Path Deviations Outperform Approximate Stability in Heterogeneous Congestion Games

Pieter Kleer[1] and Guido Schäfer[1,2](✉)

[1] Centrum Wiskunde & Informatica (CWI), Networks and Optimization Group, Amsterdam, The Netherlands
{kleer,schaefer}@cwi.nl

[2] Department of Econometrics and Operations Research, Vrije Universiteit Amsterdam, Amsterdam, The Netherlands

Abstract. We consider non-atomic network congestion games with heterogeneous players where the latencies of the paths are subject to some bounded deviations. This model encompasses several well-studied extensions of the classical Wardrop model which incorporate, for example, risk-aversion, altruism or travel time delays. Our main goal is to analyze the worst-case deterioration in social cost of a *deviated Nash flow* (i.e., for the perturbed latencies) with respect to an original Nash flow.

We show that for homogeneous players deviated Nash flows coincide with approximate Nash flows and derive tight bounds on their inefficiency. In contrast, we show that for heterogeneous populations this equivalence does not hold. We derive tight bounds on the inefficiency of both deviated and approximate Nash flows for *arbitrary* player sensitivity distributions. Intuitively, our results suggest that the negative impact of path deviations (e.g., caused by risk-averse behavior or latency perturbations) is less severe than approximate stability (e.g., caused by limited responsiveness or bounded rationality).

We also obtain a tight bound on the inefficiency of deviated Nash flows for matroid congestion games and homogeneous populations if the path deviations can be decomposed into edge deviations. In particular, this provides a tight bound on the Price of Risk-Aversion for matroid congestion games.

1 Introduction

In 1952, Wardrop [17] introduced a simple model, also known as the *Wardrop model*, to study outcomes of selfish route choices in traffic networks which are affected by congestion. In this model, there is a continuum of non-atomic players, each controlling an infinitesimally small amount of flow, whose goal is to choose paths in a given network to minimize their own travel times. The latency (or delay) of each edge is prescribed by a non-negative, non-decreasing latency function which depends on the total flow on that edge. Ever since its introduction, the Wardrop model has been used extensively, both in operations research and traffic engineering studies, to investigate various aspects of selfish routing in networks.

V. Bilò and M. Flammini (Eds.): SAGT 2017, LNCS 10504, pp. 212–224, 2017.
DOI: 10.1007/978-3-319-66700-3_17

More recently, the classical Wardrop model has been extended in various ways to capture more complex player behaviors. Examples include the incorporation of uncertainty attitudes (e.g., risk-aversion, risk-seeking), cost alterations (e.g., latency perturbations, road pricing), other-regarding dispositions (e.g., altruism, spite) and player biases (e.g., responsiveness, bounded rationality).

Several of these extensions can be viewed as defining some modified cost for each path which combines the original latency with some 'deviation' (or perturbation) along that path. Such deviations are said to be β-*bounded* if the total deviation along each path is at most β times the latency of that path. The player objective then becomes to minimize the combined cost of latency and deviation along a path (possibly using different norms). An equilibrium outcome corresponds to a β-*deviated Nash flow*, i.e., a Nash flow with respect to the combined cost. The deviations might be given explicitly (e.g., as in the altruism model of Chen et al. [1]) or be defined implicitly (e.g., as in the risk-aversion model of Nikolova and Stier-Moses [13]). Further, different fractions of players might perceive these deviations differently, i.e., players might be heterogeneous with respect to the deviations.

Another extension, which is closely related to the one above, is to incorporate different degrees of 'responsiveness' of the players. For example, each player might be willing to deviate to an alternative route only if her latency decreases by at least a certain fraction. In this context, an equilibrium outcome corresponds to an ϵ-*approximate Nash flow* for some $\epsilon \geq 0$, i.e., for each player the latency is at most $(1 + \epsilon)$ times the latency of any other path. Here, ϵ is a parameter which reflects the responsiveness of the players. An analogue definition can be given for populations with heterogeneous responsiveness parameters.

To illustrate the relation between deviated and approximate Nash flows, suppose we are given a β-deviated Nash flow f for some $\beta \geq 0$, where the latency $\ell_P(f)$ of each path P is perturbed by an arbitrary β-bounded deviation $\delta_P(f)$ satisfying $0 \leq \delta_P(f) \leq \beta l_P(f)$. Intuitively, the deviations inflate the latency on each path by at most a factor of $(1 + \beta)$. Further, assume that the population is homogeneous. From the Nash flow conditions (see Sect. 2 for formal definitions), it follows trivially that f is also an ϵ-approximate Nash flow with $\epsilon = \beta$. But does the converse also hold? That is, can every ϵ-approximate Nash flow be induced by a set of bounded path deviations? More generally, what about the relation between deviated and approximate Nash flows for heterogenous populations? Can we bound the inefficiency of these flows?

In this paper, we answer these questions by investigating the relation between the two equilibrium notions. Our main goal is to quantify the inefficiency of deviated and approximate Nash flows, both for homogeneous and heterogeneous populations. To this aim, we study the (relative) worst-case deterioration in social cost of a β-deviated Nash flow with respect to an original (unaltered) Nash flow; we use the term β-*deviation ratio* to refer to this ratio. This ratio has recently been studied in the context of risk aversion [9,13] and in the more general context of bounded path deviations [6]. Similarly, for approximate Nash flows we are interested in bounding the ϵ-*stability ratio*, i.e., the worst-case deterioration in social cost of an ϵ-approximate Nash flow with respect to an original Nash flow.

Note that these notions differ from the classical *price of anarchy* notion [8], which refers to the worst-case deterioration in social cost of a β-deviated (respectively, ε-approximate) Nash flow with respect to an *optimal* flow. While the price of anarchy typically depends on the class of latency functions (see, e.g., [1,2,6,13] for results in this context), the deviation ratio is independent of the latency functions but depends on the topology of the network (see [6,13]).

