

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

Pieter Kleer¹ and Guido Schäfer^{1,2}

¹ Centrum Wiskunde & Informatica (CWI), Networks and Optimization Group, Amsterdam, The Netherlands

² Vrije Universiteit Amsterdam, Department of Econometrics and Operations Research, Amsterdam, The Netherlands.
kleer@cwi.nl, schaefer@cwi.nl

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

1 Introduction

In the classical selfish routing game introduced by Wardrop [12], 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 given deterministically. 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 which are taken into account by the players.

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 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 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 [7, 9].

In order to model bounded deviations, we extend an idea previously put forward by Bonifaci, Salek and Schäfer [1] in the context of the *restricted network toll problem*: We assume that for every arc $a \in A$ we are given 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 [9]. 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).

Our contributions. The main contributions presented in this paper are as follows:

1. *Upper bounds:* 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 3).

In order to prove this upper bound, we first generalize a result by Bonifaci et al. [1] 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 plays a crucial role in the work by Nikolova and Stier-Moses [9] and was first used by Lin, Roughgarden, Tardos and Walkover [6] in the context of the *network design problem*.

We then specialize our bound to the case of so-called (α, β) -*deviations*, where $\theta_a^{\min} = \alpha l_a$ and $\theta_a^{\max} = \beta l_a$ with $-1 < \alpha \leq 0 \leq \beta$. We prove that the Deviation Ratio is at most $1 + (\beta - \alpha)/(1 + \alpha) \lceil (n - 1)/2 \rceil r$, where n is the number of nodes of the network and r is the sum of the demands of the commodities (Theorem 3). 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 [9] in two ways: (i) it holds for multi-commodity networks with a common source and (ii) it allows for negative risk-aversion parameters (i.e., capturing risk-taking players as well). Further, we show that our result can be used to bound the relative error in social cost incurred by small latency perturbations (Theorem 7), which is of independent interest.

2. *Lower bounds:* We prove that our bound on the Deviation Ratio for (α, β) -deviations is best possible. 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* [10] (Example 1). In particular, this complements a recent result by Lianas, Nikolova and Stier-Moses [5] who show that their bound on the Price of Risk Aversion 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 4). Finally, for general multi-commodity graphs we establish a lower bound showing that the Deviation Ratio can be exponential in n (Theorem 5). 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, Roughgarden, Tardos and Walkover [6] in their lower bound construction for the network design problem on multi-commodity networks (see also [10]).

3. *Smoothness bounds:* We improve (and slightly generalize) recent smoothness bounds on the Price of Risk Aversion given by Meir and Parkes [7] and independently by Lianas et al. [5]. In particular, we derive tight bounds for the *Biased Price of Anarchy (BPoA)* [7], 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 8).³ Note that the Biased Price of Anarchy yields an upper bound on the Deviation Ratio/Price of Risk Aversion. We also derive smoothness results for general path deviations (which are not representable by arc deviations). As a result, we obtain bounds on the Price of Risk Aversion (Theorem 9) under the non-linear *mean-std* model [5, 9] (see Section 2).

It is interesting to note that the smoothness bounds on the Biased Price of Anarchy [7] and the Price of Risk Aversion [5] are independent of the network structure (but dependent on the class of latency functions). In contrast, the bound on the Deviation Ratio depends on certain parameters of the network.⁴

Our results answer a question posed in the work by Nikolova and Stier-Moses [9] regarding possible relations between their Price of Risk Aversion model [9], the restricted network toll problem [1], and the network design problem [10]. In particular, our results also show that the analysis in [9] is not inherent to the used variance function, but rather depends on the restrictions imposed on the feasible deviations.

Related work. The modeling and studying of uncertainties in routing games has received a lot of attention in recent years. An extensive survey on this topic is given by Cominetti [2].

³ We remark that for certain types of $(0, \beta)$ -deviations, e.g., *scaled marginal tolls*, better bounds can be obtained (see, e.g., [7]).

⁴ For example, there are parallel-arc networks for which the Biased Price of Anarchy is unbounded, whereas the Deviation Ratio is a constant.

As mentioned above, our investigations are inspired by the study of the Price of Risk Aversion by Nikolova and Stier-Moses [9]. 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 elaborate in more detail on the connections to their work in Section 2.

There are several papers that study the problem of imposing tolls (which can be viewed as latency deviations) on the arcs of a network to reduce the cost of the resulting Nash flow. Conceptually, our model is related to the *restricted network toll problem* by Bonifaci et al. [1]. 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 1).

