

ON THE ASYMPTOTIC BEHAVIOR OF THE PRICE OF ANARCHY: IS SELFISH ROUTING BAD IN HIGHLY CONGESTED NETWORKS?

RICCARDO COLINI-BALDESCHI, ROBERTO COMINETTI,
PANAYOTIS MERTIKOPOULOS, AND MARCO SCARSINI

ABSTRACT. This paper examines the asymptotic behavior of the price of anarchy as a function of the total traffic inflow in nonatomic congestion games with multiple origin-destination pairs. We first show that the price of anarchy may remain bounded away from 1, even in simple three-link parallel networks with convex cost functions. On the other hand, empirical studies show that the price of anarchy is close to 1 in highly congested real-world networks, thus begging the question: under what assumptions can this behavior be justified analytically? To that end, we prove a general result showing that for a large class of cost functions (defined in terms of regular variation and including all polynomials), the price of anarchy converges to 1 in the high congestion limit. In particular, specializing to networks with polynomial costs, we show that this convergence follows a power law whose degree can be computed explicitly.

1. INTRODUCTION

Almost every commuter in a major metropolitan area has experienced the frustration of being stuck in traffic. At best, this might mean being late for dinner; at worst, it means more accidents and altercations, not to mention the vastly increased damage to the environment caused by huge numbers of idling engines.

To name but an infamous example, China's G110 traffic jam in August 2010 brought to a standstill thousands of vehicles for 100 kilometers between Hebei and Inner Mongolia. The snarl-up lasted twelve days and resulted in drivers being unable to move for more than 1 kilometer per day, reportedly spending up to five days trapped in the jam. Not caused by weather or a natural disaster, this massive tie-up was instead laid at the feet of a bevy of trucks swarming on the shortest route to Beijing, thus clogging the G110 highway to a halt (while ironically carrying supplies for construction work to ease congestion). This, therefore, raises the following question: how much better would things have been if all traffic had been routed by a benevolent social planner who could calculate (and enforce) the optimum traffic assignment?

2010 *Mathematics Subject Classification*. Primary 91A13; secondary 91A43.

Key words and phrases. nonatomic congestion games; price of anarchy; high congestion; regular variation.

Riccardo Colini-Baldeschi is a member of GNAMPA-INdAM. Roberto Cominetti and Panayotis Mertikopoulos gratefully acknowledge the support and hospitality of LUISS during a visit in which this research was initiated. Roberto Cominetti's research is also supported by FONDECYT 1130564 and Núcleo Milenio ICM/FIC RC130003 "*Información y Coordinación en Redes*". Panayotis Mertikopoulos was partially supported by the French National Research Agency (ANR) project ORACLESS (ANR-16-CE33-0004-01) and the ECOS/CONICYT Grant C15E03. Marco Scarsini is a member of GNAMPA-INdAM. He gratefully acknowledges the support and hospitality of FONDECYT 1130564 and Núcleo Milenio "*Información y Coordinación en Redes*".

In game-theoretic terms, this question boils down to the inefficiency of Nash equilibria that are not Pareto optimal. The most widely used quantitative measure of this inefficiency is the so-called *price of anarchy* (PoA): introduced by Koutsoupias and Papadimitriou (1999) and so dubbed by Papadimitriou (2001), the price of anarchy is the ratio of the social cost of the least efficient Nash equilibrium divided by the minimum achievable social cost. By virtue of this straightforward definition, deriving worst-case bounds for the price of anarchy has given rise to a vigorous literature at the interface of computer science and economics, with many surprising results.

In the context of network congestion, Pigou (1920) was probably the first to study the inefficiency of selfish routing, and his elementary two-road example with a price of anarchy equal to $4/3$ is one of the two prototypical examples thereof. The other is due to Braess (1968), and consists of a four-road network where the addition of a zero-cost segment makes things just as bad as in Pigou’s original example. These two examples were the starting point for the seminal work of Roughgarden and Tardos (2002), who showed that the price of anarchy in (nonatomic) routing games with affine costs may *never* exceed $4/3$. On the other hand, if the network’s cost functions are polynomials of degree d or less, the worst-case value of the price of anarchy grows as $\Theta(d/\log d)$, Roughgarden (2003), implying that selfish routing can be arbitrarily bad in networks with polynomial costs.

By this token, and given the typically nonlinear relation between traffic loads and travel times (the Internet being a prime example of this), the intervention of a central planner seems necessary in order to regain some measure of efficiency. At the same time however, these worst-case instances are typically realized in networks with delicately tuned traffic loads and costs. Hence, if the traffic inflow in a given network grows large (the so-called “high congestion” regime), it is not clear whether the price of anarchy also remains high in the limit.

In view of all this, our aim in this paper is to examine the asymptotic behavior of the price of anarchy in highly congested networks. By default, networks of this kind are usually modeled as nonatomic games with a continuum of players, each carrying a negligible amount of the overall traffic. In this setting, using both analytical and numerical methods, a recent study by O’Hare et al. (2016) suggests that the price of anarchy is usually close to 1 for very high and very low traffic, and it fluctuates in the intermediate regime. Colini-Baldeschi et al. (2016) further showed that the price of anarchy indeed converges to 1 in highly congested parallel-link networks with a single origin-destination (O/D) pair and costs that satisfy some regularity conditions; however, they also provided several two-link examples where the price of anarchy oscillates as a function of the traffic inflow.

All this leads to the following natural questions:

- (1) Under which conditions does the price of anarchy converge to 1 in highly congested networks? And, when it does, how fast is this convergence?
- (2) Do these conditions depend on the network topology, its cost functions, or both?
- (3) Can general results be obtained only for networks with one O/D pair or do they extend to networks with multiple such pairs?

Informally, our main result states that, for networks with multiple O/D pairs and for a large class of cost functions (including all polynomials), the price of anarchy indeed converges to 1 as the traffic inflow increases, independently of the network

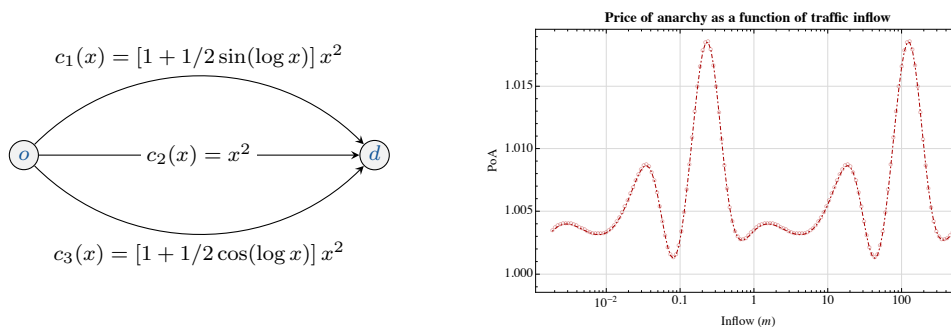


Figure 1. A simple network where selfish routing remains inefficient in the high congestion regime.

topology – and at a rate that follows a power law in the case of polynomial costs. In other words, under heavy traffic, selfish routing becomes just as efficient as the optimum traffic assignment: in this setting, myopically taking the fastest path isn’t the root of the problem, the amount of traffic is.

1.1. Our results. We start, in [Section 3](#), by constructing a 3-link network where the price of anarchy oscillates between two values strictly greater than 1 (cf. [Fig. 1](#)). As mentioned before, [Colini-Baldeschi et al. \(2016\)](#) provided several examples where the PoA oscillates in the high congestion limit, but the lower limit of this oscillation was always 1; by contrast, in the network of [Fig. 1](#), the PoA is *bounded away* from 1. Since the cost functions in this example are convex, we deduce that convexity does not guarantee that selfish routing becomes optimal infinitely often in the high congestion limit. Moreover, this counterexample involves a very simple parallel network with just three edges; hence, the asymptotically bad behavior of the PoA is not due to the network’s topology, either.

The key feature of this example is that, as the load on each of the nonlinear edges grows large, the corresponding cost function $c_e(x)$ undergoes periodic oscillations at a logarithmic scale so that

$$\liminf_{x \rightarrow \infty} \frac{c_e(x)}{c_{e'}(x)} \neq \limsup_{x \rightarrow \infty} \frac{c_e(x)}{c_{e'}(x)} \quad (1.1)$$

for any two edges e, e' in the network. To dispense with such pathologies, we focus in [Section 4](#) on general networks with multiple O/D pairs that admit a regularly varying *gauge function* $c(x)$ that classifies edges into *fast*, *slow* or *tight*, depending on the growth rate of the cost along each edge.¹ Paths are likewise fast, slow or tight, based on their slowest edge; and an O/D pair is tight if its fastest path is tight.² Our main result may then be stated as follows: *if every O/D pair has a path which is not slow and there is a tight O/D pair that carries a nonnegligible fraction of the traffic, the price of anarchy converges to 1 in the high congestion limit.*

¹Regular variation means here that $\lim_{t \rightarrow \infty} c(tx)/c(t)$ is finite and nonzero for all $x > 0$. The class of regularly varying functions therefore contains all affine functions, polynomials, (poly)logarithms and many other types of functions.

²For instance, if all the network’s cost functions are polynomials of degree d , all edges, paths and O/D pairs are tight with respect to the gauge function $c(x) = x^d$.

Put differently, in a highly congested network as above, a benevolent central authority with full control of traffic assignment would do no better than what myopic agents would do by themselves. However, this leaves open the question of *when* a network enters this “highly congested” regime. We address this question in [Section 5](#), where we examine the speed at which the price of anarchy converges to 1 as a function of the traffic inflow. Specializing to networks with polynomial costs, we find that this convergence follows a power law with respect to the total traffic inflow, and we derive an explicit estimate of the degree of this law.

Interestingly, if the network can only serve a finite amount of traffic, the situation can be quite different. Although intuition might suggest otherwise, when the total traffic inflow approaches the network’s capacity, the price of anarchy may converge to a value that is strictly greater than 1, even in simple two-link networks with $M/M/1$ cost functions. This suggests an important distinction between the two limits (high congestion vs. capacity), which also changes the impact that a social planner might have in one case or the other. We discuss this issue in [Section 6](#).

1.2. Related work. Much of the literature on congestion games is devoted to the study of bounds for the price of anarchy under different conditions. [Roughgarden and Tardos \(2002\)](#) proved a bound of $4/3$ for the price of anarchy in the case of affine cost functions, independently of the topology of the network. This bound is sharp in that, for every $M > 0$, there exists a network with traffic inflow M with price of anarchy equal to $4/3$. Importantly, our analysis in [Section 4](#) shows that the order of the quantifiers *cannot be exchanged*: in any network as above, the price of anarchy gets arbitrarily close to 1 if the traffic inflow is sufficiently large.