Our Contributions. The main contributions of this paper are as follows:

1. We show that for homogeneous populations the set of β-deviated Nash flows coincides with the set of ϵ-approximate Nash flows for $\beta = \epsilon$. Further, we derive an upper bound on the ϵ-stability ratio (and thus also on the ϵ-deviation ratio) which is at most $(1+\epsilon)/(1-\epsilon n)$, where n is the number of nodes, for single-commodity networks. We also prove that the upper bound we obtain is tight for *generalized Braess graphs*. These results are presented in Sect. 4.
2. We prove that for heterogenous populations the above equivalence does not hold. We derive tight bounds for both the β-deviation ratio and the ϵ-stability ratio for single-commodity instances on series-parallel graphs and arbitrary sensitivity distributions of the players. To the best of our knowledge, these are the first inefficiency results in the context of heterogenous populations which are tight for *arbitrary* sensitivity distributions. Our bounds show that both ratios depend on the demands and sensitivity distribution γ of the heterogenous players (besides the respective parameters β and ϵ). Further, it turns out that the β-deviation ratio is always at most the ϵ-stability ratio for $\epsilon = \beta\gamma$. These results are given in Sect. 3.
3. We also derive a tight bound on the β-deviation ratio for single-commodity matroid congestion games and homogeneous populations if the path deviations can be decomposed into edge deviations. To the best of our knowledge, this is the first result in this context which goes beyond network congestion games. In particular, this gives a tight bound on the Price of Risk-Aversion [13] for matroid congestion games. This result is of independent interest and presented in Sect. 4.

In a nutshell, our results reveal that for homogeneous populations there is no quantitative difference between the inefficiency of deviated and approximate Nash flows in the worst case. In contrast, for heterogenous populations the β-deviation ratio is always at least as good as the ϵ-stability ratio with $\epsilon = \beta\gamma$. Intuitively, our results suggest that the negative impact of path deviations (e.g., caused by risk-averse behavior or latency perturbations) is less severe than approximate stability (e.g., caused by limited responsiveness or bounded rationality).

Related Work. We give a brief overview of the works which are most related to our results. Christodoulou et al. [2] study the inefficiency of approximate equilibria in terms of the price of anarchy and price of stability (for homogeneous populations). Generalized Braess graphs were introduced by Roughgarden [14] and are used in many other lower bound constructions (see, e.g., [3,6,14]). Chen et al. [1] study an altruistic extension of the Wardrop model and, in particular, also consider heterogeneous altruistic populations. They obtain an upper bound on the ratio between an altruistic Nash flow and a social optimum for parallel graphs, which is tight for two sensitivity classes. It is mentioned that this bound is most likely not tight in general. Meir and Parkes [11] study player-specific cost functions in a smoothness framework [15]. Some of their inefficiency results are tight, although none of their bounds seems to be tight for arbitrary sensitivity distributions. Matroids have also received some attention in the Wardrop model. In particular, Fujishige et al. [5] show that matroid congestion games are immune against the Braess paradox (and their analysis is tight in a certain sense). We refer the reader to [6] for additional references and relations of other models to the bounded path deviation model considered here.

2 Preliminaries

Let $\mathcal{I} = (E, (l_e)_{e\in E}, (\mathcal{S}_i)_{i\in[k]}, (r_i)_{i\in[k]})$ be an instance of a non-atomic congestion game. Here, E is the set of resources (or edges, or arcs) that are equipped with a non-negative, non-decreasing, continuous latency function $l_e : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$. Each commodity $i \in [k]$ has a strategy set $\mathcal{S}_i \subseteq 2^E$ and demand $r_i \in \mathbb{R}_{>0}$. Note that in general the strategy set $\mathcal{S}_i$ of player i is defined by arbitrary resource subsets. If each strategy $P \in \mathcal{S}_i$ corresponds to an s_i, t_i-path in a given directed graph, then the corresponding game is called a *network* congestion game.[1] We slightly abuse terminology and use the term *path* also to refer to a strategy $P \in \mathcal{S}_i$ of player i (which does not necessarily correspond to a path in a graph); no confusion shall arise. We denote by $\mathcal{S} = \cup_i \mathcal{S}_i$ the set of all paths.

An outcome of the game is a (feasible) flow $f^i : \mathcal{S}_i \to \mathbb{R}_{\geq 0}$ satisfying $\sum_{P\in\mathcal{S}_i} f^i_P = r_i$ for every $i \in [k]$. We use $\mathcal{F}(\mathcal{S})$ to denote the set of all feasible flows $f = (f^1, \dots, f^k)$. Given a flow $f = (f^i)_{i\in[k]} \in \mathcal{F}(\mathcal{S})$, we use f^i_e to denote the total flow on resource $e \in E$ of commodity $i \in [k]$, i.e., $f^i_e = \sum_{P\in\mathcal{S}_i : e\in P} f^i_P$. The total flow on edge $e \in E$ is defined as $f_e = \sum_{i\in[k]} f^i_e$.

The latency of a path $P \in \mathcal{S}$ with respect to f is defined as $l_P(f) := \sum_{e\in P} l_e(f_e)$. The cost of commodity i with respect to f is $C_i(f) = \sum_{P\in\mathcal{S}_i} f_P l_P(f)$. The *social cost* $C(f)$ of a flow f is given by its total average latency, i.e., $C(f) = \sum_{i\in[k]} C_i(f) = \sum_{e\in E} f_e l_e(f_e)$. A flow that minimizes $C(\cdot)$ is called *(socially) optimal*.