Roughgarden [10] studies the *network design problem* of finding a subnetwork that minimizes the latency of all flow-carrying paths of the resulting Nash flow. He proves 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 $P = NP$). Later, Lin et al. [6] 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 [6, 10].

Meir and Parkes [7] and independently Lianas et al. [5] show that for non-atomic network routing games with $(1, \mu)$ -smooth⁵ 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 [9] which hold for arbitrary non-negative and non-decreasing latency functions.

2 Preliminaries

Bounded deviation model. Let $\mathcal{I} = (G = (V, A), (l_a)_{a \in A}, (s_i, t_i)_{i \in [k]}, (r_i)_{i \in [k]})$ be an instance of a non-atomic network routing game. 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 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 $r = \sum_{i \in [k]} r_i$.

⁵ Meir and Parkes [7] define a function l to be $(1, \mu)$ -smooth if $xl(y) \leq \mu y l(y) + xl(x)$ for all $x, y \geq 0$ (which is slightly different from Roughgarden's original smoothness definition [11]). Lianas et al. [5] only require *local smoothness* where y is taken fixed.

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 $f : \mathcal{P} \rightarrow \mathbb{R}_{\geq 0}$, i.e., $\sum_{P \in \mathcal{P}_i} f_P = r_i$ 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 average 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 .

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 $\delta = (\delta_a)_{a \in A}$. 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 $q_a(f_a) = l_a(f_a) + \delta_a(f_a)$; similarly, the deviated latency on path $P \in \mathcal{P}$ is given by $q_P(f) = l_P(f) + \delta_P(f)$. 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.,

$$\forall i \in [k], \forall P \in \mathcal{P}_i, f_P > 0 : \quad q_P(f) \leq q_{P'}(f) \quad \forall P' \in \mathcal{P}_i. \quad (1)$$

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 continuous threshold functions satisfying $\theta_a^{\min}(x) \leq 0 \leq \theta_a^{\max}(x)$ 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 $\theta = (\alpha, \beta)$. Throughout the paper, we assume that the deviated latencies are always non-negative, i.e., $l_a(x) + \theta_a^{\min}(x) \geq 0$ for all $x \geq 0$ and $a \in A$.

We (implicitly) assume that only deviations δ are considered for which a Nash flow exists. We briefly 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 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 $\delta' \in \Delta(\theta)$ defined by $\delta'_a = \lambda_a \theta_a^{\max}$ 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 $a \in A$. As a consequence, it follows that $q_a = l_a + \delta_a$ is a non-negative, non-decreasing and continuous function for all $a \in A$. It is well-known that for these types of functions, the existence of a Nash flow is guaranteed.

Deviation Ratio. Given an instance \mathcal{I} and threshold functions $\theta = (\theta^{\min}, \theta^{\max})$, we define the *Deviation Ratio* $DR(\mathcal{I}, \theta) = \sup_{\delta \in \Delta(\theta)} C(f^\delta)/C(f^0)$ as the worst-case ratio of the cost of a θ -inducible flow and the cost of a 0-inducible flow. Intuitively, $DR(\mathcal{I}, \theta)$ measures the worst-case deterioration of the social cost of a Nash flow due to (feasible) latency deviations.

We emphasize that the social cost function C is defined as above, i.e., with respect to the latencies (not taking into account the deviations). Note that for fixed deviations $\delta \in \Delta(\theta)$, there might be multiple Nash flows that are δ -inducible. In this case, we adopt the convention that $C(f^\delta)$ refers to the social cost of the worst Nash flow that is δ -inducible.

Our main focus in this paper is on establishing (tight) bounds on the Deviation Ratio. As a side-result, we prove that the problem of determining worst-case deviations is NP-hard.

Theorem 1. *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.*

Related notions. Nikolova and Stier-Moses [9] (see also [5,8]) 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: $v_P(f) = \sum_{a \in P} v_a(f_a)$, called the *mean-var* objective, and $v_P(f) = (\sum_{a \in P} v_a(f_a))^{1/2}$, called the *mean-std* objective. 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* [9] 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 .⁶ 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 $\text{PRA}(\mathcal{I}, \gamma, \kappa)$ of single-commodity instances \mathcal{I} with non-negative and non-decreasing latency functions is at most $1 + \gamma\kappa[(n-1)/2]$, where n is the number of nodes.

We now elaborate on the relation to our Deviation Ratio. The main technical difference is that in [9] 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 [7]. 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 $\text{DR}(\mathcal{I}, \theta) \leq \text{BPoA}(\mathcal{I}, \theta)$.

