



Network congestion games are robust to variable demand

José Correa^a, Ruben Hoeksma^b, Marc Schröder^{c,*}

^a Universidad de Chile, Santiago, Chile

^b University of Bremen, Bremen, Germany

^c RWTH Aachen University, Aachen, Germany



ARTICLE INFO

Article history:

Received 23 May 2018

Revised 6 November 2018

Accepted 19 November 2018

ABSTRACT

We consider a non-atomic network congestion game with incomplete information in which nature decides which commodities travel. The users of a commodity do not know which other commodities travel and only have distributional information about their presence. Our main result is that the price of anarchy bounds known for the deterministic demand game also apply to the Bayesian game with random demand, even if the travel probabilities of different commodities are arbitrarily correlated. Moreover, the extension result of price of anarchy bounds for complete information games to incomplete information games in which the set of players is randomly determined can be generalized to the class of smooth games.

© 2018 Elsevier Ltd. All rights reserved.

1. Introduction

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. Variability in traffic demands is the main motivation for our work.

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 they takes from their origin to their destination. The amount of traffic that they encounters (the congestion) on this route then influences their 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 (Pigou, 1920). This phenomenon has captured the attention of researchers and practitioners for quite a long time,

* Corresponding author.

E-mail addresses: correa@uchile.cl (J. Correa), hoeksma@uni-bremen.de (R. Hoeksma), marc.schroeder@oms.rwth-aachen.de (M. Schröder).

see, e.g., [Dubey \(1986\)](#). 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 their objective by unilaterally deviating from their choice. These equilibria are also known as Nash equilibria ([Nash, 1950](#)) or Wardrop equilibria ([Wardrop, 1952](#)), 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 \(1999\)](#) and [Koutsoupias and Papadimitriou \(2009\)](#) 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 a socially optimized solution ([Papadimitriou, 2001](#)).

Since its introduction, the price of anarchy has been widely studied in the context of congestion games. [Roughgarden and Tardos \(2002\)](#) and [Roughgarden \(2003\)](#) 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. \(2004, 2008\)](#) 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 ([Awerbuch et al., 2005](#); [Christodoulou and Koutsoupias, 2005](#)), with atomic splittable players ([Cominetti et al., 2009](#); [Harks, 2008](#); [Roughgarden and Schoppmann, 2011](#)), and with large number of players ([Baladeschi et al., 2016](#)), 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 \(1971\)](#) and [Sheffi \(1985\)](#) 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. [Guo et al. \(2009\)](#) show that the price of anarchy increases with the users' perception error. [Abdel-Aty et al. \(1995\)](#) 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. \(2005\)](#) and [Nie and Wu \(2009\)](#) use the expected probability of arriving on time as an objective for the users. [Ordóñez and Stier-Moses \(2010\)](#) assume that users minimize the α quantile of the experienced travel time. Whereas [Nikolova and Stier-Moses \(2014\)](#) let users minimize expected delay plus a safety margin, approximated by the standard deviation of the distribution. [Qi et al. \(2016\)](#) 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 ([Cominetti and Torrico, 2016](#)). 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 ([Nikolova and Stier Moses, 2015](#); [Lianes et al., 2016](#)).

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 ([Cominetti and Torrico, 2016](#)). 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. Therefore, we consider a basic network congestion game with incomplete information where the demand is random and users simply minimize their expected travel time. This was first considered by [Gairing et al. \(2008\)](#) who study a Bayesian routing game on parallel links in which the players' weights are unknown. More recently, [Wang et al. \(2014\)](#) found that the price of anarchy can become arbitrarily large when the variability of the demand increases. However, the unboundedness of the price of anarchy in Wang et al. finds its roots in the way equilibria are defined. In their definition, a user bases their decisions on the full knowledge of the demand distribution (even though they may have never experienced the highly congested scenarios). Here we provide an alternative view. We define a *Bayesian Nash equilibrium* in which each user bases their decisions on the demand distribution conditioned on the fact that their commodity is traveling. The concept has a solid micro foundation and allows us to prove that 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 ([Roughgarden, 2003](#); [Correa et al., 2004](#)).

