Network Congestion Games are Robust to Variable Demand*

José Correa† Ruben Hoeksma ‡ Marc Schröder §

Abstract

Network congestion games have provided a fertile ground for the algorithmic game theory community. Indeed, many of the pioneering works on bounding the efficiency of equilibria use this framework as their starting point. In recent years, there has been an increased interest in studying randomness in this context though the efforts have been mostly devoted to understanding what happens when link latencies are subject to random shocks. Although this is an important practical consideration, it is not the only source of randomness in network congestion games. Another important source is the inherent variability of the demand that most practical networks suffer from. Therefore in this paper we look at the basic non-atomic network congestion game with the additional feature that demand is random. Our main result in this paper is that under a very natural equilibrium notion, in which the basic behavioral assumption is that users evaluate their expected cost according to the demand they experience in the system, the price of anarchy of the game is actually the same as that in the deterministic demand game. Moreover, the result can be extended to the more general class of smooth games.

1 Introduction

A traffic network consists of a network on which many different entities make autonomous decisions about how to use that network. In a physical traffic network, this would be the users choosing the route that they take from their origin to their destination. Non-atomic network congestion games are a fundamental model in game theory used to model traffic flows in both physical and digital traffic networks. The model assumes a very large population of users of a network that are selfish and minimize some objective function, most often their own travel time. The choice that a user gets to make is then simply the route that she takes from her origin to her destination. The amount of traffic that she encounters (the congestion) on this route then influences her travel time, and, in turn, the routes that all the users choose influence the congestion that they all observe.

It is well known that when users of systems make autonomous choices the resulting usage of the system may be inefficient [24]. This phenomenon has captured the attention of researchers and practitioners for quite a long time, see, e.g., Dubey [10]. To assess this efficiency loss, assumptions are made on what the outcome of such a game of autonomous users is. In network congestion games this is often assumed to be an equilibrium outcome, in which no single user can improve

*This paper was presented at the 13th Conference on Web and Internet Economics (WINE 2017) and a short abstract appeared in the conference proceedings. Research was supported by the Millennium Nucleus Information and Coordination in Networks, ICM/FIC RC130003, and by the Millennium Institute for Research in Market Imperfections and Public Policy, MIPP, ICM IS130002.

†Universidad de Chile, Santiago, Chile, correa@uchile.cl

‡University of Bremen, Bremen, Germany, hoeksma@uni-bremen.de

§RWTH Aachen University, Aachen, Germany, marc.schroeder@oms.rwth-aachen.de

their objective by unilaterally deviating from their choice. These equilibria are also known as Nash equilibria [18] or Wardrop equilibria [34], depending on the specifics of the game. When we talk about efficiency in the network usage, we need a concept of social cost. The most prominent concept that has been used is the utilitarian social objective, which in the case of network congestion games is the total travel time of the users. In the last decades, we have seen an effort to put quantification to this loss in efficiency. Koutsoupias and Papadimitriou [16, 15] introduced, what was later named the *price of anarchy* as the worst case ratio between the social objective in a worst case equilibrium outcome and the social objective in an socially optimized solution [23].

Since its introduction, the price of anarchy has been widely studied in the context of congestion games. Roughgarden and Tardos [30, 26] bounded the price of anarchy of the non-atomic network congestion game and showed that the bound only depends on the class of latency functions. In particular, they bound the price of anarchy by 4/3 in networks with affine latencies. Later, Correa et al. [7, 8] extend this result to the capacitated context and provide a geometric proof allowing more general cost functions. Similar results were also obtained for a number of variants of the game, namely with atomic players [2, 4], with atomic splittable players [5, 13, 29], and with large number of players [3], among many others. All these results however, apply to deterministic situations.

More recently there has been a growing interest in having a theoretical understanding of variability in travel times in network congestion games. The issue of uncertainty is certainly of high practical relevance and has been an object of study in the transportation science community for many decades now. The work by Dial [9] and Sheffi [31] can be considered as a first approach of capturing uncertainty in network congestion games. They study a model where the perception of travel time among users varies, but the travel time is in fact deterministic. On the other hand, Abdel-Aty et al. [1] show empirically that the variability of travel times is one of the most important factors in making routing decisions. This is indeed natural, as in physical traffic networks travel times are rarely constant, because of numerous causes, like unexpected crashes, bad weather, construction works and other irregular events. Uncertainty in more recent models is often associated with a random variable with a known probability distribution and the choice of the users depends on their risk aversion. Fan et al. [11] and Nie and Wu [19] use the expected probability of arriving on time as an objective for the users. Ordóñez and Stier-Moses [22] assume that users minimize the α quantile of the experienced travel time. Whereas Nikolova and Stier-Moses [20] let users minimize expected delay plus a safety margin, approximated by the standard deviation of the distribution. Qi et al. [25] assume that the probability distribution is not fully known and thus analyze the effect of risk and ambiguous attitudes on path choices. Finally, Cominetti and Torrico take an axiomatic approach to risk aversion [6]. In the context of risk aversion, quantifying the price of anarchy becomes more difficult (it is often unbounded) and actually it is less clear what the right concept of social cost should be [21, 17].

Somewhat surprisingly, most of the work on travel time uncertainty in network congestion games focuses on variability on the links in the network. One drawback of this approach is that if users simply minimize their expected cost, this link variability has no effect on the flow reducing the problem immediately to the deterministic link cost case. Therefore risk aversion has to come into play making the analysis much more complicated and introducing some unintuitive consistency issues [6]. However, another fundamental source of travel time variability and congestion comes from the inherent randomness of the demand pattern. In this paper, we are interested in the latter source of variability so we consider a basic network congestion game with incomplete information where the demand is random and users simply minimize their expected travel time. So far, this has been considered only by Wang, Doan, and Chen [33], who found that the price of anarchy can become arbitrarily large, depending on the variability of the demand. However, as we show in this paper, the negative result of Wang et al. finds its roots in the way equilibria are defined. In their definition, a user bases her decisions on the full knowledge of the demand distribution (even though she may have never experienced the highly congested scenarios). Here we provide an alternative view. We define an equilibrium where the knowledge of users is limited to the demand they actually observe. In other words, each user bases her decisions on the demand distribution conditioned on the fact that she is traveling. The concept has a solid micro foundation that is in line with that of the Perfect Bayesian equilibrium [12] and allows us to prove that for this *observed equilibrium* the price of anarchy does not depend on the variability in the demand, but only on the class of latency functions found in the network. A result in the same spirit as that for network congestion games without variable demand [26, 7].

