ with $m \in \mathbb{N}$. The example shows tightness for $n = 2m$. The tightness for $n = 2m + 1$ then follows trivially by adding a dummy node.

graph. Furthermore, let $y_m : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ be a non-decreasing, continuous function¹⁰ with $y_m(1/m) = 0$ and $y_m(1/(m - 1)) = \beta$. We define

$$l_a^m(g) = \begin{cases} (m - j) \cdot y_m(g) & \text{for } a \in \{(s, v_j) : 1 \leq j \leq m - 1\} \\ j \cdot y_m(g) & \text{for } a \in \{(w_j, t) : 1 \leq j \leq m - 1\} \\ 1 & \text{otherwise.} \end{cases}$$

Furthermore, we define $\delta_a^m(g) = \beta$ for $a \in E_2^m$, and $\delta_a^m(g) = 0$ otherwise. Note that $0 \leq \delta_a^m(g) \leq \beta l_a^m(g)$ for all $a \in A$ and $g \geq 0$ (see Fig. 2).

A Nash flow $z = f^0$ is given by routing $1/m$ units of flow over the paths (s, w_{m-1}, t) , (s, v_1, t) and the paths in $\{(s, v_j, w_{j-1}, t) : 2 \leq j \leq m - 1\}$. Note that all these paths have latency one, and the path (s, v_j, w_j, t) , for some $1 \leq m \leq j$, also has latency one. We conclude that $C(z) = 1$.

A Nash flow $x = f^\delta$, with δ as defined above, is given by routing $1/(m - 1)$ units of flow over the paths in $\{(s, v_j, w_j, t) : 1 \leq j \leq m - 1\}$. Each such path P then has a latency of $l_P(x) = 1 + \beta m$. It follows that $C(x) = 1 + \beta m$. Note that the deviated latency of path P is $q_P(x) = 1 + \beta m$ because all deviations along this path are zero. Each path $P' = (s, v_j, w_{j-1}, t)$, for $2 \leq j \leq m - 1$, has a deviated latency of $q_{P'}(x) = 1 + \beta + (m - 1)y_m(1/(m - 1)) = 1 + \beta + (m - 1)\beta = 1 + \beta m$. The same argument holds for the paths (s, w_{m-1}, t) and (s, v_1, t) . We conclude that x is δ -inducible. It follows that $C(x)/C(z) = 1 + \beta m = 1 + \beta n/2$.

4.2 Common-Source Instances

By adapting the construction in Example 1, we obtain the following result.

Theorem 3 *There exist common source two-commodity instances \mathcal{I} such that*

$$DR(\mathcal{I}, (\alpha, \beta)) \geq \begin{cases} 1 + \frac{\beta - \alpha}{1 + \alpha} \cdot \frac{n - 1}{2} r & \text{for } n = 2m + 1 \in \mathbb{N}_{\geq 5} \\ 1 + \frac{\beta - \alpha}{1 + \alpha} \cdot \left[\left(\frac{n}{2} - 1 \right) r + 1 \right] & \text{for } n = 2m \in \mathbb{N}_{\geq 4}. \end{cases}$$

Proof We first prove the claim for n odd. Let $r \in \mathbb{R}_{\geq 1}$ and $n = 2m + 1 \in \mathbb{N}_{\geq 5}$. We modify the graph G^m by adding one extra node t_2 (the node t will be referred to as t_1 from here on). We add the arcs (s, t_2) and (t_2, t_1) (see the dotted arcs in Fig. 2). We take one commodity with sink t_1 and $r_1 = 1$, and one commodity with sink t_2 and demand $r_2 = r - 1$. Note that the latter commodity only has one (s, t_2) -path.

The pairs $(l_a^m(g), \delta_a^m(g))$, for all a except (s, t_2) and (t_2, t_1) , are defined as in Example 1, but with y a non-decreasing, non-negative, continuous function satisfying $y_m(1/m) = 0$ and $y_m((1 - \epsilon_m)/(m - 1)) = \beta$, where we choose $0 < \epsilon_m < 1/m$ so that $1/m < (1 - \epsilon_m)/(m - 1)$. For $a = (s, t_2)$, we take $(l_a^m(g), \delta_a^m(g)) = (y_m^*(x'), 0)$, where y^* is a non-decreasing, non-negative, continuous function satisfying $y_m^*(r -$

¹⁰For example $y_m(g) = m(m - 1)\beta \max\{0, (g - \frac{1}{m})\}$. That is, we define y_m to be zero for $0 \leq g \leq 1/m$ and we let it increase with constant rate to β in $1/(m - 1)$.

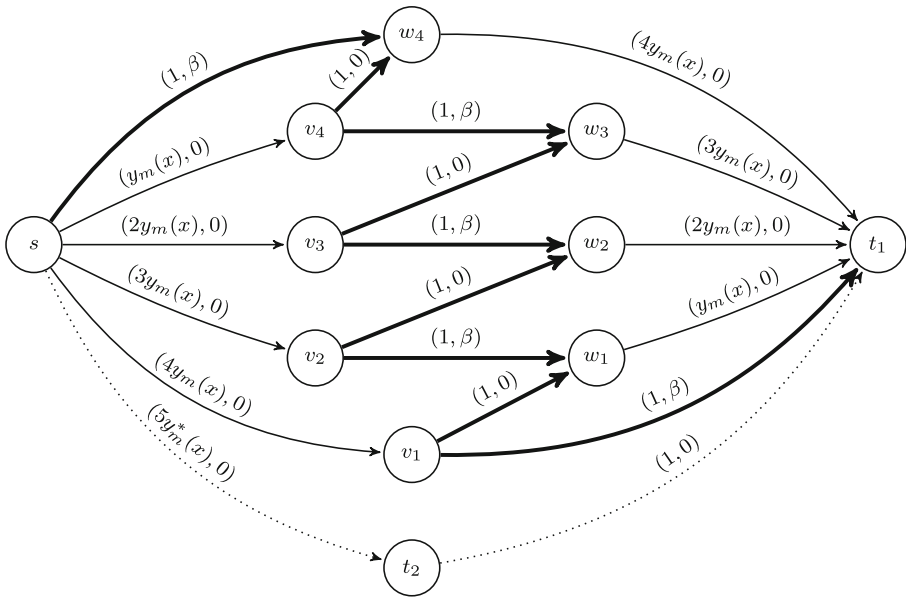


Fig. 3 The fifth (odd) Braess graph with (l_a^5, δ_a^5) on the arcs as defined above, where $t = t_1$. The thick edges indicate the alternating path π_1

$1) = 0$ and $y_m^*(r - 1 + \epsilon_m) = \beta$. For $a = (t_2, t_1)$ we take $(l_a^m(g), \delta_a^m(g)) = (1, 0)$. See Fig. 3 for an example.

