FULL LENGTH PAPER

Series B

The price of anarchy in routing games as a function of the demand

Roberto Cominetti[1] · **Valerio Dose**[2] · **Marco Scarsini**[3]

Abstract

The price of anarchy has become a standard measure of the efficiency of equilibria in games. Most of the literature in this area has focused on establishing worst-case bounds for specific classes of games, such as routing games or more general congestion games. Recently, the price of anarchy in routing games has been studied as a function of the traffic demand, providing asymptotic results in light and heavy traffic. The aim of this paper is to study the price of anarchy in nonatomic routing games in the intermediate region of the demand. To achieve this goal, we begin by establishing some smoothness properties of Wardrop equilibria and social optima for general smooth costs. In the case of affine costs we show that the equilibrium is piecewise linear, with break points at the demand levels at which the set of active paths changes. We prove that the number of such break points is finite, although it can be exponential in the size of the network. Exploiting a scaling law between the equilibrium and the social optimum, we derive a similar behavior for the optimal flows. We then prove that in any interval between break points the price of anarchy is smooth and it is either monotone (decreasing or increasing) over the full interval, or it decreases up to a certain minimum point in the

✉ Marco Scarsini
  marco.scarsini@luiss.it

  Roberto Cominetti
  roberto.cominetti@uai.cl

  Valerio Dose
  dose@diag.uniroma1.it

1 Facultad de Ingeniería y Ciencias, Universidad Adolfo Ibáñez, Santiago, Chile

2 Dipartimento di Ingegneria Informatica Automatica e Gestionale "Antonio Ruberti", "Sapienza" Università di Roma, Rome, Italy

3 Dipartimento di Economia e Finanza, Luiss University, Rome, Italy

∰ Springer

interior of the interval and increases afterwards. We deduce that for affine costs the maximum of the price of anarchy can only occur at the break points. For general costs we provide counterexamples showing that the set of break points is not always finite.

List of symbols

a_e	Coefficient of the monomial in an affine cost function
A	$Z^\top \Gamma Z$, defined in (4.7)
b_e	Constant in an affine cost function
$\boldsymbol{b}$	$(b_1, \ldots, b_m)^\top$
c_e	Cost of edge e
$\widetilde{c}_e(x_e)$	$\widetilde{C}'_e(x_e)$, defined in (3.4)
C_e	Primitive of c_e, defined in (2.5)
$\widetilde{C}_e(x_e)$	$x_e c_e(x_e)$
c_p	Cost of path p, defined in (2.2)
$\boldsymbol{d}$	$Z^\top \boldsymbol{b}$, defined in (4.7)
D	Destination of the network
D_μ	Dual problem
e	Edge
$\mathcal{E}$	Set of edges
$\widehat{\mathcal{E}}(\mu)$	Active network at demand μ, defined in Definition 3.4
$\widehat{\mathcal{E}}_0$	$\widehat{\mathcal{E}}(\mu_0)$
$\mathcal{E}(v, w)$	Set of all edges $e \in \mathcal{E}$ that go from v to w
$\widehat{\mathcal{E}}_0(v, w)$	Set of all edges $e \in \widehat{\mathcal{E}}_0$ that go from v to w
$\boldsymbol{f}$	Flow profile
$\boldsymbol{f}^*$	Equilibrium flow profile
$\bar{\boldsymbol{f}}$	Optimum flow profile
f_p	Flow of path p
$\mathcal{F}_\mu$	Set of flows of total demand μ, defined in (2.1)
$\mathcal{G}$	Graph
$\mathcal{N}^+(v)$	Out-edges of vertex v
$\mathcal{N}^-(v)$	In-edges of vertex v
O	Origin of the network
p	Path
$\mathcal{P}$	Set of paths
P_μ	Primal problem
PoA	Price of anarchy
S	Solution set
SC	Social cost, defined in (2.8)
s_e	Increment of τ_e, defined in the proof of Proposition 3.8

T_v	Equilibrium cost of shortest path to v, defined in (3.10)
u_e	Increment of x_e, defined in the proof of Proposition 3.8
$\mathcal{V}$	Set of vertices
$\mathcal{V}_0$	Set of vertices in the active network
$v_\mu(z)$	$\inf_f \varphi_\mu(\boldsymbol{f}, z)$, defined in (3.2)
$V(\mu)$	Solution of the equilibrium minimization problem, defined in (2.4)
$\widetilde{V}(\mu)$	Solution of the optimum minimization problem, defined in (2.10)
$\boldsymbol{x}$	Load profile
$\boldsymbol{x}^*$	Equilibrium load profile
x_e	Load of edge e
$\mathcal{X}_\mu$	Set of loads of total demand μ
Z	Edge-path incidence matrix
α	Element of $\mathbb{R}_+$, first used in Proposition 4.2
α_k	Element of $\mathbb{R}_+$, first used in Proposition 4.4
β	Element of $\mathbb{R}_+$, first used in Proposition 4.2
β_k	Element of $\mathbb{R}_+$, first used in Proposition 4.4
γ	Element of $\mathbb{R}_-$, first used in Proposition 4.5
γ_k	Element of $\mathbb{R}_-$, first used in Proposition 4.4
Γ	$\mathrm{diag}[(a_e)_{e\in\mathcal{E}}]$
δ	Element of $\mathbb{R}_+$, first used in Proposition 4.5
δ_v	Increment of T_v, defined in the proof of Proposition 3.8
$\Delta_{v,w}$	Constant travel time from v to w
ζ	Element of $\mathbb{R}_+$, first used in the proof of Proposition 5.5
η	Element of $\mathbb{R}_+$, first used in Proposition 4.5
λ	Equilibrium cost, defined in (2.3)
μ	Demand
μ_0	$\widehat{\mathcal{E}}$-break point for the equilibrium, defined in Definition 3.4
$\widetilde{\mu}_0$	$\widehat{\mathcal{E}}$-break point for the optimum
τ_e	Equilibrium cost of edge e
φ_μ	Perturbation function, defined in (3.1)

1 Introduction

Nonatomic routing games provide a model for the distribution of traffic over networks with a large number of drivers, each one representing a negligible fraction of the total demand. The model is based on a directed graph with one or more (OD) pairs, and the costs are identified with the delays incurred to go from origin to destination. The delay on an edge is a nondecreasing function of the load of players on that edge, and the delay of a path is additive over its edges. The standard solution concept for such nonatomic games is the Wardrop equilibrium, according to which the traffic in each OD pair travels along paths of minimum delay. The aggregate social cost experienced by the whole traffic is therefore the product of these minimal delays multiplied by the corresponding traffic demands, summed over all OD pairs.

Equilibria are known to be inefficient, so that a social planner would be able to reduce the social cost by redirecting flows along the network. The most common
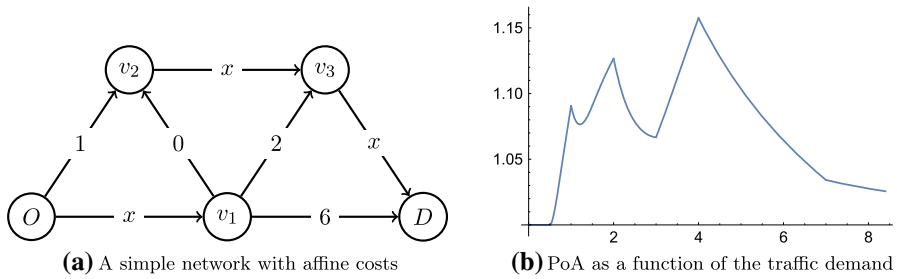
(a) A simple network with affine costs

(b) PoA as a function of the traffic demand

Fig. 1 Oscillations of the price of anarchy

measure of inefficiency is the *price of anarchy* (PoA), that is, the ratio of the social cost at equilibrium over the minimum social cost. For nonatomic congestion games with affine costs, the value of the PoA is bounded above by 4/3, and this bound is known to be sharp [32]. On the other hand, for a large class of cost functions, including all polynomials, the PoA converges to 1 as the traffic demand goes either to 0 or to infinity. In other words, equilibria tend to perfect efficiency both in light and heavy traffic [3,39].

Empirical studies have shown that in real networks, and for intermediate levels of the demand, the PoA tends to oscillate and often does not reach the worst case bounds. Figure 1 shows a typical profile of the PoA as a function of the traffic demand. It starts at 1 for low levels of traffic, then it exhibits some oscillations with a number of nonsmooth spikes, and eventually it decreases smoothly back to 1 in the highly congested regime. The shape and number of these oscillations and spikes is the object of this paper.

1.1 Our contribution

We consider nonatomic routing games over a network with a single OD, and we study the behavior of the price of anarchy as a function of the traffic demand.

To achieve our goal, we need some general results on the continuity and monotonicity of the equilibrium costs and flows. We resort to the classical result in [2] according to which a Wardrop equilibrium is a solution of a convex optimization program. We show that the optimal value of this program is convex and continuously differentiable as a function of the demand, and its derivative is precisely the equilibrium cost. A similar result is established for the minimum social cost.

When the costs have a strictly positive derivative, we show that the equilibrium loads and the PoA are in fact C^1 at each demand level that is *regular*, in the sense that all the optimal paths carry a strictly positive flow. This regularity fails in particular at the so-called $\widehat{\mathcal{E}}$-break points, which are demand levels at which the set of shortest paths at equilibrium changes.

For affine costs we bypass regularity and we show directly that the equilibrium cost is piecewise linear and differentiable except at $\widehat{\mathcal{E}}$-break points. The crucial property is that if the set of shortest paths is the same at two demand levels, then this set remains optimal in between. From here it follows that the number of $\widehat{\mathcal{E}}$-break points is finite,

though it can be exponentially large. It also follows that between $\widehat{\mathcal{E}}$-break points the PoA is differentiable and either is monotone, or is first decreasing and then increasing, so it has a unique minimum in the interior of the interval. From this we conclude that the maximum of the PoA is attained at an $\widehat{\mathcal{E}}$-break point.

We finally present several examples showing how these properties might fail for general costs.

1.2 Related work

The standard solution concept in nonatomic routing games is due to Wardrop [37]. Its mathematical properties were first studied in [2]; early algorithms for computing equilibria were proposed in [36] for affine costs, and in [12] for general convex costs. For historical surveys on the topic we refer to [10,16].

Several papers have considered the sensitivity of the equilibrium flows and costs with respect to variations of the traffic demand. Hall [19] showed that an increase in the demand of one OD pair always increases the equilibrium cost for that pair, whereas [15] observed that the cost on a different OD could be reduced. These questions were further explored in [11]. In a different direction, Patriksson [28] characterized the existence of directional derivatives for the equilibrium and [20] proved that equilibrium costs are always directionally differentiable, whereas equilibrium edge loads not always are. Takalloo and Kwon [14] showed that there exist single OD network games where an ε-increase in the traffic demand produces a global migration of traffic from one set of equilibrium paths to a disjoint set of paths, but, nevertheless, the load on each edge changes at most by ε. Moreover, if the cost functions are polynomials of degree at most d, then the equilibrium costs increase at most of a multiplicative factor $(1 + \varepsilon)^d$. [35] extended this last result to games with multiple OD pairs.