Network congestion games are an example of smooth games [27]. Smoothness is a general technique that turns out to be fruitful to prove inefficiency results. These inefficiency proofs do not only apply to Nash equilibria, but also to many other reasonable solution concepts, like mixed nash equilibria, correlated equilibria and coarse-correlated equilibria. Recently, [32] and [28] provided a set of conditions under which full-information price of anarchy bounds extend to mixed Bayesian-Nash equilibria of every corresponding game of incomplete information with a product prior distribution. We prove a similar result for the setting in which there is uncertainty about which players play the game.

Our contribution. The main contribution of this paper is of a conceptual nature. We consider the basic network congestion game with variable demand. This variable demand is modeled as several groups of users that either use the network or do not. The microscopic modeling of the situation requires an underlying probability measure p_D over the subsets of users, so that the subset of users S travel with probability $p_D(S)$. However, this measure is unknown to the users of the network who only know their own marginals. Note that we do not make any assumptions on the probability measure underlying the usage of the network. Thus, e.g., whether two different groups use the network may be highly correlated, or not at all. For technical reasons we consider probability measures over a finite number of subsets, through larger and larger numbers of subsets we can model arbitrary situations. Writing the limit model (i.e., with uncountable many possible subsets) requires significant measure theoretic technicalities that go beyond the scope of this paper.

Within the above setting, in Subsection 2.2, we introduce a new equilibrium notion for systems with variable demand, the observed equilibrium. Naturally, for this notion a user evaluates the expected cost of different paths only using her knowledge of the conditional probability measure. Subsection 2.3 uses this new equilibrium notion to show that in network congestion games equilibrium solutions do not deteriorate more when variability in the demand is introduced. More precisely, we prove that the price of anarchy in network congestion games with variable demand is bounded by $1/(1-\beta)$, where β is the standard latency-function class dependent parameter [26, 7]. With this result, we give an alternative to the pessimistic conclusion by Wang et al. [33]. Not only do we show a bound on the price of anarchy that is independent on the distribution of the demand, we show that there is no increase at all compared to the system with no variable demand.

In light of the previous result one may think that the price of anarchy actually decreases with demand variability, since already the deterministic case provides worst case instances. We prove that this intuition is wrong and find situations with variable demand in which our bound is tight. In Subsection 2.4, we consider the performance of our equilibrium concept when compared to a stronger notion of optimal flow, namely the adaptive optimum. The situation here is more mixed and the price of anarchy depends on the underlying probability measure. Finally, in Subsection 2.5 we exhibit instances showing that the observed equilibrium can have higher expected cost than Wang et al.'s user equilibrium. However, when latency functions consist of only sums of a monomial and a constant, and the demand is regular enough, we prove that the observed equilibrium always performs better than the user equilibrium.

Our equilibrium concept applies almost unchanged to the smoothness framework [27, 28, 32], see Section 3. Particular interesting classes of games that belong to this framework are routing games with atomic or atomic splittable players, as well as several scheduling games. Known price of anarchy bounds for these games [2, 4, 5, 13, 29] also apply to the model in which the set of players playing is determined stochastically.

2 Routing Games with Variable Demand

2.1 Preliminaries

We consider the following non-atomic congestion game with stochastic demand. Given are a directed network (V, A), a finite set of commodities K, and a random variable D. Here, V denotes the set of vertices and A denotes the set of (directed) arcs. For each arc $a \in A$, we are given a nondecreasing differentiable convex latency function $\ell_a : \mathbb{R}_+ \to \mathbb{R}_+$ that represents the delay experienced by users traversing this link as a function of the total flow on the link. Each commodity $k \in K$ is defined by a triple (s^k, t^k, d^k), where (s^k, t^k) is the source-sink pair and d^k is the demand. The random variable D has probability mass function $p_D : 2^K \to [0, 1]$, such that $p_D(S)$ represents the probability that only users from commodities $k \in S$ travel. For a realization $D = S$, we say that the commodities $k \in S$ are *active*. Note that within the set of commodities K, the source-sink pairs are not necessarily unique, therefore we can model, for example, the situation with several different groups with the same source-sink pair and different travel behaviors.

For each $k \in K$, let $\mathcal{P}^k$ denote the set of directed $s^k - t^k$ paths, where each path $P \in \mathcal{P}^k$ is given as a subset of the arcs, $P \subseteq A$. A flow f^k for commodity k is a nonnegative vector $f^k = (f^k_P)_{P \in \mathcal{P}^k}$ such that $\sum_{P \in \mathcal{P}^k} f^k_P = d^k$. For arc a we denote by

$$f^k_a = \sum_{\substack{P \in \mathcal{P}^k \\ P \ni a}} f^k_P$$

the amount of flow of commodity k on arc a. Similarly, for arc a let

$$f_a(S) = \sum_{k \in S} f^k_a$$

be the amount of flow on arc a if the set of active commodities is exactly S. A flow f is a vector $(f^k)_{k \in K}$, where each f^k is a feasible flow for commodity k.

Given a flow f, the corresponding total expected congestion cost is

$$C(f) = \sum_{a \in A} \sum_{S \in 2^K} p_D(S) \cdot \ell_a\left(f_a(S)\right) \cdot f_a(S)\,.$$

Definition 2.1. *A flow f^* is said to be an optimal flow if it solves the following minimization problem*

$$\min\{C(f) \mid f \text{ is a flow}\}\,.$$

2.2 Observed Equilibria

In this section we introduce observed equilibria as an equilibrium concept. We base this equilibrium concept on the following assumptions.