Worst-case bounds for the price of anarchy have been obtained for larger classes of cost functions. For polynomial costs with degree at most d , [Roughgarden \(2003\)](#) showed that the worst possible instance grows as $\Theta(d/\log d)$. [Dumrauf and Gairing \(2006\)](#) provided sharper bounds for polynomial costs that consist of monomials of maximum degree d and minimum degree s . [Roughgarden and Tardos \(2004\)](#) provided a unifying result for costs that are differentiable with $xc(x)$ convex, while [Correa et al. \(2004, 2008\)](#) considered less regular classes of cost functions; in particular, [Correa et al. \(2004\)](#) studied networks with a finite capacity. [Correa et al. \(2007\)](#) also studied the price of anarchy when the goal is to minimize the maximum latency rather than the average latency. For a survey, the reader is referred to [Roughgarden and Tardos \(2007\)](#) and [Roughgarden \(2007\)](#).

The difference between the mean value of the price of anarchy and its worst value has been studied in the context of cognitive radio networks by [Law et al. \(2012\)](#). [Youn et al. \(2008\)](#) studied the difference between optimal and actual system performance in real transportation networks, focusing in particular on Boston’s road network. They observed that the price of anarchy depends crucially on the total traffic inflow: it starts at 1, it then grows with some oscillations, and then it returns to 1 as the flow increases. [González Vayá et al. \(2015\)](#) studied optimal scheduling for the electricity demand of a fleet of plug-in electric vehicles: without using the term, they showed that the price of anarchy goes to one as the number of vehicles grows. [Cole and Tao \(2016\)](#) showed that in large Walrasian auctions and in large Fisher markets the PoA goes to one as the market size increases. Finally, [Piliouras et al. \(2013\)](#) examined how the price of anarchy is affected by uncertainty in the game.

Friedman (2004) considered the behavior of the price of anarchy when the inflow is slightly perturbed with respect to the worst case scenario. They showed that, for any value of the inflow that produces large losses for selfish routing, the cost of equilibrium routing falls rapidly. Patriksson (2004) and Josefsson and Patriksson (2007) performed a sensitivity analysis of Wardrop equilibria with respect to a range of parameters. Englert et al. (2010) examined the variation of an equilibrium in a congestion game when either the total mass of players is increased by ε or an edge that carries an ε -fraction of the mass is removed. More recently, Colini-Baldeschi et al. (2016) examined the asymptotic behavior of the price of anarchy in nonatomic congestion games when the underlying network consists of parallel edges linking a single O/D pair. Finally, Feldman et al. (2016) took a different asymptotic approach and considered atomic games where the number of players grows to infinity. Applying the notion of (λ, μ) -smoothness to the resulting sequence of atomic games, they showed that the price of anarchy converges to the corresponding nonatomic limit.

2. MODEL AND PRELIMINARIES

In this section, we present some basic definitions and results that are used throughout our paper. Specifically, in Section 2.1, we describe the class of nonatomic congestion games under study, while, in Section 2.2, we give a precise definition of the price of anarchy and the various optimality/equilibrium notions involved.

2.1. Network model. Following Beckmann et al. (1956) and Roughgarden and Tardos (2002), the basic component of our model is a finite directed multi-graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with vertex set \mathcal{V} and edge set \mathcal{E} . We further assume there is a finite set of *origin-destination* (O/D) pairs $i \in \mathcal{I}$, each with an individual *traffic demand* $\mu^i \geq 0$ which has to be routed from an origin $o^i \in \mathcal{V}$ to a destination $d^i \in \mathcal{V}$ via \mathcal{G} .

To route this traffic, the i -th O/D pair employs a set \mathcal{P}^i of (simple) paths joining o^i to d^i , each path $p \in \mathcal{P}^i$ comprising a sequence of edges that meet head-to-tail in the usual way.³ For bookkeeping reasons, we also assume that the sets \mathcal{P}^i are disjoint (which holds in particular when all pairs (o^i, d^i) are different). Then, writing $\mathcal{P} \equiv \bigcup_{i \in \mathcal{I}} \mathcal{P}^i$ for the union of all such paths, the set of feasible *routing flows* $f = (f_p)_{p \in \mathcal{P}}$ in the network is defined as

$$\mathcal{F} = \left\{ f \in \mathbb{R}_+^{\mathcal{P}} : \sum_{p \in \mathcal{P}^i} f_p = \mu^i \text{ for all } i \in \mathcal{I} \right\}. \quad (2.1)$$

In turn, a routing flow $f \in \mathcal{F}$ induces a *load* on each edge $e \in \mathcal{E}$ as

$$x_e = \sum_{p \ni e} f_p, \quad (2.2)$$

and we write $x = (x_e)_{e \in \mathcal{E}}$ for the corresponding *load profile* on the network.

Given all this, the delay (or latency) experienced by an infinitesimal traffic element in order to traverse edge e is determined by a nondecreasing, continuous *cost function* $c_e: [0, \infty) \rightarrow (0, \infty)$. Specifically, if $x = (x_e)_{e \in \mathcal{E}}$ is the load profile induced by a feasible routing flow $f = (f_p)_{p \in \mathcal{P}}$, the incurred delay on edge $e \in \mathcal{E}$ is $c_e(x_e)$.

³To be clear, we do not assume here that \mathcal{P}^i is the set of all paths joining o^i to d^i , but only some (nonempty) subset thereof. This distinction is important for Internet-like network and software-defined networking where only paths with a low hop count are used.

Hence, with a slight abuse of notation, the associated cost of path $p \in \mathcal{P}$ is given by the expression

$$c_p(f) \equiv \sum_{e \in p} c_e(x_e). \quad (2.3)$$

Putting together all of the above, the tuple $\Gamma = (\mathcal{G}, \mathcal{I}, \{\mu^i\}_{i \in \mathcal{I}}, \{\mathcal{P}^i\}_{i \in \mathcal{I}}, \{c_e\}_{e \in \mathcal{E}})$ will be referred to as a (*nonatomic*) *routing game*.

2.2. Equilibrium, optimality, and the price of anarchy. In this nonatomic context, the notion of Nash equilibrium is captured by *Wardrop's first principle*: at equilibrium, the delays along all utilized paths are equal and no higher than those that would be experienced by an infinitesimal traffic element going through an unused route (Wardrop, 1952). Formally, a routing flow f^* is said to be a *Wardrop equilibrium* (WE) of Γ if, for all $i \in \mathcal{I}$, we have

$$c_p(f^*) \leq c_{p'}(f^*) \quad \text{for all } p, p' \in \mathcal{P}^i \text{ such that } f_p^* > 0. \quad (2.4)$$

By the work of Beckmann et al. (1956), it is well known that Wardrop equilibria can be characterized equivalently as solutions of the (convex) minimization problem:

$$\begin{aligned} & \text{minimize} && \sum_{e \in \mathcal{E}} C_e(x_e), \\ & \text{subject to} && x_e = \sum_{p \ni e} f_p, \quad f \in \mathcal{F}, \end{aligned} \quad (\text{WE})$$

where $C_e(x_e) = \int_0^{x_e} c_e(w) dw$ denotes the primitive of c_e .

On the other hand, a *socially optimum* (SO) flow is defined as a solution to the total cost minimization problem:

$$\begin{aligned} & \text{minimize} && L(f) = \sum_{p \in \mathcal{P}} f_p c_p(f), \\ & \text{subject to} && f \in \mathcal{F}. \end{aligned} \quad (\text{SO})$$

To quantify the gap between solutions to (WE) and (SO), we write

$$\text{Eq}(\Gamma) = L(f^*) \quad \text{and} \quad \text{Opt}(\Gamma) = \min_{f \in \mathcal{F}} L(f), \quad (2.5)$$

where f^* is a Wardrop equilibrium of Γ .⁴ The game's *price of anarchy* (PoA) is then defined as

$$\text{PoA}(\Gamma) = \frac{\text{Eq}(\Gamma)}{\text{Opt}(\Gamma)}. \quad (2.6)$$

Obviously, $\text{PoA}(\Gamma) \geq 1$ with equality if and only if Wardrop equilibria are also socially efficient. Our main objective in the rest of our paper will be to study the asymptotic behavior of this ratio in the limit where the traffic inflow grows large.

3. A NETWORK WHERE SELFISH ROUTING IS ALWAYS INEFFICIENT

In this section, we begin our study of heavily congested networks by constructing a simple three-link network where the price of anarchy oscillates between two values strictly greater than 1.

⁴That Wardrop equilibria have the same total cost is due to Beckmann et al. (1956).

To begin, let Γ_M be a nonatomic routing game consisting of a single O/D pair with traffic inflow M . This traffic is to be routed over the three-link parallel graph of Fig. 1 with cost functions

$$c_1(x_1) = x_1^d \left[1 + \frac{1}{2} \sin(\log x_1)\right], \quad (3.1a)$$

$$c_2(x_2) = x_2^d, \quad (3.1b)$$

$$c_3(x_3) = x_3^d \left[1 + \frac{1}{2} \cos(\log x_3)\right]. \quad (3.1c)$$

The functions c_e , $e \in \{1, 2, 3\}$, are convex for $d \geq 2$. Moreover, the functions $x_e c_e(x_e)$ are strictly convex, so the problem (SO) admits a unique optimum traffic distribution. Hence, the only way for the game's price of anarchy to be equal to 1 is if the game's (unique) Wardrop equilibrium coincides with the network's socially optimum flow.

For a given value of the total traffic inflow $M = x_1 + x_2 + x_3$, the load profile $x = (x_1, x_2, x_3)$ is a Wardrop equilibrium if and only if $c_1(x_1) = c_2(x_2) = c_3(x_3)$,⁵ i.e. if the normalized profile $z = x/M$ satisfies

$$z_1^d \left[1 + \frac{1}{2} \sin(\log M z_1)\right] = z_2^d = z_3^d \left[1 + \frac{1}{2} \cos(\log M z_3)\right]. \quad (3.2)$$

Likewise, after differentiating and rearranging, the corresponding conditions for the network's socially optimum flow are

$$\begin{aligned} z_1^d \left[1 + \frac{1}{2} \sin(\log M z_1) + \frac{1}{2(d+1)} \cos(\log M z_1)\right] &= z_2^d \\ &= z_3^d \left[1 + \frac{1}{2} \cos(\log M z_3) - \frac{1}{2(d+1)} \sin(\log M z_3)\right]. \end{aligned} \quad (3.3)$$

Thus, the solutions to (3.2) and (3.3) coincide if and only if

$$\cos(\log M z_1) = 0 = \sin(\log M z_3), \quad (3.4)$$

that is, if and only if there exist integers $k_1, k_3 \in \mathbb{Z}$ such that

$$\begin{aligned} \log M z_1 &= k_1 \pi + \frac{\pi}{2}, \\ \log M z_3 &= k_3 \pi. \end{aligned} \quad (3.5)$$

This implies that $\sin(\log M z_1) = \pm 1$ and $\cos(\log M z_3) = \pm 1$, leading to the following cases:

Case 1: $\sin(\log M z_1) = 1$, $\cos(\log M z_3) = 1$. Substituting in (3.2) we get $z_1^d = z_3^d$ so that $\log z_1 = \log z_3$, and then (3.5) gives

$$k_1 \pi + \frac{\pi}{2} = k_3 \pi. \quad (3.6)$$

This amounts to $\frac{1}{2} = k_3 - k_1$, which cannot hold for integer values of k_1 and k_3 .