The recognition that selfish behavior produces social inefficiency goes back at least to Pigou [29]. A measure to quantify this inefficiency was proposed in [23], by considering the ratio of the social cost of the worst equilibrium over the optimum social cost. It was termed price of anarchy in [27]. A PoA close to one indicates efficiency of the equilibria of the game, whereas a high PoA implies that, in the worst scenario, strategic behavior can lead to significant social inefficiency. Most of the subsequent literature established sharp bounds for the PoA for specific classes of games, notably for congestion games and in particular for routing games. Roughgarden and Tardos [32] showed that in every nonatomic congestion game with affine costs the PoA is bounded above by 4/3. Moreover, this bound is sharp and attained in a traffic game with a simple two-edge parallel network. Roughgarden [31] generalized this result to polynomial functions of maximum degree d showing that the PoA grows as $\Theta(d/\log d)$. Dumrauf and Gairing [13] refined the result when the cost functions are sums of monomials whose degrees are between q and d. Roughgarden and Tardos [33] extended the analysis to all differentiable cost function c such that $xc(x)$ is convex. Less regular costs and different optimizing criteria for the social cost were studied in [7–9].

Some papers took a more applied view and studied the actual value of the PoA in real networks. Youn et al. [40,41] dealt with traffic in Boston, London, and New York, and noted that the PoA exhibits a similar pattern in the three cities: it is 1 when traffic

is light, it oscillates in the central region and then goes back to 1 when traffic increases. A similar behavior was observed by O'Hare et al. [26], who—taking an approach that relates to the one in this paper—showed how an expansion and/or retraction of the routes used at equilibrium affects the behavior of the PoA. Experimentally, [25] studied the commuting behavior of a large number of Singaporean students and concluded that the PoA is overall low and far from the worst case scenarios.

An analytical justification for the asymptotic efficiency of the PoA in light and heavy traffic was presented in [3,4,39]. Colini-Baldeschi et al. [4] considered the case of single OD parallel networks and proved that, in heavy traffic, the PoA converges to one when the cost functions are regularly varying. These results were extended in various directions in [3], considering general networks and analyzing both the light and heavy traffic asymptotics. A different technique, called *scalability*, was used in [39] to study the case of heavy traffic.

Wu and Möhring [38] considered issues that are quite close to the ones examined here. They defined a metric on the space of nonatomic congestion games that share the same network, commodities, and strategies, but differ in terms of demands and cost functions. Using this metric they showed that the PoA is a continuous function of both the demand and the costs. Then they performed a sensitivity analysis of the PoA with respect to variations of the game in terms of this metric.

In a very interesting recent paper Klimm and Warode [22], [21, see also the conference version] considered nonatomic routing games with piecewise linear costs and provided algorithms to track the full path of Wardrop equilibria when the demands vary proportionally along a fixed direction. These algorithms are based on (positive or negative) electrical flows on undirected graphs and are then suitably adapted to positive flows on directed graphs. The connection between their paper and ours will be discussed in Sect. 4.

The behavior of the PoA as a function of a different parameter was studied in [6]. In that case the parameter of interest is the probability that players actually take part in the game. Colini-Baldeschi et al. [5] studied the possibility of achieving efficiency in routing games via the use of tolls, when the demand can vary, whereas [18] analyzed the income inequality effects of reducing the PoA via tolls.

1.3 Organization of the paper

In Sect. 2 we recall the model of nonatomic routing games, and we set the notations and the standing assumptions. Section 3 investigates the smoothness of the equilibrium costs and of the PoA as a function of the demand, for general nondecreasing smooth costs. The behavior of the PoA for affine costs is studied in Sect. 4. Section 5 presents various examples.

2 The nonatomic congestion model

We consider a nonatomic routing game with a single origin-destination pair. The network is described by a directed multigraph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with vertex set $\mathcal{V}$, edge set

$\mathcal{E}$, an origin $O \in \mathcal{V}$, and a destination $D \in \mathcal{V}$. The traffic demand is given by a positive real number $\mu > 0$, interpreted as vehicles per hour, which has to be routed along a set $\mathcal{P}$ of simple paths from O to D. The nonatomic hypothesis means that each vehicle controls a negligible fraction of the total traffic, and consequently the traffic flows are treated as continuous variables.

The traffic flow on path p is denoted by f_p and the set of *feasible flow profiles* is

$$\mathcal{F}_\mu = \left\{ \boldsymbol{f} = (f_p)_{p \in \mathcal{P}} : f_p \geq 0 \text{ and } \sum_{p \in \mathcal{P}} f_p = \mu \right\}. \tag{2.1}$$

Each flow profile $\boldsymbol{f} \in \mathcal{F}_\mu$ induces a *load profile* $\boldsymbol{x} = (x_e)_{e \in \mathcal{E}}$ with $x_e = \sum_{p \ni e} f_p$ representing the aggregate traffic over the edge e. We call $\mathcal{X}_\mu$ the set of all such load profiles. Note that different flow profiles may induce the same edge loads so this correspondence is not bijective.

Every edge $e \in \mathcal{E}$ has a continuous nondecreasing *cost function* $c_e \colon [0, +\infty) \to [0, +\infty)$, where $c_e(x_e)$ represents the travel time (or unit cost) of traversing the edge when the load is x_e. When the traffic is distributed according to $\boldsymbol{f} = (f_p)_{p \in \mathcal{P}}$ with induced load profile $\boldsymbol{x} = (x_e)_{e \in \mathcal{E}}$, the cost experienced by traveling on a path $p \in \mathcal{P}$ is given by

$$c_p(\boldsymbol{f}) := \sum_{e \in p} c_e \left(\sum_{p' \ni e} f_{p'} \right) = \sum_{e \in p} c_e(x_e). \tag{2.2}$$

With a slight abuse of notation, we use the same symbol c for the cost function over paths and over edges. The meaning should be clear from the context.

2.1 Wardrop equilibrium

A feasible flow profile is called a *Wardrop equilibrium* if the paths that are actually used have minimum cost. Formally, $\boldsymbol{f}^* \in \mathcal{F}_\mu$ is an equilibrium iff there exists $\lambda \in \mathbb{R}$ such that

$$\begin{cases} c_p(\boldsymbol{f}^*) = \lambda & \text{for each } p \in \mathcal{P} \text{ such that } f_p^* > 0, \\ c_p(\boldsymbol{f}^*) \geq \lambda & \text{for each } p \in \mathcal{P} \text{ such that } f_p^* = 0. \end{cases} \tag{2.3}$$

The quantity λ is called the *equilibrium cost*, and is a function of μ.

As noted in [2], Wardrop equilibria coincide with the optimal solutions of the convex minimization problem

$$V(\mu) := \min_{\boldsymbol{f} \in \mathcal{F}_\mu} \sum_{e \in \mathcal{E}} C_e \left(\sum_{p' \ni e} f_{p'} \right) = \min_{\boldsymbol{x} \in \mathcal{X}_\mu} \sum_{e \in \mathcal{E}} C_e(x_e), \tag{2.4}$$

where $C_e(\cdot)$ is the primitive of the edge cost $c_e(\cdot)$, that is,

$$C_e(x_e) = \int_0^{x_e} c_e(z)\, dz. \tag{2.5}$$

This follows by noting that (2.3) are the optimality conditions for $V(\mu)$, with the equilibrium cost λ playing the role of a Lagrange multiplier for the constraint $\sum_p f_p = \mu$. It follows that, for each fixed demand μ an equilibrium flow f^* always exists.

Although Wardrop equilibria are not always unique, all of them induce the same edge costs $\tau_e(\mu) := c_e(x_e^*)$. In particular they have the same equilibrium cost $\lambda = \lambda(\mu)$, which is simply the shortest O-D distance:

$$\lambda(\mu) = \min_{p \in \mathcal{P}} \sum_{e \in p} \tau_e(\mu). \tag{2.6}$$

As a matter of fact, as proved in [17], the equilibrium edge costs $\tau_e = \tau_e(\mu)$ are the unique optimal solution of the strictly convex dual program

$$\min_{\tau} \sum_{e \in \mathcal{E}} C_e^*(\tau_e) - \mu \min_{p \in \mathcal{P}} \sum_{e \in p} \tau_e \tag{2.7}$$

where $C_e^*(\cdot)$ is the Fenchel conjugate of $C_e(\cdot)$, which is strictly convex.

2.2 Social optimum and efficiency of equilibria

The total cost experienced by all users traveling across the network is called the *social cost* and is denoted by

$$SC(f) := \sum_{p \in \mathcal{P}} f_p\, c_p(f) = \sum_{e \in \mathcal{E}} x_e\, c_e(x_e). \tag{2.8}$$

Since in equilibrium all the paths that carry flow have the same cost $\lambda(\mu)$, it follows that all equilibria have the same social cost, that is,

$$SC(f^*) = \sum_{p \in \mathcal{P}} f_p^*\, c_p(f^*) = \sum_{p \in \mathcal{P}} f_p^*\, \lambda(\mu) = \mu\, \lambda(\mu). \tag{2.9}$$

A feasible flow $\tilde{f} \in \mathcal{F}_\mu$ is called an *optimum flow* if it minimizes the social cost, that is, $\tilde{f}$ is an optimal solution of

$$\tilde{V}(\mu) := \min_{f \in \mathcal{F}_\mu} SC(f) = \min_{f \in \mathcal{F}_\mu} \sum_{e \in \mathcal{E}} \tilde{C}_e \left(\sum_{p' \ni e} f_{p'} \right), \tag{2.10}$$

where $\widetilde{C}_e(x_e) = x_e \, c_e(x_e)$. The *price of anarchy* (PoA) is then defined as the ratio between the social cost at equilibrium $\mu \, \lambda(\mu)$ and the minimum social cost $\widetilde{V}(\mu)$:

$$\mathsf{PoA}(\mu) = \frac{\mu \, \lambda(\mu)}{\widetilde{V}(\mu)}. \tag{2.11}$$

Our main goal is to investigate the smoothness of the function $\mu \mapsto \mathsf{PoA}(\mu)$ and to understand the kinks and monotonicity properties observed in the example of Fig. 1b. To this end, Sect. 3 presents some preliminary results on the differentiability of equilibria as a function of the demand μ. More precise results will be discussed in Sect. 4 for the case with affine costs.

3 Differentiability of equilibria and price of anarchy

In order to study the smoothness of the PoA, we begin by establishing some preliminary facts on the differentiability of the optimal value function $\mu \mapsto V(\mu)$ and the smoothness of the equilibrium loads $\mu \mapsto x_e^*(\mu)$. These results follow from general convex duality and sensitivity analysis of parametric optimization problems. The following property does not seem to have been stated earlier in the literature, at least in this generality.