- Each user aims to minimize her own expected congestion cost.
- Users have no a priori knowledge of the distribution of the demand.
- Users observe an approximate distribution of the demand by experience. Thus, they only observe the demand when they themselves use the network.

When we use these assumptions we arrive at the following conclusion. Users minimize their own expected cost, given some approximate demand distribution, that they know from observing when using the network. Formally, this means that users use conditional probabilities given that the user travels, instead of the exact distribution, to calculate expected costs. This yields the concept of observed equilibrium.

Definition 2.2 (Observed equilibrium)**.** *A flow f is said to be an observed equilibrium if for all $k \in K$ and all $P, Q \in \mathcal{P}$ with $f_P^k > 0$*

$$\sum_{a\in P}\sum_{\substack{S\in 2^K\\ S\ni k}} \frac{p_D(S)}{\sum\limits_{\substack{S\in 2^K\\ S\ni k}} p_D(S)} \cdot \ell_a\left(f_a(S)\right) \le \sum_{a\in Q}\sum_{\substack{S\in 2^K\\ S\ni k}} \frac{p_D(S)}{\sum\limits_{\substack{S\in 2^K\\ S\ni k}} p_D(S)} \cdot \ell_a\left(f_a(S)\right) . \tag{1}$$

Since observed equilibria are in fact Nash equilibria with an adapted objective for the players, all the usual properties of Nash equilibria hold. In particular, there always exists an observed equilibrium.

The following example illustrates the concept of observed equilibrium.

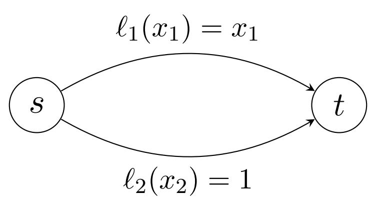

Figure 1: A two-link parallel network.

Example 2.3. Consider the network in Fig. 1. Let $K = \{k^1, k^2\}$, where both k^1 and k^2 are defined by $(s, t, 1)$. Assume that $p_D(\{k^1\}) = \frac{1}{2}$ and $p_D(K) = \frac{1}{2}$. Notice that even though both commodities are defined by the same triple, the two groups of users have asymmetric information: commodity k^1 faces a game of incomplete information, whereas k^2 knows (from past experience) that k^1 also travels when k^2 is selected for travel. The observed equilibrium is $f^1 = (1, 0)$ and $f^2 = (0, 1)$, yielding an expected congestion cost of $C(f) = 3/2$. The optimal flow is $f^{*1} = (1/2, 1/2)$ and $f^{*2} = (0, 1)$, yielding an expected congestion cost of $C(f^*) = 5/4$.

Similar as in the deterministic demand model, we can reformulate the equilibrium condition as a variational inequality.

Lemma 2.4. *A flow f is an observed equilibrium if and only if for all flows g*

$$\sum_{k\in K}\sum_{a\in A}\sum_{\substack{S\in 2^K\\ S\ni k}} p_D(S)\cdot \ell_a\left(f_a(S)\right)\cdot (f_a^k - g_a^k) \le 0\,. \tag{2}$$

Proof. Assume (2) is satisfied. We show that (1) holds. Let $k \in K$ and $P, Q \in \mathcal{P}$ with $f_P^k > 0$. Define g as follows

$$g_a^{k'} = \begin{cases} f_a^{k'} & \text{if } k' \neq k\text{, or } a \in P \cap Q\text{, or } a \notin P \cup Q\,,\\ f_a^{k'} - f_P^k & \text{if } k' = k \text{ and } a \in P \setminus Q\,,\\ f_a^{k'} + f_P^k & \text{if } k' = k \text{ and } a \in Q \setminus P\,. \end{cases}$$

By construction, g is a feasible flow. We get that

$$\begin{aligned}
&\sum_{k\in K}\sum_{a\in A}\sum_{\substack{S\in 2^K\\ S\ni k}} p_D(S)\cdot \ell_a\left(f_a(S)\right)\cdot (f_a^k - g_a^k)\\
&= \sum_{a\in P\setminus Q}\sum_{\substack{S\in 2^K\\ S\ni k}} p_D(S)\cdot \ell_a\left(f_a(S)\right)\cdot f_P^k - \sum_{a\in Q\setminus P}\sum_{\substack{S\in 2^K\\ S\ni k}} p_D(S)\cdot \ell_a\left(f_a(S)\right)\cdot f_P^k \le 0\,.
\end{aligned}$$

Dividing by $f_P^k > 0$ on both sides, and adding and subtracting

$$\sum_{a\in P\cup Q}\sum_{\substack{S\in 2^K\\ S\ni k}} p_D(S)\cdot \ell_a\left(f_a(S)\right)$$

yields an inequality equivalent to the observed equilibrium condition.

Now, assume that (1) is satisfied. We show that (2) holds. For all $k \in K$, there exists an π^k such that for all $P \in \mathcal{P}$ with $f_P^k > 0$ we have

$$\sum_{a\in P}\sum_{\substack{S\in 2^K\\ S\ni k}} p_D(S)\cdot \ell_a\left(f_a(S)\right) = \pi^k$$

and for all $Q \in \mathcal{P}$ with $f_Q = 0$ we have

$$\sum_{a\in Q}\sum_{\substack{S\in 2^K\\ S\ni k}} p_D(S)\cdot \ell_a\left(f_a(S)\right) \ge \pi^k\,.$$