Network congestion games are an example of smooth games ([Roughgarden, 2015a](#)). 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 ([Aumann, 1974](#)) and coarse-correlated equilibria ([Hannan, 1957](#); [Moulin and Vial, 1978](#)). Recently, [Syrgkanis \(2012\)](#) and [Roughgarden \(2015b\)](#) 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, i.e. games for which players' types are drawn independently from prior distributions. [Roughgarden \(2015b\)](#) also shows that there is no general extension theorem for incomplete information games with correlated types. Our main result is an extension result for the setting in which there is uncertainty about which players play the game, and allows for correlations.

Our contribution. The main contribution of this paper is the extension of bounds on the price of anarchy for smooth games with incomplete information, where the incomplete information comes from randomness in the set of players participating in the game. In order to derive this result, we first consider a 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 over the subsets of users, so that the subset of users S travels with probability $p(S)$. 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 Section 3, we consider the Bayesian Nash equilibrium of the corresponding incomplete information game. Naturally, for this notion a user evaluates the expected cost of different paths only using their knowledge of the conditional probability measure. Section 4 uses this 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 (Roughgarden, 2003; Correa et al., 2004). With this result, we give an alternative to the pessimistic conclusion by Wang et al. (2014). 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.

Our extension result applies almost unchanged to the smoothness framework (Roughgarden, 2015a; 2015b), see Section 5. 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 (Awerbuch et al., 2005; Christodoulou and Koutsoupias, 2005; Cominetti et al., 2009; Harks, 2008; Roughgarden and Schoppmann, 2011) also apply to the model in which the set of players playing is determined stochastically.

2. Routing games with random demand

We consider the following non-atomic network congestion game with random demand. Given is a directed network (V, A) , a finite set of commodities K , and a probability space $(K, 2^K, p)$. 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 non-decreasing differentiable convex latency function $\ell_a : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ that represents the delay experienced by users traversing this link as a function of the total flow on the link. All users face the same latency function on an arc.

Each commodity $k \in K$ has a source s^k , a sink t^k , and a demand d^k . A commodity consists of infinitely many negligible small users that are always selected for travel together. Each commodity is selected for travel at random according to a *travel distribution* p . The travel distribution is a joint distribution over the set of all subsets of commodities. For each $S \subseteq K$, $p(S)$ represents the probability that exactly the users from commodities $k \in S$ travel. For a realization S , we say that the commodities $k \in S$ are *active*.

Note that we intentionally allow different commodities to have the same source-sink pairs. This permits modeling situations in which two or more different groups of users share the same source-sink pair, but have different travel frequencies.

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 non-negative vector $f^k = (f_p^k)_{P \in \mathcal{P}^k}$ such that $\sum_{P \in \mathcal{P}^k} f_p^k = d^k$. For arc a we denote by

$$f_a^k = \sum_{\substack{P \in \mathcal{P}^k \\ P \ni a}} f_p^k$$

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

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

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 travel behavior of a commodity is not only determined by its routing choices as described by the flow f , but also by the travel distribution that determines which commodities active. That means that for a given flow f the corresponding total expected congestion costs are equal to

$$C(f) = \sum_{a \in A} \sum_{S \subseteq K} p(S) \cdot \ell_a(f_a(S)) \cdot f_a(S).$$




Fig. 1. A two-link parallel network.

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

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

Note that Definition 2.1 defines the optimal flow as that flow that minimizes the total expected congestion cost.

3. Bayesian Nash equilibrium

In this section we introduce the equilibrium concept that we use. In a Bayesian Nash equilibrium each user minimizes their expected congestion costs with respect to their beliefs about the presence of other commodities and their choices. We assume that the beliefs about the presence of other commodities are derived from the distribution of the random demand variable. Whenever users travel the network, they observe the congestion on the arcs and they base their beliefs on that. In particular, since users of one commodity either are all active or all inactive, each user within that commodity bases their decision on the same information. This implies that the expected costs of a path for a user are calculated using their conditional probability, conditioned on their own commodity being active. This yields the following solution concept.