Case 2: $\sin(\log M z_1) = 1$, $\cos(\log M z_3) = -1$. Similarly as above, from (3.2) we get $3z_1^d = z_3^d$ so that $\frac{1}{d} \log 3 + \log z_1 = \log z_3$, and then (3.5) gives

$$\frac{1}{d} \log 3 + k_1 \pi + \frac{\pi}{2} = k_3 \pi. \quad (3.7)$$

This yields $\log 3/\pi = d(k_3 - k_1 - \frac{1}{2})$, which again cannot hold for $k_1, k_3, d \in \mathbb{Z}$.

The remaining two cases lead to a contradiction in the same way, implying that the game's Wardrop equilibrium and socially optimum flow cannot coincide for any value of M . Since Eqs. (3.2) and (3.3) are periodic in $\log M$, it also follows that

⁵Since an unused edge always has a cost of zero, all paths are used at equilibrium.

the game's price of anarchy is periodic at a logarithmic scale. Thus, by focusing on the period $1 \leq M \leq e^{2\pi}$, we conclude that

$$\inf_{M \geq 0} \text{PoA}(\Gamma_M) = \min_{1 \leq M \leq e^{2\pi}} \text{PoA}(\Gamma_M) > 1, \quad (3.8)$$

i.e., Wardrop equilibria in the network of Fig. 1 *remain strictly inefficient, even as* $M \rightarrow \infty$. Notice also that, given the logarithmic periodicity of $\text{PoA}(\Gamma_M)$, the same holds when $M \rightarrow 0$ as well.

Remark. The network discussed above complements a recent set of two-link counterexamples provided by Colini-Baldeschi et al. (2016), who showed that the price of anarchy may oscillate between 1 and an arbitrarily large value in the high congestion limit. From an efficiency standpoint, the saving grace of the latter class of instances is that the price of anarchy gets arbitrarily close (or even equal) to 1 infinitely often; by contrast, the instance presented here shows that even this weak guarantee may fail to hold in general.

4. THE ASYMPTOTIC BEHAVIOR OF THE PRICE OF ANARCHY

The example of the previous section shows that, in complete generality, the price of anarchy can be bounded away from 1 for all values of the traffic inflow. However, given the irregular behavior of the cost model (3.1), the question remains: *does selfish routing remain inefficient in highly congested networks under reasonable assumptions for the network's cost functions?*

To answer this, we must first give precise meaning to what “highly congested” means. If there is a single O/D pair in the network, this is straightforward: as in Section 3, the network becomes highly congested when the inflow of said pair grows to infinity. However, if there are several O/D pairs in the network, the traffic inflow of each pair could be growing at very different rates, with some pairs possibly contributing a negligible amount of traffic in the limit. In particular, the inflow of some O/D pairs could remain finite (or even vanish), but the network may still become heavily congested if the aggregation of these traffic demands grows large.

To account for this, we define the *total traffic inflow* in the network as

$$M = \sum_{i \in \mathcal{I}} \mu^i, \quad (4.1)$$

and we write

$$\lambda^i = \frac{\mu^i}{M} \quad (4.2)$$

for the fraction of the traffic generated by the i -th O/D pair. Then, somewhat informally, the high congestion regime refers to the limit $M \rightarrow \infty$, without making any assumptions on the behavior of the relative inflow vector $\lambda = (\lambda^i)_{i \in \mathcal{I}}$ in this limit.

To make the statement of our main result (and its proof) less cumbersome, it will be more convenient to consider a sequence of inflow vectors $\mu_n = M_n \lambda_n$, $n \in \mathbb{N}$. In this way, for each n , we get a routing game $\Gamma_n \equiv (\mathcal{G}, \mathcal{I}, \mu_n, \{\mathcal{P}^i\}_{i \in \mathcal{I}}, \{c_e\}_{e \in \mathcal{E}})$ over the same underlying network, but with different total inflow M_n . The high congestion limit is then defined as follows:

Definition 4.1. We say that Γ_n becomes *highly congested* if $M_n \rightarrow \infty$ as $n \rightarrow \infty$. In this case, we will also refer to Γ_n as a *highly congested network*.

To distinguish between O/D pairs that consistently contribute to congestion, we also say that $i \in \mathcal{I}$ has *heavy traffic* (or is *heavy*) whenever $\liminf_{n \rightarrow \infty} \lambda_n^i > 0$; similarly, a subset $\mathcal{I}' \subseteq \mathcal{I}$ of O/D pairs has *heavy traffic* if $\liminf_{n \rightarrow \infty} \sum_{i \in \mathcal{I}'} \lambda_n^i > 0$.

Going back to the example of Section 3, the main source of inefficiency under heavy traffic is that two of the cost functions (3.1) exhibit (logarithmic) oscillations as $x \rightarrow \infty$. To dispense with this kind of pathologies, it is necessary to impose some regularity on the asymptotic growth of the network's cost functions. We do so by means of the following regularity notion:

Definition 4.2. A function $g: [0, \infty) \rightarrow (0, \infty)$ is said to be *regularly varying* if

$$\lim_{t \rightarrow \infty} \frac{g(tx)}{g(t)} \text{ is finite and nonzero for all } x \geq 0. \quad (4.3)$$

Standard examples of regularly varying functions include all affine, polynomial and logarithmic/polylogarithmic functions. The notion itself dates back to the work of Karamata (1933, 1930) and has been used extensively in probability and large deviations theory (see, e.g. de Haan and Ferreira, 2006; Jessen and Mikosch, 2006; Resnick, 2007). Economic applications can also be found in Alcalá (2014); for a comprehensive survey, the reader is referred to Bingham et al. (1989).

With all this at hand, we will discard irregular growth phenomena (such as those observed in Section 3) by positing that each cost function $c_e(x)$ can be compared asymptotically to some regularly varying function $c(x)$. Specifically, given an ensemble of cost functions $\mathcal{C} = \{c_e\}_{e \in \mathcal{E}}$, a regularly varying function $c: [0, \infty) \rightarrow (0, \infty)$ is called a *benchmark function* for \mathcal{C} if the (possibly infinite) limit

$$\alpha_e = \lim_{x \rightarrow \infty} \frac{c_e(x)}{c(x)} \quad (4.4)$$

exists for all $e \in \mathcal{E}$. When it exists, this limit will be called the *c-index* of edge e (or, more simply, its *index*), and e will be called *fast*, *slow*, or *tight* (relative to c) if α_e is respectively 0, ∞ , or in-between. Since bottlenecks are caused by the slowest edges, we also define the *c-index* of a path $p \in \mathcal{P}$ as

$$\alpha_p = \max_{e \in p} \alpha_e, \quad (4.5)$$

and the path is said to be *fast*, *slow*, or *tight* based on whether α_p is 0, ∞ , or between these two extremes. Finally, extending this terminology to O/D pairs, we say that a pair $i \in \mathcal{I}$ is *tight* if $0 < \alpha^i < \infty$, where the *c-index* α^i of a pair is defined as the minimal index of its corresponding paths, namely

$$\alpha^i = \min_{p \in \mathcal{P}^i} \alpha_p. \quad (4.6)$$

In other words, a path is fast (resp. tight, resp. slow) if its slowest edge is fast (resp. tight, resp. slow), and an O/D pair is tight if its fastest path is tight.

Finding a benchmark is trivial (for instance, we can always take $c(x) = 1$), but finding a benchmark that yields a meaningful classification in the high congestion limit is less so. To that end, we introduce the following refinement:

Definition 4.3. A highly congested network is called *gaugeable* if it admits a benchmark function $c: [0, \infty) \rightarrow (0, \infty)$ such that:

- i) Every O/D pair has a path which is not slow: $\alpha^i < \infty$ for all $i \in \mathcal{I}$.
- ii) The set of tight O/D pairs has heavy traffic: $\liminf_{n \rightarrow \infty} \sum_{i: 0 < \alpha^i < \infty} \lambda_n^i > 0$.

In this case, the function c will be called the network's *gauge function*.

Essentially, the existence of a gauge function $c(x)$ guarantees that (i) every O/D pair can route its traffic along a path whose cost grows as $\mathcal{O}(c(x))$; and (ii) a non-negligible fraction of the total traffic incurs a cost that scales as $\Theta(c(x))$. Both of these conditions are fairly general and ensure that the network's cost functions are mutually comparable in the high congestion limit (as opposed to the oscillatory example of Section 3). For instance, (i) is satisfied automatically if the network's cost functions are bounded or if they are polynomials (cf. Corollaries 4.5 and 4.7 below); likewise, (ii) is satisfied if every O/D pair is tight or if some O/D pair is both tight and heavy. Finally, in networks with a single O/D pair, the situation is even simpler as it suffices for the pair to be tight.

We are now in a position to state our main result for highly congested networks. This result extends Corollary 3 in Colini-Baldeschi et al. (2016), who considered only the single O/D case on a parallel link network.

Theorem 4.4. *If a highly congested network Γ_n is gaugeable, then $\text{PoA}(\Gamma_n) \rightarrow 1$.*

Before discussing the proof of Theorem 4.4, we first present some immediate corollaries:

Corollary 4.5. *Suppose that each O/D pair $i \in \mathcal{I}$ has a path $p \in \mathcal{P}^i$ such that all edges $e \in p$ have bounded costs, i.e. $\lim_{x \rightarrow \infty} c_e(x) < \infty$. Then, $\text{PoA}(\Gamma_n) \rightarrow 1$.*

Proof. Taking $c(x) = 1$, we get $\alpha_e = \lim_{x \rightarrow \infty} c_e(x) \in (0, \infty]$ for all $e \in \mathcal{E}$. By assumption, every O/D pair is tight, so our claim follows from Theorem 4.4. ■

Corollary 4.6. *Suppose there exists a regularly varying function $c(x)$ such that $c(x) \leq c_e(x) \leq (1 + o(1))c(x)$ for all $e \in \mathcal{E}$. Then, $\text{PoA}(\Gamma_n) \rightarrow 1$.*

Proof. Taking $c(x)$ as a benchmark, we get $\alpha_e = 1$ for all $e \in \mathcal{E}$. This implies that every O/D pair is tight, so our claim follows from Theorem 4.4. ■

Corollary 4.7. *If the network's costs are polynomial and every O/D pair $i \in \mathcal{I}$ has heavy traffic, then $\text{PoA}(\Gamma_n) \rightarrow 1$.*

Proof. Let d_e be the degree of c_e and set $d_p = \max_{e \in p} d_e$ for all $p \in \mathcal{P}$. Furthermore, let $d^i = \min_{p \in \mathcal{P}^i} d_p$ and set $d = \max_{i \in \mathcal{I}} d^i$. It is then easy to see that $c(x) = x^d$ is a gauge function for the network, so the result follows from Theorem 4.4. ■

Corollary 4.8. *Suppose that the network's cost functions c_e are regularly varying and the (possibly infinite) limits $\alpha_{e,e'} = \lim_{x \rightarrow \infty} c_e(x)/c_{e'}(x)$ exist for all $e, e' \in \mathcal{E}$. Assume further that every O/D pair $i \in \mathcal{I}$ has heavy traffic. Then, $\text{PoA}(\Gamma_n) \rightarrow 1$.*

Proof. Define a total order among the network's edges by setting $e \preceq e'$ if and only if $\alpha_{e,e'} < \infty$. For each path $p \in \mathcal{P}$, choose a maximal element e_p of p , i.e. an edge $e_p \in p$ such that $e \preceq e_p$ for all $e \in p$. Also, for each O/D pair $i \in \mathcal{I}$, choose a path p^i for which e_{p^i} is minimal, i.e. $e_{p^i} \preceq e_{p'}$ for all $p' \in \mathcal{P}^i$. Finally, pick an O/D pair $j \in \mathcal{I}$ such that e_{p^j} is maximal, i.e. $e_{p^i} \preceq e_{p^j}$ for all $i \in \mathcal{I}$. Letting $e^* = e_{p^j}$, it follows that $c(x) = c_{e^*}(x)$ is a gauge function for the network, so our claim follows from Theorem 4.4. ■

Remark 4.1. Similar conclusions for networks with cost functions that are not regularly varying are not hard to obtain. For instance, an extension of the proof of Theorem 4.4 shows that Corollary 4.7 still holds if some of the network's cost functions are superpolynomial, provided there exists a path with polynomial costs.

