

Methods

Network Pricing: How to Induce Optimal Flows Under Strategic Link Operators

 José Correa,^a Cristóbal Guzmán,^b Thanasis Lianas,^c Evdokia Nikolova,^d Marc Schröder^e

^a Department of Industrial Engineering, Universidad de Chile, Santiago, Chile; ^b Institute for Mathematical and Computational Engineering, Facultad de Matemáticas y Escuela de Ingeniería, Pontificia Universidad Católica de Chile, Santiago, Chile; ^c School of Electrical and Computer Engineering, National Technical University of Athens, Athens, Greece; ^d Department of Electrical and Computer Engineering, Cockrell School of Engineering, University of Texas at Austin, Austin, Texas 78712; ^e School of Business and Economics, Maastricht University, 6211 LM Maastricht, Netherlands

Contact: correa@uchile.cl,  <https://orcid.org/0000-0002-3012-7622> (JC); crguzmanp@mat.uc.cl,

 <https://orcid.org/0000-0002-1498-2055> (CG); lianas@corelab.ntua.gr (TL); nikolova@austin.utexas.edu,

 <https://orcid.org/0000-0002-8801-0320> (EN); marc.schroeder@oms.rwth-aachen.de,  <https://orcid.org/0000-0002-0048-2826> (MS)

Received: June 22, 2018

Revised: December 20, 2019; June 7, 2020

Accepted: June 26, 2020

Published Online in Articles in Advance: March 12, 2021

Area of Review: Optimization

<https://doi.org/10.1287/opre.2020.2067>

Copyright: © 2021 INFORMS

Abstract. Network pricing games provide a framework for modeling real-world settings with two types of strategic agents: operators of a network and users of the network. Operators of the network post a price so as to attract users and maximize profit; users of the network select routes based on these prices and congestion from other users. Motivated by the fact that equilibrium in these games may not exist, may not be unique, and may induce an inefficient network performance, our main result is to observe that a simple regulation on the network owners' market solves these three issues. Specifically, if an authority could set appropriate caps (upper bounds) on the tolls (prices) operators can charge, then the game among the link operators has a unique and strong Nash equilibrium and the users' game results in a Wardrop equilibrium that achieves the optimal total delay. We call any price vector with these properties a great set of tolls and investigate the efficiency of great tolls with respect to the users' surplus. We derive a bicriteria bound that compares the users' surplus under great tolls with the users' surplus under optimal tolls. Finally, we consider two different extensions of the model. First, we assume that operators face operating costs that depend on the amount of flow on the link, for which we prove existence of great tolls. Second, we allow operators to own more than one link. In this case, we prove that, when operators own complementary links (i.e., links for which an increase in toll value may only increase the flow on the other owned links), any toll vector that induces the optimal flow and that is upper bounded by the marginal tolls is a great set of tolls, and furthermore, we show that, when all links in the network are complementary, then the aforementioned toll vector is also a strong cap equilibrium.

Funding: This work was supported by the Hellenic Foundation for Research and Innovation (HFRI) under the "First Call for HFRI Research Projects to Support Faculty Members and Researchers and the Procurement of High-Cost Research Equipment Grant," project BALSAM [Grant HFRI-FM17-1424]; Agencia Nacional de Investigación y Desarrollo (ANID, Chile) [Grants FONDECYT 11160939, FONDECYT 1190043, AFB170001, and AFB180003]; the National Science Foundation [Grants CCF-1216103, CCF-1331863, CCF-1350823, and CCF-1733832]; and by the Millennium Institute for Research in Market Imperfections and Public Policy [Grants ICM IS130002, RC130003].

Keywords: network pricing games • competition regulation • price caps • selfish routing

1. Introduction

Network pricing games provide a framework for modeling real-world settings with two types of strategic agents: operators of a network and users of the network. Operators of the network post a price for usage of the links they own so as to attract users and maximize profit, and users of the network select routes based on these prices and congestion from other users. A landmark example of this type of interaction, which sets the ground in which we describe our work, occurs in transportation networks. Here, link owners

are road operators and may toll the usage of their road. Users are travelers that seek to minimize their travel time plus payments. The challenge in these games is that there are two levels of competition: one among the owners to attract users to their link so as to maximize profit and second among users of the network to select routes that are cheap yet not too congested.

In the absence of self-interested link owners, these games reduce to the well-studied network congestion games—a widely accepted and practically useful model for selfish resource allocation in transportation

and communication networks (see, e.g., Beckmann et al. 1956, Roughgarden 2005, Correa and Stier-Moses 2011, and references therein). In congestion games, self-interested users travel in a network from their origin to their destination with the goal of minimizing their own travel cost. The natural solution concept of the game is the so-called Wardrop equilibrium flow, under which all users route along the shortest paths given the strategic choices of other users. We assume that the total amount of traffic is dependent on the disutility the users experience—a model also known as congestion games with elastic demand.

Because selfish behavior usually drives systems to socially inefficient situations, a central authority is typically interested in optimizing the social welfare for the network users—a task that can be implemented by setting appropriate *marginal tolls* (Beckmann et al. 1956), which simply charge each user the negative externality the user imposes on the system. However, the goal of the link owners is to set tolls so as to attract users and maximize their own profit. Imagine a link owner increasing the price. Clearly, some of the users will move to alternative routes, thereby increasing the congestion on these routes and making them less attractive. This implies that link owners have an incentive to set a price that is higher than what is socially desired and thereby introduce new inefficiencies. For instance, under marginal tolls, some operators may want to increase their toll in order to make a higher profit. In this regard, an equilibrium for the link owners is a toll vector such that a change in a single toll does not increase the profit of the corresponding toll operator (under the implied user Wardrop equilibrium flow).