Definition 3.1. A flow f is said to be a Bayesian Nash equilibrium if for all $k \in K$ and all $P, Q \in \mathcal{P}$ with $f_P^k > 0$

$$\sum_{a \in P} \sum_{\substack{S \subseteq K \\ S \ni k}} p(S) \cdot \ell_a(f_a(S)) \leq \sum_{a \in Q} \sum_{\substack{S \subseteq K \\ S \ni k}} p(S) \cdot \ell_a(f_a(S)). \quad (1)$$

Note that both the left hand side and the right hand side of (1) sum only over sets S that contain commodity k . Thus (1) says that if commodity k uses the path P , any other path cannot have a lower expected congestion cost, conditioned on the fact that commodity k is part of the set of active commodities. Note that we do not claim that this equilibrium notion is the result of a multi-stage game dynamics where users observe the network every time they travel. Yet, if the users are in a Bayesian Nash equilibrium and they observe the congestion in the way described, they have no incentive to change their strategy.

Bayesian Nash equilibria are in fact Nash/Wardrop equilibria with an adapted objective for the players. After we illustrate the concept of a Bayesian Nash equilibrium in an example, we show that the usual properties, i.e. existence of an equilibrium, also hold.

The following example illustrates the concept of a Bayesian Nash equilibrium.

Example 3.2. Consider the network in Fig. 1. Let $K = \{k^1, k^2\}$, where both k^1 and k^2 have source s , sink t , and demand 1. The travel distribution is given by $p(\{k^1\}) = \frac{1}{2}$ and $p(K) = \frac{1}{2}$. In this example, with probability $\frac{1}{2}$ only commodity k^1 travels, whereas with probability $\frac{1}{2}$ both commodities travel. Notice that even though both commodities have the same source-sink pair and demand, due to the nature of the travel distribution the two groups of users have asymmetric information: commodity k^1 always travels, while k^2 travels half the time. Therefore, commodity k^1 faces a game of incomplete information, whereas k^2 knows that k^1 travels whenever k^2 travels. Let us now argue that $f = ((1, 0), (0, 1))$ is a Bayesian Nash equilibrium.¹ First, consider condition (1) for commodity k^1 ,

$$\begin{aligned} p(\{k^1\}) \cdot \ell_1(f_1^1) + p(\{k^1, k^2\}) \cdot \ell_1(f_1^1 + f_1^2) &= \frac{1}{2} \cdot f_1^1 + \frac{1}{2} \cdot (f_1^1 + f_1^2) = \frac{1}{2} \cdot 1 + \frac{1}{2} \cdot 1 \\ &\leq 1 = p(\{k^1\}) \cdot \ell_2(f_1^1) + p(\{k^1, k^2\}) \cdot \ell_2(f_1^1 + f_1^2), \end{aligned}$$

where the left-hand side is the expected congestion on the top arc and the right-hand side is the expected congestion cost on the bottom arc. Second, consider condition (1) for commodity k^2 ,

$$\begin{aligned} p(\{k^1, k^2\}) \cdot \ell_2(f_2^1 + f_2^2) &= \frac{1}{2} \cdot 1 \\ &\leq \frac{1}{2} = p(\{k^1, k^2\}) \cdot \ell_1(f_1^1 + f_1^2), \end{aligned}$$

¹ Finding Bayesian Nash equilibria can be done by solving (3).

where the left-hand side is the expected congestion on the bottom arc and the right-hand side is the expected congestion cost on the top arc. Hence (1) is satisfied for both commodities and thus f is a Bayesian Nash equilibrium. The corresponding expected total congestion costs are

$$C(f) = p(\{k^1\}) \cdot (\ell_1(f_1^1) + \ell_2(f_2^1)) + p(\{k^1, k^2\}) \cdot (\ell_1(f_1^1 + f_1^2) + \ell_2(f_2^1 + f_2^2)) = \frac{1}{2} \cdot 1 + \frac{1}{2} \cdot 2 = \frac{3}{2}.$$

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

Lemma 3.3. *A flow f is a Bayesian Nash equilibrium if and only if, for all flows g , we have*

$$\sum_{k \in K} \sum_{a \in A} \sum_{\substack{S \subseteq K \\ S \ni k}} p(S) \cdot \ell_a(f_a(S)) \cdot (f_a^k - g_a^k) \leq 0. \tag{2}$$