Hence,

$$\begin{aligned}
&\sum_{k\in K}\sum_{a\in A}\sum_{\substack{S\in 2^K\\ S\ni k}} p_D(S)\cdot \ell_a\left(f_a(S)\right)\cdot f_a^k\\
&= \sum_{k\in K}\sum_{P\in \mathcal{P}^k} f_P^k \cdot \sum_{a\in P}\sum_{\substack{S\in 2^K\\ S\ni k}} p_D(S)\cdot \ell_a\left(f_a(S)\right)
\end{aligned}$$

$$
\begin{aligned}
&= \sum_{k\in K} \pi^k \cdot \sum_{P\in\mathcal{P}^k} f_P^k = \sum_{k\in K} \pi^k \cdot \sum_{P\in\mathcal{P}^k} g_P^k \\
&\leq \sum_{k\in K} \sum_{P\in\mathcal{P}^k} g_P^k \cdot \sum_{a\in P} \sum_{\substack{S\in 2^K \\ S\ni k}} p_D(S) \cdot \ell_a\left(f_a(S)\right) \\
&= \sum_{k\in K} \sum_{a\in A} \sum_{\substack{S\in 2^K \\ S\ni k}} p_D(S) \cdot \ell_a\left(f_a(S)\right) \cdot g_a^k .
\end{aligned}
$$

□

Again, just like in the deterministic demand setting, we can formulate observed equilibria as solutions to a minimization problem.

Proposition 2.5. *A flow f is an observed equilibrium if and only if it solves the following minimization problem*

$$
\min_f \sum_{a\in A} \sum_{S\in 2^K} p_D(S) \cdot \int_0^{f_a(S)} \ell_a(x)dx.
$$

Proof. Observe that the objective functions is convex and the feasible region is convex and compact. Thus the minimization problem is a convex optimization problem. So it is necessary and sufficient for f to satisfy the first order optimality conditions, for all flows g

$$
\sum_{k\in K} \sum_{a\in A} \sum_{\substack{S\in 2^K \\ S\ni k}} p_D(S) \cdot \ell_a\left(f_a(S)\right) \cdot (g_a^k - f_a^k) \geq 0,
$$

which is equivalent to (2).

□

2.3 Price of Anarchy

In this section, we analyze the price of anarchy for observed equilibria in network congestion games.

Let $\mathcal{L}$ be a class of latency functions. For the latency function $\ell \in \mathcal{L}$, define

$$
\beta(\ell) = \sup_{f,x\geq 0} \frac{(\ell(f) - \ell(x)) \cdot x}{\ell(f) \cdot f} ,
$$

where by convention $0/0 = 1$. In addition, define $\beta(\mathcal{L}) = \sup_{\ell\in\mathcal{L}} \beta(\ell)$.

Theorem 2.6. *Let f and f^* be the observed equilibrium and optimal flow, respectively, with $\ell_a \in \mathcal{L}$ for all $a \in A$. Then*

$$
C(f) \leq \frac{1}{1-\beta(\mathcal{L})} \cdot C(f^*) .
$$

Proof.

$$
C(f) = \sum_{a\in A} \sum_{S\in 2^K} p_D(S) \cdot \ell_a\left(f_a(S)\right) \cdot f_a(S)
$$

$$
\begin{aligned}
&= \sum_{k\in K}\sum_{a\in A}\sum_{\substack{S\in 2^K\\ S\ni k}} p_D(S)\cdot \ell_a\left(f_a(S)\right)\cdot f_a^k \\
&\le \sum_{k\in K}\sum_{a\in A}\sum_{\substack{S\in 2^K\\ S\ni k}} p_D(S)\cdot \ell_a\left(f_a(S)\right)\cdot f_a^{*k} \\
&= \sum_{a\in A}\sum_{S\in 2^K} p_D(S)\cdot \ell_a\left(f_a(S)\right)\cdot f_a^{*}(S) \\
&\le \sum_{a\in A}\sum_{S\in 2^K} p_D(S)\cdot \left(\beta(\mathcal{L})\cdot \ell_a\left(f_a(S)\right)\cdot f_a(S) + \ell_a\left(f_a^{*}(S)\right)\cdot f_a^{*}(S)\right) \\
&= \beta(\mathcal{L})\cdot C(f) + C(f^*)\,,
\end{aligned}
$$

where the first inequality follows by Lemma 2.4 and the second inequality by definition of $\beta(\mathcal{L})$. $\square$

Note that this bound on the price of anarchy is independent of the random demand variable. This is in contrast to the result of Wang et al. [33], for which the price of anarchy depends on the demand distribution. Moreover, our bound is equal to the bound on the price of anarchy for the fixed demand model [7, 30]. Since the bound is tight for that model, and observed equilibria coincide with Nash equilibria when the demand is deterministic, the bound is also tight for the model with variable demand. While the latter may suggest that variability in the demand could even improve the price of anarchy, the following example shows that the bound for polynomials with nonnegative coefficients is also tight if the distribution of the demand is non-trivial.

Example 2.7. Consider the network in Fig. 2. Let $K = \{k^1, k^2\}$, where both k^1 and k^2 are defined by $(s, t, 1)$.

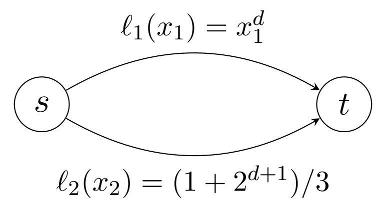

Figure 2: A two-link parallel network.

Assume that $p_D(\{k^1\}) = \frac{1}{4}$, $p_D(\{k^2\}) = \frac{1}{4}$ and $p_D(K) = \frac{1}{2}$. The observed equilibrium is $f^1 = f^2 = (1, 0)$, yielding an expected congestion cost of

$$C(f) = \frac{1+2^{d+1}}{2}\,.$$

The optimal flow is $f^{*1} = f^{*2} = \left(\frac{1}{(d+1)^{1/d}}, \frac{(d+1)^{1/d}-1}{(d+1)^{1/d}}\right)$, yielding an expected congestion cost of

$$C(f^*) = \left(1 - \frac{d}{(d+1)^{(d+1)/d}}\right)\frac{1+2^{d+1}}{2}\,.$$

Hence,

$$C(f) = \left(1 - \frac{d}{(d+1)^{(d+1)/d}}\right)^{-1} C(f^*) = \beta(\mathcal{L}_d)\cdot C(f^*)\,,$$

where $\mathcal{L}_d$ is the set of polynomial latency functions with nonnegative coefficients and degree at most d.