Proposition 3.1 *The map $\mu \mapsto V(\mu)$ is convex and C^1 on $(0, \infty)$ with $V'(\mu) = \lambda(\mu)$ continuous and nondecreasing. Moreover, the equilibrium costs $\tau_e(\mu)$ are uniquely defined and continuous.*

Proof This is a consequence of the convex duality theorem. Indeed, consider the perturbation function $\varphi_\mu : \mathbb{R}^{|\mathcal{P}|} \times \mathbb{R} \to \mathbb{R} \cup \{+\infty\}$, given by

$$\varphi_\mu(\boldsymbol{f}, z) = \begin{cases} \sum_{e \in \mathcal{E}} C_e(\sum_{p' \ni e} f_{p'}) & \text{if } \boldsymbol{f} \geq \boldsymbol{0}, \ \sum_p f_p = \mu + z \\ +\infty & \text{otherwise.} \end{cases} \tag{3.1}$$

Clearly φ_μ is a proper closed convex function [30, page 24]. Calling

$$v_\mu(z) = \inf_{\boldsymbol{f}} \varphi_\mu(\boldsymbol{f}, z), \tag{3.2}$$

we have $V(\mu + z) = v_\mu(z)$ and in particular $V(\mu) = v_\mu(0)$ which we consider as the primal problem (P_μ). From general convex duality, we have that $z \mapsto v_\mu(z) = V(\mu + z)$ is a convex function, from which we deduce that $V(\cdot)$ is convex. Moreover, the perturbation function φ_μ yields a corresponding dual

$$(\mathsf{D}_\mu) \quad \min_{\lambda \in \mathbb{R}} \varphi_\mu^*(\boldsymbol{0}, \lambda), \tag{3.3}$$

where φ_μ^* is the Fenchel conjugate function, that is,

$$\varphi_\mu^*(\boldsymbol{0}, \lambda) = \sup_{\boldsymbol{f}, z} \ \langle \boldsymbol{0}, \boldsymbol{f} \rangle + \lambda z - \varphi_\mu(\boldsymbol{f}, z)$$

$$= \sup_{f \geq 0} \lambda \left(\sum_p f_p - \mu \right) - \sum_{e \in \mathcal{E}} C_e \left(\sum_{p' \ni e} f_{p'} \right).$$

Since $V(\mu')$ is finite for all $\mu' \in (0, \infty)$, it follows that $v_\mu(z) = V(\mu + z)$ is finite for z in some interval around 0, and then the convex duality theorem implies that there is no duality gap and the subdifferential $\partial v_\mu(0)$ at $z = 0$ coincides with the optimal solution set $S(D_\mu)$ of the dual problem, that is, $\partial V(\mu) = \partial v_\mu(0) = S(D_\mu)$.

We claim that the dual problem has a unique solution, which is exactly the equilibrium cost $\lambda(\mu)$. Indeed, fix an optimal solution f^* for $v_\mu(0) = V(\mu)$ and recall that this is just a Wardrop equilibrium. The dual optimal solutions are precisely the λ's such that

$$\varphi_\mu(f^*, 0) + \varphi_\mu^*(\mathbf{0}, \lambda) = 0.$$

This equation can be written explicitly as

$$\sum_{e \in \mathcal{E}} C_e \left(\sum_{p \ni e} f_p^* \right) + \sup_{f \geq 0} \lambda \left(\sum_p f_p - \mu \right) - \sum_{e \in \mathcal{E}} C_e \left(\sum_{p' \ni e} f_{p'} \right) = 0,$$

from which it follows that $f = f^*$ is an optimal solution in the latter supremum. The corresponding optimality conditions are

$$\lambda - \sum_{e \in p} c_e \left(\sum_{p' \ni e} f_{p'}^* \right) = 0, \quad \text{if } f_p^* > 0,$$

$$\lambda - \sum_{e \in p} c_e \left(\sum_{p' \ni e} f_{p'}^* \right) \leq 0, \quad \text{if } f_p^* = 0,$$

which imply that λ is the equilibrium cost for the Wardrop equilibrium, that is, $\lambda = \lambda(\mu)$. It follows that $\partial V(\mu) = \{\lambda(\mu)\}$ so that $\mu \mapsto V(\mu)$ is not only convex but also differentiable with $V'(\mu) = \lambda(\mu)$. The conclusion follows by noting that every convex differentiable function is automatically of class C^1, with $V'(\mu)$ nondecreasing.

The continuity of the equilibrium edge costs $\tau_e = \tau_e(\mu)$ is a consequence of Berge's maximum theorem [see, e.g., [1], Section 17.5]. Indeed, the equilibrium edge costs are optimal solutions for the dual program in (2.7). Since the objective function is jointly continuous in (τ, μ), Berge's theorem implies that the optimal solution correspondence is upper-semicontinous. However, in this case the optimal solution is unique, so that the optimal correspondence is single-valued, and, as a consequence, the equilibrium edge costs $\tau_e(\mu)$ are continuous. $\qquad \square$

As an immediate consequence of Proposition 3.1 we obtain the following result:

Corollary 3.2 *If the costs $c_e(\cdot)$ are strictly increasing and continuous, the equilibrium loads $x_e^*(\mu) = c_e^{-1}(\tau_e(\mu))$ are unique and continuous.*

For multiple origin-destination networks with continuous and strictly increasing costs, the continuity of the *equilibrium loads* x_e^* as a function of the demands was already proved in [19, Theorem 1]. On the other hand, [19, Theorem 2] proved the continuity of the *equilibrium costs* provided that all paths carry a strictly positive flow, whereas [19, Theorem 3] showed that the equilibrium cost of each OD increases with the corresponding demand. As shown in Proposition 3.1, for a single origin-destination the continuity and monotonicity of the equilibrium cost requires neither that costs be strictly increasing nor that all paths carry a strictly positive flow. Although this might be considered a minor improvement, allowing for nondecreasing costs and particularly constant costs is a convenient extension.

For the analysis of the PoA, the most relevant part of Proposition 3.1 is the smoothness of $V(\cdot)$ and the characterization of its derivative $V'(\mu) = \lambda(\mu)$. In particular, considering the social optimum problem (2.10) we get the following direct consequence:

Proposition 3.3 *Let the costs $c_e(\cdot)$ be C^1 and nondecreasing with $x_e \mapsto x_e\, c_e(x_e)$ convex. Then the optimal social cost $\mu \mapsto \widetilde{V}(\mu)$ is convex and C^1. Moreover, $\mu \mapsto PoA(\mu)$ is continuous in $(0, +\infty)$ and differentiable at every μ where the equilibrium cost $\lambda(\mu)$ is differentiable.*

Proof The assumptions on $c_e(\cdot)$ imply that the *marginal costs*

$$\widetilde{c}_e(x_e) := \widetilde{C}_e'(x_e) = c_e(x_e) + x_e\, c_e'(x_e) \tag{3.4}$$

are continuous and nondecreasing. It follows that the optimal flows are the Wardrop equilibria for these marginal costs, and the smoothness of $\widetilde{V}(\mu)$ follows from Proposition 3.1. For the PoA it suffices to observe that $\lambda(\mu)$ is continuous and then use the equality (2.11). □

Wu and Mohring [38] recently established a very general result on the continuity of the PoA with respect to the demands and also with respect to perturbations of the cost functions. However, differentiability was not addressed in their paper.

3.1 Differentiability of equilibrium costs

In order to use Proposition 3.3 it is convenient to find conditions that ensure the differentiability of the equilibrium cost $\lambda(\mu)$. In this section we present one such result, which also guarantees the differentiability of the resource loads $x_e^*(\mu)$. This follows from the implicit function theorem applied to the system of first order optimality conditions for (2.4). Given a vertex $v \in \mathcal{V}$, call $\mathcal{N}^+(v)$ and $\mathcal{N}^-(v)$ the sets of out-edges and in-edges of v, respectively, and $\mathcal{P}(\mathsf{O}, v)$ the set of all paths from O to v.

Moreover, call $\mathcal{E}(v, w)$ the set of all edges $e \in \mathcal{E}$ that go from v to w. Then, an equilibrium load profile $(x_e^*)_{e \in \mathcal{E}}$ for a total demand μ is characterized as a solution of

$$x_e \geq 0 \quad \text{for all } e \in \mathcal{E}, \tag{3.5}$$

$$\sum_{e \in \mathcal{N}^+(v)} x_e - \sum_{e \in \mathcal{N}^-(v)} x_e = \begin{cases} 0 & \text{if } v \neq \mathsf{O}, \mathsf{D}, \\ \mu & \text{if } v = \mathsf{O}, \\ -\mu & \text{if } v = \mathsf{D}, \end{cases} \quad \text{for all } v \in \mathcal{V}, \quad (3.6)$$

$$T_{\mathsf{O}} = 0 \text{ and } T_w \leq T_v + \tau_e \text{ for all } e \in \mathcal{E}(v, w), \quad (3.7)$$

$$x_e \cdot (T_w - T_v - \tau_e) = 0 \text{ for all } e \in \mathcal{E}(v, w), \quad (3.8)$$

$$\tau_e = c_e(x_e) \text{ for all } e \in \mathcal{E}, \quad (3.9)$$

where τ_e is the equilibrium cost of the edge e and

$$T_v := \min_{p \in \mathcal{P}(\mathsf{O}, v)} \sum_{e \in p} \tau_e \quad (3.10)$$

is the equilibrium cost of a shortest path from the origin O to vertex v.

For the subsequent analysis we define the *active network* as the set of all edges that lie on some shortest path. We also consider the demand levels at which this set changes.

Definition 3.4 For each $\mu > 0$ we let $\mathcal{P}(\mu)$ be the set of all shortest paths from the source O to the sink D with cost at equilibrium equal to $\lambda(\mu)$, and we define the *active network* $\widehat{\mathcal{E}}(\mu)$ as the union of the edges on all these paths $p \in \mathcal{P}(\mu)$.

The active network is said to be *locally constant* at μ_0 if there exists $\varepsilon > 0$ such that $\widehat{\mathcal{E}}(\mu)$ is the same for all $\mu \in [\mu_0 - \varepsilon, \mu_0 + \varepsilon]$.

The demand μ_0 is called an $\widehat{\mathcal{E}}$-*break point* if there exists $\varepsilon > 0$ such that $\widehat{\mathcal{E}}(\mu)$ is constant over each of the intervals $[\mu_0 - \varepsilon, \mu_0)$ and $(\mu_0, \mu_0 + \varepsilon]$, with $\widehat{\mathcal{E}}(\mu_0 - \varepsilon) \neq \widehat{\mathcal{E}}(\mu_0 + \varepsilon)$.

Remark 3.5 Since the equilibrium costs $\tau_e = \tau_e(\mu)$ and $\lambda(\mu)$ are unique for each μ, it follows that both $\mathcal{P}(\mu)$ and $\widehat{\mathcal{E}}(\mu)$ are also uniquely determined. Moreover, the continuity of $\tau_e(\mu)$ implies that an edge $e \notin \widehat{\mathcal{E}}(\mu_0)$ that is inactive at μ_0 remains inactive for μ near μ_0, that is to say $\widehat{\mathcal{E}}(\mu_0)^c \subseteq \widehat{\mathcal{E}}(\mu)^c$.

Figure 2 shows the evolution of the active network at different demand levels for the game in Fig. 1a, with five $\widehat{\mathcal{E}}$-break points at $\mu = 1, 2, 3, 4, 7$. Notice the correspondence with the break points in the price of anarchy in Fig. 1b.

Although in general there can be infinitely many $\widehat{\mathcal{E}}$-break points (see Proposition 5.5 and Remark 5.6), their number is finite for series-parallel networks (*cf.* Proposition 3.12) and also for general networks with affine costs (*cf.* Proposition 4.2). In both cases, once an active network changes, it may never occur again at higher demand levels.

An edge carrying a strictly positive flow at equilibrium must be on some optimal path, and, as a consequence, belongs to the active network. However, the converse may fail when a path becomes active but carries no flow. In order to prove the smoothness of the equilibrium flows we need to avoid this situation, which leads to the following definition of a *regular demand*.