If the population is heterogenous, then each commodity $i \in [k]$ is further partitioned in h_i *sensitivity classes*, where class $j \in [h_i]$ has demand r_{ij} such that $r_i = \sum_{j\in[h_i]} r_{ij}$. Given a path $P \in \mathcal{S}_i$, we use $f_{P,j}$ to refer to the amount of flow on path P of sensitivity class j (so that $\sum_{j\in[h_i]} f_{P,j} = f_P$).

[1] If a network congestion game with a single commodity is considered (i.e., $k = 1$), we omit the commodity index for ease of notation.

Deviated Nash flows. We consider a *bounded deviation model* similar to the one introduced in [6].[2] We use $\delta = (\delta_P)_{P\in\mathcal{S}}$ to denote some arbitrary path deviations, where $\delta_P : \mathcal{F}(\mathcal{S}) \to \mathbb{R}_{\geq 0}$ for all $P \in \mathcal{S}$. Let $\beta \geq 0$ be fixed. Define the set of *β-bounded path deviations* as $\Delta(\beta) = \{(\delta_P)_{P\in\mathcal{S}} \mid 0 \leq \delta_P(f) \leq \beta l_P(f)$ for all $f \in \mathcal{F}(\mathcal{S})\}$.

Every commodity $i \in [k]$ and sensitivity class $j \in [h_i]$ has a non-negative sensitivity γ_{ij} with respect to the path deviations. The population is *homogeneous* if $\gamma_{ij} = \gamma$ for all $i \in [k]$, $j \in [h_i]$ and some $\gamma \geq 0$; otherwise, it is *heterogeneous*. Define the *deviated latency* of a path $P \in \mathcal{S}_i$ for sensitivity class $j \in [h_i]$ as $q_P^j(f) = l_P(f) + \gamma_{ij}\delta_P(f)$.

We say that a flow f is a *β-deviated Nash flow* if there exist some β-bounded path deviations $\delta \in \Delta(\beta)$ such that

$$\forall i \in [k], \forall j \in [h_i], \forall P \in \mathcal{S}_i, f_{P,j} > 0: \qquad q_P^j(f) \leq q_{P'}^j(f) \ \ \forall P' \in \mathcal{S}_i. \tag{1}$$

We define the *β-deviation ratio* β-DR$(\mathcal{I})$ as the maximum ratio $C(f^\beta)/C(f^0)$ of an β-deviated Nash flow f^β and an original Nash flow f^0. Intuitively, the deviation ratio measures the worst-case deterioration in social cost as a result of (bounded) deviations in the path latencies. Note that here the comparison is done with respect to an *unaltered* Nash flow to measure the impact of these deviations.

The set $\Delta(\beta)$ can also be restricted to path deviations which are defined as a function of edge deviations along that path. Suppose every edge $e \in E$ has a deviation $\delta_e : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ satisfying $0 \leq \delta_e(x) \leq \beta l_e(x)$ for all $x \geq 0$. For example, feasible path deviations can then be defined by the L_1-norm objective $\delta_P(f) = \sum_{e\in P}\delta_e(x)$ (as in [6,13]) or the L_2-norm objective $\delta_P(f) = \sqrt{\sum_{e\in P}\delta_e(x)^2)}$ (as in [9,13]). The *Price of Risk-Aversion* introduced by Nikolova and Stier-Moses [13] is technically the same ratio as the deviation ratio for the L_1- and L_2-norm (see [6] for details).

Approximate Nash Flows. We introduce the notion of an approximate Nash flow. Also here, each commodity $i \in [k]$ and sensitivity class $j \in [h_i]$ has a non-negative sensitivity ϵ_{ij}. We say that the population is *homogeneous* if $\epsilon_{ij} = \epsilon$ for all $i \in [k]$, $j \in [h_i]$ and some $\epsilon \geq 0$; otherwise, it is *heterogeneous*.

A flow f is an *ϵ-approximate Nash flow* with respect to sensitivities $\epsilon = (\epsilon_{ij})_{i\in[k],j\in[h_i]}$ if

$$\forall i \in [k],\ \forall j \in [h_i],\ \forall P \in \mathcal{S}_i, f_{P,j} > 0: \qquad l_P(f) \leq (1+\epsilon_{ij})l_{P'}(f) \ \ \forall P' \in \mathcal{S}_i \tag{2}$$

Note that a 0-approximate Nash flow is simply a Nash flow. We define the *ϵ-stability ratio* ϵ-SR$(\mathcal{I})$ as the maximum ratio $C(f^\epsilon)/C(f^0)$ of an ϵ-approximate Nash flow f^ϵ and an original Nash flow f^0.

Some of the proofs are missing in the main text below and can be found in [7].

[2] In fact, in [6] more general path deviations are introduced; the path deviations considered here correspond to $(0, \beta)$-*path deviations* in [6].

3 Heterogeneous Populations

We first elaborate on the relation between deviated and approximate Nash flows for general congestion games with heterogeneous populations.

Proposition 1. *Let $\mathcal{I}$ be a congestion game with heterogeneous players. If f is a β-deviated Nash flow for $\mathcal{I}$, then f is an ϵ-approximate Nash flow for $\mathcal{I}$ with $\epsilon_{ij} = \beta\gamma_{ij}$ for all $i \in [k]$ and $j \in [h_i]$ (for the same demand distribution r).*

Discrete Sensitivity Distributions. Subsequently, we show that the reverse of Proposition 1 does not hold. We do this by providing tight bounds on the β-deviation ratio and the ϵ-stability ratio for instances on (single-commodity) series-parallel graphs and arbitrary discrete sensitivity distributions.