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

$$\sum_{k \in K} \sum_{a \in A} \sum_{\substack{S \subseteq K \\ S \ni k}} p(S) \cdot \ell_a(f_a(S)) \cdot (f_a^k - g_a^k) = \sum_{a \in P \setminus Q} \sum_{\substack{S \subseteq K \\ S \ni k}} p(S) \cdot \ell_a(f_a(S)) \cdot f_P^k - \sum_{a \in Q \setminus P} \sum_{\substack{S \subseteq K \\ S \ni k}} p(S) \cdot \ell_a(f_a(S)) \cdot f_P^k \leq 0.$$

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

$$\sum_{a \in P \cup Q} \sum_{\substack{S \subseteq K \\ S \ni k}} p(S) \cdot \ell_a(f_a(S))$$

yields an inequality equivalent to the Bayesian Nash equilibrium condition.

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

$$\sum_{a \in P} \sum_{\substack{S \subseteq K \\ S \ni k}} p(S) \cdot \ell_a(f_a(S)) = \pi^k$$

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

$$\sum_{a \in Q} \sum_{\substack{S \subseteq K \\ S \ni k}} p(S) \cdot \ell_a(f_a(S)) \geq \pi^k.$$

Hence,

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

□

Just as for the deterministic demand setting, we have that Bayesian Nash equilibria are solutions of a convex minimization problem and hence equilibrium flows exist and have equal total expected costs.

Proposition 3.4. A Bayesian Nash equilibrium flow f exists. Moreover, if f and g are Bayesian Nash equilibrium flows, then $C(f) = C(g)$.

Proof. Consider the following convex minimization problem

$$\min \left\{ \sum_{a \in A} \sum_{S \subseteq K} p(S) \cdot \int_0^{f_a(S)} \ell_a(x) dx \mid f \text{ is a flow.} \right\}. \quad (3)$$

For all flows g , we have that $f + (g - f) = g$ is a feasible flow, and thus $g - f$ is a feasible direction and therefore the first order optimality conditions are

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

which is equivalent to (2). In other words, an optimal solution of (3) is a Bayesian Nash equilibrium flow. Since the objective function is continuous, and the feasible region is convex and compact, there always exists a solution. By convexity of the objective function, we have that whenever f and g are solutions, $\sum_{S \subseteq K} p(S) \cdot \ell_a(f_a(S)) = \sum_{S \subseteq K} p(S) \cdot \ell_a(g_a(S))$ for all $a \in A$ and hence $C(f) = C(g)$. \square

4. Price of anarchy

In this section, we analyze the price of anarchy for Bayesian Nash equilibria in network congestion games with random demand.

Example 4.1. Consider the instance of Example 3.2, where $K = \{k^1, k^2\}$, $p(\{k^1\}) = \frac{1}{2}$ and $p(K) = \frac{1}{2}$. The Bayesian Nash equilibrium is $f^1 = (1, 0)$ and $f^2 = (0, 1)$ with corresponding expected total congestion costs of $C(f) = \frac{1}{2} \cdot 1 + \frac{1}{2} \cdot 2 = \frac{3}{2}$.

The optimal flow is $f^{*1} = (\frac{1}{2}, \frac{1}{2})$ and $f^{*2} = (0, 1)$ with corresponding expected total congestion costs of $C(f) = \frac{1}{2} \cdot \frac{3}{4} + \frac{1}{2} \cdot \frac{7}{4} = \frac{5}{4}$. Hence, the Bayesian Nash equilibrium is not efficient.

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 $\beta(\mathcal{L}) = \sup_{\ell \in \mathcal{L}} \beta(\ell)$.

The main result of this section shows that the loss in efficiency for Bayesian Nash equilibria in the incomplete information game is at most the loss in efficiency for Nash/Wardrop equilibria in the complete information game.

Theorem 4.2. Let f and f^* be the Bayesian Nash 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.

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

where the first inequality follows by Lemma 3.3 and the second inequality by definition of $\beta(\mathcal{L})$. \square

Corollary 4.3. Let f and f^* be the Bayesian Nash equilibrium and optimal flow, respectively, with ℓ_a affine for all $a \in A$. Then

$$C(f) \leq \frac{4}{3} \cdot C(f^*).$$




Fig. 2. A two-link parallel network.