Definition 3.6 A demand $\mu > 0$ is called *regular* if there is an equilibrium with $x_e > 0$ for all $e \in \widehat{\mathcal{E}}(\mu)$.
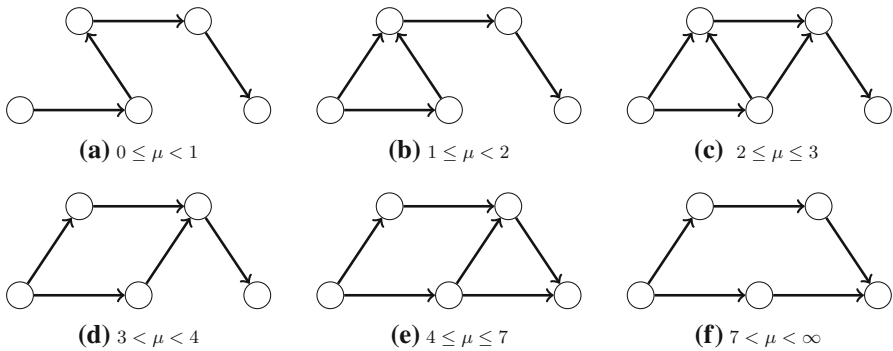
Fig. 2 The active network for the graph on Fig. 1, with $\widehat{\mathcal{E}}$-break points at $\mu = 1, 2, 3, 4, 7$

Regularity is just *strict complementarity*. Indeed, the complementarity condition (3.8) imposes that for each $e \in \mathcal{E}$ either x_e or $(T_w - T_v - \tau_e)$ is zero, whereas strict complementarity requires exactly one of these expressions to be zero. As shown next, when the costs are strictly increasing this implies that the active network is locally constant, so that a regular demand cannot be an $\widehat{\mathcal{E}}$-break point. We note however that there can be nonregular demands at which the active network is still locally constant (see Example 5.2).

Lemma 3.7 *Suppose that the costs $c_e(\cdot)$ are strictly increasing and continuous. If μ_0 is a regular demand then the active network $\widehat{\mathcal{E}}(\mu)$ is locally constant at μ_0.*

Proof From regularity all active edges $e \in \widehat{\mathcal{E}}(\mu_0)$ satisfy $x_e^*(\mu_0) > 0$. By Proposition 3.1 the maps $x_e^*(\cdot)$ are continuous, so these strict inequalities are preserved for μ near μ_0, and therefore $\widehat{\mathcal{E}}(\mu_0) \subseteq \widehat{\mathcal{E}}(\mu)$. This, combined with Remark 3.5, yields $\widehat{\mathcal{E}}(\mu) = \widehat{\mathcal{E}}(\mu_0)$ for μ close to μ_0. □

We are now ready to establish the smoothness of the equilibrium.

Proposition 3.8 *Assume that the costs $c_e(\cdot)$ are C^1 with strictly positive derivative. If μ_0 is a regular demand then $\mu \mapsto (x^*(\mu), \tau(\mu), T(\mu))$ is continuously differentiable in a neighborhood of μ_0. In particular the equilibrium cost $\mu \mapsto \lambda(\mu)$ is C^1 near μ_0.*

Proof From Proposition 3.1 and Corollary 3.2, the equilibrium costs $\tau_e(\mu)$ and loads $x_e^*(\mu)$ are uniquely defined and continuous in μ. Hence, the equilibrium cost $T_v(\mu)$ of a shortest path to any vertex v is also continuous. These functions x_e^*, τ_e and T_v satisfy in particular (3.6), (3.8), and (3.9).

Now, since μ_0 is regular, the active network is locally constant. Let $\widehat{\mathcal{E}}_0 = \widehat{\mathcal{E}}(\mu_0)$ be this active network and $\mathcal{V}_0$ the corresponding vertices. Moreover, call $\widehat{\mathcal{E}}_0(v, w)$ the set of all edges $e \in \widehat{\mathcal{E}}_0$ that go from v to w. For μ near μ_0 we have $x_e^*(\mu) = 0$ for all $e \notin \widehat{\mathcal{E}}_0$; hence, these functions are trivially differentiable. Also for $w \notin \mathcal{V}_0$ we can take $v \in \mathcal{V}_0$ the last vertex on a shortest path from O to w, so that $T_w(\mu) = T_v(\mu) + \Delta_{v,w}$ where $\Delta_{v,w}$ is a constant travel time from v to w. Hence it suffices to establish the smoothness of $x_e^*(\mu)$ for $e \in \widehat{\mathcal{E}}_0$ and $T_v(\mu)$ for $v \in \mathcal{V}_0$. To this end, consider (3.6),

(3.8), and (3.9) restricted to the edges in $\widehat{\mathcal{E}}_0$, together with the equation $T_O = 0$, which gives the following system:

$$\sum_{e \in \mathcal{N}^+(v) \cap \widehat{\mathcal{E}}_0} x_e - \sum_{e \in \mathcal{N}^-(v) \cap \widehat{\mathcal{E}}_0} x_e = \begin{cases} 0 \text{ if } v \neq O, D, \\ \mu \text{ if } v = O, \end{cases} \quad \text{for all } v \in \mathcal{V}_0, \quad (3.11)$$

$$x_e \cdot (T_w - T_v - \tau_e) = 0 \quad \text{for all } e \in \widehat{\mathcal{E}}_0(v, w), \quad (3.12)$$

$$\tau_e = c_e(x_e), \quad (3.13)$$

$$T_O = 0. \quad (3.14)$$

To apply the implicit function theorem to this reduced system, we must check that the associated linearized system has a unique solution. Let u_e, δ_v and s_e be respectively the increments in the variables x_e, T_v and τ_e for each $e \in \widehat{\mathcal{E}}_0$ and $v \in \mathcal{V}_0$. The homogeneous linear system obtained from (3.11), (3.12), (3.13), and (3.14) is:

$$\sum_{e \in \mathcal{N}^+(v) \cap \widehat{\mathcal{E}}_0} u_e - \sum_{e \in \mathcal{N}^-(v) \cap \widehat{\mathcal{E}}_0} u_e = 0 \quad \text{for all } v \in \mathcal{V}_0, \quad (3.15)$$

$$x_e(\delta_w - \delta_v - s_e) + u_e(T_w - T_v - \tau_e) = 0 \quad \text{for all } e \in \widehat{\mathcal{E}}_0(v, w), \quad (3.16)$$

$$s_e = c'_e(x_e)u_e, \quad (3.17)$$

$$\delta_O = 0. \quad (3.18)$$

Strict complementarity on an active link implies that (3.16) is equivalent to

$$\delta_w = \delta_v + s_e \quad \text{for all } e \in \widehat{\mathcal{E}}_0(v, w), \quad (3.19)$$

which, together with (3.17), gives

$$\delta_w = \delta_v + c'_e(x_e)u_e \quad \text{for all } e \in \widehat{\mathcal{E}}_0(v, w). \quad (3.20)$$

These equations are the stationarity conditions for the strongly convex (since $c'_e(x_e) > 0$) quadratic program

$$\min_{u} \sum_{e \in \widehat{\mathcal{E}}_0} \frac{1}{2} c'_e(x_e) \cdot u_e^2 \quad (P)$$

under the constraints (3.15). Indeed, associating a Lagrange multiplier to each of those constraints, we get the Lagrangian

$$\mathcal{L}(u, \delta) = \sum_{e \in \widehat{\mathcal{E}}_0} \frac{1}{2} c'_e(x_e)u_e^2 + \sum_{v \in \mathcal{V}_0} \delta_v \left(\sum_{e \in \mathcal{N}^+(v) \cap \widehat{\mathcal{E}}_0} u_e - \sum_{e \in \mathcal{N}^-(v) \cap \widehat{\mathcal{E}}_0} u_e \right)$$

and the equation $\partial \mathcal{L}/\partial u_e = 0$ is precisely equivalent to (3.20). Hence, every solution of (3.15), (3.16), (3.17), and (3.18) corresponds to an optimal solution of (P). Since

$u_e = 0$ for all $e \in \widehat{\mathcal{E}}_0$ is feasible, it is also the unique optimal solution. Then (3.17) yields $s_e = 0$ for all $e \in \widehat{\mathcal{E}}_0$, and from (3.19) and (3.18) we also get $\delta_v = 0$ for all $v \in \mathcal{V}_0$.

Since the linear system (3.15), (3.16), (3.17), and (3.18) has only the trivial solution, the Jacobian of (3.11), (3.12), (3.13), and (3.14) with respect to x, τ, T is invertible and the implicit function theorem implies the smoothness of the solution. In particular $\lambda(\mu) = T_D(\mu)$ is continuously differentiable. $\qquad\qquad\square$

Remark 3.9 The directional differentiability of the equilibrium loads with respect to parameters was investigated by [20,28], though it can also be derived from [34, Theorem 5.1]. Theorem 10 in [28] shows that differentiability holds provided that the edge loads in the linearized system are unique (this is what is done in the proof above) and assuming in addition that all the route flows corresponding to the solutions of the linearized system satisfy an additional vanishing condition. Our Proposition 3.8 is more straightforward as it follows directly from the implicit function theorem, which gives in addition the continuity of the derivatives. It is also easier to apply and to interpret: it just requires to check that all active links carry a positive flow. In other words, smoothness can only fail at critical values of the demand where a new link becomes active but is not yet carrying flow. This simpler sufficient condition is all that is needed hereafter.

Combining Propositions 3.3 and 3.8 we obtain the following result on the differentiability of the price of anarchy.

Theorem 3.10 *Suppose that $c_e(\cdot)$ are C^1 with strictly positive derivative and $x_e\, c_e(x_e)$ convex. Then PoA $(\cdot)$ is continuously differentiable at each regular demand level μ_0.*

While all $\widehat{\mathcal{E}}$-break points are nonregular, there might exist other nonregular points that are not $\widehat{\mathcal{E}}$-break points (see Example 5.2). We do not know if differentiability of PoA$(\cdot)$ can fail at such additional nonregular points. In the next section we will show that for affine costs, nonsmoothness can only occur at $\widehat{\mathcal{E}}$-break points and that there are finitely many of them. In contrast, for general networks and nonlinear costs the number of $\widehat{\mathcal{E}}$-break points can be unbounded. In this regard, it is worth noting that for networks with a series-parallel topology (which excludes the Wheatstone network), the active network increases monotonically with the demand, which yields a sharp bound on the number of different active networks and $\widehat{\mathcal{E}}$-break points that can occur as the demand grows from 0 to $+\infty$.

Definition 3.11 The class of *series-parallel* (SP) networks can be constructed as follows:

- A network with two vertices O, D and one edge (O, D) connecting them is SP.
- A network obtained by joining in series two SP networks by merging D_1 with O_2 is SP.
- A network obtained by joining in parallel two SP networks by merging O_1 with O_2 and D_1 with D_2 is SP.

Proposition 3.12 *Let $\mathcal{G}$ be a series-parallel network. Then there exist equilibrium load profiles $x(\mu)$ whose components $x_e(\mu)$ are nondecreasing functions of the demand μ.*

Moreover, the active network $\widehat{\mathcal{E}}(\mu)$ is also nondecreasing with respect to inclusion so that the number of $\widehat{\mathcal{E}}$-break points is bounded by the minimum between the number of paths and the number of edges.