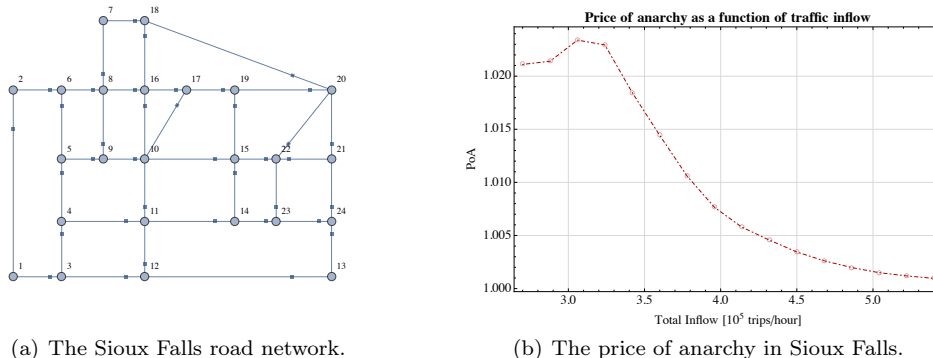


Figure 2. The price of anarchy in Sioux Fall as a function of overall traffic.

For an illustration of [Theorem 4.4](#), [Fig. 2](#) shows the behavior of the price of anarchy as a function of traffic inflow in the road network of the Sioux Falls metropolitan area, a standard case study in the transportation literature (see, e.g., [O’Hare et al., 2016](#)). The network’s (two-way) arterial roads are shown in [Fig. 2\(a\)](#), and the delay function of each road is taken to be of the BPR (Bureau of Public Roads) type

$$c_e(x) = a_e + b_e x^4, \quad (4.7)$$

with constants a_e and b_e taken from the standard reference paper of [LeBlanc et al. \(1975, Table 1\)](#).⁶ Traffic data for the number of trips per hour between each pair of nodes in the network were taken from the same source; [LeBlanc et al. \(1975\)](#) report an average value of $M_{\text{avg}} \approx 3.6 \times 10^5$ trips/hour for the network’s total traffic inflow. The network’s price of anarchy for various values of M (both above and below M_{avg}) is then plotted in [Fig. 2\(b\)](#): as can be seen, the price of anarchy decreases rapidly, even when the total inflow is close to the reported average M_{avg} .

Proof of [Theorem 4.4](#). We now turn to the proof of [Theorem 4.4](#). The main idea is to use the gauge function $c(x)$ to capture the asymptotic behavior of the convex optimization problems (WE) and (SO), and then show that their solutions (Wardrop equilibria and socially optimums flows respectively) coincide in the limit.

We begin with some required notation. First, let $\mu = (\mu^i)_{i \in \mathcal{I}}$ denote a traffic inflow vector with total inflow $M = \sum_{i \in \mathcal{I}} \mu^i$ and relative inflows given by the vector $\lambda = (\lambda^i)_{i \in \mathcal{I}}$. Instead of working directly with the flow variables $f \in \mathcal{F}$, it will be more convenient to introduce the rescaled *traffic allocation* variables $y^i = (y_p^i)_{p \in \mathcal{P}^i}$ defined as

$$y_p^i = f_p / \mu^i \quad \text{for all } p \in \mathcal{P}^i, i \in \mathcal{I}. \quad (4.8)$$

We clearly have $\sum_{p \in \mathcal{P}^i} y_p^i = 1$ for all $i \in \mathcal{I}$; we will also write $\mathcal{Y}^i = \Delta(\mathcal{P}^i)$ for the simplex of traffic allocations of $i \in \mathcal{I}$ and $\mathcal{Y} = \prod_{i \in \mathcal{I}} \mathcal{Y}^i$ for the product ensemble thereof. Moreover, descending to the edge level, we define the rescaled load per destination induced on $e \in \mathcal{E}$ as

$$z_e^i(y) = \sum_{p \in \mathcal{P}^i, p \ni e} y_p^i \quad (4.9)$$

⁶See also the online repository <https://github.com/bstabler/TransportationNetworks>.

and we denote respectively the *rescaled* and *total load* on $e \in \mathcal{E}$ as

$$\zeta_e(y, \lambda) = \sum_{i \in \mathcal{I}} \lambda^i z_e^i(y), \quad (4.10)$$

$$x_e(y, \mu) = \sum_{i \in \mathcal{I}} \mu^i z_e^i(y). \quad (4.11)$$

Clearly $x_e(y, \mu) = M \zeta_e(y, \lambda)$, and the objective function of (SO) may be expressed in terms of y as

$$L_\mu(y) = \sum_{e \in \mathcal{E}} x_e(y, \mu) c_e(x_e(y, \mu)). \quad (4.12)$$

With these transformed variables at hand, the next important component of our proof is the following characterization of regularly varying functions:

Lemma 4.9 (Karamata). *If g is regularly varying, there exists some $\rho \in \mathbb{R}$ such that*

$$\lim_{t \rightarrow \infty} \frac{g(tx)}{g(x)} = x^\rho \quad \text{for all } x > 0. \quad (4.13)$$

When this is the case, g is called ρ -regularly varying.

A proof of Karamata's lemma above can be found in Bingham et al. (1989). In what follows, we will often need to compare the asymptotic behavior of the cost along an edge to that of the network's gauge, so we require a slightly stronger result for the asymptotic behavior of functions that are only asymptotically equal to a regularly varying function:

Lemma 4.10. *Consider two functions $f, g: [0, \infty) \rightarrow (0, \infty)$ such that:*

- (1) f is nondecreasing.
- (2) g is ρ -regularly varying for some $\rho > 0$.
- (3) $\lim_{x \rightarrow \infty} f(x)/g(x) = \alpha \in [0, \infty)$.

If $M_n \rightarrow \infty$ and $z_n \rightarrow z \in [0, \infty)$, then

$$\lim_{n \rightarrow \infty} \frac{f(M_n z_n)}{g(M_n)} = \alpha z^\rho. \quad (4.14)$$

Proof. If $z > 0$, the sequence $x_n = M_n z_n$ diverges to infinity, so our claim follows from Theorem 1.5.2 in Bingham et al. (1989) by writing

$$\frac{f(M_n z_n)}{g(M_n)} = \frac{f(x_n)}{g(x_n)} \cdot \frac{g(M_n z_n)}{g(M_n)} \rightarrow \alpha z^\rho. \quad (4.15)$$

If $z = 0$, then, for all $\varepsilon > 0$, we have $z_n \leq \varepsilon$ if n is sufficiently large. Then, using the monotonicity of f and the previous argument, we get

$$0 \leq \limsup_{n \rightarrow \infty} \frac{f(M_n z_n)}{g(M_n)} \leq \limsup_{n \rightarrow \infty} \frac{f(M_n \varepsilon)}{g(M_n)} = \alpha \varepsilon^\rho. \quad (4.16)$$

Taking $\varepsilon \rightarrow 0$, we conclude that $f(M_n z_n)/g(M_n) \rightarrow 0 = \alpha z^\rho$, as claimed. \blacksquare

Having established this auxiliary lemma, the heavy lifting for the proof of Theorem 4.4 is provided by the following asymptotic approximation result:

Lemma 4.11. *Consider an ensemble of nondecreasing functions $g_e: [0, \infty) \rightarrow (0, \infty)$, $e \in \mathcal{E}$, and suppose that it admits a ρ -regularly varying benchmark g with $\rho > 0$. Consider also a sequence of inflow vectors $\mu_n = M_n \lambda_n$ such that:*

- (a) The total inflow M_n grows to infinity and the vector of relative inflows λ_n converges to some $\lambda \in \Delta(\mathcal{I})$.
- (b) Every O/D pair has a path which is not slow (relative to g), i.e. $\alpha^i < \infty$ for all $i \in \mathcal{I}$.
- (c) There exists a heavy O/D pair which is also tight (relative to g), i.e. $\lambda^i > 0$ and $0 < \alpha^i < \infty$ for some $i \in \mathcal{I}$.

Then, the optimal allocation problem

$$G_n = \min_{y \in \mathcal{Y}} \sum_{e \in \mathcal{E}} g_e(x_e(y, \mu_n)) \quad (4.17)$$

is such that

$$\lim_{n \rightarrow \infty} \frac{G_n}{g(M_n)} = V_\rho(\lambda), \quad (4.18)$$

where $V_\rho(\lambda) \in (0, \infty)$ is the value of the gauge problem⁷

$$V_\rho(\lambda) = \min_{y \in \mathcal{Y}} \sum_{e \in \mathcal{E}} \alpha_e \zeta_e(y, \lambda)^\rho \quad (4.19)$$

Finally, if \hat{y}_n is a sequence of solutions of G_n , every limit point of \hat{y}_n solves $V_\rho(\lambda)$.

Proof. The arguments in the proof are similar (though not exactly the same) as those based on epi-convergence (Attouch, 1984). To streamline the presentation, we break up the proof in five steps as follows:

Step 1: $V_\rho(\lambda) < \infty$. Since each O/D pair admits a path that is not slow, routing all traffic through said path gives a finite value for the objective of (4.19), implying in turn that $V_\rho(\lambda) < \infty$. More precisely, for every $i \in \mathcal{I}$, take a traffic allocation $y^i \in \mathcal{Y}^i$ that assigns zero weight to the slow paths of i . Then, for every e with $\alpha_e = \infty$ we have $z_e^i(y) = 0$ and, a fortiori, $\zeta_e(y, \lambda) = 0$; hence

$$V_\rho(\lambda) \leq \sum_{e: \alpha_e < \infty} \alpha_e \zeta_e(y, \lambda)^\rho < \infty. \quad (4.20)$$

Step 2: $V_\rho(\lambda) > 0$. Let $i \in \mathcal{I}$ be a heavy pair which is also tight, so $\lambda^i > 0$ and $\alpha^i = \min_{p \in \mathcal{P}^i} \alpha_p \in (0, \infty)$. For every $y \in \mathcal{Y}$ we have $\sum_{p \in \mathcal{P}^i} y_p = 1$, so there exists some route $p \in \mathcal{P}^i$ with $y_p^i \geq 1/|\mathcal{P}^i|$. This gives $z_e^i(y) \geq 1/|\mathcal{P}^i|$ for all $e \in p$, and hence:

$$\sum_{e \in \mathcal{E}} \alpha_e \zeta_e(y, \lambda)^\rho \geq \sum_{e \in \mathcal{E}} \alpha_e (\lambda^i z_e^i(y))^\rho \geq \sum_{e \in p} \alpha_e (\lambda^i / |\mathcal{P}^i|)^\rho \geq \alpha^i (\lambda^i / |\mathcal{P}^i|)^\rho > 0. \quad (4.21)$$

Minimizing over $y \in \mathcal{Y}$ then yields $V_\rho(\lambda) > 0$, as claimed.