A Nash flow z for this instance is given by routing $1/m$ units of flow over the paths (s, w_{m-1}, t_1) , (s, v_1, t_1) and the paths in $\{(s, v_j, w_{j-1}, t_1) : 2 \leq j \leq m - 1\}$ for the first commodity, and $r - 1$ units of flow over (s, t_2, t_1) for the second commodity. This claim is true since all the paths for the first commodity have latency one, as well as the paths (s, v_j, w_j, t) , for $1 \leq m \leq j$. This is also true for (s, t_2, t_1) . The latency for the other commodity is zero. We may conclude that $C(z) = 1$.

A Nash flow x under deviation δ , as defined here, is given by, for the first commodity, routing $(1 - \epsilon_m)/(m - 1)$ units of flow over the paths in $\{(s, v_j, w_j, t) : 1 \leq j \leq m - 1\}$, and ϵ_m units of flow over the path (s, t_2, t_1) . Note that the perceived latency on all these paths p is $q_p(x) = 1 + \beta m$ (which is also the true latency, since all the deviations are zero on the arcs of these paths). Using the same reasoning as in Example 1 it can be seen that the perceived latency on the paths $P' = (s, v_j, w_{j-1}, t)$, for $2 \leq j \leq m - 1$, is also $q_{P'}(x) = 1 + \beta m$, from which we may conclude that x is indeed a Nash flow under the deviation δ . We have $C(x) = 1 + \beta m + (r - 1)\beta m = 1 + \beta r m$, since for the first commodity the (true) latency along every path is $1 + \beta m$, and for the other commodity the latency along (s, t_2) is βm .

We next prove the claim for n even. Let $r \in \mathbb{R}_{\geq 1}$ and $n = 2m \in \mathbb{N}_{\geq 4}$. We use the same Braess graphs as in Example 1, without modifications. We introduce another commodity with demand $r_2 = r - 1$, for which we choose $t_2 = v_1$. We replace the pair $((m - 1)y_m(x'), 0)$ on $a = (s, v_1)$ by the pair $((m - 1)y'_m(g), 0)$ where y'_m

satisfies $y'_m(1/m + r - 1) = 0$ and $y'_m(1/(m - 1) + r - 1) = \beta$. Note that the flows x and z , as defined in Example 1 with the extension that the second commodity uses the arc (s, v_1) in both cases, still form feasible Nash flows for their respective deviations. We obtain

$$\begin{aligned} C(x) &= \sum_i \sum_{q \in \mathcal{P}_i} x_q^i l_q(x) = 1 + \beta m + (r - 1)(m - 1)\beta \\ &= 1 + \beta m + \beta(r - 1)(m - 1) = (1 + \beta r m) - \beta(r - 1). \end{aligned}$$

This completes the proof. \square

Remark 1 For two-commodity instances with n even, we can actually improve the upper bound in Theorem 1 to the lower bound stated in Theorem 3: Suppose the upper bound of Theorem 1 is tight. Then we need to have $\eta_1 = \eta_2 = n/2$. This means that the alternating path tree is actually a path, in the sense that all nodes are adjacent to at most two arcs of the alternating path tree, that alternates between arcs in X and Z , starting and ending with an arc in Z (see Fig. 2). However, because $t_1 \neq t_2$ this means that at least one of the two commodities has no more than $n/2 - 1$ arcs in Z , which is a contradiction.

4.3 Multi-Commodity Instances

For general multi-commodity instances we establish the following exponential lower bound on the Deviation Ratio. In particular, this proves that there is an exponential gap between the cases of multi-commodity networks with and without a common source.

Theorem 4 *For every $p = 2q + 1 \in \mathbb{N}$, there exists a two-commodity instance \mathcal{I} whose size is polynomially bounded in p such that*

$$DR(\mathcal{I}, (\alpha, \beta)) \geq 1 + \beta F_{p+1} \approx 1 + 0.45\beta \cdot \phi^{p+1},$$

where F_p is the p -th Fibonacci number and $\phi \approx 1.618$ is the golden ratio.

The instance used in the proof of Theorem 4 is based on the following graph introduced by Lin et al. [12].

Definition 2 ([12]) For $p = 2q + 1 \in \mathbb{N}$, the graph $G^p = (V^p, A^p)$ is defined by

$$V^p = \{s_1, s_2, t_1, t_2, e, w_0, \dots, w_p, v_1, \dots, v_p\},$$

and $A^p = A(P_1^p) \cup A(P_2^p) \cup A_1^p \cup A_2^p \cup \{s_1, w_0\}$ where

$$P_1^p = (s_1, e, w_1, v_1, v_2, \dots, v_p, t_1) \text{ and } P_2^p = (s_2, w_0, w_1, \dots, w_p, t_2)$$




Fig. 4 The graph G^p for $p = 7$ (this is a reproduction of Figure 4 in [12]). The arc $a = (s_1, e)$ has $\delta_a = \beta$, whereas all the other arcs have $\delta_a = 0$

are the horizontal (s_1, t_1) -path and vertical (s_2, t_2) -path, respectively; see Fig. 4. Further,

$$A_1^p = \{(s_2, v_i) : i = 1, 3, 5, 7, \dots, p - 2\} \cup \{(e, w_i) : i = 2, 4, 6, 8, \dots, p - 1\}$$

and

$$A_2^p = \{(w_i, v_i) : i = 3, 5, 7, \dots, p\} \cup \{(v_i, w_i) : i = 2, 4, 6, 8, \dots, p - 1\}.$$

Lastly, the paths T_i are denoted by

$$T_i = \begin{cases} (s_1, w_0, w_1, v_1, \dots, v_p, t_1) & i = 0 \\ (s_1, e, w_i, w_{i+1}, v_{i+1}, \dots, v_p, t_1) & i = 2, 4, 6, \dots, p - 1 \\ (s_2, v_1, v_{i+1}, w_{i+1}, \dots, w_p, t_2) & i = 1, 3, 5, \dots, p. \end{cases}$$

These paths can be seen as ‘shortcuts’ for the paths P_1 and P_2 .

Proof (Theorem 4) We consider instances $(G^p, l^p, \delta^p, r^p)_{p=1,3,5,7,\dots}$ with G^p as in Definition 2. It is not hard to see that $|V^p|, |A^p| \in \mathcal{O}(p)$. The latency functions l^p are given as follows:

$$l_a^p(x') = \begin{cases} \beta g_\delta^i(x') & \text{for } a \in \{(v_i, v_{i+1}) : i = 1, 3, 5, \dots, p - 2\} \\ \beta g_\delta^i(x') & \text{for } a \in \{(w_i, w_{i+1}) : i = 0, 2, 4, 6, \dots, p - 1\} \\ 1 & \text{for } a \in \{(s_1, e), (s_1, w_0)\} \\ 0 & \text{otherwise.} \end{cases}$$

Here

$$g_\delta^i(x') = \begin{cases} 0 & x' \leq 1 \\ h_\delta^i(x') & 1 \leq x' \leq 1 + \delta \\ F_i & x' \geq 1 + \delta, \end{cases}$$

where F_i is the i -th Fibonacci number, and $h_\delta^i(x')$ is some non-decreasing, non-negative, continuous function satisfying $h_\delta^i(1) = 0$ and $h_\delta^i(1 + \delta) = F_i$ (so that $g_\delta^i(x')$ is also non-decreasing, non-negative and continuous). Furthermore, we take $\delta_a = \beta$ for $a = (s_1, e)$ and $\delta_a = 0$ for all $a \in A \setminus \{(s_1, e)\}$. Finally, we have $r_1^p = r_2^p = 1$.