Proof The result clearly holds for the network with only two vertices and a single link.

Let $\mathcal{G}^1$ and $\mathcal{G}^2$ be two series-parallel networks for which the result is true, and fix two nondecreasing equilibrium loads $\boldsymbol{x}^1(\mu)$ and $\boldsymbol{x}^2(\mu)$, with their corresponding equilibrium costs $\lambda^1(\mu)$ and $\lambda^2(\mu)$ and active networks $\mathcal{E}^1(\mu)$ and $\mathcal{E}^2(\mu)$.

If $\mathcal{G}^1$ and $\mathcal{G}^2$ are connected in series, then an equilibrium is given by the coupling $\boldsymbol{x}(\mu) = (\boldsymbol{x}^1(\mu), \boldsymbol{x}^2(\mu))$ with active network $\widehat{\mathcal{E}}(\mu) = \mathcal{E}^1(\mu) \cup \mathcal{E}^2(\mu)$, all of which are nondecreasing with μ.

The case in which $\mathcal{G}^1$ and $\mathcal{G}^2$ are joined in parallel, is slightly more involved. Here an equilibrium splits as $\boldsymbol{x}(\mu) = (\boldsymbol{x}^1(\mu^1), \boldsymbol{x}^2(\mu^2))$ where $\mu^1 + \mu^2 = \mu$ with $\mu^2 = 0$ if $\lambda^1(\mu) < \lambda^2(0)$, $\mu^1 = 0$ if $\lambda^2(\mu) < \lambda^1(0)$, and $\lambda^1(\mu^1) = \lambda^2(\mu^2)$ otherwise. More explicitly, if we let $g(\mu) = \inf\{x \in [0, \mu] : \lambda^1(x) \geq \lambda^2(\mu - x)\}$, with $g(\mu) = \mu$ when the latter set is empty, then both $\mu^1 = g(\mu)$ and $\mu^2 = \mu - g(\mu)$ turn out to be nondecreasing and therefore $\boldsymbol{x}(\mu) = (\boldsymbol{x}^1(g(\mu)), \boldsymbol{x}^2(\mu - g(\mu)))$ is a nondecreasing equilibrium. The monotonicity of the active network $\widehat{\mathcal{E}}(\mu)$ is similar. We have $\widehat{\mathcal{E}}(\mu) = \mathcal{E}^1(\mu)$ if $\lambda^1(\mu) < \lambda^2(0), \widehat{\mathcal{E}}(\mu) = \mathcal{E}^2(\mu)$ when $\lambda^2(\mu) < \lambda^1(0)$, and $\widehat{\mathcal{E}}(\mu) = \mathcal{E}^1(g(\mu)) \cup \mathcal{E}^2(\mu - g(\mu))$ otherwise, and in all three cases the active network is nondecreasing. $\square$

Remark 3.13 A related result was obtained in [24] for undirected networks. His Lemma 2 shows that in a series-parallel network *there exists some path* whose edge loads are increasing in the total traffic demand. Proposition 3.12 proves that there exists an equilibrium in which this monotonicity holds for all edges and paths.

Concerned about complexity of computing parametric mincost flows, [22, Corollary 4] also proved monotonicity of the output flows on the edges when costs are piecewise linear.

4 Networks with affine cost functions

4.1 Equilibrium flows

In this section we consider the case of affine cost functions

$$c_e(x) = a_e \cdot x + b_e, \tag{4.1}$$

with $a_e, b_e \geq 0$ for each $e \in \mathcal{E}$. We recall that in this case we have a scaling law that relates the equilibrium and optimum flows.

Lemma 4.1 *(Roughgarden and Tardos [32, Lemma 2.3]) Suppose c_e is affine for all $e \in \mathcal{E}$. Let $\boldsymbol{f}^*(\mu)$ be an equilibrium flow with corresponding load $\boldsymbol{x}^*(\mu)$. Then $\widetilde{\boldsymbol{f}}(\mu) = \frac{1}{2}\boldsymbol{f}^*(2\mu)$ is a socially optimal flow with load $\widetilde{\boldsymbol{x}}(\mu) = \frac{1}{2}\boldsymbol{x}^*(2\mu)$.*

From this it follows directly that for affine costs the $\widehat{\mathcal{E}}$-break points for the optimum are in one-to-one correspondence with the $\widehat{\mathcal{E}}$-break points for the equilibrium, that is,

$$\tilde{\mu}_0 = \frac{1}{2}\mu_0. \tag{4.2}$$

[26, Theorem 3.5] established a similar scaling law when all the edge costs are BPR functions of the same degree.

We now describe the behavior of the equilibrium and the price of anarchy. We first show that any given subset of edges $\widehat{\mathcal{E}}_0 \subseteq \mathcal{E}$ can only be an active network over an interval. In other words, once a given active network changes, it will never become active again. This implies that the number of $\widehat{\mathcal{E}}$-break points is finite. Moreover, we show that between $\widehat{\mathcal{E}}$-break points the equilibrium cost is affine with nonnegative slope and intersect.

Proposition 4.2 *Suppose that the costs* $c_e(\cdot)$ *are affine. Let* $\mathcal{E}_0 \subseteq \mathcal{E}$ *and suppose that* $\mu^1 < \mu^2$ *are such that* $\widehat{\mathcal{E}}(\mu^1) = \widehat{\mathcal{E}}(\mu^2) = \mathcal{E}_0$. *Then, for every* $\mu \in [\mu^1, \mu^2]$ *we have* $\widehat{\mathcal{E}}(\mu) = \mathcal{E}_0$ *and we can select an equilibrium flow* $\boldsymbol{f}(\mu)$ *that is affine in* μ, *so that the equilibrium cost* $\lambda(\mu)$ *is also affine on* $[\mu^1, \mu^2]$. *More precisely* $\lambda(\mu) = \alpha + \beta\mu$ *with nonnegative coefficients* $\alpha \geq 0$ *and* $\beta \geq 0$.

Proof Let $\boldsymbol{f}^1$ and $\boldsymbol{f}^2$ be equilibrium profiles for μ^1 and μ^2, and consider the following affine interpolation with $\theta = \theta(\mu) := (\mu - \mu^1)/(\mu^2 - \mu^1) \in [0, 1]$

$$\boldsymbol{f}(\mu) = (1 - \theta)\boldsymbol{f}^1 + \theta\boldsymbol{f}^2 \in \mathcal{F}_\mu. \tag{4.3}$$

The corresponding load profile is given by $\boldsymbol{x}(\mu) = (1 - \theta)\boldsymbol{x}^1 + \theta\boldsymbol{x}^2$ and, since the costs are affine, it follows that $c_e(x_e(\mu)) = (1 - \theta)c_e(x_e^1) + \theta c_e(x_e^2)$. Hence, the same affine behavior holds for the path costs $c_p(\boldsymbol{f}(\mu)) = (1 - \theta)c_p(\boldsymbol{f}^1) + \theta c_p(\boldsymbol{f}^2)$. Now, since $\widehat{\mathcal{E}}(\mu^1) = \widehat{\mathcal{E}}(\mu^2) = \mathcal{E}_0$, the optimal paths are the same for μ^1 and μ^2. Thus, if p is an optimal path and p' is not optimal, we have

$$c_p(\boldsymbol{f}^1) = \lambda(\mu^1) < c_{p'}(\boldsymbol{f}^1)$$
$$c_p(\boldsymbol{f}^2) = \lambda(\mu^2) < c_{p'}(\boldsymbol{f}^2)$$

and taking a convex combination of these inequalities we get

$$c_p(\boldsymbol{f}(\mu)) = (1 - \theta)\lambda(\mu^1) + \theta\lambda(\mu^2) < c_{p'}(\boldsymbol{f}(\mu)).$$

This implies that the paths p and p' remain respectively optimal and nonoptimal for μ. It follows that $\widehat{\mathcal{E}}(\mu) = \mathcal{E}_0$ and also that $\boldsymbol{f}(\mu)$ is an equilibrium with $\lambda(\mu) = (1 - \theta)\lambda(\mu^1) + \theta\lambda(\mu^2)$.

This shows that the equilibrium cost is affine over the interval $[\mu^1, \mu^2]$, that is, $\lambda(\mu) = \alpha + \beta\mu$. From Proposition 3.1 we know that $\lambda(\mu)$ is nondecreasing so that $\beta \geq 0$, and therefore it remains to show that $\alpha \geq 0$. We rewrite the interpolated flow

as $f(\mu) = \mu w + z$ with

$$w = \frac{f^2 - f^1}{\mu^2 - \mu^1},\tag{4.4}$$

$$z = \frac{\mu^2 f^1 - \mu^1 f^2}{\mu^2 - \mu^1}.\tag{4.5}$$

Let $\mathcal{P}_0$ be the set of all the O-D shortest paths included in $\mathcal{E}_0$. We observe that $z_p = 0$ for every $p \notin \mathcal{P}_0$, and also that $\sum_{p \in \mathcal{P}} z_e = 0$. On the other hand, since $f(\mu)$ is an equilibrium we have $c_p(f(\mu)) = \lambda(\mu)$ for all $p \in \mathcal{P}_0$, so that $\langle c(f(\mu)), z \rangle = 0$ and, as a consequence,

$$\mu\lambda(\mu) = \langle c(f(\mu)), f(\mu) \rangle = \langle c(f(\mu)), \mu w + z \rangle = \mu\langle c(f(\mu)), w \rangle.\tag{4.6}$$

Defining

$$A = Z^\top \Gamma Z \quad \text{and} \quad d = Z^\top b,\tag{4.7}$$

with Z the edge-path incidence matrix, $\Gamma = \mathrm{diag}[(a_e)_{e \in \mathcal{E}}]$, and $b = (b_e)_{e \in \mathcal{E}}$, the vector of path costs can be expressed as

$$c(f(\mu)) = Af(\mu) + d = \mu Aw + Az + d,\tag{4.8}$$

so that

$$\lambda(\mu) = \langle c(f(\mu)), w \rangle = \mu\langle Aw, w \rangle + \langle Az + d, w \rangle,\tag{4.9}$$

which yields $\beta = \langle Aw, w \rangle$ and $\alpha = \langle Az + d, w \rangle$. Since A is positive semidefinite, we get again $\beta \geq 0$. Now, since $0 = \langle c(f(\mu)), z \rangle = \mu\langle Aw, z \rangle + \langle Az + d, z \rangle$ for all μ, it follows that $\langle Aw, z \rangle = 0$ so that $\alpha = \langle d, w \rangle$, and also $\langle Az + d, z \rangle = 0$ so that $\langle d, z \rangle = -\langle Az, z \rangle \leq 0$. To conclude, we note that all the entries in $f(\mu)$ and in d are nonnegative, and therefore

$$0 \leq \langle f(\mu), d \rangle = \mu\langle w, d \rangle + \langle z, d \rangle,\tag{4.10}$$

from which we deduce that $\alpha = \langle w, d \rangle \geq 0$ as claimed. $\qquad\square$

Remark 4.3 The paper by Klimm and Warode [22] developed a homotopy method for computing the full path of Wardrop equilibria as a function of the traffic demand. The method is designed to work with piecewise linear costs and produces a piecewise linear path of equilibrium loads. The algorithm first determines the equilibrium costs and then recovers the loads by inverting the link costs. This requires the costs to be strictly increasing. In contrast, we work directly in the space of flows so we can handle nondecreasing and constant costs. However, we restrict to affine costs, which is essential for Proposition 4.2. Indeed, beyond the piecewise affine character of the equilibrium, the most relevant part of this result is the identification of the break points

as the demand levels at which the active network changes. As shown in Proposition 4.2, each particular subnetwork can be active on a demand interval, and, on each of these intervals, the equilibrium varies linearly. Once an active network is abandoned, it will never occur again at higher demand levels. This property fails to hold for nonlinear or even piecewise linear costs, as shown in Proposition 5.5.