Note that this bound on the price of anarchy is independent of the random demand variable, and equal to the bound on the price of anarchy for the fixed demand model (Correa et al., 2004; Roughgarden and Tardos, 2002). Interestingly, this is in contrast to the result of Wang et al. (2014), who conclude that the price of anarchy depends on the demand distribution. The key difference is that in their model they assume that users of the network use the full distribution of demand to calculate expected costs, whereas in our setting users only use the marginal distribution (i.e., the distribution of demand conditioned on the fact that their commodity is active). Evidently, this modeling difference makes a significant difference for the resulting equilibrium. Our result is in line with a recent stream of literature that shows that price of anarchy bounds of Nash equilibria in complete information games extend to Bayesian–Nash equilibria in games of incomplete information. We refer to Roughgarden (2015b) for an extensive list of references. Moreover, recent empirical work shows that the price of anarchy in real-life traffic networks is not too large. See, for example, Monnot et al. (2017) for a recent study of travel behavior in Singapore.

Since the price of anarchy bound is tight for the deterministic model, and Bayesian Nash equilibria coincide with Nash equilibria when the demand is deterministic, the bound is also tight for the model with random 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 non-negative coefficients is also tight if the distribution of the demand is non-trivial.

Example 4.4. Consider the network in Fig. 2. Let $K = \{k^1, k^2\}$, where both k^1 and k^2 have source s , sink t , and demand 1 and let the travel distribution be given by $p(\{k^1\}) = \frac{1}{4}$, $p(\{k^2\}) = \frac{1}{4}$ and $p(K) = \frac{1}{2}$. The Bayesian Nash equilibrium solves

$$\begin{aligned} \frac{1}{4} \cdot (f_1^1)^d + \frac{1}{2} \cdot (f_1^1 + f_1^2)^d &= \frac{3}{4} \cdot \frac{1 + 2^{d+1}}{3} \\ \frac{1}{4} \cdot (f_1^2)^d + \frac{1}{2} \cdot (f_1^1 + f_1^2)^d &= \frac{3}{4} \cdot \frac{1 + 2^{d+1}}{3}, \end{aligned}$$

yielding $f^1 = f^2 = (1, 0)$. The corresponding expected total congestion costs are

$$C(f) = \frac{1}{4} \cdot 1 + \frac{1}{4} \cdot 1 + \frac{1}{2} \cdot 2^{d+1} = \frac{1 + 2^{d+1}}{2}.$$

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 total 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 non-negative coefficients and degree at most d .

Remark 4.5. We can also compare the performance of a Bayesian Nash equilibrium to a stronger notion of social optimum: the adaptive optimal flow. An adaptive flow f is a vector $(f^S)_{S \subseteq K}$, where each f^S is a vector $(f_a^{k,S})_{k \in S}$ where $f_a^{k,S}$ is a flow for commodity k . An adaptive flow f^{**} is said to be an *adaptive optimal flow* if it solves the following minimization problem for all $S \subseteq K$

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

The total expected cost of the adaptive optimal flow is

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

The adaptive optimal flow minimizes the congestion cost for each possible demand separately, i.e. for each $S \subseteq K$ a possibly different flow is chosen. Therefore, its total expected congestion cost is in general lower than the cost for the optimal flow as defined in Definition 2.1, that minimizes the total expected congestion cost with a single flow. The following example shows that the Bayesian Nash equilibrium can perform arbitrarily worse than the adaptive optimal flow.

Example 4.6. 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 has source s , sink t , and demand $1/m$, for each $i = 1, \dots, m^2$. Assume that

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

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 Bayesian Nash 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

$$C(f) = \frac{m^3 + m^2 - 1}{m^4} \text{ and } C(f^{**}) = \frac{8m^3 - m^2 - 4}{4m^5}.$$

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

5. Smooth games

The main result of Section 2, that the price of anarchy for non-atomic network congestion games with random demand is the same as for deterministic games, also applies to atomic congestion games. The extension result even holds for the more general class of smooth games, as defined by Roughgarden (2015a). Roughgarden (2015b) provides an extension theorem of price of anarchy bounds for independent player types. However, his definition of smooth games with incomplete information is slightly more restrictive than ours and, therefore, yields a stronger result as he compares to the adaptive optimal costs. In this section, we first introduce smooth games, and then we state the result.