Step 3: $\limsup_{n \rightarrow \infty} G_n/g(M_n) \leq V_\rho(\lambda)$. Fix an optimal solution $\hat{y} \in \mathcal{Y}$ of (4.19). By the finiteness of $V_\rho(\lambda)$, we have $\zeta_e(\hat{y}, \lambda) = 0$ for every edge $e \in \mathcal{E}$ with $\alpha_e = \infty$. If $\lambda^i > 0$, this implies that $z_e^i(\hat{y}) = 0$. Otherwise, if $\lambda^i = 0$, the objective function of (4.19) does not depend on y^i , so every y^i with $z_e^i(y) = 0$ is also optimal. This means that we can choose the solution \hat{y} of (4.19) so that all traffic is routed along edges that are not slow.

⁷By convention, we take $\alpha_e z_e^\rho = 0$ if $\alpha_e = \infty$ and $z_e = 0$.

Now, from optimality we have

$$\frac{G_n}{g(M_n)} \leq \frac{1}{g(M_n)} \sum_{e \in \mathcal{E}} g_e(M_n \zeta_e(\hat{y}, \lambda_n)). \quad (4.22)$$

Using Lemma 4.10, for every edge $e \in \mathcal{E}$ with $\alpha_e < \infty$ we get

$$\frac{g_e(M_n \zeta_e(\hat{y}, \lambda_n))}{g(M_n)} \rightarrow \alpha_e \zeta_e(\hat{y}, \lambda)^\rho. \quad (4.23)$$

On the other hand, if e is slow (i.e. $\alpha_e = \infty$), we have $\zeta_e(\hat{y}, \lambda_n) = 0$ and $g(M_n) \rightarrow \infty$ (Bingham et al., 1989, Proposition 1.5.1), so

$$\lim_{n \rightarrow \infty} \frac{g_e(M_n \zeta_e(\hat{y}, \lambda_n))}{g(M_n)} = \lim_{n \rightarrow \infty} \frac{g_e(0)}{g(M_n)} = 0 = \alpha_e \zeta_e(\hat{y}, \lambda)^\rho. \quad (4.24)$$

Combining the previous three displayed equations, we obtain

$$\limsup_{n \rightarrow \infty} \frac{G_n}{g(M_n)} \leq \sum_{e \in \mathcal{E}} \lim_{n \rightarrow \infty} \frac{g_e(M_n \zeta_e(\hat{y}, \lambda_n))}{g(M_n)} = \sum_{e \in \mathcal{E}} \alpha_e \zeta_e(\hat{y}, \lambda)^\rho = V_\rho(\lambda). \quad (4.25)$$

Step 4: $\liminf_{n \rightarrow \infty} G_n/g(M_n) \geq V_\rho(\lambda)$. Passing to a subsequence if necessary, we may assume that $\liminf_{n \rightarrow \infty} G_n/g(M_n)$ is attained as a limit. Let \hat{y}_n be a sequence of solutions of G_n . Taking a further subsequence we may assume that \hat{y}_n converges to some $\hat{y} \in \mathcal{Y}$. Then, ignoring the network's slow edges, we have

$$\frac{G_n}{g(M_n)} \geq \sum_{e: \alpha_e < \infty} \frac{g_e(M_n \zeta_e(\hat{y}_n, \lambda_n))}{g(M_n)}, \quad (4.26)$$

and hence, by Lemma 4.10, we obtain

$$\liminf_{n \rightarrow \infty} \frac{G_n}{g(M_n)} \geq \sum_{e: \alpha_e < \infty} \alpha_e \zeta_e(\hat{y}, \lambda)^\rho. \quad (4.27)$$

To proceed, we next show that $\zeta_e(\hat{y}, \lambda) = 0$ for every slow edge. Indeed, if this were not the case, we could find some $\varepsilon > 0$ such that $\zeta_e(\hat{y}_n, \lambda_n) > \varepsilon$ for all sufficiently large n . With g_e nondecreasing, we then get

$$\frac{G_n}{g(M_n)} \geq \frac{g_e(M_n \zeta_e(\hat{y}_n, \lambda_n))}{g(M_n)} \geq \frac{g_e(M_n \varepsilon)}{g(M_n)} = \frac{g_e(M_n \varepsilon)}{g(M_n \varepsilon)} \frac{g(M_n \varepsilon)}{g(M_n)} \rightarrow \alpha_e \varepsilon^\rho = \infty, \quad (4.28)$$

in contradiction to Steps 1 and 3 above. From all this, it follows that

$$\liminf_{n \rightarrow \infty} \frac{G_n}{g(M_n)} \geq \sum_{e \in \mathcal{E}} \alpha_e \zeta_e(\hat{y}, \lambda)^\rho \geq V_\rho(\lambda). \quad (4.29)$$

Step 5: Optimality of limit points. As above, let \hat{y}_n be a sequence of optimal solutions of (4.17) and, by descending to a subsequence if necessary, assume that it converges to some $\hat{y} \in \mathcal{Y}$. From the previous steps we have $G_n/g(M_n) \rightarrow V_\rho(\lambda)$ so, by proceeding as in Step 4, we get

$$V_\rho(\lambda) = \lim_{n \rightarrow \infty} \frac{G_n}{g(M_n)} \geq \sum_{e \in \mathcal{E}} \alpha_e \zeta_e(\hat{y}, \lambda)^\rho \geq V_\rho(\lambda), \quad (4.30)$$

showing that \hat{y} solves (4.19). ■

With all this at hand, we are finally in a position to prove our main result:

Proof of Theorem 4.4. Let $L_n \equiv L_{\mu_n}$ denote the objective (4.12), and let y_n^* , \tilde{y}_n respectively denote the traffic allocation vectors of a Wardrop equilibrium and a socially optimum flow respectively. Then, the network's price of anarchy may be expressed as

$$\text{PoA}(\Gamma_n) = \frac{\text{Eq}(\Gamma_n)}{\text{Opt}(\Gamma_n)} = \frac{L_n(y_n^*)}{L_n(\tilde{y}_n)}, \quad (4.31)$$

so it suffices to show that $\limsup_{n \rightarrow \infty} \text{PoA}(\Gamma_n) \leq 1$.

Descending to a subsequence if necessary, we may assume without loss of generality that (i) $\lim_{n \rightarrow \infty} \text{PoA}(\Gamma_n)$ exists; (ii) the sequence λ_n of relative inflows converges to some $\lambda \in \Delta(\mathcal{I})$; and (iii) the sequences y_n^* and \tilde{y}_n respectively converge to some y^* , $\tilde{y} \in \mathcal{Y}$.

With this in mind, we will derive the asymptotic behavior of $\text{Opt}(\Gamma_n)$ by means of Lemma 4.11. First, by Lemma 4.9, Proposition 1.5.1 of Bingham et al. (1989), and the fact that the network's cost functions are nondecreasing, it follows that the network's gauge function c is β -regularly varying for some $\beta \geq 0$. Then, letting $g_e(x) = xc_e(x)$ and $g(x) = xc(x)$, we also get that g is ρ -regularly varying with $\rho = 1 + \beta > 0$ and $\lim_{x \rightarrow \infty} g_e(x)/g(x) = \lim_{x \rightarrow \infty} c_e(x)/c(x) = \alpha_e$. This means that the requirements of Lemma 4.11 are all satisfied, so we obtain:⁸

$$L_n(\tilde{y}_n) = \text{Opt}(\Gamma_n) \sim V_\rho(\lambda) g(M_n) \quad \text{as } n \rightarrow \infty. \quad (4.32)$$

In view of this, and since $0 < V_\rho(\lambda) < \infty$, it suffices to show that $L_n(y_n^*) \sim V_\rho(\lambda) g(M_n)$. To this end, we first analyze the asymptotic behavior of

$$V(\Gamma_n) = \min_{y \in \mathcal{Y}} \sum_{e \in \mathcal{E}} C_e(x_e(y, \mu_n)) \quad (4.33)$$

by applying Lemma 4.11 to the primitives C_e and C of c_e and c respectively. By a standard result (Bingham et al., 1989, Theorem 1.5.11), C is also ρ -regularly varying with $\rho = 1 + \beta$; moreover, by L'Hôpital's rule we also have $\lim_{x \rightarrow \infty} C_e(x)/C(x) = \lim_{x \rightarrow \infty} c_e(x)/c(x) = \alpha_e$. By Lemma 4.11, it follows that $V(\Gamma_n)/C(M_n) \rightarrow V_\rho(\lambda)$. In addition, since the Wardrop equilibrium traffic allocations y_n^* are solutions of $V(\Gamma_n)$, the limit y^* of y_n^* is optimal for $V_\rho(\lambda)$ by Lemma 4.11.