Let z be defined by sending one unit of flow over the paths P_1 and P_2 . We claim that z is a Nash flow with respect to the latencies l^p and $C(z) = 1$. By construction, the latency along the path P_1 is $l_{P_1}(z) = 1$. It is not hard to see that any (s_1, t_1) -path has latency greater or equal than one (because every path for commodity 1 uses either (s_1, e) or (s_1, w_0)). For commodity 2 the latency along P_2 is $l_{P_2}(z) = 0$, which is clearly a shortest path. This proves that z is a Nash flow. Further, $C(z) = 1$.

We use Lemma 4 (given below) to describe a Nash flow x with respect to the deviated latencies $l^p + \delta^p$. It follows that $C(x) = C(x)/C(z) \geq 1 + \beta F_{p-1} + \beta F_p = 1 + \beta F_{p+1}$. This concludes the proof (since $F_p \approx c \cdot \phi^p$ where $c \approx 0.4472$ and $\phi \approx 1.618$). \square

The following lemma is similar to Lemma 5.4, Lemma 5.5 and Lemma 5.6 in [12].

Lemma 4 *There exists a $\delta > 0$ and a feasible flow x satisfying the following properties:*

- (i) $x_a \geq 1 + \delta$ for all $a \in \{(v_i, v_{i+1}) : i = 1, 3, 5, \dots, p - 2\} \cup \{(w_i, w_{i+1}) : i = 0, 2, 4, 6, \dots, p - 1\}$.
- (ii) $l_P(x) \geq 1 + \beta F_{p-1}$ for all $P \in \mathcal{P}_1$, with equality if and only if $P = T_i$ for some $i = 2, 4, 6, \dots, p - 1$.
- (iii) $l_P(x) \geq \beta F_p$ for all $P \in \mathcal{P}_2$, with equality if and only if $P = T_i$ for some $i = 1, 3, 5, \dots, p$.
- (iv) x is a Nash flow under the perceived latencies $l^p + \delta^p$.

Proof The statements (i)–(iii) follow from Lemmas 5.4, 5.5 and 5.6 in [12]. The last statement is clearly true for commodity 2 (since this commodity is not affected by

the deviation on arc (s_1, e)). For commodity 1, all the flow-carrying paths T_i have a perceived latency of $Q_{T_i}(x) = 1 + \beta(F_p + 1)$, and the perceived latency along any other (s_1, t_1) -path is greater or equal than that. The actual latencies along these paths are $l_{T_i}(x) = 1 + \beta F_{p-1}$ for $i = 2, 4, 6, \dots, p - 1$, and $l_{T_0}(x) = 1 + \beta(F_{p-1} + 1)$. \square

5 Smoothness Bounds on the Biased Price of Anarchy

We derive tight smoothness bounds on the Biased Price of Anarchy for $(0, \beta)$ -deviations. Our approach is a generalization of the framework of Correa, Schulz and Stier-Moses [6] (which we obtain for $\beta = 0$).

Let \mathcal{L} be a given set of latency functions and $\beta \geq 0$ fixed. For $l \in \mathcal{L}$, define

$$\hat{\mu}(l, \beta) = \sup_{x, z \geq 0} \left\{ \frac{z[l(x) - (1 + \beta)l(z)]}{xl(x)} \right\} \quad \text{and} \quad \hat{\mu}(\mathcal{L}, \beta) = \sup_{l \in \mathcal{L}} \hat{\mu}(l, \beta). \quad (8)$$

Theorem 5 *Let \mathcal{L} be a set of non-negative, non-decreasing and continuous functions. Let \mathcal{I} be a general multi-commodity instance with $(l_a)_{a \in A} \in \mathcal{L}^A$. Let x be δ -inducible for some $(0, \beta)$ -deviation δ and let z be an arbitrary feasible flow. If $\hat{\mu}(\mathcal{L}, \beta) < 1$, then*

$$C(x) \leq \frac{1 + \beta}{1 - \hat{\mu}(\mathcal{L}, \beta)} C(z).$$

Moreover, this bound is tight if \mathcal{L} contains all constant functions and is closed under scalar multiplication, i.e., for every $l \in \mathcal{L}$ and $\gamma \geq 0, \gamma l \in \mathcal{L}$.

Proof We use a similar approach as Correa et al. [6]. Since x is a deviated Nash flow with respect to $l + \delta$, the following variational inequality holds:


$$\sum_{a \in A} x_a(l_a(x_a) + \delta_a(x_a)) \leq \sum_{a \in A} z_a(l_a(x_a) + \delta_a(x_a)).$$

We have

$$\begin{aligned} C(x) &= \sum_{a \in A} x_a l_a(x_a) \leq \sum_{a \in A} z_a l_a(x_a) + (z_a - x_a) \delta_a(x_a) \\ &\leq \sum_{x_a > z_a} z_a l_a(x_a) + \sum_{z_a \geq x_a} z_a (l_a(x_a) + \delta_a(x_a)) \\ &\leq \sum_{x_a > z_a} z_a l_a(x_a) + (1 + \beta) \sum_{z_a \geq x_a} z_a l_a(x_a) \\ &\leq \sum_{x_a > z_a} z_a l_a(x_a) + (1 + \beta) \sum_{z_a \geq x_a} z_a l_a(z_a), \end{aligned}$$

where the third inequality holds because δ is a $(0, \beta)$ -deviation and the last inequality holds because the latency functions are non-decreasing.

Fig. 5 Example used in the proof of Theorem 5. The arcs are labeled by their respective (l_a, δ_a) functions. Note that $\delta \in \Delta(0, \beta)$



We obtain

$$\begin{aligned}
 C(x) &\leq \sum_{x_a > z_a} z_a l_a(x_a) + (1 + \beta) \sum_{z_a \geq x_a} z_a l_a(z_a) \\
 &= \sum_{x_a > z_a} z_a [l_a(x_a) - (1 + \beta)l_a(z_a) + (1 + \beta)l_a(z_a)] + (1 + \beta) \sum_{z_a \geq x_a} z_a l_a(z_a) \\
 &= (1 + \beta)C(z) + \sum_{x_a > z_a} z_a [l_a(x_a) - (1 + \beta)l_a(z_a)] \\
 &\leq (1 + \beta)C(z) + \hat{\mu}(\mathcal{L}, \beta) \sum_{x_a > z_a} x_a l_a(x_a) \\
 &\leq (1 + \beta)C(z) + \hat{\mu}(\mathcal{L}, \beta)C(x).
 \end{aligned}$$

Thus, for $\hat{\mu}(\mathcal{L}, \beta) < 1$, we obtain $C(x) \leq (1 + \beta)/(1 - \hat{\mu}(\mathcal{L}, \beta))C(z)$.