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

Definition 5.1. A cost-minimization game is called (λ, μ) -smooth, with $\lambda > 0$ and $\mu < 1$, for player set $S \subseteq N$, if for all corresponding strategy profiles s, s^* , we have

$$\sum_{i \in S} C_i^S(s_i^*, s_{-i}) \leq \lambda \cdot C^S(s^*) + \mu \cdot C^S(s).$$

We assume that the subset of players that participate in the game is determined stochastically. To that end, let $(N, 2^N, p)$ be a probability space. For each $S \subseteq N$, $p(S)$ represents the probability that only players from S play the game.

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

$$C(s) = \sum_{S \subseteq N} p(S) \cdot C^S(s).$$

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

We define a Bayesian Nash equilibrium in these games as follows.

Definition 5.2. A strategy profile s is a Bayesian Nash equilibrium if for all $i \in N$ and all $s'_i \in S_i$,

$$\sum_{\substack{S \subseteq N \\ S \ni i}} p(S) \cdot C_i^S(s) \leq \sum_{\substack{S \subseteq N \\ S \ni i}} p(S) \cdot C_i^S(s'_i, s_{-i}).$$

The main result shows that the price of anarchy bounds for smooth games also extend to smooth games with random player sets.

Theorem 5.3. *Let s be an Bayesian Nash equilibrium strategy profile and let s^* be an optimal solution. If a cost-minimization game is (λ, μ) -smooth for all player sets $S \subseteq N$, then*

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

Proof.

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

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 $S \subseteq N$. \square

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.

6. Discussion and conclusion

We consider a non-atomic routing game with incomplete information in which the travel frequency of each commodity is determined by a random variable. We show existence and essential uniqueness of Bayesian Nash equilibria and derive the following main result regarding the price of anarchy. Every routing game with incomplete information, in which the incomplete information comes from the randomness by which commodities are selected for travel, has the same upper bound on the price of anarchy as the corresponding deterministic demand game. Moreover, the extension result also holds for the more general class of smooth games. It would be interesting to see if there are different settings in which price of anarchy bounds carry over when allowing for correlations. Another interesting research direction is to study situations in which incomplete information on the demand side can in fact improve (in terms of worst case guarantees) the associated price of anarchy bounds.

Acknowledgments

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.

References

- Abdel-Aty, M., Kitamura, R., Jovanis, P., 1995. Investigating effect of travel time variability on route choice using repeated-measurement stated preference data. *Transp. Res. Rec.* 1493, 39–45.
- Aumann, R., 1974. Subjectivity and correlation in randomized strategies. *J. Math. Econ.* 1 (1), 67–96.
- Awerbuch, B., Azar, Y., Epstein, A., 2005. The price of routing unsplittable flow. In: *Proceedings of the STOC*.
- Baldeschi, R., Cominetti, R., Mertikopoulos, P., Scarsini, M., 2016. On the price of anarchy of highly congested non-atomic network games. In: *Proceedings of the SAGT*.
- Christodoulou, G., Koutsoupias, E., 2005. The price of anarchy of finite congestion games. In: *Proceedings of the STOC*.
- Cominetti, R., Correa, J., Stier-Moses, N., 2009. The impact of oligopolistic competition in networks. *Oper. Res.* 57 (6), 1421–1437.
- Cominetti, R., Torrico, A., 2016. Additive consistency of risk measures and its application to risk-averse routing in networks. *Math. Oper. Res.* 41 (4), 1510–1521.
- Correa, J., Schulz, A., Stier-Moses, N., 2004. Selfish routing in capacitated networks. *Math. Oper. Res.* 29 (4), 961–976.
- Correa, J., Schulz, A., Stier-Moses, N., 2008. A geometric approach to the price of anarchy in non-atomic congestion games. *Games Econ. Behav.* 64 (2), 457–469.
- Dial, R., 1971. A probabilistic multi-path traffic assignment algorithm which obviates path enumeration. *Transp. Res.* 5 (2), 83–111.
- Dubey, P., 1986. Inefficiency of Nash equilibria. *Math. Oper. Res.* 11 (1), 1–8.
- Fan, Y., Kalaba, R., Moore, J., 2005. Arriving on time. *J. Optim. Theory Appl.* 127 (3), 497–513.
- Gairing, M., Monien, B., Tiemann, K., 2008. Selfish routing with incomplete information. *Theory Comput. Syst.* 42 (1), 91–130.
- Guo, X., Yang, H., Liu, T.L., 2009. Bounding the inefficiency of Logit-based stochastic user equilibrium. *Eur. J. Oper. Res.* 201 (2), 463–469.
- Hannan, J., 1957. Approximation to Bayes risk in repeated play. *Contrib. Theory Games* 3, 97–139.
- Harks, T., 2008. Stackelberg strategies and collusion in network games with splittable flow. In: *Proceedings of the WAOA*.
- Koutsoupias, E., Papadimitriou, C., 2009. Worst-case equilibria. *Comput. Sci. Rev.* 3 (2), 65–69.