Theorem 1. *Let $\mathcal{I}$ be a single-commodity network congestion game on a series-parallel graph with heterogeneous players, demand distribution $r = (r_i)_{i\in[h]}$ normalized to 1, i.e., $\sum_{j\in[h]} r_i = 1$, and sensitivity distribution $\gamma = (\gamma_i)_{i\in[h]}$, with $\gamma_1 < \gamma_2 < \cdots < \gamma_h$. Let $\beta \geq 0$ be fixed and define $\epsilon = (\beta\gamma_i)_{i\in[h]}$. Then the ϵ-stability ratio and the β-deviation ratio are bounded by:*

$$\epsilon\text{-}SR(\mathcal{I}) \leq 1 + \beta \sum_{j=1}^{h} r_j\gamma_j \quad \text{and} \quad \beta\text{-}DR(\mathcal{I}) \leq 1 + \beta \cdot \max_{j\in[h]} \left\{ \gamma_j \left(\sum_{p=j}^{h} r_p \right) \right\}. \tag{3}$$

Further, both bounds are tight for all distributions r and γ.

It is not hard to see that the bound on the β-deviation ratio is always smaller than the bound on the ϵ-stability ratio.[3] Our bound on the β-deviation ratio also yields tight bounds on the *Price of Risk-Aversion* [13] for series-parallel graphs and arbitrary heterogeneous risk-averse populations, both for the L_1-norm and L_2-norm objective.[4]

We need the following technical lemma for the proof of the β-deviation ratio.

Lemma 1. *Let $0 \leq \tau_{k-1} \leq \cdots \leq \tau_1 \leq \tau_0$ and $c_i \geq 0$ for $i = 1, \ldots, k$ be given. We have $c_1\tau_0 + \sum_{i=1}^{k-1}(c_{i+1} - c_i)\tau_i \leq \tau_0 \cdot \max_{i=1,\ldots,k}\{c_i\}$.*

Proof (Theorem 1, β-deviation ratio). Let $x = f^\beta$ be a β-deviated Nash flow with path deviations $(\delta_P)_{P\in\mathcal{S}} \in \Delta(\beta)$ and let $z = f^0$ be an original Nash flow. Let $X = \{a \in A : x_a > z_a\}$ and $Z = \{a \in A : z_a \geq x_a \text{ and } z_a > 0\}$ (arcs with $x_a = z_a = 0$ may be removed without loss of generality).

In order to analyze the ratio $C(x)/C(z)$ we first argue that we can assume without loss of generality that the latency function $l_a(y)$ is constant for values $y \geq x_a$ for all arcs $a \in Z$. To see this, note that we can replace the function $l_a(\cdot)$ with the function $\hat{l}_a$ defined by $\hat{l}_a(y) = l_a(x_a)$ for all $y \geq x_a$ and $\hat{l}_a(y) = l_a(y)$

[3] This follows from Markov's inequality: for a random variable Y, $P(Y \geq t) \leq E(Y)/t$.

[4] Observe that we show tightness of the bound on parallel arcs, in which case these objectives coincide.

for $y \leq x_a$. In particular, this implies that the flow x is still a β-deviated Nash flow for the same path deviations as before. This holds since for any path P the latency $l_P(x)$ remains unchanged if we replace the function l_a by $\hat{l}_a$.

By definition of arcs in Z, we have $x_a \leq z_a$ and therefore $\hat{l}_a(z_a) = l_a(x_a) \leq l_a(z_a)$. Let z' be an original Nash flow for the instance with l_a replaced by $\hat{l}_a$. Then we have $C(z') \leq C(z)$ using the fact that series-parallel graphs are immune to the Braess paradox, see Milchtaich [12, Lemma 4]. Note that, in particular, we find $C(x)/C(z) \leq C(x)/C(z')$. By repeating this argument, we may without loss of generality assume that all latency functions l_a are constant between x_a and z_a for $a \in Z$. Afterwards, we can even replace the function $\hat{l}_a$ by a function that has the constant value of $l_a(x_a)$ everywhere.

In the remainder of the proof, we will denote P_j as a flow-carrying arc for sensitivity class $j \in [h]$ that maximizes the path latency amongst all flow-carrying path for sensitivity class $j \in [h]$, i.e., $P_j = \text{argmax}_{P \in \mathcal{P}: x_{P,j} > 0}\{l_P(x)\}$. Moreover, there also exists a path P_0 with the property that $z_a \geq x_a$ and $z_a > 0$ for all arcs $a \in P_0$ (see, e.g., Lemma 2 [12]).

For fixed $a < b \in \{1, \dots, h\}$, the Nash conditions imply that (these steps are of a similar nature as Lemma 1 [4])

$$l_{P_a}(x) + \gamma_a \cdot \delta_{P_a}(x) \leq l_{P_b}(x) + \gamma_a \cdot \delta_{P_b}(x)$$
$$l_{P_b}(x) + \gamma_b \cdot \delta_{P_b}(x) \leq l_{P_a}(x) + \gamma_b \cdot \delta_{P_a}(x).$$

Adding up these inequalities implies that $(\gamma_b - \gamma_a)\delta_{P_b}(x) \leq (\gamma_b - \gamma_a)\delta_{P_a}(x)$, which in turn yields that $\delta_{P_b}(x) \leq \delta_{P_a}(x)$ (using that $\gamma_a < \gamma_b$ if $a < b$). Furthermore, we also have

$$l_{P_1}(x) + \gamma_1 \delta_{P_1}(x) \leq l_{P_0}(x) + \gamma_1 \delta_{P_0}(x), \tag{4}$$

and $l_{P_0}(x) = l_{P_0}(z) \leq l_{P_1}(z) \leq l_{P_1}(x)$, which can be seen as follows. The equality follows from the fact that l_a is constant for all $a \in Z$ and, by choice, P_0 only consists of arcs in Z. The first inequality follows from the Nash conditions of the original Nash flow z, since there exists a flow-decomposition in which the path P_0 is used (since the flow on all arcs of P_0 is strictly positive in z). The second inequality follows from the fact that