We will now prove the tightness of the obtained bound if \mathcal{L} contains all constant functions and is closed under scalar multiplication. For arbitrary $c \in \mathcal{L}$ and demand r , consider the parallel-arc instance in Fig. 5.

Clearly, a deviated Nash flow is given by $x = (x_1, x_2) = (r, 0)$, since then $l_1(x_1) + \delta_1(x_1) = l_2(x_2) + \delta_2(x_2) = (1 + \beta)/r$. We have $C(x) = (1 + \beta)$.

For a feasible flow $z = (\epsilon, r - \epsilon)$. We have

$$C(z) = \frac{(1 + \beta)\epsilon c(\epsilon) + (r - \epsilon)c(r)}{rc(r)} = \frac{rc(r) - \epsilon[c(r) - (1 + \beta)c(\epsilon)]}{rc(r)}$$

which implies that, with z^* a socially optimal flow,

$$\frac{C(x)}{C(z^*)} \geq \frac{C(x)}{C(z)} = (1 + \beta) \left(1 - \frac{\epsilon[c(r) - (1 + \beta)c(\epsilon)]}{r \cdot c(r)} \right)^{-1}.$$

In order to claim tightness we can choose $c \in \mathcal{L}$, and $r \geq \epsilon \geq 0$, arbitrary close to $\hat{\mu}(\mathcal{L}, \beta)$. □

6 Applications

We use our results on the Deviation Ratio and the Biased Price of Anarchy obtained in the previous sections to derive several new results below.

6.1 Price of Risk Aversion

We obtain the following bound on the Price of Risk Aversion for multi-commodity networks with a common source.

Theorem 6 *The Price of Risk Aversion for a common source multi-commodity instance \mathcal{I} with non-negative and non-decreasing latency functions, variance-to-mean-ratio $\kappa > 0$ and risk-aversion parameter $\gamma \geq -1/\kappa$ is at most*

$$PRA(\mathcal{I}, \gamma, \kappa) \leq \begin{cases} 1 - \gamma\kappa/(1 + \gamma\kappa)\lceil(n - 1)/2\rceil r & \text{for } -1/\kappa < \gamma \leq 0 \\ 1 + \gamma\kappa\lceil(n - 1)/2\rceil r & \text{for } \gamma \geq 0. \end{cases}$$

Moreover, these bounds are tight in all its parameters if $n = 2m + 1$ and almost tight if $n = 2m$. In particular, for single-commodity instances we obtain tightness for all $n \in \mathbb{N}$.

Note that Theorem 6 generalizes the result in [16] to multi-commodity networks with a common source and to negative risk-aversion parameters. Further, it establishes that the bound is tight in all its parameters.

Proof (Theorem 6) Recall from the discussion in Section 2 that the deviations $\delta_a = \gamma v_a$ can be interpreted as θ -deviations with

$$\theta_a^{\min} = \begin{cases} 0 & \text{if } \gamma \geq 0 \\ \gamma\kappa l_a & \text{if } -1/\kappa < \gamma \leq 0 \end{cases} \quad \text{and} \quad \theta_a^{\max} = \begin{cases} \gamma\kappa l_a & \text{if } \gamma \geq 0 \\ 0 & \text{if } -1/\kappa < \gamma \leq 0. \end{cases}$$

Here, the restriction $\gamma > -1/\kappa$ is necessary to satisfy Assumption 1. The theorem now follows directly from Theorem 1, Example 1 and Theorem 3. □

6.2 Stability of Nash Flows Under Small Perturbations

We next show that our results can be used to bound the relative error in social cost incurred by small latency perturbations.

We introduce some more notation. We say that $(\tilde{l}_a)_{a \in A}$ are ϵ -perturbed latency functions with respect to $(l_a)_{a \in A}$ if

$$\sup_{a \in A, x \geq 0} \left| \frac{l_a(x) - \tilde{l}_a(x)}{l_a(x)} \right| \leq \epsilon$$

for some small $\epsilon > 0$. We are interested in bounding the relative error in social cost due to ϵ -perturbations of the latency functions. More precisely, the *relative error in social cost* is defined as the ratio

$$\frac{C(\tilde{f}) - C(f)}{C(f)},$$

where f is a Nash flow with respect to $(l_a)_{a \in A}$ and \tilde{f} is a Nash flow with respect to ϵ -perturbed latency functions $(\tilde{l}_a)_{a \in A}$. To the best of our knowledge, this notion has not been studied in the literature before.

The theorem below establishes an upper bound on the relative error in social cost. In particular, for small ϵ -perturbations the theorem implies that the relative error is asymptotically $O(\epsilon n)$.

Theorem 7 *Let \mathcal{I} be a common source multi-commodity instance with non-negative and non-decreasing latency functions $(l_a)_{a \in A}$. Let f be a Nash flow with respect to $(l_a)_{a \in A}$ and let \tilde{f} be a Nash flow with respect to ϵ -perturbed latency functions $(\tilde{l}_a)_{a \in A}$ for some $0 < \epsilon < 1$. Then the relative error in social cost satisfies*

$$\frac{C(\tilde{f}) - C(f)}{C(f)} \leq \frac{2\epsilon}{1 - \epsilon} \cdot \left\lceil \frac{n - 1}{2} \right\rceil r.$$

Proof Note that the ϵ -perturbation $l - \tilde{l}$ can be seen as a $(-\epsilon, \epsilon)$ -deviation. Using Theorem 1, we obtain

$$\frac{C(\tilde{f})}{C(f)} \leq 1 + \frac{2\epsilon}{1 - \epsilon} \cdot \left\lceil \frac{n - 1}{2} \right\rceil r.$$

The claim follows. \square

6.3 Biased Price of Anarchy

Our smoothness bound on the Biased Price of Anarchy derived in Theorem 5 improves upon the bounds of $(1 + \beta)/(1 - \mu)$ recently obtained by Meir and Parkes [13] and Lineas et al. [11] for $(1, \mu)$ -smooth latency functions. To see this, note that the bound stated in Theorem 5 is not worse than the bound $(1 + \beta)/(1 - \mu)$ because for $(1, \mu)$ -smooth latency functions it holds that

$$\hat{\mu}(\mathcal{L}, \beta) \leq \hat{\mu}(\mathcal{L}, 0) \leq \mu.$$

As a direct consequence, we also obtain better smoothness bounds on the Price of Risk Aversion.

We exemplify the increased strength of our general smoothness bound by deriving a closed form expression on the Biased Price of Anarchy for affine latency functions.