Due to space limitations, some material is omitted from this extended abstract and can be found in the full version of the paper (see [4]).

⁶ The existence of a risk-averse Nash flow is proven in [8].

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.

We first derive a characterization result for the inducibility of a given flow f . This generalizes the characterization in [1] to common source multi-commodity instances and negative deviations. We define an *auxiliary graph* $\hat{G} = \hat{G}(f) = (V, \hat{A})$ with $\hat{A} = A \cup \bar{A}$, where $\bar{A} = \{(v, u) : a = (u, v) \in A^+\}$. That is, \hat{A} consists of the set of arcs in A , which we call *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:

$$c_a = \begin{cases} l_{(u,v)}(f_a) + \theta_{(u,v)}^{\max}(f_a) & \text{for } a = (u, v) \in A \\ -l_{(u,v)}(f_a) - \theta_{(u,v)}^{\min}(f_a) & \text{for } a = (v, u) \in \bar{A}. \end{cases} \quad (2)$$

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

Theorem 2 does not hold for general multi-commodity instances. The proof of Lemma 1 follows directly from Theorem 2.

Lemma 1. *Let x be θ -inducible and let X_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 $\hat{G}(x)$, respectively. Then*

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

The following notion of alternating paths turns out to be crucial. It was first introduced by Lin et al. [6] and is also used by Nikolova and Stier-Moses [9].

Definition 1 (Alternating path [6, 9]). *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 .*

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 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*).

The following lemma proves the existence of an *alternating path tree*, i.e., a spanning tree of alternating paths, rooted at the common source node s . It is a direct generalization of Lemma 4.6 in [6] and Lemma 4.5 in [9].

Lemma 2. *Let z and x be feasible flows and let Z and X be a partition of A as in Definition 1. Then there exists an alternating path tree.*

We now have all the ingredients to prove the following main result.

Theorem 3. *Let x be θ -inducible and let z be 0-inducible. Further, let $A = X \cup Z$ be a partition 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 $P \in \mathcal{P}_i$.⁷ Then*

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

(ii) *Suppose $\theta = (\alpha, \beta)$ with $-1 < \alpha \leq 0 \leq \beta$. Let η_i is the number of disjoint segments of consecutive arcs in Z on the alternating s, t_i -path π_i for $i \in [k]$.⁸ Then*

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

Proof (i). We have $C(x) = \sum_i \sum_{P \in \mathcal{P}_i} x_P^i l_P(x) \leq \sum_i r_i \sum_{a \in X_i} l_a(x_a)$ by the choice of X_i . By applying the first inequality of Lemma 1 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

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

Let Z_i be an arbitrary flow-carrying path of commodity $i \in [k]$ with respect to z . By applying the second inequality of Lemma 1 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

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

Combining these inequalities and exploiting the definition of X and Z , we obtain

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

The claim now follows by multiplying the above inequality with r_i and summing over all commodities $i \in [k]$. Note that $C(z) = \sum_i r_i \sum_{a \in Z_i} l_a(z_a)$. \square

⁷ Note 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 $l_P(x)$.

⁸ Note that $\eta_i \leq \lceil (n-1)/2 \rceil$.




Fig. 1. 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 for (α, β) -deviations

We show that the bound in Theorem 3 is tight in all its parameters for (α, β) -deviations. We start with single-commodity instances.

Our instance is based on the generalized Braess graph [10]. 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 $E_3^m = \{(v_1, t) \cup \{(s, w_{m-1})\}\}$.

Example 1. We can assume without loss of generality that $\alpha = 0$ (see [4]). 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 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 Figure 1).

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 $C(z) = 1$.

⁹ Note 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.

¹⁰ For 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 $1/(m-1)$.

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 $C(x)/C(z) = 1 + \beta m = 1 + \beta n/2$. \square

By adapting the construction above, we obtain the following result.

Theorem 4. *There exist common source two-commodity instances \mathcal{I} such that*

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

For two-commodity instances and n even, we can actually improve the upper bound in Theorem 3 to the lower bound stated in Theorem 4 (see [4]).

For general multi-commodity instances the situation is much worse. In particular, we establish an exponential lower bound on the Deviation Ratio. The instance used in proof of Theorem 5 is similar to the one used by Lin et al. [6].

Theorem 5. *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.*

5 Applications

By using our bounds on the Deviation Ratio, we obtain the following results.

Price of Risk Aversion.

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 $n = 2m$ (see [4]). In particular, for single-commodity instances we obtain tightness for all $n \in \mathbb{N}$.

Stability of Nash flows under small perturbations.