[22, Theorem 10] showed that, even for affine costs, the number of $\widehat{\mathcal{E}}$-break points can be exponential in the number of paths. In Sect. 5 we will use their example to show an interesting behavior of the price of anarchy.

4.2 Behavior of the price of anarchy

We now prove that the social cost at the equilibrium and at the optimum have a very similar quadratic form, from which we deduce that between $\widehat{\mathcal{E}}$-break points the function $\mu \mapsto \mathsf{PoA}(\mu)$ has a unique minimum and its maximum must be attained at some of the $\widehat{\mathcal{E}}$-break points.

Proposition 4.4 *Let μ_k and μ_{k+1} be two consecutive $\widehat{\mathcal{E}}$-break points for the equilibrium. Then, there exist $\alpha_k \geq 0$, $\beta_k \geq 0$, and $\gamma_k \leq 0$, such that*

$$\mathsf{SC}(\boldsymbol{f}^*(\mu)) = \alpha_k \mu + \beta_k \mu^2 \quad when \ \mu \in (\mu_k, \mu_{k+1}), \tag{4.11}$$

$$\mathsf{SC}(\widetilde{\boldsymbol{f}}(\mu)) = \gamma_k + \alpha_k \mu + \beta_k \mu^2 \quad when \ 2\mu \in (\mu_k, \mu_{k+1}), \ i.e., \ \mu \in (\widetilde{\mu}_k, \widetilde{\mu}_{k+1}). \tag{4.12}$$

Proof Since $\mathsf{SC}(\boldsymbol{f}^*(\mu)) = \mu \, \lambda(\mu)$, the equality (4.11) follows directly from Proposition 4.2 with $\alpha_k = \langle A\boldsymbol{z} + \boldsymbol{d}, \boldsymbol{w} \rangle$ and $\beta_k = \langle A\boldsymbol{w}, \boldsymbol{w} \rangle$, where $\boldsymbol{w}, \boldsymbol{z}, A, \boldsymbol{d}$ are defined as in (4.4), (4.5), and (4.7).

In order to prove (4.12), let $\boldsymbol{f}(\mu) = \mu \boldsymbol{w} + \boldsymbol{z}$ be the affine interpolated equilibria as in the proof of Proposition 4.2. When $2\mu \in (\mu_k, \mu_{k+1})$, Proposition 4.1 implies that an optimum flow is

$$\widetilde{\boldsymbol{f}}(\mu) = \frac{1}{2} \boldsymbol{f}(2\mu) = \mu \, \boldsymbol{w} + \frac{1}{2}\boldsymbol{z}.$$

Then the social cost at optimum will be a quadratic function in μ with the same linear coefficient (α_k) and quadratic coefficient (β_k) of the social cost at equilibrium, and constant coefficient

$$\gamma_k = \frac{1}{4}\langle A\boldsymbol{z}, \boldsymbol{z} \rangle + \frac{1}{2}\langle \boldsymbol{d}, \boldsymbol{z} \rangle,$$

which is less or equal to $\langle A\boldsymbol{z}, \boldsymbol{z} \rangle + \langle \boldsymbol{d}, \boldsymbol{z} \rangle = 0$ since $\langle A\boldsymbol{z}, \boldsymbol{z} \rangle \geq 0$ and $\langle \boldsymbol{d}, \boldsymbol{z} \rangle \leq 0$. □

From this result it follows that the PoA between $\widehat{\mathcal{E}}$-break points is a quotient of quadratics. Moreover, the specific signs of the coefficients of these quadratics imply that PoA has a unique minimum between break points and that its local maxima can only occur at these break points.

Theorem 4.5 *Let μ_k and μ_{k+1} be two consecutive $\widehat{\mathcal{E}}$-break points for the equilibrium. Then on the interval (μ_k, μ_{k+1}) the function $\mu \mapsto \mathsf{PoA}(\mu)$ is smooth and it is either decreasing, or increasing, or first decreasing and then increasing with a local minimum in the interior of the interval. In particular $\mathsf{PoA}(\mu)$ does not attain a local maximum on (μ_k, μ_{k+1}).*

Proof From Proposition 3.3 the optimal social cost is a C^1 function of μ, so that (4.11) implies that $\mathsf{PoA}(\mu)$ is smooth over the full interval (μ_k, μ_{k+1}). Consider first the case when there is no $\widehat{\mathcal{E}}$-break point $\widetilde{\mu}_\ell$ for the optimum in this interval. Using (4.11), and (4.12) we can express the PoA in the form

$$\mathsf{PoA}(\mu) = \frac{\mathsf{SC}(f^*(\mu))}{\mathsf{SC}(\widetilde{f}(\mu))} = \frac{\alpha\mu + \beta\mu^2}{\gamma + \delta\mu + \eta\mu^2},$$

with $\alpha, \beta, \delta, \eta, \geq 0$ and $\gamma \leq 0$. The derivative is given by

$$\mathsf{PoA}'(\mu) = \frac{1}{\mathsf{SC}(\widetilde{f}(\mu))^2} \cdot [\alpha\gamma + (2\beta\gamma)\mu + (\beta\delta - \alpha\eta)\mu^2]$$

and, since $\alpha\gamma \leq 0$ and $2\beta\gamma \leq 0$, it can have at most one positive zero, and only if $\beta\delta - \alpha\eta > 0$. Hence, either $\mathsf{PoA}'(\mu)$ has a constant sign over (μ_k, μ_{k+1}), or it changes from negative to positive if the zero lies on (μ_k, μ_{k+1}), in which case we have a local minimum at this zero.

If the optimum has $\widehat{\mathcal{E}}$-break points in (μ_k, μ_{k+1}), we can repeat the argument on each subinterval, noting that $\mathsf{PoA}(\mu)$ is C^1 so that the sign $\mathsf{PoA}'(\mu)$ does not change at these $\widehat{\mathcal{E}}$-break points. $\qquad\qquad\square$

Proposition 4.5 shows that the typical profile of the Price of Anarchy for networks with affine costs is similar to the one shown in the example Fig. 1b. In particular, it implies:

Corollary 4.6 *For networks with affine costs the maximum of the Price of Anarchy is attained at some $\widehat{\mathcal{E}}$-break point.*

5 Examples and counter examples

In this section we present a set of examples that illustrate the results of the previous sections.

The first example shows the difference between the set of $\widehat{\mathcal{E}}$-break points and the demands at which the set of paths used at equilibrium changes.

Example 5.1 Given a selection of equilibrium flows $f^*(\mu)$, the values μ at which the set of used paths changes may be different from the $\widehat{\mathcal{E}}$-break points. For instance, in the network of Fig. 3, for all $\mu \geq 2$ and for any choice of $\eta \in [0, 1]$ the flow in the flow in the following table is an equilibrium.

.

Path	{O, v_1, v_2, D}	{O, v_1, D}	{O, v_2, D}	{O, D}
Flow	$1 - \eta$	η	η	$\mu - 1 - \eta$

Fig. 3 The set of paths used in equilibrium may change arbitrarily often
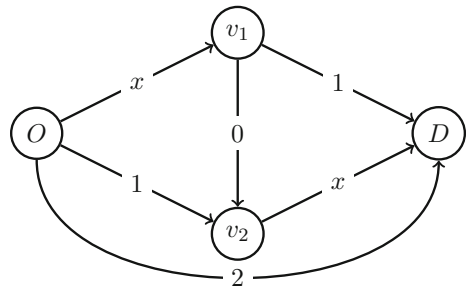
By letting η oscillate between 0 and 1 arbitrarily often, the set of used paths might change an arbitrary number of times, whereas the active network is always the set of edges in all four paths, each of them with an equilibrium cost equal to 2.

Our second example shows that in general the set of nonregular demands may be strictly larger than the set of $\widehat{\mathcal{E}}$-break points.

Example 5.2 Consider again Fig. 3 with the cost 2 in the lower link replaced by $2 + x$, with $\widehat{\mathcal{E}}$-break points at $\mu = 1$ and $\mu = 2$. Note that all the demands $\mu < 1$ and $\mu > 2$ are regular. However, for $\mu \in (1, 2)$ the active network $\widehat{\mathcal{E}}(\mu)$ is constant and comprises all the links, so that there are no $\widehat{\mathcal{E}}$-break points, while the unique equilibrium sends a zero flow on the lower link and hence μ is not regular. It is worth noting that, although the loss of regularity in the interval $(1, 2)$ precludes the use of Proposition 3.8, in this case the costs are affine so that the equilibrium flows are piecewise affine and differentiable except at the $\widehat{\mathcal{E}}$-break points $\mu = 1$ and $\mu = 2$.

The next example deals with the fact that the PoA can attain the value 1 several times.

Example 5.3 In the example of Fig. 1, the price of anarchy shows an initial phase in which it is identically equal to 1, after which it oscillates and eventually goes back to 1 but only asymptotically. We will use an example taken from [22, Section 6.1.2] to show that the PoA can oscillate and go back to 1 more than once. The idea is to nest several Wheatstone networks and choose constant costs that increase exponentially as we go from the inner to the outer networks.

Figure 4a shows the version where the network is obtained by nesting with two Wheatstone networks. Figure 4b shows the graph of the corresponding PoA. The PoA is equal to 1 for small demand ($\mu \leq 1/2$), then increases and reaches a local maximum, then it decreases back to 1 and it remains equal to 1 for the entire interval [2, 3] of demand, then reaches its maximum and, after that, decreases back to 1, where it remains indefinitely.