Theorem 8 *Let \mathcal{I} be a general multi-commodity instance with affine latency functions $(l_a)_{a \in A}$. Then*

$$BPoA(\mathcal{I}, \beta) \leq \frac{(1 + \beta)^2}{\frac{3}{4} + \beta}.$$

Note that the upper bound of $4(1 + \beta)/3$ on the Biased Price of Anarchy for affine latency functions given in [11, 13] is inferior to our bound.

Proof (Theorem 8) Let \mathcal{L} be the set of all affine latency functions with non-negative coefficients. The claim follows from Theorem 5 by showing that $\hat{\mu}(\mathcal{L}, \beta) = \frac{1}{4(1 + \beta)}$.

Let $l_a(y) = c_a y + d_a$ be an arbitrary affine latency function with $c_a, d_a \geq 0$. We need to show that

$$z_a [c_a x_a + d_a - (1 + \beta)(c_a z_a + d_a)] \leq \frac{1}{4(1 + \beta)} x_a [c_a x_a + d_a],$$

or, equivalently,

$$c_a [z_a x_a - (1 + \beta)z_a^2] + d_a [z_a - z_a(1 + \beta)] \leq c_a \left[\frac{1}{4(1 + \beta)} x_a^2 \right] + d_a \left[\frac{1}{4(1 + \beta)} x_a \right].$$

It suffices to show that

$$z_a x_a - (1 + \beta)z_a^2 \leq \frac{1}{4(1 + \beta)} x_a^2 \quad \text{and} \quad z_a - z_a(1 + \beta) \leq \frac{1}{4(1 + \beta)} x_a.$$

The second inequality is always true, using the non-negativity of z_a, x_a and β . For the first inequality, we have

$$0 \leq \left(\frac{x_a}{2} - (1 + \beta)z_a \right)^2 = (1 + \beta)^2 z_a^2 + \frac{x_a^2}{4} - (1 + \beta)x_a z_a,$$

which implies that

$$[1 + \beta] \left(x_a z_a - (1 + \beta)z_a^2 \right) \leq \frac{x_a^2}{4}.$$

Dividing this inequality by $(1 + \beta)$ gives the desired result. Further, we have tightness for $(x_a, z_a) = (1, \frac{1}{2(1+\beta)})$. □

6.4 Absolute Gap Between the BPoA and the Deviation Ratio

Finally, we derive an upper bound on the absolute gap between the Biased Price of Anarchy and the Deviation Ratio.

Theorem 9 *Let \mathcal{L} be a set of non-negative, non-decreasing and continuous functions (containing constants and closed under scalar multiplication). Let \mathcal{G} be the set of all instances with $(l_a)_{a \in A} \in \mathcal{L}^A$. If $\hat{\mu}(\mathcal{L}, \beta) < 1$, then*

$$|BPoA(\mathcal{G}, (0, \beta)) - DR(\mathcal{G}, (0, \beta))| \leq (1 + \beta) \frac{\hat{\mu}(\mathcal{L}, \beta)}{1 - \hat{\mu}(\mathcal{L}, \beta)}.$$

For example, for affine latencies we have $\hat{\mu}(\mathcal{L}, \beta) = \frac{1}{4(1+\beta)}$ as shown in the proof of Theorem 8. As a result,

$$|BPoA(\mathcal{G}, (0, \beta)) - DR(\mathcal{G}, (0, \beta))| \leq \frac{1 + \beta}{3 + 4\beta} \leq \frac{1}{3}$$

for all $\beta \geq 0$, i.e., the gap is bounded by a constant. In particular, this suggests that for large β the Biased Price of Anarchy provides a good approximation for the Deviation Ratio (or the Price of Risk Aversion). Note that this does not follow from the bound $4(1 + \beta)/3$ for affine latencies obtained in [11, 13].

Proof (Theorem 9) Consider the instance dedicated in Fig. 5. No matter how the flow splits over the two arcs in the unaltered Nash flow z with respect to $\delta = 0$, we always

have $C(z) = 1$. Further, as argued in the proof of Theorem 5, the deviated Nash flow x has social cost $C(x) = 1 + \beta$. Thus,

$$\frac{C(x)}{C(z)} = 1 + \beta \leq \text{DR}(\mathcal{G}, (0, \beta)) \leq \text{BPoA}(\mathcal{G}, (0, \beta)) \leq \frac{1 + \beta}{1 - \hat{\mu}(\mathcal{L}, \beta)}.$$

This implies that

$$\begin{aligned} |\text{BPoA}(\mathcal{G}, (0, \beta)) - \text{DR}(\mathcal{G}, (0, \beta))| &\leq \frac{1 + \beta}{1 - \hat{\mu}(\mathcal{L}, \beta)} - (1 + \beta) \\ &= (1 + \beta) \left(\frac{1}{1 - \hat{\mu}(\mathcal{L}, \beta)} - 1 \right). \end{aligned}$$

□

7 Generalizations of our Model

In this section, we consider two natural generalizations of our model and derive some additional results. In the first generalization, we consider *general path deviations* which are more expressive than the arc deviations studied above. In the second generalization, we consider *heterogenous players* where players have different attitudes towards general path deviations.

7.1 General Path Deviations

We consider general path deviations which are not necessarily decomposable into arc deviations. The main motivation for investigating such deviations is that we can apply such bounds to the *mean-std* objective of the Price of Risk Aversion model by Nikolova and Stier-Moses [16] (see Section 2).

First, we need to adjust some definitions of Section 2. Throughout this section, we assume that we are given non-positive and non-negative, respectively, continuous threshold functions $\theta^{\min} = (\theta_P^{\min})_{P \in \mathcal{P}}$ and $\theta^{\max} = (\theta_P^{\max})_{P \in \mathcal{P}}$. The set of *feasible path deviations* is defined as

$$\Delta(\theta) = \{(\delta_P)_{P \in \mathcal{P}} \mid \theta_P^{\min}(f) \leq \delta_P(f) \leq \theta_P^{\max}(f) \text{ for all feasible flows } f\}.$$

In particular, (α, β) -path deviations are deviations $\delta \in \Delta(\theta)$ with $\theta_P^{\min} = \alpha l_P$ and $\theta_P^{\max} = \beta l_P$ for all $P \in \mathcal{P}$. Given (α, β) -path deviations δ , a flow f is δ -inducible if f is a Nash flow with respect to $l + \delta$ (as defined in (2)).

We obtain the following theorem for general $(0, \beta)$ -path deviations.