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 slightly perturbed latency functions $(\tilde{l}_a)_{a \in A}$ satisfying $\sup_{a \in A, x \geq 0} |(l_a(x) - \tilde{l}_a(x))/l_a(x)| \leq \epsilon$ for some small $\epsilon > 0$. Then the relative error in social cost is $(C(\tilde{f}) - C(f))/C(f) \leq 2\epsilon/(1 - \epsilon) \lceil (n - 1)/2 \rceil \cdot r = \mathcal{O}(\epsilon rn)$.*

6 Smoothness based approaches

We derive tight smoothness bounds on the Biased Price of Anarchy for $(0, \beta)$ -deviations. Our bounds improve upon the bounds of $(1 + \beta)/(1 - \mu)$ recently obtained by Meir and Parkes [7] and Lianas et al. [5] for $(1, \mu)$ -smooth latency functions. As a direct consequence, we also obtain better smoothness bounds on the Price of Risk Aversion. Our approach is a generalization of the framework of Correa, Schulz and Stier-Moses [3] (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

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

Theorem 8. *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. Then $C(x)/C(z) \leq (1 + \beta)/(1 - \hat{\mu}(\mathcal{L}, \beta))$ if $\hat{\mu}(\mathcal{L}, \beta) < 1$. 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 $\gamma \geq 0$, $\gamma l \in \mathcal{L}$.*

For example, for affine latencies $\hat{\mu}(\mathcal{L}, \beta) = 1/(4(1 + \beta))$ (see [4]) and we obtain a bound of $(1 + \beta)^2/(\frac{3}{4} + \beta)$ on the Biased Price of Anarchy, which is strictly better than the bound $4(1 + \beta)/3$ obtained in [5, 7].

We also provide an upper bound on the absolute gap between the Biased Price of Anarchy and the Deviation Ratio (see [4]).

As a final result we derive smoothness bounds for general path deviations, which are not necessarily decomposable into arc deviations (see [4] for formal definitions). 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 [9] (see Section 2).

Theorem 9. *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 an arbitrary feasible flow. If $\hat{\mu}(\mathcal{L}, 0) < 1/(1 + \beta)$, then $C(x)/C(z) \leq (1 + \beta)/(1 - (1 + \beta)\hat{\mu}(\mathcal{L}, 0))$.*

7 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 (see Section 3) 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.

References

1. V. Bonifaci, M. Salek, and G. Schäfer. Efficiency of restricted tolls in non-atomic network routing games. In *Proc. of the International Symposium on Algorithmic Game Theory*, pages 302–313. 2011.
2. R. Cominetti. Equilibrium routing under uncertainty. *Mathematical Programming*, 151(1):117–151, 2015.
3. J. R. Correa, A. S. Schulz, and N. E. Stier-Moses. A geometric approach to the price of anarchy in nonatomic congestion games. *Games and Economic Behavior*, 64(2):457–469, 2008.
4. P. Kleer and G. Schäfer. The impact of worst-case deviations in non-atomic network routing games. *CoRR*, abs/1605.01510, 2016.
5. T. Lianeas, E. Nikolova, and N. E. Stier-Moses. Asymptotically tight bounds for inefficiency in risk-averse selfish routing. *CoRR*, abs/1510.02067, 2015.
6. H. Lin, T. Roughgarden, É. Tardos, and A. Walkover. Stronger bounds on Braess’s paradox and the maximum latency of selfish routing. *SIAM Journal on Discrete Mathematics*, 25(4):1667–1686, 2011.
7. R. Meir and D. Parkes. Playing the wrong game: Smoothness bounds for congestion games with behavioral biases. *SIGMETRICS Performance Evaluation Review*, 43(3):67–70, 2015.
8. E. Nikolova and N. E. Stier-Moses. A mean-risk model for the traffic assignment problem with stochastic travel times. *Operations Research*, 62(2):366–382, 2014.
9. E. Nikolova and N. E. Stier-Moses. The burden of risk aversion in mean-risk selfish routing. In *Proc. of the ACM Conference on Economics and Computation*, pages 489–506, 2015.
10. T. Roughgarden. On the severity of Braess’s paradox: Designing networks for selfish users is hard. *Journal of Computer and System Sciences*, 72(5):922–953, 2006.
11. T. Roughgarden. Intrinsic robustness of the price of anarchy. *Journal of the ACM*, 62(5):32, 2015.
12. J. G. Wardrop. Some theoretical aspects of road traffic research. *Proc. of the Institution of Civil Engineers*, 1:325–378, 1952.