2.4 Adaptive Optimum

In this section, we compare the performance of the observed equilibrium to a stronger notion of social optimum: the adaptive optimal flow. An adaptive flow f is a vector $(f^S)_{S\in 2^K}$, where each f^S is a vector $(f^{k,S})_{k\in S}$ where $f^{k,S}$ is a flow for commodity k.

Definition 2.8. *An adaptive flow f^{**} is said to be an adaptive optimal flow if it solves the following minimization problem for all $S \in 2^K$*

$$\min_{f^S} \sum_{a\in A} \ell_a \left(\sum_{k\in S} f_a^{k,S} \right) \cdot \left(\sum_{k\in S} f_a^{k,S} \right) .$$

The total expected costs of the adaptive optimal flow are

$$C(f^{**}) = \sum_{a\in A} \sum_{S\in 2^K} p_D(S) \cdot \ell_a \left(\sum_{k\in S} f_a^{k,S} \right) \cdot \left(\sum_{k\in S} f_a^{k,S} \right) .$$

The following example shows that the observed equilibrium can perform arbitrarily worse than the adaptive optimal flow.

Example 2.9. Consider the network in Fig. 1. Let $m \in \mathbb{N}$, with $m > 1$, and let $K = \{k^1, \dots, k^{m^2}\}$, where k^i , for each $i = 1, \dots, m^2$, is defined by $(s, t, 1/m)$. Assume that

$$p_D(\{k^i\}) = \frac{1 - 1/m^3}{m^2} \text{ for each } i = 1, \dots, m^2, \text{ and } p_D(K) = \frac{1}{m^3} .$$

This implies that each commodity observes the following random variable: with probability $\frac{1-1/m^3}{m^2}$ only their own commodity travels and with probability $\frac{1}{m^3}$ all commodities travel. Let f and f^{**} denote the observed equilibrium and the adaptive optimum, respectively. Then

$$f^{k^i} = \left(\frac{m^4 + m^3 - m}{m^5 + m^4 - m}, \frac{m - 1}{m^5 + m^4 - m} \right) \text{ for each } i = 1, \dots, m^2 \text{ and}$$

$$f^{**} = \begin{cases} (\frac{1}{m}, 0) & \text{if the demand is } \frac{1}{m}, \\ (\frac{1}{2}, \frac{2m-1}{2}) & \text{if the demand is } m\,. \end{cases}$$

This yields respective expected congestion costs of

$$C(f) = \frac{m^3 + m^2 - 1}{m^4}$$

and

$$C(f^{**}) = \frac{8m^3 - m^2 - 4}{4m^5} .$$

Hence $C(f)/C(f^{**}) \to \infty$ as $m \to \infty$.

In order to obtain a constant upper bound, we have to restrict the set of allowable random demand variables. To this end, let $K = \{k^1, \dots, k^m\}$ be the set of commodities. For all $S \in 2^K$, we assume that $p_D(S) > 0$ if and only if $S = \{k^1, \dots, k^i\}$ for some $i \in \{1, \dots, m\}$. In particular, the sequence of sets of commodities with positive probability is a nested sequence of sets. Let $\mathcal{L}$ be a class of latency functions. For the latency function $\ell \in \mathcal{L}$, define

$$\omega(\ell) = \sup_{m \in \mathbb{N}, (x^i)_{i=1}^m, (f^i)_{i=1}^m} \frac{\left(\ell\left(\sum_{i=1}^m x^i\right) - \ell\left(\sum_{i=1}^m f^i\right)\right)\left(\sum_{i=1}^m f^i\right) + \ell\left(\sum_{i=1}^m x^i\right)\left(\sum_{i=1}^m x^i\right) - \sum_{i=1}^m \ell\left(\sum_{j=1}^i x^j\right) x^i}{\ell\left(\sum_{i=1}^m x^i\right)\left(\sum_{i=1}^m x^i\right)}$$

where, by convention, $0/0 = 1$. In addition, define $\omega(\mathcal{L}) = \sup_{\ell \in \mathcal{L}} \omega(\ell)$. Notice that if we would fix $m = 1$ in the definition of $\omega(\ell)$, then $\omega(\ell) = \beta(\ell)$. The extra additive term accounts for the adaptive setting. For more details on $\omega(\mathcal{L})$, see [14].

Theorem 2.10. *Let f and f^{**} be the observed equilibrium and adaptive optimal flow, respectively, with $\ell_a \in \mathcal{L}$ for all $a \in A$. Then*

$$C(f) \leq \frac{1}{1 - \beta(\mathcal{L})} \cdot \frac{1}{1 - \omega(\mathcal{L})} \cdot C(f^{**}) .$$

Proof. Let $x = (x^1, \dots, x^m)$ be defined recursively as the Wardrop equilibrium for demand level d^i given fixed flows $(x^1, \dots, x^{i-1})$ for each $i \in \{1, \dots, m\}$. See [14] for more details on this online Wardrop equilibrium. As x is a feasible flow, by Theorem 2.6

$$C(f) \leq \frac{1}{1 - \beta(\mathcal{L})} \cdot C(x) .$$

But then by [14], we have

$$\begin{aligned} C(f) &\leq \frac{1}{1 - \beta(\mathcal{L})} \cdot C(x) \\ &\leq \frac{1}{1 - \beta(\mathcal{L})} \cdot \frac{1}{1 - \omega(\mathcal{L})} \cdot C(f^{**}) . \end{aligned}$$

□

Proposition 2.11. *Let $\mathcal{L}$ be the class of affine latency functions. Then*

$$\frac{1}{1 - \beta(\mathcal{L})} \cdot \frac{1}{1 - \omega(\mathcal{L})} \leq \frac{16}{3} .$$

2.5 Continuous demand

The results discussed in this paper also hold for a single commodity with a continuous demand distribution, where we assume that every user in the commodity is equaly likely to be part of the demand. Note that this setting is still very different from the one where a user knows the exact distribution since the conditional distribution "favors" the higher demands.