$$\sum_{e \in P_1} l_e(z_e) = \sum_{e \in P_1 \cap X} l_e(z_e) + \sum_{e \in P_1 \cap Z} l_e(z_e) \leq \sum_{e \in P_1 \cap X} l_e(x_e) + \sum_{e \in P_1 \cap Z} l_e(x_e)$$

using that $z_e \leq x_e$ for $e \in X$ and the fact that latency functions for $e \in Z$ are constant. In particular, we find that $l_{P_0}(x) \leq l_{P_1}(x)$. Adding this inequality to (4), we obtain $\gamma_1 \delta_{P_1}(x) \leq \gamma_1 \delta_{P_0}(x)$ and therefore $\delta_{P_1}(x) \leq \delta_{P_0}(x)$. Thus $\delta_{P_h}(x) \leq \delta_{P_{h-1}}(x) \leq \dots \leq \delta_{P_1}(x) \leq \delta_{P_0}(x)$. Moreover, by using induction it can be shown that

$$l_{P_j}(x) \leq l_{P_0}(x) + \gamma_1 \delta_{P_0}(x) + \left[\sum_{g=1}^{j-1} (\gamma_{g+1} - \gamma_g)\delta_{P_g}(x)\right] - \gamma_j \delta_{P_j}(x). \tag{5}$$

Using (5), we then have

$$\begin{aligned}
C(x) &\le \sum_{j=1}^{h} r_j l_{P_j}(x) \quad \text{(by choice of the paths } P_j\text{)} \\
&\le \sum_{j=1}^{h} r_j \left(l_{P_0}(x) + \gamma_1 \delta_{P_0}(x) + \left[\sum_{g=1}^{j-1} (\gamma_{g+1} - \gamma_g) \delta_{P_g}(x) \right] - \gamma_j \delta_{P_j}(x) \right) \\
&= l_{P_0}(x) + \gamma_1 \delta_{P_0}(x) + \sum_{j=1}^{h} (r_{j+1} + \cdots + r_h)(\gamma_{j+1} - \gamma_j) \delta_{P_j}(x) - r_j \gamma_j \delta_{P_j}(x) \\
&\le l_{P_0}(x) + \gamma_1 \delta_{P_0}(x) \\
&\quad + \sum_{j=1}^{h-1} \Big[(r_{j+1} + \cdots + r_h)\gamma_{j+1} - (r_j + r_{j+1} + \cdots + r_h)\gamma_j \Big] \delta_{P_j}(x)
\end{aligned}$$

In the last inequality, we leave out the last negative term $-r_h \gamma_h \delta_{P_h}(x)$. Note that $\gamma_1 = (r_1 + \cdots + r_h)\gamma_1$ since we have normalized the demand to 1. We can then apply Lemma 1 with $\tau_i = \delta_{P_i}(x)$ for $i = 0, \ldots, h-1$ and $c_i = \gamma_i \cdot \sum_{p=i}^{h} r_p$ for $i = 1, \ldots, k$. Continuing the estimate, we get

$$C(x) \le l_{P_0}(x) + \max_{j \in [h]} \left\{ \gamma_j \cdot \sum_{p=j}^{h} r_p \right\} \cdot \delta_{P_0}(x) \le \left[1 + \beta \cdot \max_{j \in [h]} \left\{ \gamma_j \left(\sum_{p=j}^{h} r_p \right) \right\} \right] C(z)$$

where for the second inequality we use that $\delta_{P_0}(x) \le \beta l_{P_0}(x)$, which holds by definition, and $l_{P_0}(x) = l_{P_0}(z) = C(z)$, which holds because z is an original Nash flow and all arcs in P_0 have strictly positive flow in z (and because of the fact that all arcs in P_0 have a constant latency functions).

To prove tightness, fix $j \in [h]$ and consider the following instance on two arcs. We take $(l_1(y), \delta_1(y)) = (1, \beta)$ and $(l_2(y), \delta_2(y))$ with $\delta_2(y) = 0$ and $l_2(y)$ a strictly increasing function satisfying $l_2(0) = 1 + \epsilon$ and $l_2(r_j + r_{j+1} + \cdots + r_h) = 1 + \gamma_j \beta$, where $\epsilon < \gamma_j \beta$. The (unique) original Nash flow is given by $z = (z_1, z_2) = (1, 0)$ with $C(z) = 1$. The (unique) β-deviated Nash flow x is given by $x = (x_1, x_2) = (r_1 + r_2 + \cdots + r_{j-1}, r_j + r_{j+1} + \cdots + r_h)$ with $C(x) = 1 + \beta \cdot \gamma_j (r_j + \cdots + r_h)$. Since this construction holds for all $j \in [h]$, we find the desired lower bound. □

Continuous Sensitivity Distributions. We obtain a similar result for more general (not necessarily discrete) sensitivity distributions. That is, we are given a Lebesgue integrable *sensitivity density function* $\psi : \mathbb{R}_{\ge 0} \to \mathbb{R}_{\ge 0}$ over the total demand. Since we can normalize the demand to 1, we have the condition that $\int_0^\infty \psi(y)dy = 1$. We then find the following natural generalizations of our upper bounds:

1. ϵ-SR$(\mathcal{I}) \le 1 + \beta \int_0^\infty y \cdot \psi(y)dy$, and
2. β-DR$(\mathcal{I}) \le 1 + \beta \cdot \sup_{t \in \mathbb{R}_{\ge 0}} \left\{ t \cdot \int_t^\infty \psi(y)dy \right\}$.

These bounds are both asymptotically tight for all distributions. Details are given the full version [7].