Noting that $x_e(y_n^*, \mu_n) = M_n \zeta_e(y_n^*, \lambda_n)$ we obtain

$$\frac{L_n(y_n^*)}{g(M_n)} = \sum_{e \in \mathcal{E}} \frac{g_e(M_n \zeta_e(y_n^*, \lambda_n))}{g(M_n)}. \quad (4.34)$$

Lemma 4.10 yields the following limit for every edge $e \in \mathcal{E}$ with $\alpha_e < \infty$

$$\frac{g_e(M_n \zeta_e(y_n^*, \lambda_n))}{g(M_n)} \rightarrow \alpha_e \zeta_e(y^*, \lambda)^\rho. \quad (4.35)$$

To establish a similar limiting result when $e \in \mathcal{E}$ is slow ($\alpha_e = \infty$), we first claim that there exists a constant $B \geq 0$ such that

$$g_e(M_n \zeta_e(y_n^*, \lambda_n)) \leq B \zeta_e(y_n^*, \lambda_n) g(M_n) \quad (4.36)$$

or, equivalently:

$$c_e(M_n \zeta_e(y_n^*, \lambda_n)) \leq B c(M_n) \quad (4.37)$$

This is trivial when $\zeta_e(y_n^*, \lambda_n) = 0$, so it suffices to consider the case $\zeta_e(y_n^*, \lambda_n) > 0$. In this case, e receives some equilibrium traffic from at least one O/D pair $i \in \mathcal{I}$,

⁸The asymptotic equality notation " $f_n \sim g_n$ " means here that $\lim_{n \rightarrow \infty} f_n/g_n = 1$.

so it must belong to a path $p \in \mathcal{P}^i$ with minimal cost. Then, if we consider an alternative path $p' \in \mathcal{P}^i$ all of whose edges are tight or fast, we have

$$c_e(M_n \zeta_e(y_n^*, \lambda_n)) \leq \sum_{e' \in p} c_{e'}(M_n \zeta_{e'}(y_n^*, \lambda_n)) \leq \sum_{e' \in p'} c_{e'}(M_n \zeta_{e'}(y_n^*, \lambda_n)). \quad (4.38)$$

Now, using the trivial bound $M_n \zeta_{e'}(y_n^*, \lambda_n) \leq M_n$, we further get

$$c_e(M_n \zeta_e(y_n^*, \lambda_n)) \leq \sum_{e' \in p'} c_{e'}(M_n) \leq \sum_{e': \alpha_{e'} < \infty} c_{e'}(M_n). \quad (4.39)$$

However, for all $e' \in \mathcal{E}$ with $\alpha_{e'} < \infty$, we can find a constant $B_{e'}$ such that $c_{e'}(x)/c(x) \leq B_{e'}$; consequently, (4.37) follows by taking $B = \sum_{e': \alpha_{e'} < \infty} B_{e'}$. Thus, given that y^* is optimal for $V_\rho(\lambda)$, we get $\zeta_e(y_n^*, \lambda_n) \rightarrow \zeta_e(y^*, \lambda) = 0$, and hence

$$\frac{g_e(M_n \zeta_e(y_n^*, \lambda_n))}{g(M_n)} \leq B \zeta_e(y_n^*, \lambda_n) \rightarrow 0 = \alpha_e \zeta_e(y^*, \lambda)^\rho. \quad (4.40)$$

Combining (4.35) and (4.40) into (4.34) we then obtain

$$\lim_{n \rightarrow \infty} \frac{L_n(y_n^*)}{g(M_n)} = \sum_{e \in \mathcal{E}} \alpha_e \zeta_e(y^*, \lambda)^\rho = V_\rho(\lambda), \quad (4.41)$$

and our proof is complete. \blacksquare

5. CONVERGENCE RATE ANALYSIS

Theorem 4.4 guarantees that $\text{PoA}(\Gamma_n) \rightarrow 1$, but it does not provide any indication for the rate of convergence thereof. This rate can be computed to leading order in the special case where all edges have polynomial costs of the form

$$c_e(x) = \sum_{k=0}^{d_e} c_{e,k} x^k \quad (5.1)$$

with $c_{e,k} \geq 0$ and a strictly positive leading coefficient $c_{e,d_e} > 0$.⁹ To state our result, we define the *degree* of an edge $e \in \mathcal{E}$ to be d_e , that of a path $p \in \mathcal{P}$ as $d_p = \max_{e \in p} d_e$, of an O/D pair $i \in \mathcal{I}$ as $d^i = \min_{p \in \mathcal{P}^i} d_p$, and of the network as $d = \max_{i \in \mathcal{I}} d^i$.¹⁰ Moreover, we write \mathcal{E}_d for the set of edges with $d_e \leq d$, and we let

$$g = \min_{e \notin \mathcal{E}_d} d_e - d. \quad (5.2)$$

We call g the *gap* between d and the lowest edge degree greater than d ; we also write $a = g/(d+g)$.¹¹ Our main result for the rate of convergence of the price of anarchy may then be stated as follows:

Theorem 5.1. *With notation as above, let Γ_n be a highly congested network, and assume there exists a heavy O/D pair $i \in \mathcal{I}$ with $d^i = d$. Then, there exist nonnegative constants $K_1, K_a \geq 0$ such that*

$$\text{PoA}(\Gamma_n) \leq 1 + \frac{K_1}{M_n} + \frac{K_a}{M_n^a}. \quad (5.3)$$

If, in addition, there are no edges with degree greater than d , we have $K_a = 0$.

⁹This covers in particular the case where all edges have BPR cost functions of the form (4.7).

¹⁰Note here the similarity between the above definitions and those for the index of an edge, path, and O/D pair in Section 4.

¹¹We employ here the standard convention that $\min \emptyset = \infty$, so $g = \infty$ and $a = 1$ when $\mathcal{E} = \mathcal{E}_d$; obviously, $g \geq 1$ always.

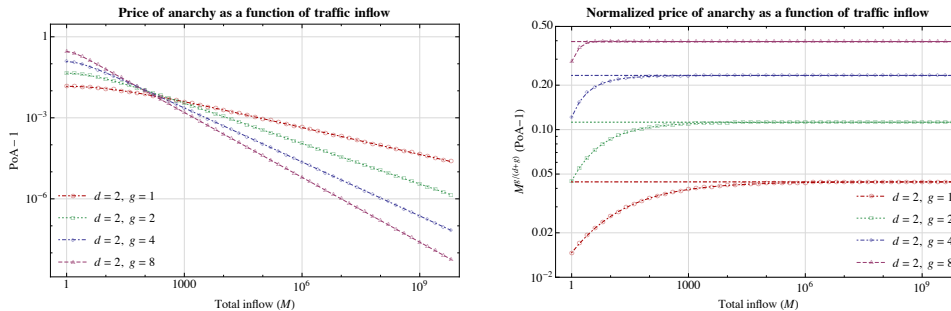


Figure 3. The rate of convergence of the price of anarchy in a two-edge Pigou network with cost functions x^d and x^{d+g} , $g \geq 1$. To the left, we provide a log-log plot of the price of anarchy as a function of the inflow. To the right, the price of anarchy is rescaled by M^a , showing that $\text{PoA}(\Gamma_M) \sim 1 + b/M^a$ with b given by (5.5); the horizontal lines correspond precisely to this value.

Remark 5.1. If the costs are monomials of the same degree, then $K_1 = K_a = 0$ and $\text{PoA}(\Gamma_n) \equiv 1$ for all n .

According to this result (which we prove in [Appendix A](#)), $\text{PoA}(\Gamma_n)$ converges to 1 as $\mathcal{O}(1/M_n^a)$. In particular, if $\mathcal{E}_d = \mathcal{E}$, the rate of convergence is $\mathcal{O}(1/M_n)$; otherwise, in the presence of a single edge e with $d_e > d$, the rate of convergence drops abruptly to $\mathcal{O}(1/M_n^a)$. In this case, the exponent a may be as small as $1/(d+1)$, but it gets closer to 1 as the gap g increases.

This reduction in the speed of convergence is not an artifact of the proof but an essential feature of the problem. This can be seen by considering a two-edge Pigou network which exhibits a $\Theta(M_n^{-a})$ rate of convergence (see [Fig. 3](#)); as an added bonus, in this case we can also compute explicitly the proportionality constants in the high congestion limit. For simplicity, since there is only one O/D pair in this network, we state our result directly in terms of the total traffic inflow M (as opposed to a sequence of inflow vectors):

Proposition 5.2. *Consider a routing game Γ_M with a single O/D pair on a two-link parallel network with cost functions $c_1(x) = x^d$ and $c_2(x) = x^{d+g}$ for some $g \geq 1$. Then:*

$$\lim_{M \rightarrow \infty} (\text{PoA}(\Gamma_M) - 1) M^{g/(d+g)} = b, \quad (5.4)$$

where

$$b = (d+g) \left(\frac{d+1}{d+g+1} \right)^{1+1/(d+g)} - d > 0. \quad (5.5)$$

We prove [Proposition 5.2](#) in [Appendix A](#); [Fig. 3](#) illustrates numerically the convergence rate predicted by [Eqs. \(5.4\) and \(5.5\)](#) for different values of d and g .

6. DISCUSSION

Most of the literature on the price of anarchy – for congestion games and not only – has traditionally focused on establishing worst-case upper bounds for different classes of networks, cost functions, types of players (atomic vs. nonatomic),

etc. Several of these results have become milestones in the field and have had a significant impact in real-world considerations for traffic networks; however, many real-world situations involve a fixed network and heavy traffic flows that are not necessarily close to these worst-case scenarios. Thus, in addition to determining how bad can selfish routing be in general, it is also important to determine how bad it is under heavy traffic.

Our goal in this paper was to provide an answer to this question by determining the asymptotic behavior of the price of anarchy in networks where the total traffic inflow grows large. Under fairly light assumptions (namely that the network’s cost functions can be compared asymptotically to a regularly varying gauge function (for instance, a monomial of the form $c(x) = x^\rho$), we found that the price of anarchy goes to 1 in the high congestion limit, independently of the network’s topology, and even when there are multiple O/D pairs. Moreover, focusing on the baseline case of polynomial costs, a sharper asymptotic analysis reveals that this convergence follows a power law whose degree can be computed explicitly as a function of the degrees of the network’s cost functions.

What we find appealing about this result is that it is essentially independent of the underlying graph and/or the way with which the traffic inflow of different O/D pairs might grow to infinity. From the social planner’s point of view, this means that selfish routing is not the real cause of increased delays in highly congested networks. Instead, our findings suggest that if a network consistently exhibits very heavy traffic and increased delays, sophisticated tolling/rerouting schemes that target the optimum traffic assignment will not yield considerable gains over a “laissez-faire” approach where each traffic element takes the shortest path available.

That said, a word of caution is due if this high congestion regime cannot be reached because the network can only serve a finite amount of traffic – for instance, as in data networks with latencies modeled by M/M/1 functions (Bertsekas and Gallager, 1992). Intuitively, one might expect a certain rescaling to take place, mapping the high congestion limit to the so-called *capacity limit* (i.e. when the total inflow approaches the network’s capacity). However, this is *not* so, even in simple two-link networks with a single O/D pair and M/M/1 latencies of the form $c_e(x) = (m_e - x)^{-1}$, $e = 1, 2$, with m_e denoting the capacity of each link.

As can be seen by a straightforward calculation (which we carry out in Appendix B), the capacity limit of the price of anarchy in this example is

$$\lim_{M \rightarrow m_1 + m_2} \text{PoA}(\Gamma_M) = 1 + \left(\frac{1 - \sqrt{m_2/m_1}}{1 + \sqrt{m_2/m_1}} \right)^2. \quad (6.1)$$

Therefore, if there is a strong disparity between the capacity of each link, the price of anarchy could become as high as 2 when the traffic inflow approaches capacity. In this case, myopic routing decisions can become problematic, so the intervention of a social planner is warranted to reinstate efficiency. We intend to explore this direction in more detail in future work.