To illustrate this idea, we restrict our attention to single-source, single-sink networks. In order to compare our results to the results of Wang et al. [33], we make the following assumptions. Given are a directed single-source, single-sink network (V, A), and a random variable D with probability density function $p_D : \mathbb{R}_+ \to [0, 1]$, such that $p_D(x)$ represents the probability that a subset of x units of users travels. Note that D is a random variable over the real numbers and not over all possible subsets of players. We assume that for each realization x, a subset of x units of users is drawn uniformly at random for travel.

Let $\mathcal{P}$ denote the set of directed $s-t$ paths. A unit flow f is a nonnegative vector $f = (f_P)_{P\in\mathcal{P}}$ such that $\sum_{P\in\mathcal{P}} f_P = 1$. For a unit flow f, let $f_a = \sum_{P\in\mathcal{P}} f_P$ be the amount of flow on arc a. Since all players are symmetric and have the same information, we assume that each player plays the same mixed strategy in equilibrium. Let f denote the unit flow of each player.

A unit flow f is an *observed equilibrium with uniform selection* if for all $P, Q \in \mathcal{P}$ with $f_P > 0$, we have

$$\sum_{a\in P} \int_{x\in\mathbb{R}_+} \frac{p_D(x)\cdot x}{\int_{x\in\mathbb{R}_+} p_D(x)\cdot x\, dx} \cdot \ell_a(f_a \cdot x)\, dx \le \sum_{a\in Q} \int_{x\in\mathbb{R}_+} \frac{p_D(x)\cdot x}{\int_{x\in\mathbb{R}_+} p_D(x)\cdot x\, dx} \cdot \ell_a(f_a \cdot x)\, dx\,.$$

A unit flow g is a *user equilibrium* as defined by Wang et al. [33], if for all $P, Q \in \mathcal{P}$ with $g_P > 0$, we have

$$\sum_{a\in P} \int_{x\in\mathbb{R}_+} p_D(x)\cdot \ell_a(g_a \cdot x)\, dx \le \sum_{a\in Q} \int_{x\in\mathbb{R}_+} p_D(x)\cdot \ell_a(g_a \cdot x)\, dx\,. \tag{3}$$

The next example shows that the observed equilibrium with uniform selection can perform arbitrarily better than the user equilibrium of [33].

Example 2.12. Let $m \in \mathbb{N}$, with $m > 1$. Define the probability, p, as $p(\frac{1}{m}) = \frac{m}{m+1}$ and $p(m) = \frac{1}{m+1}$. Consider the network in Fig. 1 . Let f and g be the observed equilibrium and the user equilibrium, respectively. Then

$$f = \left(\frac{m}{m^2-m+1}, \frac{m^2-2m+1}{m^2-m+1}\right) \text{ and } g = (1, 0)\,.$$

This yields respective expected congestion costs of

$$C(f) = 1$$

and

$$C(g) = m - 1 + 1/m\,.$$

Hence $C(g)/C(f) \to \infty$ as $m \to \infty$.

The following result shows that the observed equilibrium with uniform selection performs better than the user equilibrium (3) when latency functions are of the form $a_a \cdot x^p + b_a$ with $a_a, b_a \ge 0$.

Theorem 2.13. *Let f and g be the observed equilibrium and user equilibrium, respectively, with $\ell_a(x) = a_a \cdot x^p + b_a$ for some $a_a, b_a \ge 0$ for all $a \in A$. Then*

$$C(f) \le C(g)\,.$$

Proof. By the variational inequality of the observed equilibrium, we have

$$\begin{aligned}
&\sum_{a\in A}\int_{x\in\mathbb{R}_+} p_D(x)\cdot x\cdot(a_a\cdot f_a^p\cdot x^p+b_a)\cdot(f_a-g_a)\;dx\\
&=\sum_{a\in A}(a_a\cdot f_a^p\cdot\mathbb{E}(D^{p+1})+b_a\cdot\mathbb{E}(D))\cdot(f_a-g_a)\le 0\,.
\end{aligned}$$

By the variational inequality of the user equilibrium, we have

$$\begin{aligned}
&\mathbb{E}(D)\cdot\sum_{a\in A}\int_{x\in\mathbb{R}_+} p_D(x)\cdot(a_a\cdot g_a^p\cdot x^p+b_a)\cdot(g_a-f_a)\;dx\\
&=\sum_{a\in A}(a_a\cdot g_a^p\cdot\mathbb{E}(D)\cdot\mathbb{E}(D^p)+b_a\cdot\mathbb{E}(D))\cdot(g_a-f_a)\le 0\,.
\end{aligned}$$

Adding the above two inequalities yields

$$\begin{aligned}
&\sum_{a\in A} a_a\cdot(f_a^p\cdot\mathbb{E}(D^{p+1})-g_a^p\cdot\mathbb{E}(D)\cdot\mathbb{E}(D^p))\cdot(f_a-g_a)\\
&=\sum_{a\in A} a_a(f_a^p\cdot(\mathbb{E}(D^{p+1})-\mathbb{E}(D)\cdot\mathbb{E}(D^p))+(f_a^p-g_a^p)\mathbb{E}(D)\cdot\mathbb{E}(D^p))(f_a-g_a)\\
&=(\mathbb{E}(D^{p+1})-\mathbb{E}(D)\cdot\mathbb{E}(D^p))\cdot\sum_{a\in A} a_a\cdot f_a^p\cdot(f_a-g_a)\\
&+\mathbb{E}(D)\cdot\mathbb{E}(D^p)\cdot\sum_{a\in A} a_a\cdot(\sum_{i=1}^{p} f_a^{p-i}\cdot x_a^{i-1})\cdot(f_a-g_a)^2\le 0\,.
\end{aligned}$$