Theorem 10 *Let \mathcal{I} be a general multi-commodity instance with $(l_a)_{a \in A} \in \mathcal{L}^A$. Let x be δ -inducible with respect to some $(0, \beta)$ -path deviation δ and let z be an arbitrary feasible flow. If $\hat{\mu}(\mathcal{L}, 0) < 1/(1 + \beta)$, then*

$$C(x) \leq \frac{1 + \beta}{1 - (1 + \beta)\hat{\mu}(\mathcal{L}, 0)} C(z).$$

Proof We know that the flow x satisfies the variational inequality

$$\sum_{P \in \mathcal{P}} x_P [l_P(x) + \delta_P(x)] \leq \sum_{P \in \mathcal{P}} z_P [l_P(x) + \delta_P(x)].$$

It follows that

$$C(x) \leq \sum_{P \in \mathcal{P}} x_P [l_P(x) + \delta_P(x)] \leq \sum_{P \in \mathcal{P}} z_P [l_P(x) + \delta_P(x)] \leq (1 + \beta) \sum_{P \in \mathcal{P}} z_P l_P(x)$$

using the non-negativity of the flow and the deviations. Using the definition of the smoothness parameter $\hat{\mu} := \hat{\mu}(\mathcal{L}, 0)$ in (8), we find that

$$\sum_{P \in \mathcal{P}} z_P l_P(x) = \sum_{a \in A} z_a l_a(x_a) \leq \sum_{a \in A} z_a l_a(z_a) + \sum_{a \in A} \hat{\mu} x_a l_a(x_a) = C(z) + \hat{\mu} C(x).$$

The claim follows by rearranging terms and exploiting that $\hat{\mu} < 1/(1 + \beta)$. □

As a final observation, we draw a connection between general (α, β) -path deviations and approximate Nash flows. Suppose f is δ -inducible with respect to some (α, β) -path deviation δ . The Nash flow conditions in (2) then imply that for every commodity $i \in [k]$ and for every path $P \in \mathcal{P}_i$ with $f_P > 0$, we have

$$(1 + \alpha)l_P(f) \leq l_P(f) + \delta_P(f) \leq l_{P'}(f) + \delta_{P'}(f) \leq (1 + \beta)l_{P'}(f) \quad \forall P' \in \mathcal{P}_i.$$

In particular, the above inequalities imply that f is an $(1 + \beta)/(1 + \alpha)$ -approximate Nash flow (see [4]). As a consequence, the bounds by Christodoulou et al. [4] on the Price of Anarchy for approximate Nash flows in non-atomic routing games with polynomial latency functions, yield upper bounds on the BPoA and DR of instances with polynomial latency functions.

7.2 Heterogeneous Players

As a second generalization, we consider a heterogeneous population of players in which different fractions of players have different attitudes towards the path deviations (as introduced above).

We consider k different *player types* in a single-commodity network (i.e., all player types share the same source and destination). For each type $i \in [k]$ we have a demand r_i and an attitude τ_i towards the deviations. We assume without loss of generality that the demands are normalized such that $\sum_{i \in [k]} r_i = 1$.

In this context, a feasible flow $f = (f_P^i)_{i \in [k], P \in \mathcal{P}}$ is δ -inducible if:

$$\forall i \in [k], \forall P \in \mathcal{P}, f_P^i > 0: \quad l_P(f) + \tau_i \delta_P(f) \leq l_{P'}(f) + \tau_i \delta_{P'}(f) \quad \forall P' \in \mathcal{P}.$$

We show in the next theorem that the deviation ratio is upper bounded by a function that is linear in the weighted average of the sensitivities of the different player types with respect to the deviations.

Theorem 11 *Let \mathcal{I} be a single-commodity instance. Let x be δ -inducible with respect to some $(0, \beta)$ -path deviation δ and let z be a 0-inducible Nash flow. If there is an alternating (s, t) -path π consisting only of arcs in Z , then*

$$C(x) \leq \left(1 + \beta \sum_{i \in [k]} \tau_i r_i \right) C(z).$$

Note that the condition of the alternating path π to consist of arcs in Z only is equivalent to having $\eta = 1$, i.e., π is an actual (s, t) -path in the underlying graph. In particular, this condition is satisfied for *series-parallel* graphs (see, e.g., Corollary 4.8 [16]). This implies that the bound derived above holds for all instances with series-parallel graphs. It would be interesting to see if this bound extends to arbitrary alternating paths. In a recent work [10], we further improved upon the result of Theorem 11.

In the proof below, we use some ideas that have been used in the proof of Lemma 4 [11].

Proof (Theorem 11) For $i \in [k]$, let \bar{P}_i be a path maximizing $l_P(x)$ over all flow-carrying paths $P \in \mathcal{P}$ of type i . We have

$$l_{\bar{P}_i}(x) \leq l_{\bar{P}_i}(x) + \tau_i \delta_{\bar{P}_i}(x) \leq l_{\pi}(x) + \tau_i \delta_{\pi}(x) \leq (1 + \beta \tau_i) l_{\pi}(x) = (1 + \beta \tau_i) \sum_{a \in \pi} l_a(x_a).$$

Note that by definition of the alternating path π , we have $x_a \leq z_a$ for all $a \in \pi$. Continuing with the estimate, we find $l_{\bar{P}_i}(x) \leq (1 + \beta \tau_i) \sum_{a \in \pi} l_a(z_a)$ and thus

$$C(x) \leq \sum_{i \in [k]} r_i l_{\bar{P}_i}(x) \leq \sum_{i \in [k]} r_i \left((1 + \beta \tau_i) \sum_{a \in \pi} l_a(z_a) \right).$$

Note that we can write $C(z) = \sum_{a \in \pi} l_a(z_a)$ because there exists a flow-decomposition of z in which π is flow-carrying (here we use $z_a > 0$ for all $a \in \pi$). We thus obtain

$$C(x) \leq \sum_{i \in [k]} r_i \left((1 + \beta \tau_i) \sum_{a \in \pi} l_a(z_a) \right) = \left(\sum_{i \in [k]} r_i (1 + \beta \tau_i) \right) C(z).$$

The claim follows because $\sum_{i \in [k]} r_i = 1$. □

8 Conclusions

We introduced a unifying model to study the impact of (bounded) worst-case latency deviations in non-atomic selfish routing games. We demonstrated that the Deviation Ratio is a useful measure to assess the cost deterioration caused by such deviations. Among potentially other applications, we showed that the Deviation Ratio provides bounds on the Price of Risk Aversion and the relative error in social cost if the latency functions are subject to small perturbations.

Our approach to bound the Deviation Ratio is quite generic and, albeit considering a rather general setting, enables us to obtain tight bounds. We believe that this approach will turn out to be useful to derive bounds on the Deviation Ratio of other games (e.g., network cost sharing games).

In general, studying the impact of (bounded) worst-case deviations of the input data of more general classes of games (e.g., congestion games) is an interesting and challenging direction for future work.