APPENDIX A. SPEED OF CONVERGENCE

In this appendix, we provide the proofs of the results presented in Section 5

Proof of Theorem 5.1. Let $y_n^*, \tilde{y}_n \in \mathcal{Y}$ be an equilibrium and an optimum flow of Γ_n respectively. The social cost of y_n^* can then be estimated as

$$\begin{aligned}
L_n(y_n^*) &= \sum_{e \in \mathcal{E}} \sum_{k=0}^{d_e} c_{e,k} \cdot (M_n \zeta_e(y_n^*, \lambda_n))^{k+1} \\
&= \sum_{e \in \mathcal{E}} \sum_{k=0}^{d_e} \left[\frac{d+1}{k+1} + \frac{k-d}{k+1} \right] c_{e,k} \cdot (M_n \zeta_e(y_n^*, \lambda_n))^{k+1} \\
&= (d+1) \sum_{e \in \mathcal{E}} C_e(M_n \zeta_e(y_n^*, \lambda_n)) + \sum_{e \in \mathcal{E}} \sum_{k=0}^{d_e} \frac{k-d}{k+1} c_{e,k} \cdot (M_n \zeta_e(y_n^*, \lambda_n))^{k+1} \\
&\leq (d+1) \sum_{e \in \mathcal{E}} C_e(M_n \zeta_e(\tilde{y}_n, \lambda_n)) + \sum_{e \notin \mathcal{E}_d} (M_n \zeta_e(y_n^*, \lambda_n)) \cdot c_e(M_n \zeta_e(y_n^*, \lambda_n)),
\end{aligned} \tag{A.1}$$

where we used the fact that $\zeta_e(y_n^*, \lambda_n)$ minimizes the first sum, while in the double sum we dropped the edges $e \in \mathcal{E}_d$ since $(k-d)/(k+1) \leq 0$ for all $k \leq d_e$, and we used the inequality $(k-d)/(k+1) \leq 1$ to bound the remaining terms $e \notin \mathcal{E}_d$ by factoring out $M_n \zeta_e(y_n^*, \lambda_n)$ and using the definition (5.1) of c_e .

Call L_n^I the first sum above and L_n^{II} the second. We assume in what follows that n is large enough so that $M_n \geq 1$. This assumption is done for convenience and it only affects the value of the constants K_1 and K_a : by redefining them appropriately, the bound (5.3) can be assumed to hold for all n . Now, denoting

$$G = \sum_{e \in \mathcal{E}} \sum_{k=0}^{d-1} \frac{d-k}{k+1} c_{e,k} \tag{A.2}$$

and using the fact that $\tilde{z}_{e,n} \leq 1$, we can bound L_n^I as

$$\begin{aligned}
L_n^I &= \sum_{e \in \mathcal{E}} \sum_{k=0}^{d_e} \frac{d+1}{k+1} c_{e,k} \cdot (M_n \zeta_e(\tilde{y}_n, \lambda_n))^{k+1} \\
&= \sum_{e \in \mathcal{E}} \sum_{k=0}^{d_e} c_{e,k} \cdot (M_n \zeta_e(\tilde{y}_n, \lambda_n))^{k+1} + \sum_{e \in \mathcal{E}} \sum_{k=0}^{d_e} \frac{d-k}{k+1} c_{e,k} \cdot (M_n \zeta_e(\tilde{y}_n, \lambda_n))^{k+1} \\
&= \text{Opt}(\Gamma_n) + \sum_{e \in \mathcal{E}} \sum_{k=0}^{d_e} \frac{d-k}{k+1} c_{e,k} \cdot (M_n \zeta_e(\tilde{y}_n, \lambda_n))^{k+1} \\
&\leq \text{Opt}(\Gamma_n) + \sum_{e \in \mathcal{E}} \sum_{k=0}^{d-1} \frac{d-k}{k+1} c_{e,k} \cdot M_n^{k+1} \\
&\leq \text{Opt}(\Gamma_n) + GM_n^d.
\end{aligned} \tag{A.3}$$

In order to bound L_n^{II} consider any edge $e \notin \mathcal{E}_d$ that contributes to the sum with $\zeta_e(y_n^*, \lambda_n) > 0$. Since y_n^* is an equilibrium, this edge e belongs to a path $p \in \mathcal{P}^i$ with minimal cost for some $i \in \mathcal{I}$. Hence, taking an alternative path $p' \in \mathcal{P}^i$ with all its edges in \mathcal{E}_d and denoting

$$B = \sum_{e' \in \mathcal{E}_d} \sum_{k=0}^{d_{e'}} c_{e',k}, \tag{A.4}$$

we get the bound

$$\begin{aligned} c_e(M_n \zeta_e(y_n^*, \lambda_n)) &\leq \sum_{e' \in p} c_{e'}(M_n \zeta_{e'}(y_n^*, \lambda_n)) \\ &\leq \sum_{e' \in p'} c_{e'}(M_n \zeta_{e'}(y_n^*, \lambda_n)) \leq \sum_{e' \in p'} c_{e'}(M_n) \leq BM_n^d. \end{aligned} \quad (\text{A.5})$$

In particular, we have

$$c_{e,d_e} \cdot (M_n \zeta_e(y_n^*, \lambda_n))^{d_e} \leq c_e(M_n \zeta_e(y_n^*, \lambda_n)) \leq BM_n^d, \quad (\text{A.6})$$

so, letting $c_0 = \min_{e \notin \mathcal{E}_d} c_{e,d_e}$, we get $M_n \zeta_e(y_n^*, \lambda_n) \leq (BM_n^d/c_0)^{1/d_e}$, and hence

$$L_n^{\text{II}} \leq \sum_{e \notin \mathcal{E}_d} (BM_n^d/c_0)^{1/d_e} BM_n^d \leq DM_n^{d+d/(d+g)} \quad (\text{A.7})$$

where we set¹² $D = B \sum_{e \notin \mathcal{E}_d} (B/c_0)^{1/d_e}$. Hence, combining (A.3) and (A.7) we get

$$\begin{aligned} \text{PoA}(\Gamma_n) &= \frac{L_n(y_n^*)}{\text{Opt}(\Gamma_n)} \leq \frac{\text{Opt}(\Gamma_n) + GM_n^d + DM_n^{d+d/(d+g)}}{\text{Opt}(\Gamma_n)} \\ &= 1 + \frac{GM_n^d + DM_n^{d+d/(d+g)}}{\text{Opt}(\Gamma_n)}. \end{aligned} \quad (\text{A.8})$$

Now, if we set $H = \min_{e \in \mathcal{E}} c_{e,d_e}$, we have the following lower bound for $\text{Opt}(\Gamma_n)$.

$$\begin{aligned} \text{Opt}(\Gamma_n) &= \min_{y \in \mathcal{Y}} \sum_{e \in \mathcal{E}} M_n \zeta_e(y, \lambda_n) \cdot c_e(M_n \zeta_e(y, \lambda_n)) \\ &\geq H \min_{y \in \mathcal{Y}} \sum_{e \in \mathcal{E}} (M_n \zeta_e(y, \lambda_n))^{d_e+1}. \end{aligned} \quad (\text{A.9})$$

Take a heavy O/D pair $j \in \mathcal{I}$ with $d^j = d$, and select $\varepsilon > 0$ so that $\lambda_n^j \geq \varepsilon$ for $n \in \mathbb{N}$ large. For each feasible $y \in \mathcal{Y}$ there exists a path $p \in \mathcal{P}^j$ with $y_p \geq 1/|\mathcal{P}^j|$ so $\zeta_e(y, \lambda_n) \geq \varepsilon/|\mathcal{P}^j|$ for all $e \in p$. For n large we may assume that $M_n \varepsilon/|\mathcal{P}^j| \geq 1$ and since the path p contains at least one edge $e \in p$ with $d_e \geq d$, setting $\bar{H} = H(\varepsilon/|\mathcal{P}^j|)^{d+1}$ we get

$$\text{Opt}(\Gamma_n) \geq H (M_n \varepsilon/|\mathcal{P}^j|)^{d_e+1} \geq \bar{H} M_n^{d+1}, \quad (\text{A.10})$$

which, combined with (A.8), yields (5.3) with $K_1 = G/\bar{H}$ and $K_a = D/\bar{H}$. We conclude by observing that when $\mathcal{E}_d = \mathcal{E}$ we can take $D = 0$ and hence $K_a = 0$. ■

Proof of Proposition 5.2. Let x denote the flow on edge e_2 . At equilibrium, the costs on both edges must be equal so that $x^{d+g} = (M-x)^d$, which is equivalent to $x + x^\kappa = M$ with $\kappa = (d+g)/d$. Since $\kappa > 1$ it follows that

$$x = M^{1/\kappa}(1 + o(1)) \quad (\text{A.11})$$

so the equilibrium cost $\text{Eq}(\Gamma_M) = M \cdot c_2(x) = M \cdot c_1(M-x)$ scales as

$$\text{Eq}(\Gamma_M) = M \cdot \left[M - M^{1/\kappa}(1 + o(1)) \right]^d = M^{d+1} - dM^{d+1/\kappa} + o(M^{d+1/\kappa}). \quad (\text{A.12})$$

Similarly, if \tilde{x} is the optimal flow on edge e_2 , both edges have the same marginal cost

$$(d+g+1)\tilde{x}^{d+g} = (d+1)(M-\tilde{x})^d. \quad (\text{A.13})$$

¹²Note that the above sum is zero if $\mathcal{E} = \mathcal{E}_d$, so the bound also holds trivially in that case.

Therefore, if we let

$$\theta = \left(\frac{d+1}{d+g+1} \right)^{1/d}, \quad (\text{A.14})$$

we get $\theta\tilde{x} + \tilde{x}^\kappa = \theta M$ as before, and hence

$$\tilde{x} = (\theta M)^{1/\kappa}(1 + o(1)). \quad (\text{A.15})$$

Observing that $(d+g+1)/\kappa = d+1/\kappa$ it follows that the optimal cost scales as

$$\begin{aligned} \text{Opt}(\Gamma_M) &= (M - \tilde{x}) \cdot c_1(M - \tilde{x}) + \tilde{x} \cdot c_2(\tilde{x}) \\ &= \left[M - (\theta M)^{1/\kappa}(1 + o(1)) \right]^{d+1} + \left[(\theta M)^{1/\kappa}(1 + o(1)) \right]^{d+g+1} \\ &= M^{d+1} - (d+1)M^d(\theta M)^{1/\kappa} + (\theta M)^{d+1/\kappa} + o(M^{d+1/\kappa}) \\ &= M^{d+1} - \theta^{1/\kappa} [(d+1) - \theta^d] M^{d+1/\kappa} + o(M^{d+1/\kappa}) \\ &= M^{d+1} - (b+d)M^{d+1/\kappa} + o(M^{d+1/\kappa}) \end{aligned} \quad (\text{A.16})$$

where the last equality follows from the identity $\theta^{1/\kappa} [(d+1) - \theta^d] = b+d$.

Combining the previous expressions we get

$$\begin{aligned} \text{PoA}(\Gamma_M) &= \frac{\text{Eq}(\Gamma_M)}{\text{Opt}(\Gamma_M)} = \frac{\text{Opt}(\Gamma_M) + bM^{d+1/\kappa} + o(M^{d+1/\kappa})}{\text{Opt}(\Gamma_M)} \\ &= 1 + \frac{bM^{d+1/\kappa} + o(M^{d+1/\kappa})}{M^{d+1} + o(M^{d+1})} \\ &= 1 + bM^{-g/(d+g)} + o(M^{-g/(d+g)}). \end{aligned} \quad (\text{A.17})$$

It remains to show that $b > 0$, which is equivalent to

$$\left(\frac{d+1}{d+g+1} \right)^{d+g+1} > \left(\frac{d}{d+g} \right)^{d+g} \quad (\text{A.18})$$

and which itself follows from the fact that $(1 - g/x)^x$ is increasing in x . \blacksquare

APPENDIX B. THE CAPACITY LIMIT

In this appendix, we discuss an example of a capacitated network where the price of anarchy does not converge to 1 as the total inflow approaches capacity. In particular, we consider a routing game Γ_M over a parallel network with a single O/D pair and two edges with M/M/1 cost functions of the form

$$c_e(x) = \frac{1}{m_e - x} \quad 0 \leq x < m_e, \quad e = 1, 2. \quad (\text{B.1})$$

To ensure that the inflow does not exceed the network's capacity, we assume throughout that $M < m_1 + m_2$; moreover, for convenience (and without loss of generality), we assume that $m_1 \leq m_2$.