Notice that if $\mathbb{E}(D^{p+1})=\mathbb{E}(D)\cdot\mathbb{E}(D^p)$, then $f=g$ and thus $C(f)=C(g)$. So in the remainder of the proof we assume that $\mathbb{E}(D^{p+1})>\mathbb{E}(D)\cdot\mathbb{E}(D^p)$, which implies that $\sum_{a\in A} a_a\cdot f_a^p\cdot(f_a-g_a)\le 0$, as the second term is nonnegative. Since

$$\begin{aligned}
&\sum_{a\in A} a_a\cdot(f_a^p-g_a^p)\cdot g_a\\
&=p\cdot\sum_{a\in A} a_a\cdot f_a^p\cdot(f_a-g_a)-\sum_{a\in A} a_a\cdot(\sum_{i=1}^{p}(p+1-i)\cdot f_a^{p-i}\cdot g_a^{i-1})\cdot(f_a-g_a)^2\le 0\,,
\end{aligned}$$

we have

$$\begin{aligned}
C(f)&=\sum_{a\in A}(a_a\cdot f_a^p\cdot\mathbb{E}(D^{p+1})+b_a\cdot\mathbb{E}(D))\cdot f_a\\
&\le\sum_{a\in A}(a_a\cdot f_a^p\cdot\mathbb{E}(D^{p+1})+b_a\cdot\mathbb{E}(D))\cdot g_a\\
&\le\sum_{a\in A}(a_a\cdot g_a^p\cdot\mathbb{E}(D^{p+1})+b_a\cdot\mathbb{E}(D))\cdot g_a=C(g)\,,
\end{aligned}$$

where the first inequality follows from the variational inequality and the second inequality from the observation above. □

The last example shows that the previous result cannot be extended to more general polynomial latency functions.

Example 2.14. Assume that $p_D(1) = \frac{1}{2}$ and $p_D(2) = \frac{1}{2}$. Consider the network in Figure 3. Let f and g be the observed equilibrium and the user equilibrium, respectively. Then $f = (\frac{115-\sqrt{4045}}{90}, \frac{-25+\sqrt{4045}}{90})$ and $g = (\frac{13-\sqrt{53}}{10}, \frac{-3+\sqrt{53}}{10})$. Hence

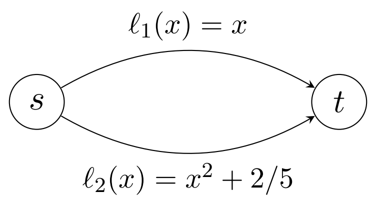

Figure 3: A two-link parallel network.

$$C(f) = \frac{115 - \sqrt{4045}}{36} > \frac{1551 - 115\sqrt{53}}{500} = C(g)\,.$$

3 Smooth games

The main result of Section 2, the fact that the price of anarchy for non-atomic routing games with variable demand is the same as for deterministic games, also applies to atomic routing games. It even holds for the more general class of smooth games, as defined by Roughgarden [27]. In this section, we first introduce smooth games in a slightly different, yet equivalent, way as compared to Roughgarden. Then, we state the more general result.

Let $K = \{1, \ldots, k\}$ denote the set of players. Each player k selects a strategy s_k from a set S_k. For each subset of players $S \subseteq K$ and corresponding strategy profile $s = (s_k)_{k\in S}$, player k incurs a costs $C_k^S(s)$. Note that the strategies are independent of the subset S. Let $C^S(s) = \sum_{k\in S} C_k^S(s)$ denote the total cost of a subset of players S and corresponding strategy profile s. We denote by s_{-k} the strategy profile of all players except player k, such that $s = (s_k, s_{-k})$.

Definition 3.1 (Smooth games)**.** *A cost-minimization game is called* (λ, μ)*-smooth, with* $\lambda > 0$ *and* $\mu < 1$*, for player set* $S \subseteq K$*, if for all corresponding strategy profiles* $s, s^\star$*, we have*

$$\sum_{k\in S} C_k^S(s_k^\star, s_{-k}) \le \lambda \cdot C^S(s^\star) + \mu \cdot C^S(s)\,.$$

We assume that the subset of players that participate in the game is determined stochastically. To that end, let D be a random variable with probability mass function $p_D : 2^K \to [0, 1]$. For each $S \subseteq K$, $p_D(S)$ represents the probability that only players from S play the game.

Given a strategy profile s, define the expected total costs as

$$C(s) = \sum_{k\in K} \sum_{\substack{S\in 2^K \\ S\ni k}} p_D(S) \cdot C_i^S(s)\,.$$

The strategy profile s^* minimizing the expected total costs is called *optimal.*

We propose the following definition of the observed equilibrium for cost-minimization games with stochastic demand, following the same reasoning as described in Section 2.4. The definition is in line with the Perfect Bayesian equilibrium in these games.

Definition 3.2 (Observed equilibrium)**.** *A strategy profile s is an observed equilibrium if for all $k \in K$ and all $s'_i \in S_i$,*

$$\sum_{\substack{S \in 2^K \\ S \ni k}} \frac{p_D(S)}{\sum_{\substack{S \in 2^K \\ S \ni k}} p_D(S)} \cdot C_i^S(s) \leq \sum_{\substack{S \in 2^K \\ S \ni k}} \frac{p_D(S)}{\sum_{\substack{S \in 2^K \\ S \ni k}} p_D(S)} \cdot C_i^S(s'_i, s_{-i}) \, .$$

Theorem 3.3. *Let s an observed equilibrium strategy profile and let s^* be an optimal solution. If a cost-minimization game is (λ, μ)-smooth for all player sets $S \subseteq K$, then*

$$C(s) \leq \frac{\lambda}{1-\mu} \cdot C(s^*) \, .$$

Proof.