Appendix A: Proof of Lemma 1

Lemma 1 *Let $-1 < \alpha \leq 0 \leq \beta$ be fixed. Then f is inducible with an (α, β) -deviation if and only if it is inducible with a $(0, \frac{\beta-\alpha}{1+\alpha})$ -deviation.*

Proof Let f be δ -inducible for some $\alpha l \leq \delta \leq \beta l$, and for $a \in A$, write $\delta_a(f_a) = d_a l_a(f_a)$. Without loss of generality we may assume that $\delta_a(x) = d_a l_a(x)$ (since by definition $d_a l_a(x)$ also induces f). From the equilibrium conditions (2), we know that

$$\forall i \in [k], \forall P \in \mathcal{P}_i, f_P > 0 : \sum_{a \in P} l_a(f_a) + \delta_a(f_a) \leq \sum_{a \in P'} l_a(f_a) + \delta_a(f_a) \quad \forall P' \in \mathcal{P}_i.$$

This is equivalent to $\forall i \in [k], \forall P \in \mathcal{P}_i, f_P > 0 :$

$$\sum_{a \in P} \left(1 + \frac{d_a - \alpha}{1 + \alpha} \right) l_a(f_a) \leq \sum_{a \in P'} \left(1 + \frac{d_a - \alpha}{1 + \alpha} \right) l_a(f_a) \quad \forall P' \in \mathcal{P}_i$$

which can be seen by writing

$$l_a(f_a) + \delta_a(f_a) = (1 + d_a)l_a(f_a) = (1 + \alpha + d_a - \alpha)l_a(f_a),$$

and then dividing the inequality by $1 + \alpha$. We then see that δ' , defined by $\delta'_a(x) = \frac{d_a - \alpha}{1 + \alpha} l_a(x)$ for all $a \in A$ and $x \geq 0$, also induces f since

$$\alpha l_a(x) \leq d_a l_a(x) \leq \beta l_a(x) \quad \Leftrightarrow \quad 0 \leq \frac{d_a - \alpha}{1 + \alpha} l_a(x) \leq \frac{\beta - \alpha}{1 + \alpha} l_a(x).$$

□

Appendix B: Computing Optimal Deviations

The bounded deviation model introduced in Section 2.2 naturally gives rise to the following two optimization problems:¹¹

1. *Best deviation problem:* compute a deviation $\delta^* \in \Delta(\theta)$ such that

$$\delta^* = \arg \inf_{\delta \in \Delta(\theta)} \inf \{ C(f^\delta) \mid f^\delta \text{ is } \delta\text{-inducible} \}.$$

¹¹We assume that the infimum and supremum are attained in the set $\Delta(\theta)$; in particular, this is true for (α, β) -deviations.

2. *Worst deviation problem*: compute a deviation $\delta^* \in \Delta(\theta)$ such that

$$\delta^* = \arg \sup_{\delta \in \Delta(\theta)} \sup\{C(f^\delta) \mid f^\delta \text{ is } \delta\text{-inducible}\}.$$

Recall that the social cost function C only takes into account the latencies but not the deviations. A somewhat subtle point here is that for a fixed deviation $\delta \in \Delta(\theta)$, the social cost of a δ -inducible flow might not be unique. In particular, in the best deviation problem we seek a feasible deviation δ such that the social cost of the best Nash flow that is δ -inducible is minimized (similar as in [2]). In contrast, in the worst deviation problem we want to determine a feasible deviation δ such that the social cost of the worst Nash flow that is δ -inducible is maximized.

Below we elaborate on relations between the best deviation problem and various network toll problems. As a side result, we also show that the worst deviation problem is NP-hard, even for single-commodity instances with linear latencies (Theorem 12).

B.1 Relations to Network Toll Problems

The best deviation problem is a direct generalization of the *restricted network toll problem* introduced by Bonifaci et al. [2]. We obtain this model for $\theta^{\min} = 0$. The deviations are interpreted as non-negative tolls on the arcs. The objective minimized in [2] is measured against the social optimum, i.e., the authors are interested in the ratio $C(f^\delta)/C(f^*)$, where f^* is an optimal flow for the instance \mathcal{I} . Also, our definition of $(0, \beta)$ -deviations is equivalent to the definition of β -restricted tolls in [2].

The work by Fotakis et al. [8] can technically be seen as a variant of the restricted network toll problem in which the tolls are interpreted as risk-averse behavior of players. Here, we have $\theta_a^{\min} = 0$ and $\theta_a^{\max} = \gamma l_a$ for all $a \in A$. The authors consider deviations of the form $\delta_a(x) = \gamma_a l_a(x)$ for $0 \leq \gamma_a \leq \gamma$ for all $a \in A$. In particular, deviations of this form induce an approximate Nash flow as studied by Christodoulou et al. [4]. For example, if all latency functions in the network are polynomials of degree at most d , then we obtain a γd -approximate Nash flow.

Hoeyer et al. [9] consider the *taxing subnetwork problem*, which is a special case of the restricted network toll problem. Here only a designated subset of the arcs can be tolled, which is equivalent to $\theta_a^{\min} = 0$ and $\theta_a^{\max} \in \{0, \infty\}$ for all $a \in A$. They show that best deviation problem is NP-complete, even for two commodities. To the best of our knowledge, the single-commodity case is still an open problem. On the positive side, Hoeyer et al. [9] and Bonifaci et al. [2] give polynomial time algorithms for parallel-arc networks, solving the best deviation problem for their respective definitions of the threshold functions.

Beckmann et al. [1] proved that the social optimum can be induced as a Nash flow using *marginal tolls*, that is, by setting $\delta_a(x) = x \cdot l'_a(x)$, where $l'_a(x)$ is the derivative of $l_a(x)$ (assuming the existence of l'_a). In particular, if these tolls are feasible, i.e., $\delta \in \Delta(\theta)$, then δ is an optimal solution for the best deviation problem.

There are several models that study deviations in the form of *scaled marginal tolls*, i.e., deviations defined by $\delta_a(x) = \rho x l'_a(x)$ for some $\rho \in \mathbb{R}$. We elaborate on two such models in more detail:

In the standard non-atomic routing model it is assumed that players are completely selfish in the sense that they want to minimize their own latencies. However, more recently researchers also considered settings where players are (*partially*) *altruistic*. Chen et al. [3], for example, model such altruistic behavior by including scaled marginal tolls in the objective of the players. In particular, they study scaled marginal tolls with $-1 \leq \rho \leq 1$.

Meir and Parkes [13] also study deviations in the form of scaled marginal tolls, which are interpreted as *behavioral biases* towards the marginal tolls. A conceptual difference here is that the parameter ρ is chosen by the players, instead of the system designer (as, for example, in the restricted network toll model). Here, the authors are also interested in the case $\rho \geq 1$ (which is less relevant in the other models). The authors also study this model in [14], where these deviations are interpreted as *distance-based strict uncertainty*.

B.2 Hardness of the Worst Deviation Problem

As a side-result, we prove that the problem of determining worst-case deviations is NP-hard.

Theorem 12 *It is NP-hard to compute deviations $\delta \in \Delta(\theta)$ such that $C(f^\delta)$ is maximized, even for single-commodity networks with linear latencies.*

Proof We give a reduction from the DIRECTED HAMILTONIAN s, t -PATH problem: We are given a directed graph $G = (V, A)$, and fixed $s, t \in V$, and the goal is to decide whether or not there exists a simple directed s, t -path in G that visits every node exactly once. Let \mathcal{J} be an instance of DIRECTED HAMILTONIAN s, t -PATH problem.