4 Homogeneous Population

The reverse of Proposition 1 also holds for homogeneous players in single-commodity instances. As a consequence, the set of β-deviated Nash flows and the set of ϵ-approximate Nash flows with $\epsilon = \beta\gamma$ coincide in this case.

Recall that for homogeneous players we have $\gamma_{ij} = \gamma$ for all $i \in [k]$, $j \in [h_i]$ and some $\gamma \geq 0$.

Proposition 2. *Let $\mathcal{I}$ be a single-commodity congestion game with homogeneous players. f is an ϵ-approximate Nash flow for $\mathcal{I}$ if and only if f is a β-deviated Nash flow for $\mathcal{I}$ with $\epsilon = \beta\gamma$.*

Upper Bound on the Stability Ratio. Our main result in this section is an upper bound on the ϵ-stability ratio. Given the above equivalence, this bound also applies to the β-deviation ratio with $\epsilon = \beta\gamma$.

The following concept of alternating paths is crucial. For single-commodity instances an alternating path always exists (see, e.g., [13]) (Fig. 1).

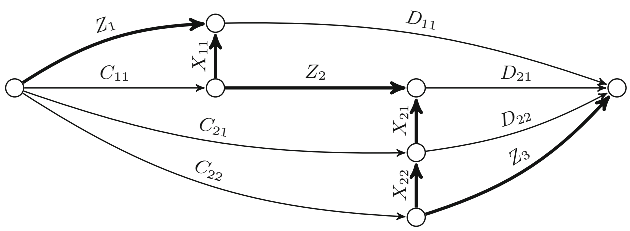

Fig. 1. Sketch of the situation in the proof of Theorem 2 with $q_1 = 1$ and $q_2 = 2$.

Definition 1 (Alternating path [10,13]). *Let $\mathcal{I}$ be a single-commodity network congestion game and let x and z be feasible flows. We partition the edges $E = X \cup Z$ such that $Z = \{a \in E : z_a \geq x_a \text{ and } z_a > 0\}$ and $X = \{a \in E : z_a < x_a \text{ or } z_a = x_a = 0\}$. We say that π is an alternating s,t-path if the arcs in $\pi \cap Z$ are oriented in the direction of t, and the arcs in $\pi \cap X$ are oriented in the direction of s. We call the number of backward arcs on π the* backward length *of π and refer to it by $q(\pi) = |\pi \cap X|$.*

Theorem 2. *Let $\mathcal{I}$ be a single-commodity network congestion game. Let $\epsilon \geq 0$ be fixed and consider an arbitrary alternating path π with backward length $q = q(\pi)$. If $\epsilon < 1/q$, then the ϵ-stability ratio is bounded by*

$$\epsilon\text{-}SR(\mathcal{I}) \leq \frac{1+\epsilon}{1-\epsilon \cdot q} \leq \frac{1+\epsilon}{1-\epsilon \cdot n}.$$

Note that the restriction on ϵ stated in the theorem always holds if $\epsilon < 1/n$. In particular, for $\epsilon \ll 1/n$ we roughly get ϵ-SR$(\mathcal{I}) \leq 1+\epsilon n$. The proof of Theorem 2 is inspired by a technique of Nikolova and Stier-Moses [13], but technically more involved.

Proof. Let $x = f^\epsilon$ be an ϵ-approximate Nash flow and let $z = f^0$ an original Nash flow. Let $\pi = Z_1X_1Z_2X_2 \dots Z_{\eta-1}X_{\eta-1}Z_\eta$ be an alternating path for x and z, where Z_i and X_i are maximal sections consisting of consecutive arcs, respectively, in Z and X (i.e., $Z_i \subseteq Z$ and $X_i \subseteq X$ for all i). Furthermore, we let $q_i = |X_i|$ and write $X_i = (X_{iq_i}, \dots, X_{i2}, X_{i1})$, where X_{ij} are the arcs in the section X_i. By definition, for every arc X_{ij} there exists a path $C_{ij}X_{ij}D_{ij}$ that is flow-carrying for x.[5]

For convenience, we define $C_{01} = D_{\eta,0} = \emptyset$. Furthermore, we denote $P^{\max}$ as a path maximizing $l_P(x)$ over all paths $P \in \mathcal{S}$. For convenience, we will abuse notation, and write $Q = Q(x) = \sum_{a\in Q} l_a(x)$ for $Q \subseteq E$.

Note that for all i, j:

$$C_{ij}(x) + X_{ij}(x) + D_{ij}(x) \le P^{\max}(x). \tag{6}$$

Fix some $i \in \{1, \dots, \eta-1\}$. Then we have $C_{i1}+X_{i1}+D_{i1} \le (1+\epsilon)(C_{i-1,q_{i-1}} + Z_i + D_{i1})$ by definition of an ϵ-approximate Nash flow. This implies that (leaving out D_{i1} on both sides) $C_{i1} + X_{i1} \le (1+\epsilon)Z_i + C_{i-1,q_{i-1}} + \epsilon(C_{i-1,q_{i-1}} + D_{i1})$. Furthermore, for all $j \in \{2, \dots, q_i\}$, we have $C_{ij}+X_{ij}+D_{ij} \le (1+\epsilon)(C_{i,j-1}+D_{ij})$ which implies (again leaving out D_{ij} on both sides)

$$C_{ij} + X_{ij} \le C_{i,j-1} + \epsilon(C_{i,j-1} + D_{ij}).$$

Adding up these inequalities for $j \in \{1, \dots, q_i\}$ and subtracting $\sum_{j=1}^{q_i-1} C_{ij}$ from both sides, we obtain for all $i \in \{1, \dots, \eta-1\}$

$$C_{i,q_i} + \sum_{j=1}^{q_i} X_{ij} \le C_{i-1,q_{i-1}} + (1+\epsilon)Z_i + \epsilon\Bigg(\sum_{j=1}^{q_i} D_{ij} + C_{i-1,q_{i-1}} + \sum_{j=1}^{q_i-1} C_{ij}\Bigg). \tag{7}$$