Below we list the paths in the network in Fig. 4a:
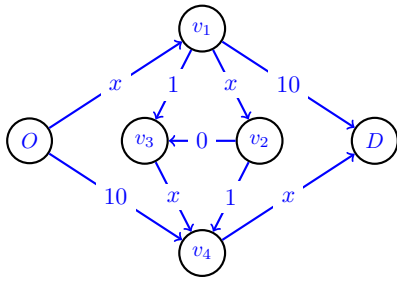$$p_1 = O \rightarrow v_1 \rightarrow D,$$

$$p_2 = O \rightarrow v_4 \rightarrow D,$$
$$p_3 = O \rightarrow v_1 \rightarrow v_2 \rightarrow v_4 \rightarrow D,$$
$$p_4 = O \rightarrow v_1 \rightarrow v_3 \rightarrow v_4 \rightarrow D,$$
$$p_5 = O \rightarrow v_1 \rightarrow v_2 \rightarrow v_3 \rightarrow v_4 \rightarrow D.$$

The equilibrium flow for $\mu \in [0, +\infty)$ is given explicitly in the following table:

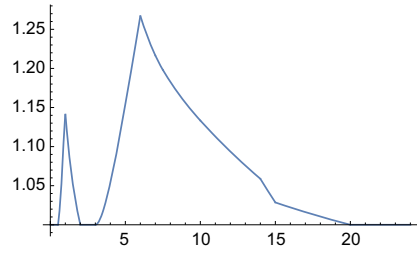| Interval	Cost $\lambda(\mu)$	p_1	p_2	p_3	p_4	p_5
$\mu \in [0, 1)$	4μ	0	0	0	0	μ
$\mu \in [1, 2)$	$2 + 2\mu$	0	0	$\mu - 1$	$\mu - 1$	$2 - \mu$
$\mu \in [2, 6)$	$1 + \dfrac{5}{2}\mu$	0	0	$\dfrac{\mu}{2}$	$\dfrac{\mu}{2}$	0
$\mu \in [6, 14)$	$\dfrac{29}{2} + \dfrac{\mu}{4}$	$\dfrac{3}{4}\mu - \dfrac{9}{2}$	$\dfrac{3}{4}\mu - \dfrac{9}{2}$	$\dfrac{9}{2} - \dfrac{\mu}{4}$	$\dfrac{9}{2} - \dfrac{\mu}{4}$	0
$\mu \in [14, 15)$	18	$\mu - 8$	$\mu - 8$	$15 - \mu$	$15 - \mu$	$\mu - 14$
$\mu \in [15, 20)$	$12 + \dfrac{2}{5}\mu$	$\dfrac{3}{5}\mu - 2$	$\dfrac{3}{5}\mu - 2$	0	0	$4 - \dfrac{\mu}{5}$
$\mu \geq 20$	$10 + \dfrac{\mu}{2}$	$\dfrac{\mu}{2}$	$\dfrac{\mu}{2}$	0	0	0

As a consequence, the PoA is

$$
\mathrm{PoA}(\mu) = \begin{cases}
1 & \text{if } \mu \in [0, \tfrac{1}{2}) \\[4pt]
\dfrac{8\mu^2}{-1 + 4\mu + 4\mu^2} & \text{if } \mu \in [\tfrac{1}{2}, 1) \\[8pt]
\dfrac{4 + 4\mu}{2 + 5\mu} & \text{if } \mu \in [1, 2) \\[8pt]
1 & \text{if } \mu \in [2, 3) \\[4pt]
\dfrac{4\mu + 10\mu^2}{-81 + 58\mu + \mu^2} & \text{if } \mu \in [3, 6) \\[8pt]
\dfrac{58\mu + \mu^2}{-81 + 58\mu + \mu^2} & \text{if } \mu \in [6, 7) \\[8pt]
\dfrac{58\mu + \mu^2}{-130 + 72\mu} & \text{if } \mu \in [7, \tfrac{15}{2}) \\[8pt]
\dfrac{290\mu + 5\mu^2}{-200 + 240\mu + 8\mu^2} & \text{if } \mu \in [\tfrac{15}{2}, 10) \\[8pt]
\dfrac{58 + \mu}{40 + 2\mu} & \text{if } \mu \in [10, 14) \\[8pt]
\dfrac{36}{20 + \mu} & \text{if } \mu \in [14, 15) \\[8pt]
\dfrac{120 + 4\mu}{100 + 5\mu} & \text{if } \mu \in [15, 20) \\[8pt]
1 & \text{if } \mu \geq 20.
\end{cases}
\tag{5.1}
$$

(a) Two nested Wheatstone networks   **(b)** Corresponding PoA

Fig. 4 An example where PoA goes back to 1 once at intermediate demands

The case of three nested Wheatstone networks can be treated similarly (see Fig. 5). Here the PoA is 1 in the intervals $[0, \frac{1}{2}]$, $[2, 3]$, $[20, 30]$ and $[200, +\infty)$.

Example 5.4 This example shows that the result in Proposition 4.5 fails for polynomials, even for the simplest network topology. Indeed, consider the two-link parallel network with cost functions

$$
\begin{aligned}
c_1(x_1) &= x_1, \\
c_2(x_2) &= 1 + x_2^2.
\end{aligned}
\tag{5.2}
$$

The equilibrium and optimum flows can be computed explicitly as in the following table:

Demand	$x_1^*(\mu)$	$x_2^*(\mu)$	$\widetilde{x}_1(\mu)$	$\widetilde{x}_2(\mu)$
$\mu \in [0, 1/2]$	μ	0	μ	0
$\mu \in [1/2, 1]$	μ	0	$\mu - \dfrac{-1 + \sqrt{6\mu - 2}}{3}$	$\dfrac{-1 + \sqrt{6\mu - 2}}{3}$
$\mu \in [1, +\infty)$	$\mu - \dfrac{-1 + \sqrt{4\mu - 3}}{2}$	$\dfrac{-1 + \sqrt{4\mu - 3}}{2}$	$\mu - \dfrac{-1 + \sqrt{6\mu - 2}}{3}$	$\dfrac{-1 + \sqrt{6\mu - 2}}{3}$

Note that the sole $\widehat{\mathcal{E}}$-break point is at $\mu = 1$. Moreover we have $x_1^*(3) = \widetilde{x}_1(3) = 2$ and $x_2^*(3) = \widetilde{x}_2(3) = 1$ so that the equilibrium and optimal flows coincide, and PoA(3) = 1. This implies that in the interval $(3, +\infty)$ the PoA has a local maximum (see Fig. 6).

When the cost functions are less regular, the set of paths used at equilibrium can have a recurring behavior, and an active network that is abandoned at some point can be reactivated at larger demands. In particular we cannot ensure that the number of $\widehat{\mathcal{E}}$-break points is finite.

Proposition 5.5 *There exist networks and nondecreasing cost functions $c_e(\cdot)$ such that a given active network $\widehat{\mathcal{E}}(\mu) = \mathcal{E}_0$ can repeat itself over disjoint demand intervals defined by $\widehat{\mathcal{E}}$-break points.*
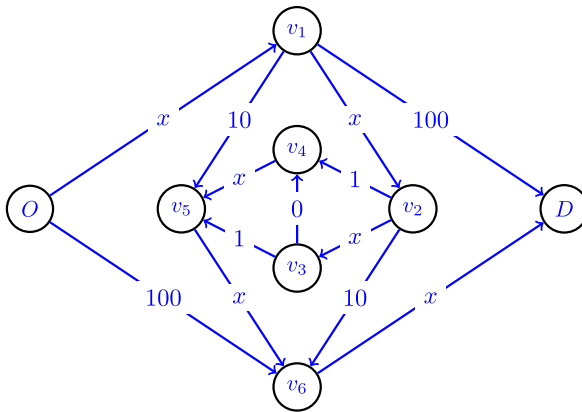
(a) Three nested Wheatstone networks
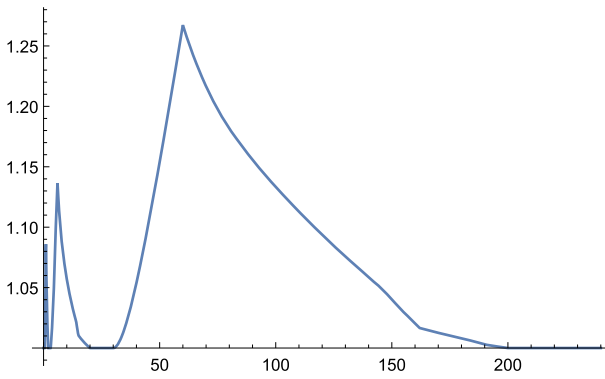
(b) Corresponding PoA

Fig. 5 An example where PoA goes back to 1 twice at intermediate demands
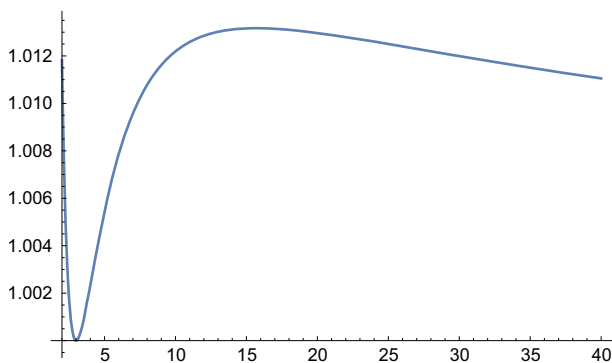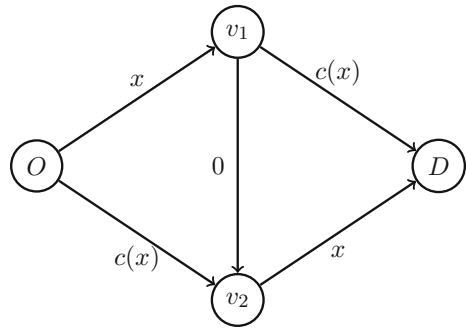

Fig. 6 Plot of the PoA on the interval $[2, 40]$ for the parallel network with costs as in (5.2)

Fig. 7 Wheatsone network with nonaffine costs

Proof Consider the network in Fig. 7 with the cost $c(x)$ defined in $[0, \zeta] \cup [\zeta + \eta, +\infty]$ as follows

$$c(x) = \begin{cases} a & \text{if } x \leq \zeta \\ b & \text{if } x \geq \zeta + \eta, \end{cases}$$

with $0 < \eta < a < \zeta < \zeta + \eta < b$, and in the interval $[\zeta, \zeta + \eta]$ we interpolate in any way that makes $c(x)$ continuous and nondecreasing in the whole $[0, +\infty]$.

Then we have the following regimes:

(a) when $\mu \in [0, a]$, the equilibrium flow uses only the path $\{O, v_1, v_2, D\}$, the load on the two edges $\{O, v_2\}$, and $\{v_1, D\}$ is zero and the equilibrium cost is $\lambda(\mu) = 2\mu$;

(b) when $\mu \in (a, 2a)$, the equilibrium flow uses all the three paths as in the following table:

Path	$\{O, v_1, v_2, D\}$	$\{O, v_1, D\}$	$\{O, v_2, D\}$
Flow	$2a - \mu$	$\mu - a$	$\mu - a$

the load on the two edges $\{O, v_2\}$, $\{v_1, D\}$ is $\mu - a < \zeta$, and the equilibrium cost is $\lambda(\mu) = 2a$;

(c) when $\mu \in [2a, 2\zeta]$, the equilibrium flow splits equally between the two paths $\{O, v_1, D\}$, $\{O, v_2, D\}$, the load on the two edges $\{O, v_2\}$, $\{v_1, D\}$ is $\mu/2 \leq \zeta$, and the equilibrium cost is $\lambda(\mu) = a + \mu/2$;

(d) the regime on the demand interval $[2\zeta, b + \zeta + \eta]$ is complicated to describe, but this is not relevant for our purpose;

(e) when $\mu \in [b + \zeta + \eta, 2b)$, the equilibrium uses all three paths as in the following table:

Path	$\{O, v_1, v_2, D\}$	$\{O, v_1, D\}$	$\{O, v_2, D\}$
Flow	$2b - \mu$	$\mu - b$	$\mu - b$

the load on the two edges $\{O, v_2\}$, $\{v_1, D\}$ is $\mu - b \geq \zeta + \eta$, and the equilibrium cost is $\lambda(\mu) = 2b$;

(f) when $\mu \geq 2b$, the equilibrium flow splits equally between the two paths $\{O, v_1, D\}$, $\{O, v_2, D\}$, and the load on the two edges $\{O, v_2\}$, $\{v_1, D\}$ is $\mu/2 > \zeta + \eta$.