- Koutsoupias, E., Papadimitriou, C., 1999. Worst-case Equilibria. In: Proceedings of the STACS.
- Lianes, T., Nikolova, E., Stier Moses, N., 2016. Asymptotically tight bounds for inefficiency in risk-averse selfish routing. In: Proceedings of the IJCAI.
- Monnot, B., Benita, F., Piliouras, G., 2017. How bad is selfish routing in practice? arXiv:1703.01599.
- Moulin, H., Vial, J., 1978. Strategically zero-sum games: the class of games whose completely mixed equilibria cannot be improved upon. *Int. J. Game Theory* 7 (3–4), 201–221.
- Nash, J., 1950. Equilibrium points in n-person games. *Proc. Natl. Acad. Sci. U.S.A.* 36 (1), 48–49.
- Nie, Y., Wu, X., 2009. Shortest path problem considering on-time arrival probability. *Transp. Res. Part B Methodol.* 43 (6), 597–613.
- Nikolova, E., Stier-Moses, N., 2014. A mean-risk model for the traffic assignment problem with stochastic travel times. *Oper. Res.* 62 (2), 366–382.
- Nikolova, E., Stier Moses, N., 2015. The burden of risk aversion in mean-risk selfish routing. In: Proceedings of the EC.
- Ordóñez, F., Stier-Moses, N., 2010. Wardrop equilibria with risk-averse users. *Transp. Sci.* 44 (1), 63–86.
- Papadimitriou, C., 2001. Algorithms, games, and the internet. In: Proceedings of the STOC.
- Pigou, A., 1920. *The Economics of Welfare*. Macmillan.
- Qi, J., Sim, M., Sun, D., Yuan, X., 2016. Preferences for travel time under risk and ambiguity: implications in path selection and network equilibrium. *Transp. Res. Part B Methodol.*
- Roughgarden, T., 2003. The price of anarchy is independent of the network topology. *J. Comput. Syst. Sci.* 67 (2), 341–364.
- Roughgarden, T., 2015a. Intrinsic robustness of the price of anarchy. *J. ACM* 62 (5), 32.
- Roughgarden, T., 2015b. The price of anarchy in games of incomplete information. *ACM Trans. Econ. Comput.* 32 (1), 6.
- Roughgarden, T., Schoppmann, F., 2011. Local smoothness and the price of anarchy in atomic splittable congestion games. In: Proceedings of the SODA.
- Roughgarden, T., Tardos, E., 2002. How bad is selfish routing? *J. ACM* 49 (2), 236–259.
- Sheffi, Y., 1985. *Urban Transportation Networks*. Prentice-Hall.
- Syrkkanis, S., 2012. Bayesian games and the smoothness framework. arXiv:1203.5155.
- Wang, C., Doan, X., Chen, B., 2014. Price of anarchy for non-atomic congestion games with stochastic demands. *Transp. Res. Part B* 70, 90–111.
- Wardrop, J., 1952. Some theoretical aspects of road traffic research. *Proc. Inst. Civil Eng.* 1 (3), 325–362.