Moreover, we also have

$$P^{\max} \le (1+\epsilon)(C_{\eta-1,\eta-1} + Z_\eta) = C_{\eta-1,\eta-1} + (1+\epsilon)Z_\eta + \epsilon C_{\eta-1,\eta-1}. \tag{8}$$

Adding up the inequalities in (7) for all $i \in \{1, \dots, \eta-1\}$, and the inequality in (8), we obtain

$$P^{\max} + \sum_{i=1}^{\eta-1} C_{i,q_i} + \sum_{i=1}^{\eta-1}\sum_{j=1}^{q_i} X_{ij} \le \sum_{i=1}^{\eta-1} C_{i,q_i} + (1+\epsilon)\sum_{i=1}^{\eta} Z_i + \epsilon\Bigg(\sum_{i=1}^{\eta-1}\sum_{j=1}^{q_i} C_{ij} + D_{ij}\Bigg)$$

which simplifies to

$$P^{\max} + \sum_{i=1}^{\eta-1}\sum_{j=1}^{q_i} X_{ij} \le (1+\epsilon)\sum_{i=1}^{\eta} Z_i + \epsilon\Bigg(\sum_{i=1}^{\eta-1}\sum_{j=1}^{q_i} C_{ij} + D_{ij}\Bigg). \tag{9}$$

[5] Note that for a Nash flow one can assume that there is a flow-carrying path traversing all arcs $X_{iq_i}, \dots, X_{i1}$; but this cannot be done for an approximate Nash flow.

Using (6), we obtain

$$\sum_{i=1}^{\eta-1}\sum_{j=1}^{q_i} C_{ij} + D_{ij} \leq \sum_{i=1}^{\eta-1}\sum_{j=1}^{q_i} P^{\max} - X_{ij} = \left(\sum_{i=1}^{\eta-1} q_i\right)P^{\max} - \sum_{i=1}^{\eta-1}\sum_{j=1}^{q_i} X_{ij}.$$

Combining this with (9), and rearranging some terms, we get

$$\begin{aligned}(1-\epsilon\cdot q)P^{\max} &\leq (1+\epsilon)\left[\sum_{i=1}^{\eta} Z_i - \sum_{i=1}^{\eta-1}\sum_{j=1}^{q_i} X_{ij}\right] \\ &= (1+\epsilon)\left[\sum_{e\in Z\cap\pi} l_e(x_e) - \sum_{e\in X\cap\pi} l_e(x_e)\right]\end{aligned}$$

where $q = q(\pi) = \sum_{i=1}^{\eta-1} q_i$ is the backward length of π.

Similarly (see also [13, Lemma 4.5]), it can be shown that

$$l_Q(z) \geq \sum_{e\in Z\cap\pi} l_e(z_e) - \sum_{e\in X\cap\pi} l_e(z_e) \tag{10}$$

for any path Q with $z_Q > 0$ (these all have the same latency, since z is an original Nash flow). Using a similar argument as in [13, Theorem 4.6], we obtain

$$\begin{aligned}(1-\epsilon\cdot q)l_{P^{\max}}(x) &\leq (1+\epsilon)\left[\sum_{e\in Z\cap\pi} l_e(x_e) - \sum_{e\in X\cap\pi} l_e(x_e)\right] \\ &\leq (1+\epsilon)\left[\sum_{e\in Z\cap\pi} l_e(z_e) - \sum_{e\in X\cap\pi} l_e(z_e)\right] \leq (1+\epsilon)l_Q(z).\end{aligned}$$

By multiplying both sides with the demand r, we obtain $(1-\epsilon\cdot q)C(x) \leq (1-\epsilon\cdot q)r\cdot l_{P^{\max}}(x) \leq (1+\epsilon)r\cdot l_Q(z) = (1+\epsilon)C(z)$ for $\epsilon < 1/q$, which proves the claim. □

Tight Bound on the Stability Ratio. In this section, we consider instances for which all backward sections of the alternating path π consist of a single arc., i.e., $q_i = 1$ for all $i = 1,\dots,\eta-1$. We then have $q = \sum_{i=1}^{\eta-1} q_i \leq \lfloor n/2\rfloor - 1$ since every arc in X must be followed directly by an arc in Z (and we can assume w.l.o.g. that the first and last arc are contained in Z). By Theorem 2, we obtain ϵ-SR$(\mathcal{I}) \leq (1+\epsilon)/(1-\epsilon\cdot(\lfloor n/2\rfloor-1))$ for all $\epsilon < 1/(\lfloor n/2\rfloor-1)$. We show that this bound is tight. Further, we show that there exist instances for which ϵ-SR$(\mathcal{I})$ is unbounded for $\epsilon \geq 1/(\lfloor n/2\rfloor - 1)$. This completely settles the case of $q_i = 1$ for all i.

Our construction is based on the *generalized Braess graph* [14]. By construction, alternating paths for these graphs satisfy $q_i = 1$ for all i. See the full version [7] for details.

Theorem 3. *Let $n = 2m$ be fixed and let $\mathcal{B}^m$ be the set of all instances on the generalized Braess graph with n nodes. Then*

$$\sup_{\mathcal{I}\in\mathcal{B}^m} \epsilon\text{-}SR(\mathcal{I}) = \begin{cases}\frac{1+\epsilon}{1-\epsilon\cdot(\lfloor n/2\rfloor-1)} & \text{if } \epsilon < \frac{1}{\lfloor n/2\rfloor-1},\\ \infty & \text{otherwise.}\end{cases}$$

Non-symmetric Matroid Congestion Games. In the previous sections, we considered (symmetric) network congestion games only. It is interesting to consider other combinatorial strategy sets as well. In this section we make a first step in this direction by focusing on the bases of matroids as strategies.