This proves our claim, as in the intervals sub (c) and (f) the equilibrium uses only the two paths $\{O, v_1, D\}$ and $\{O, v_2, D\}$, while in the interval sub (b) and (e) it uses all three paths. □

Remark 5.6 Note that in the same way one can construct examples of networks with an infinite number of $\widehat{\mathcal{E}}$-break points for the equilibrium. Furthermore, one can make the increasing sequence of such $\widehat{\mathcal{E}}$-break points to be convergent. Indeed, one could chose infinite sequences $(\eta_i)_{i \in \mathbb{N}}$, $(a_i)_{i \in \mathbb{N}}$, $(\zeta_i)_{i \in \mathbb{N}}$, $(b_i)_{i \in \mathbb{N}}$ with

$$0 < \eta_i < a_i < \zeta_i < \zeta_i + \eta_i < b_i < a_{i+1} \quad \text{for every } i \in \mathbb{N}. \tag{5.3}$$

Then we set the cost function $c(x)$ to be

$$c(x) = \begin{cases} a_1 & \text{if } x \leq \zeta_1 \\ b_i & \text{if } \zeta_i + \eta_i \leq x \leq \dfrac{a_{i+1} - b_i}{2} \\ a_{i+1} & \text{if } a_{i+1} \leq x \leq \zeta_{i+1}, \end{cases}$$

and in the intervals $[\zeta_i, \zeta_i + \eta]$, $[(a_{i+1} - b_i)/2, a_{i+1}]$ we interpolate in any way that makes $c(x)$ continuous and nondecreasing in the whole $[0, +\infty]$. With this choice of $c(x)$, the situation is similar to the one of Proposition 5.5, repeated infinitely many times. Furthermore, since the η_i can be as small as we want, we can choose the sequences $(a_i)_{i \in \mathbb{N}}$, $(\zeta_i)_{i \in \mathbb{N}}$, $(b_i)_{i \in \mathbb{N}}$ to be convergent to the same limit.

Remark 5.7 Repetition of an active network can happen for cost functions that are smooth or piecewise affine, as in [22]. We know that this cannot happen for affine cost functions, but, so far we have not been able to characterize the class of cost functions for which repetitions of active networks are impossible.

Declarations

References

1. Aliprantis, C.D., Border, K.C.: Infinite Dimensional Analysis, 3rd edn. Springer, Berlin (2006)
2. Beckmann, M.J., McGuire, C., Winsten, C.B.: Studies in the Economics of Transportation. Yale University Press, New Haven (1956)
3. Colini-Baldeschi, R., Cominetti, R., Mertikopoulos, P., Scarsini, M.: When is selfish routing bad? The price of anarchy in light and heavy traffic. Oper. Res. **68**(2), 411–434 (2020). https://doi.org/10.1287/opre.2019.1894
4. Colini-Baldeschi, R., Cominetti, R., Scarsini, M.: Price of anarchy for highly congested routing games in parallel networks. Theory Comput. Syst. **63**(1), 90–113 (2019). https://doi.org/10.1007/s00224-017-9834-1
5. Colini-Baldeschi, R., Klimm, M., Scarsini, M.: Demand-independent optimal tolls. In: 45th International Colloquium on Automata, Languages, and Programming, *LIPIcs. Leibniz Int. Proc. Inform.*, vol. 107, pp. Art. No. 151, 14. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern (2018). https://doi.org/10.4230/LIPIcs.ICALP.2018.151
6. Cominetti, R., Scarsini, M., Schröder, M., Stier-Moses, N.: Price of anarchy in stochastic atomic congestion games with affine costs. In: Proceedings of The Twentieth ACM Conference on Economics and Computation (ACM EC '19) (2019). https://doi.org/10.1145/3328526.3329579
7. Correa, J.R., Schulz, A.S., Stier-Moses, N.E.: Selfish routing in capacitated networks. Math. Oper. Res. **29**(4), 961–976 (2004). https://doi.org/10.1287/moor.1040.0098
8. Correa, J.R., Schulz, A.S., Stier-Moses, N.E.: Fast, fair, and efficient flows in networks. Oper. Res. **55**(2), 215–225 (2007). https://doi.org/10.1287/opre.1070.0383
9. Correa, J.R., Schulz, A.S., Stier-Moses, N.E.: A geometric approach to the price of anarchy in nonatomic congestion games. Games Econom. Behav. **64**(2), 457–469 (2008). https://doi.org/10.1016/j.geb.2008.01.001
10. Correa, J.R., Stier-Moses, N.E.: Wardrop equilibria. In: Cochran, J.J. (ed.) Encyclopedia of Operations Research and Management Science. Wiley, New York (2011). https://doi.org/10.1002/9780470400531.eorms0962
11. Dafermos, S., Nagurney, A.: Sensitivity analysis for the asymmetric network equilibrium problem. Math. Program. **28**(2), 174–184 (1984). https://doi.org/10.1007/BF02612357
12. Dafermos, S.C., Sparrow, F.T.: The traffic assignment problem for a general network. J. Res. Nat. Bur. Standards Sect. B **73B**, 91–118 (1969)
13. Dumrauf, D., Gairing, M.: Price of anarchy for polynomial Wardrop games. In: WINE '06: Proceedings of the 2nd Conference on Web and Internet Economics, pp. 319–330. Springer, Berlin (2006). https://doi.org/10.1007/11944874_29
14. Englert, M., Franke, T., Olbrich, L.: Sensitivity of Wardrop equilibria. Theory Comput. Syst. **47**(1), 3–14 (2010). https://doi.org/10.1007/s00224-009-9196-4
15. Fisk, C.: More paradoxes in the equilibrium assignment problem. Transport. Res. Part B **13**(4), 305–309 (1979). https://doi.org/10.1016/0191-2615(79)90023-7
16. Florian, M., Hearn, D.: Network equilibrium models and algorithms. In: M.H. et al (ed.) Handbooks in Operation Research and Management Science, vol. 8. North Holland, Amsterdam (1995)
17. Fukushima, M.: On the dual approach to the traffic assignment problem. Transport. Res. Part B **18**(3), 235–245 (1984). https://doi.org/10.1016/0191-2615(84)90034-1
18. Gemici, K., Koutsoupias, E., Monnot, B., Papadimitriou, C.H., Piliouras, G.: Wealth inequality and the price of anarchy. In: 36th International Symposium on Theoretical Aspects of Computer

Science. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern (2019). https://doi.org/10.4230/LIPIcs.STACS.2019.31

19. Hall, M.A.: Properties of the equilibrium state in transportation networks. Transport. Sci. **12**(3), 208–216 (1978). https://doi.org/10.1287/trsc.12.3.208

20. Josefsson, M., Patriksson, M.: Sensitivity analysis of separable traffic equilibrium equilibria with application to bilevel optimization in network design. Transport. Res. Part B **41**(1), 4–31 (2007). https://doi.org/10.1016/j.trb.2005.12.004

21. Klimm, M., Warode, P.: Computing all Wardrop equilibria parametrized by the flow demand. In: Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 917–934. SIAM, Philadelphia (2019). https://doi.org/10.1137/1.9781611975482.56

22. Klimm, M., Warode, P.: Parametric computation of minimum-cost flows with piecewise quadratic costs. Math. Oper. Res. (forthcoming) (2021)

23. Koutsoupias, E., Papadimitriou, C.: Worst-case equilibria. In: STACS 99 (Trier), Lecture Notes in Comput. Sci., vol. 1563, pp. 404–413. Springer, Berlin (1999). https://doi.org/10.1007/3-540-49116-3_38

24. Milchtaich, I.: Network topology and the efficiency of equilibrium. Games Econom. Behav. **57**(2), 321–346 (2006). https://doi.org/10.1016/j.geb.2005.09.005

25. Monnot, B., Benita, F., Piliouras, G.: Routing games in the wild: efficiency, equilibration and regret. In: Devanur, N.R., Lu, P. (eds.) Web and Internet Economics, pp. 340–353. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-71924-5_24

26. O'Hare, S.J., Connors, R.D., Watling, D.P.: Mechanisms that govern how the price of anarchy varies with travel demand. Transport. Res. Part B **84**, 55–80 (2016). https://doi.org/10.1016/j.trb.2015.12.005

27. Papadimitriou, C.: Algorithms, games, and the internet. In: Proceedings of the Thirty-Third Annual ACM Symposium on Theory of Computing, pp. 749–753. ACM, New York (2001). https://doi.org/10.1145/380752.380883

28. Patriksson, M.: Sensitivity analysis of traffic equilibria. Transport. Sci. **38**(3), 258–281 (2004). https://doi.org/10.1287/trsc.1030.0043

29. Pigou, A.C.: The Economics of Welfare, 1st edn. Macmillan and Co., London (1920)

30. Rockafellar, R.T.: Convex Analysis. Princeton Landmarks in Mathematics. Princeton University Press, Princeton (1997). (Reprint of the 1970 original)

31. Roughgarden, T.: The price of anarchy is independent of the network topology. J. Comput. Syst. Sci. **67**(2), 341–364 (2003). https://doi.org/10.1016/S0022-0000(03)00044-8

32. Roughgarden, T., Tardos, E.: How bad is selfish routing? J. ACM **49**(2), 236–259 (2002). https://doi.org/10.1145/506147.506153

33. Roughgarden, T., Tardos, E.: Bounding the inefficiency of equilibria in nonatomic congestion games. Games Econ. Behav. **47**(2), 389–403 (2004). https://doi.org/10.1016/j.geb.2003.06.004

34. Shapiro, A.: Sensitivity analysis of nonlinear programs and differentiability properties of metric projections. SIAM J. Control. Optim. **26**(3), 628–645 (1988). https://doi.org/10.1137/0326037

35. Takalloo, M., Kwon, C.: Sensitivity of Wardrop equilibria: revisited. Optim. Lett. **14**(3), 781–796 (2020). https://doi.org/10.1007/s11590-020-01552-9

36. Tomlin, J.A.: Minimum-cost multicommodity network flows. Oper. Res. **14**(1), 45–51 (1966). https://doi.org/10.1287/opre.14.1.45

37. Wardrop, J.G.: Some theoretical aspects of road traffic research. Proc. Inst. Civil Eng. Part II **1**, 325–378 (1952). https://doi.org/10.1680/ipeds.1952.11259

38. Wu, Z., Möhring, R.: A sensitivity analysis for the price of anarchy in non-atomic congestion games. arXiv:2007.13979

39. Wu, Z., Möhring, R.H., Chen, Y., Xu, D.: Selfishness need not be bad. Oper. Res. **69**(2), 410–435 (2021). https://doi.org/10.1287/opre.2020.2036

40. Youn, H., Gastner, M.T., Jeong, H.: Price of anarchy in transportation networks: efficiency and optimality control. Phys. Rev. Lett. **101**(12), 128701 (2008)

41. Youn, H., Gastner, M.T., Jeong, H.: Erratum: Price of anarchy in transportation networks: Efficiency and optimality control. Phys. Rev. Lett. **102**, 049905 (2009). https://doi.org/10.1103/PhysRevLett.102.049905