Now, define an instance \mathcal{I} of the bounded deviation model on the graph G by taking $l_a(x) = x$ for all $a \in A$, $\theta_a^{\min} = 0$ for all $a \in A$, and $\theta_a^{\max} = n - 1$ for all $a \in A$. Furthermore, take $r = 1$.

We claim that G has a Hamiltonian path from s to t if and only if there is a deviation $\delta \in \Delta(\theta)$ such that $C(f^\delta) \geq n - 1$. First, let G have a Hamiltonian path P from s to t , and define δ by $\delta_a = 0$ if $a \in P$, and $\delta_a = n - 1$ otherwise. We then have that f^δ is given by $f_a^\delta = 1$ if $a \in P$ and $f_a^\delta = 0$ otherwise, since the perceived latency along P is then equal to $l_P(f^\delta) = n - 1$, and any other path P' uses at least one different arc $a' \notin P$, which gives us that

$$l_{P'}(f^\delta) + \delta_{P'}(f^\delta) \geq l_{a'}(f^\delta) + \delta_{a'}(f^\delta) \geq n - 1 = l_P(f^\delta) + \delta_P(f^\delta).$$

Note that f^δ is the unique Nash flow in this case (since all the perceived latencies $l_a + \delta_a$ are strictly increasing).

Conversely, suppose there is a $\delta \in \Delta(\theta)$ such that $C(f^\delta) \geq n - 1$. For any *feasible* flow g we have that $l_P(g) \leq n - 1$, with strict inequality if $f_P^\delta < 1$ (since then there will be at least one arc $a \in P$ with $f_a^\delta < 1$). This means that

$$C(g) = \sum_{P \in \mathcal{P}} g_P l_P(g) \leq \sum_{P \in \mathcal{P}} g_P (n - 1) = n - 1,$$




Fig. 6 All the values of l_a , θ_a^{\min} and θ_a^{\max} that are not explicitly stated are zero

using that $r = 1$. Again, we have strict inequality if $0 < g_P < 1$ for some path P , i.e., if not all players use the same path. This means that for f^δ there is at most one path P^* with $f_{P^*}^\delta > 0$, which then implies that $f_{P^*}^\delta = 1$. Furthermore, we can conclude that $|A(P^*)| = l_{P^*}(f^\delta) = C(f^\delta) = n - 1$, which implies that P^* is a Hamiltonian path from s to t , since it is a simple path by assumption. \square

Appendix C: Necessity of Common Source Assumption in Theorem 2

The example below shows that Theorem 2 does not hold if the assumption that all commodities share a common source is dropped.

Example 2 Consider the graph $G = (V, A)$ in Fig. 6 and suppose that $r_1 = r_2 = 1$. Then the flow f that routes one unit of flow over both paths $(s_1, v_1, 1, 2, t_1)$ and $(s_2, v_2, 3, 4, t_2)$ is feasible and inducible (take $\delta = 0$). However, looking at the graph $\hat{G}(f)$, we obtain a negative cost cycle $(1, 4, 3, 2, 1)$ (by using the reversed arcs of $(1, 2)$ and $(3, 4)$).

References

1. Beckmann, M., McGuire, B., Winsten, C.: Studies in the Economics of Transportation. Yale University Press, New Haven (1956)
2. Bonifaci, V., Salek, M., Schäfer, G.: On the efficiency of restricted tolls in network routing games. Lecture Notes in Computer Science (2011)
3. Chen, P.-A., Kempe, D.: Altruism, selfishness, and spite in traffic routing. In: Proceedings of the 9th ACM conference on electronic commerce, pp. 140–149. ACM (2008)
4. Christodoulou, G., Koutsoupias, E., Spirakis, P.G.: On the performance of approximate equilibria in congestion games. Algorithmica **61**(1), 116–140 (2011)
5. Cominetti, R.: Equilibrium routing under uncertainty. Math. Program. **151**(1), 117–151 (2015)
6. Correa, J.R., Schulz, A.S., Stier-Moses, N.E.: A geometric approach to the price of anarchy in nonatomic congestion games. Games and Economic Behavior **64**(2), 457–469 (2008). Special Issue in Honor of Michael B. Maschler
7. Englert, M., Franke, T., Olbrich, L.: Sensitivity of Wardrop Equilibria, pp. 158–169. Springer, Berlin (2008)
8. Fotakis, D., Kalimeris, D., Lianas, T.: Improving selfish routing for risk-averse players. In: Proceedings of Web and Internet Economics - 11th International Conference, WINE 2015, Amsterdam, the netherlands, december 9–12, 2015, pp. 328–342 (2015)

9. Hoeyer, M., Olbrich, L., Skopalik, A.: Taxing subnetworks. In: Papadimitriou, C.H., Zhang, S. (eds.) WINE, volume 5385 of Lecture Notes in Computer Science, pp. 286–294. Springer (2008)
10. Kleer, P., Schäfer, G.: Path Deviations Outperform Approximate Stability in Heterogeneous Congestion Games, pp. 212–224. Springer International Publishing Cham, Berlin (2017)
11. Lianas, T., Nikolova, E., Stier-Moses, N.E.: Asymptotically tight bounds for inefficiency in risk-averse selfish routing. CoRR, arXiv:[1510.02067](https://arxiv.org/abs/1510.02067) (2015)
12. Lin, H., Roughgarden, T., Tardos, É., Walkover, A.: Stronger bounds on Braess's paradox and the maximum latency of selfish routing. *SIAM J. Discret. Math.* **25**(4), 1667–1686 (2011)
13. Meir, R., Parkes, D.: Playing the wrong game smoothness bounds for congestion games with behavioral biases. *SIGMETRICS Perform. Eval. Rev.* **43**(3), 67–70 (2015)
14. Meir, R., Parkes, D.C.: Congestion games with distance-based strict uncertainty. CoRR, arXiv:[1411.4943](https://arxiv.org/abs/1411.4943) (2014)
15. Nikolova, E., Stier-Moses, N.E.: A mean-risk model for the traffic assignment problem with stochastic travel times. *Oper. Res.* **62**(2), 366–382 (2014)
16. Nikolova, E., Stier-Moses, N.E.: The Burden of Risk Aversion in Mean-Risk Selfish Routing. In: Proceedings of the 16th ACM Conference on Economics and Computation, EC '15, pp. 489–506. ACM, New York (2015)
17. Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V.: Algorithmic Game Theory. Cambridge University Press, New York (2007)
18. Roughgarden, T.: On the severity of Braess's paradox: Designing networks for selfish users is hard. *J. Comput. Syst. Sci.* **72**(5), 922–953 (2006)
19. Roughgarden, T.: Intrinsic robustness of the price of anarchy. *J. ACM* **62**(5), 32 (2015)
20. Wardrop, J.G.: Some theoretical aspects of road traffic research. *Proc. Inst. Civ. Eng.* **1**, 325–378 (1952)