A matroid congestion game is given by $\mathcal{J} = (E, (l_e)_{e\in E}, (\mathcal{S}_i)_{i\in[k]}, (r_i)_{i\in[k]})$, and matroids $\mathcal{M}_i = (E, \mathcal{I}_i)$ over the ground set E for every $i \in [k]$.[6] The strategy set $\mathcal{S}_i$ consists of the *bases* of the matroid $\mathcal{M}_i$, which are the independent sets of maximum size, e.g., spanning trees in an undirected graph. We refer the reader to Schrijver [16] for an extensive overview of matroid theory.

As for network congestion games, it can be shown that in general the ϵ-stability ratio can be unbounded (see Theorem 5 [7] in the appendix of [7]); this also holds for general path deviations because the proof of Proposition 2 in the appendix holds for arbitrary strategy sets. However, if we consider path deviations induced by the sum of edge deviations (as in [6,13]), then we can obtain a more positive result for general matroids.

Recall that for every resource $e \in E$ we have a deviation function $\delta_e : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ satisfying $0 \leq \delta_e(x) \leq \beta l_e(x)$ for all $x \geq 0$. The deviation of a basis B is then given by $\delta_B(f) = \sum_{e\in B} \delta_e(f_e)$.

Theorem 4. *Let $\mathcal{J} = (E, (l_e)_{e\in E}, (\mathcal{S}_i)_{i\in[k]}, (r_i)_{i\in[k]})$ be a matroid congestion game with homogeneous players. Let $\beta \geq 0$ be fixed and consider β-bounded basis deviations as defined above. Then the β-deviation ratio is upper bounded by β-DR$(\mathcal{J}) \leq 1 + \beta$. Further, this bound is tight already for 1-uniform matroid congestion games.*

Acknowledgements. We thank the anonymous referees for their very useful comments, and one reviewer for pointing us to Lemma 1 [4].

References

1. Chen, P.-A., Keijzer, B.D., Kempe, D., Schäfer, G.: Altruism and its impact on the price of anarchy. ACM Trans. Econ. Comput. **2**(4), 17:1–17:45 (2014)
2. Christodoulou, G., Koutsoupias, E., Spirakis, P.G.: On the performance of approximate equilibria in congestion games. Algorithmica **61**(1), 116–140 (2011)
3. Englert, M., Franke, T., Olbrich, L.: Sensitivity of Wardrop equilibria. In: Monien, B., Schroeder, U.-P. (eds.) SAGT 2008. LNCS, vol. 4997, pp. 158–169. Springer, Heidelberg (2008). doi:10.1007/978-3-540-79309-0_15
4. Fleischer, L.: Linear tolls suffice: new bounds and algorithms for tolls in single source networks. Theoret. Comput. Sci. **348**(2), 217–225 (2005)
5. Fujishige, S., Goemans, M.X., Harks, T., Peis, B., Zenklusen, R.: Matroids are immune to braess paradox. CoRR, abs/1504.07545 (2015)
6. Kleer, P., Schäfer, G.: The impact of worst-case deviations in non-atomic network routing games. CoRR, abs/1605.01510 (2016)

[6] A matroid over E is given by a collection $\mathcal{I} \subseteq 2^E$ of subsets of E (called *independent sets*). The pair $\mathcal{M} = (E, \mathcal{I})$ is a *matroid* if the following three properties hold: (i) $\emptyset \in \mathcal{I}$; (ii) If $A \in \mathcal{I}$ and $B \subseteq A$, then $B \in \mathcal{I}$. (iii) If $A, B \in \mathcal{I}$ and $|A| > |B|$, then there exists an $a \in A \setminus B$ such that $B + a \in \mathcal{I}$.

7. Kleer, P., Schäfer, G.: Path deviations outperform approximate stability in heterogeneous congestion games. CoRR, abs/1707.01278 (2017)
8. Koutsoupias, E., Papadimitriou, C.: Worst-case equilibria. In: Meinel, C., Tison, S. (eds.) STACS 1999. LNCS, vol. 1563, pp. 404–413. Springer, Heidelberg (1999). doi:10.1007/3-540-49116-3_38
9. Lianeas, T., Nikolova, E., Stier-Moses, N.E.: Asymptotically tight bounds for inefficiency in risk-averse selfish routing. In: Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence. IJCAI 2016, NY, USA, New York, pp. 338–344 (2016)
10. Lin, H., Roughgarden, T., Tardos, É., Walkover, A.: Stronger bounds on Braess's paradox and the maximum latency of selfish routing. SIAM J. Discrete Math. **25**(4), 1667–1686 (2011)
11. Meir, R., Parkes, D.C.: Playing the wrong game: smoothness bounds for congestion games with risk averse agents. CoRR, abs/1411.1751 (2017)
12. Milchtaich, I.: Network topology and the efficiency of equilibrium. Games Econ. Behav. **57**(2), 321–346 (2006)
13. Nikolova, E., Stier-Moses, N.E.: The burden of risk aversion in mean-risk selfish routing. In: Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC 2015, pp. 489–506. ACM, New York (2015)
14. 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)
15. Roughgarden, T.: Intrinsic robustness of the price of anarchy. J. ACM **62**(5), 32 (2015)
16. Schrijver, A.: Combinatorial Optimization: Polyhedra and Efficiency Vol. B Matroids Trees Stable Sets Algorithms and Combinatorics. Springer, Heidelberg (2003). Chaps. 39–69
17. Wardrop, J.G.: Some theoretical aspects of road traffic research. In: Proceedings of the Institution of Civil Engineers, vol. 1, pp. 325–378 (1952)