$$\begin{aligned} C(s) &= \sum_{k \in K} \sum_{\substack{S \in 2^K \\ S \ni k}} p_D(S) \cdot C_i^S(s) \\ &\leq \sum_{k \in K} \sum_{\substack{S \in 2^K \\ S \ni k}} p_D(S) \cdot C_i^S(s_i^*, s_{-i}) \\ &= \sum_{S \in 2^K} \sum_{k \in S} p_D(S) \cdot C_i^S(s_i^*, s_{-i}) \\ &\leq \sum_{S \in 2^K} p_D(S) \cdot (\lambda \cdot C^S(s^*) + \mu \cdot C^S(s)) \\ &= \lambda \cdot C(s^*) + \mu \cdot C(s) \, , \end{aligned}$$

where the first inequality follows by the equilibrium condition and the second inequality by (λ, μ)-smoothness of the cost-minimization game for all player sets $S \subseteq K$. □

In particular, all known price of anarchy bounds for, for example, atomic or atomic splittable players also extend to the setting with random player sets. Moreover, one could define mixed observed equilibria, correlated observed equilibria, and coarse-correlated observed equilibria in a similar fashion as these are defined for games with deterministic demand, and the bound would still hold.

References

[1] Abdel-Aty, M., Kitamura, R., Jovanis, P.: Investigating effect of travel time variability on route choice using repeated-measurement stated preference data. Transportation research record 1493, 39–45 (1995)

[2] Awerbuch, B., Azar Y., Epstein, A.: The price of routing unsplittable flow. STOC 2005.

[3] Baldeschi, R., Cominetti, R., Mertikopoulos, P., Scarsini, M.: On the Price of Anarchy of highly congested non-atomic network games. SAGT 2016.

[4] Christodoulou, G., Koutsoupias, E.: The price of anarchy of finite congestion games. STOC 2005.

[5] Cominetti, R., Correa, J., Stier-Moses, N.: The impact of oligopolistic competition in networks. Operations Research, 57(6), 1421–1437 (2009)

[6] Cominetti, R., Torrico, A.: Additive consistency of risk measures and its application to risk-averse routing in networks. Mathematics of Operations Research 41(4), 1510–1521 (2016)

[7] Correa, J.R., Schulz, A.S., Stier Moses, N.E.: Selfish routing in capacitated networks. Mathematics of Operations Research 29(4), 961–976 (2004)

[8] Correa, J., Schulz, A., Stier-Moses, N.: A geometric approach to the price of anarchy in non-atomic congestion games. Games and Economic Behavior 64(2), 457–469 (2008)

[9] Dial, R.: A probabilistic multi-path traffic assignment algorithm which obviates path enumeration. Transportation Research 5(2), 83–111 (1971)

[10] Dubey, P.: Inefficiency of Nash equilibria. Mathematics of Operations Research 11(1), 1–8 (1986)

[11] Fan, Y., Kalaba, R., Moore, J.: Arriving on time. Journal of Optimization Theory and Applications 127(3), 497–513 (2005)

[12] Fudenberg, D., Tirole, J.: Perfect Bayesian equilibrium and sequential equilibrium. Journal of Economic Theory 53, 236–260 (1991)

[13] Harks, T.: Stackelberg strategies and collusion in network games with splittable flow. WAOA 2008.

[14] Harks, T., Végh, L.: Nonadaptive selfish routing with online demands. CAAN 2007.

[15] Koutsoupias, E., Papadimitriou, C.: Worst-case equilibria. Computer Science Review 3(2), 65–69 (2009)

[16] Koutsoupias, E., Papadimitrou, C.: Worst-case equilibria. STACS 1999.

[17] Lianeas, T., Nikolova, E., Stier Moses, N.: Asymptotically tight bounds for inefficiency in risk-averse selfish routing. IJCAI 2016.

[18] Nash, J.: Equilibrium points in n-person games. Proceedings of the National Academy of Sciences of the United States of America 36(1), 48–49 (1950)

[19] Nie, Y., Wu, X.: Shortest path problem considering on-time arrival probability. Transportation Research Part B: Methodological 43(6), 597–613 (2009)

[20] Nikolova, E., Stier-Moses, N.: A mean-risk model for the traffic assignment problem with stochastic travel times. Operations Research 62(2), 366–382 (2014)

[21] Nikolova, E., Stier Moses, N.: The burden of risk aversion in mean-risk selfish routing. EC 2015.

[22] Ordóñez, F., Stier-Moses, N.: Wardrop equilibria with risk-averse users. Transportation Science 44(1), 63–86 (2010)

[23] Papadimitriou, C.: Algorithms, games, and the internet. STOC 2001.

[24] Pigou, A.: The economics of welfare. Macmillan (1920)

[25] Qi, J., Sim, M., Sun, D., Yuan, X.: Preferences for travel time under risk and ambiguity: Implications in path selection and network equilibrium. Transportation Research Part B: Methodological (2016)

[26] Roughgarden, T.: The price of anarchy is independent of the network topology. Journal of Computer and System Sciences 67(2), 341–364 (2003)

[27] Roughgarden, T.: Intrinsic robustness of the price of anarchy. Journal of the ACM 62(5), 32 (2015)

[28] Roughgarden, T.: The price of anarchy in games of incomplete information. ACM Transactions on Economics and Computation, 32(1), 6 (2015)

[29] Roughgarden, T., Schoppmann, F.: Local smoothness and the price of anarchy in atomic splittable congestion games. SODA 2011.

[30] Roughgarden, T., Tardos, E.: How bad is selfish routing? Journal of the ACM 49(2), 236–259 (2002)

[31] Sheffi, Y.: Urban Transportation Networks. Prentice-Hall (1985)

[32] Syrgkanis, S.: Bayesian games and the smoothness framework. arXiv.cs.GT:1203.5155v1, 2012.

[33] Wang, C., Doan, X., Chen, B.: Price of anarchy for non-atomic congestion games with stochastic demands. Transportation Research Part B 70, 90–111 (2014)

[34] Wardrop, J.: Some theoretical aspects of road traffic research. Proceedings of the Institute of Civil Engineers 1(3), 325–362 (1952)