Now, with notation as in the proof of Proposition 5.2, the game's Wardrop equilibrium is

$$x^* = \begin{cases} 0 & \text{if } M \leq m_2 - m_1, \\ \frac{1}{2}(M + m_1 - m_2) & \text{if } m_2 - m_1 < M < m_1 + m_2. \end{cases} \quad (\text{B.2})$$




Figure 4. M/M/1 costs (left) and the network's price of anarchy (right).

Hence, the social cost at equilibrium is

$$\text{Eq}(\Gamma_M) = \begin{cases} \frac{M}{m_2 - M} & \text{if } M \leq m_2 - m_1, \\ \frac{2M}{m_1 + m_2 - M} & \text{if } m_2 - m_1 < M < m_1 + m_2. \end{cases} \quad (\text{B.3})$$

On the other hand, the network's socially optimum flow is given by

$$\tilde{x} = \begin{cases} 0 & \text{if } M \leq m_2 - \sqrt{m_1 m_2}, \\ \frac{M - (m_2 - \sqrt{m_1 m_2})}{1 + \sqrt{m_2/m_1}} & \text{if } m_2 - \sqrt{m_1 m_2} < M < m_1 + m_2. \end{cases} \quad (\text{B.4})$$

Hence, the minimum total cost in the network will be

$$\text{Opt}(\Gamma_M) = \begin{cases} \frac{M}{m_2 - M} & \text{if } M \leq m_2 - \sqrt{m_1 m_2}, \\ \frac{2M - (\sqrt{m_1} - \sqrt{m_2})^2}{m_1 + m_2 - M} & \text{if } m_2 - \sqrt{m_1 m_2} < M < m_1 + m_2. \end{cases} \quad (\text{B.5})$$

Therefore, we get the following expression for the price of anarchy

$$\text{PoA}(\Gamma_M) = \begin{cases} 1 & \text{if } M \leq m_2 - \sqrt{m_1 m_2}, \\ \frac{(m_1 + m_2 - M)M}{(m_2 - M)(2M - (\sqrt{m_1} - \sqrt{m_2})^2)} & \text{if } m_2 - \sqrt{m_1 m_2} < M \leq m_2 - m_1, \\ \frac{2M}{2M - (\sqrt{m_1} - \sqrt{m_2})^2} & \text{if } m_2 - m_1 < M < m_1 + m_2. \end{cases} \quad (\text{B.6})$$

Thus, taking the limit $M \rightarrow m_1 + m_2$, we obtain

$$\lim_{M \rightarrow \infty} \text{PoA}(\Gamma_M) = 1 + \left(\frac{1 - \sqrt{m_2/m_1}}{1 + \sqrt{m_2/m_1}} \right)^2. \quad (\text{B.7})$$

From the above, it follows that, if $m_1 = m_2$, the price of anarchy is identically equal to 1 for all values of the traffic inflow M . Otherwise, not only is it strictly greater than 1, it does not even converge to 1 as the inflow approaches the network's capacity (i.e. $M \rightarrow m_1 + m_2$). For illustration purposes, Fig. 4 shows the behavior of the price of anarchy as a function of the inflow for $m_1 = 1$, $m_2 = 2$.

REFERENCES

- Alcalá, L. A. (2014). On the relation between regular variation and the asymptotic elasticity of substitution. *Economics Letters*, 125(1):29–31.
- Attouch, H. (1984). *Variational Convergence for Functions and Operators*. Pitman, Boston, MA.

- Beckmann, M. J., McGuire, C., and Winsten, C. B. (1956). *Studies in the Economics of Transportation*. Yale University Press, New Haven, CT.
- Bertsekas, D. P. and Gallager, R. (1992). *Data Networks*. Prentice Hall, Englewood Cliffs, NJ, 2 edition.
- Bingham, N. H., Goldie, C. M., and Teugels, J. L. (1989). *Regular Variation*, volume 27 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, Cambridge.
- Braess, D. (1968). Über ein Paradoxon aus der Verkehrsplanung. *Unternehmensforschung*, 12:258–268.
- Cole, R. and Tao, Y. (2016). Large market games with near optimal efficiency. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 791–808. ACM.
- Colini-Baldeschi, R., Cominetti, R., and Scarsini, M. (2016). On the price of anarchy of highly congested nonatomic network games. In *Algorithmic Game Theory: 9th International Symposium, SAGT 2016*, pages 117–128.
- Correa, J. R., Schulz, A. S., and Stier-Moses, N. E. (2004). Selfish routing in capacitated networks. *Mathematics of Operations Research*, 29(4):961–976.
- Correa, J. R., Schulz, A. S., and Stier-Moses, N. E. (2007). Fast, fair, and efficient flows in networks. *Operations Research*, 55(2):215–225.
- Correa, J. R., Schulz, A. S., and Stier-Moses, N. E. (2008). A geometric approach to the price of anarchy in nonatomic congestion games. *Games and Economic Behavior*, 64(2):457–469.
- de Haan, L. and Ferreira, A. (2006). *Extreme Value Theory: An Introduction*. Springer Series in Operations Research and Financial Engineering. Springer, New York.
- Dumrauf, D. and Gairing, M. (2006). Price of anarchy for polynomial Wardrop games. In *WINE '06: Proceedings of the 2nd Conference on Web and Internet Economics*, pages 319–330. Springer Berlin Heidelberg, Berlin, Heidelberg.
- Englert, M., Franke, T., and Olbrich, L. (2010). Sensitivity of Wardrop equilibria. *Theory of Computing Systems*, 47(1):3–14.
- Feldman, M., Immorlica, N., Lucier, B., Roughgarden, T., and Syrgkanis, V. (2016). The price of anarchy in large games. In *STOC '16: Proceedings of the 48th Annual ACM Symposium on the Theory of Computing*.
- Friedman, E. J. (2004). Genericity and congestion control in selfish routing. In *CDC '04: Proceedings of the 43rd IEEE Annual Conference on Decision and Control*, volume 5, pages 4667–4672.
- González Vayá, M., Grammatico, S., Andersson, G., and Lygeros, J. (2015). On the price of being selfish in large populations of plug-in electric vehicles. In *CDC '15: Proceedings of the 53rd IEEE Annual Conference on Decision and Control*, pages 6542–6547. IEEE.
- Jessen, A. H. and Mikosch, T. (2006). Regularly varying functions. *Publ. Inst. Math. (Beograd) (N.S.)*, 80(94):171–192.
- Josefsson, M. and Patriksson, M. (2007). Sensitivity analysis of separable traffic equilibria with application to bilevel optimization in network design. *Transportation Research Part B: Methodological*, 41(1):4–31.
- Karamata, J. (1930). Sur un mode de croissance régulière des fonctions. *Mathematica (Cluj)*, 4:38–53.
- Karamata, J. (1933). Sur un mode de croissance régulière. Théorèmes fondamentaux. *Bull. Soc. Math. France*, 61:55–62.
- Koutsoupias, E. and Papadimitriou, C. H. (1999). Worst-case equilibria. In *STACS '99: Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science*.
- Law, L. M., Huang, J., and Liu, M. (2012). Price of anarchy for congestion games in cognitive radio networks. *IEEE Trans. Wireless Commun.*, 11(10):3778–3787.
- LeBlanc, L. J., Morlok, E. K., and Pierskalla, W. P. (1975). An efficient approach to solving the road network equilibrium traffic assignment problem. *Transportation Research*, 9:309–318.
- O'Hare, S. J., Connors, R. D., and Watling, D. P. (2016). Mechanisms that govern how the price of anarchy varies with travel demand. *Transportation Research Part B: Methodological*, 84:55–80.

- Papadimitriou, C. H. (2001). Algorithms, games, and the Internet. In *STOC '01: Proceedings of the 33rd Annual ACM Symposium on the Theory of Computing*.
- Patriksson, M. (2004). Sensitivity analysis of traffic equilibria. *Transportation Science*, 38(3):258–281.
- Pigou, A. C. . (1920). *The Economics of Welfare*. Macmillan and Co., London, 1st edition.
- Piliouras, G., Nikolova, E., and Shamma, J. S. (2013). Risk sensitivity of price of anarchy under uncertainty. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce*, pages 715–732.
- Resnick, S. I. (2007). *Heavy-Tail Phenomena*. Springer Series in Operations Research and Financial Engineering. Springer, New York. Probabilistic and statistical modeling.
- Roughgarden, T. (2003). The price of anarchy is independent of the network topology. *J. Comput. System Sci.*, 67(2):341–364.
- Roughgarden, T. (2007). Routing games. In *Algorithmic Game Theory*, pages 461–486. Cambridge Univ. Press, Cambridge.
- Roughgarden, T. and Tardos, É. (2002). How bad is selfish routing? *Journal of the ACM*, 49(2):236–259.
- Roughgarden, T. and Tardos, É. (2004). Bounding the inefficiency of equilibria in nonatomic congestion games. *Games and Economic Behavior*, 47(2):389–403.
- Roughgarden, T. and Tardos, É. (2007). Introduction to the inefficiency of equilibria. In *Algorithmic Game Theory*, pages 443–459. Cambridge Univ. Press, Cambridge.
- Wardrop, J. G. (1952). Some theoretical aspects of road traffic research. In *Proceedings of the Institute of Civil Engineers, Part II*, volume 1, pages 325–378.
- Youn, H., Gastner, M. T., and Jeong, H. (2008). Price of anarchy in transportation networks: Efficiency and optimality control. *Physical Review Letters*, 101(12):128701.

DIPARTIMENTO DI ECONOMIA E FINANZA, LUISS, VIALE ROMANIA 32, 00197 ROMA, ITALY,
E-mail address: `rcolini@luiss.it`

FACULTAD DE INGENIERÍA Y CIENCIAS, UNIVERSIDAD ADOLFO IBÁÑEZ, SANTIAGO, CHILE
E-mail address: `roberto.cominetti@uai.cl`

CNRS, LIG, F-38000 GRENOBLE, FRANCE, AND UNIV. GRENOBLE ALPES, LIG, F-38000
 GRENOBLE, FRANCE
E-mail address: `panayotis.mertikopoulos@imag.fr`

DIPARTIMENTO DI ECONOMIA E FINANZA, LUISS, VIALE ROMANIA 32, 00197 ROMA, ITALY,
E-mail address: `marco.scarsini@luiss.it`