In this more complex game-theoretic environment, (i) an equilibrium may fail to exist (the only case in which an equilibrium is proved to exist is in parallel link networks with affine latency functions; see, e.g., Acemoglu and Ozdaglar 2007a, Hayrapetyan et al.

2007, Harks et al. 2019); (ii) an equilibrium might not be unique (see, e.g., Acemoglu and Ozdaglar 2007b, Harks et al. 2019); and (iii) the total delay of the equilibrium flow can be arbitrarily higher than the optimal delay, implying that the network can behave arbitrarily worse than the case in which tolls are completely absent (see, e.g., Acemoglu and Ozdaglar 2007b, Harks et al. 2019).

In the face of these challenges, we set out to find ways to mitigate the effect of selfish toll-operator behavior. We introduce competition regulation by allowing a regulator to set specific price caps on the toll values that each toll operator can set on a link. Setting different price caps for different operators is consistent with the practice in some privately operated networks of highways. For example, in Santiago de Chile, there are currently 12 different operators who set tolls on different urban highways as shown in Figure 1. The current regulation sets a price cap that is unique to each highway, and the toll operators are allowed to set tolls upper bounded by the caps (Gonzalez 2016). As it turns out, introducing such regulation can resolve all of the previously mentioned issues as there are caps for which the game has a unique (strong) Nash equilibrium under which the Wardrop equilibrium is optimal.

1.1. Contribution

Because the unregulated network pricing game behaves so poorly, in Section 3, we set out to study a simple mechanism that improves it. In particular, we investigate the regulatory policy of setting upper bounds (caps) on the toll values that each operator is allowed to set. Note that there is a tension between the toll charged by the operator and the amount of flow the operator will get. It is plausible for a toll operator to gain from decreasing a toll because the link will attract more flow, which may result in an

Figure 1. (Color online) Privately Operated Highways in Santiago de Chile



Note. Each highway has its own toll cap.

overall higher profit. For large enough caps, it indeed happens that it is optimal for the operator to set a toll upper bounded by the cap. Our main result, Theorem 1, shows that, when the central planner chooses the marginal tolls as caps, then the unique Nash equilibrium for the operators is to set the tolls equal to the caps, which is known to induce the optimal flow as a Wardrop equilibrium. In what follows, any toll vector that induces the optimal flow as a Wardrop equilibrium is called an optimal toll. Under this definition, the theorem shows something stronger: any optimal cap vector upper bounded by the marginal tolls when chosen as toll caps leads to a unique Nash equilibrium in which every operator charges precisely the cap. Moreover, we show this equilibrium is robust to coalitions, a concept known as *strong Nash equilibrium*. We show that, even though the flow on a given link is a decreasing function of the toll on that link (as we show in Lemma 1), the profit of the toll operator as a function of the toll the operator charges is an increasing function up to a certain point (Lemma 2), which is a key tool to prove Theorem 1.

Inspired by Theorem 1, we study *great tolls* in more detail in Section 4. These are optimal toll vectors and, when they are set as caps, are themselves the unique Nash equilibrium. It is easy to observe that, given an optimal flow, if, for every commodity, there are some users that do not travel, then great tolls are unique. However, there are also simple examples in which great tolls are not unique. A basic question to ask is how efficient great tolls are with respect to the users' surplus. Note that the users' total delay under great tolls is fixed because all great tolls are optimal; thus, our question is equivalent to that of finding great tolls that minimize the total payments. By the results in Section 3, any optimal toll vector that is upper bounded by the marginal toll vector is a great set of tolls. As a benchmark, we use the *minimum payment tolls* (MPT), defined as those tolls that, among the optimal ones, minimize the total payments. Minimum payment tolls, just as great tolls upper bounded by the marginal tolls, can be computed by a linear program. On the negative side, we give an example that shows the users' surplus under great tolls can be arbitrarily worse than the optimal users' surplus. On the positive side, we prove that the users' surplus under great tolls is at least as high as the optimal surplus if each user has no less than half the original valuation (Theorem 2).

In Section 5, we study two different extensions of our model. First, we consider the setting in which each of the operators faces operating costs for maintaining the link. We prove that the main result, the existence of great tolls, is still valid. However, if toll caps are too low or operating costs are too high, there might be links for which the profit in the Nash equilibrium is negative.

Second, we study the setting in which operators are allowed to own more than one link. We start by showing that, for instances in which players own complementary links, that is, links for which an increase of the toll value in one of them may only increase the flow on the other (complementary) links, any optimal tolls upper bounded by the marginal tolls are great tolls (Theorem 5). Intuitively, this holds because an operator that operates a single link has an incentive to use the upper bound as a toll, and this incentive remains if, by doing so, the operator only gains more flow on the operator's other links and, thus, only gains more profit. Then, we show that, when all the links are complementary, for example, in a parallel link network, such tolls are additionally a strong Nash equilibrium for the operators (Theorem 6). Finally, we show that Theorems 5 and 6 are essentially tight by providing two examples.

1.2. Related Work

Acemoglu and Ozdaglar (2007a) introduce a model of price competition between link operators in which each user has some fixed reservation value for travel. They show that increasing competition among operators from a monopoly to an oligopoly may cause a reduction in efficiency, measured as the difference between the users' willingness to pay and the delay, and provide a (tight) bound on efficiency in pure strategy equilibria. In a follow-up work, Acemoglu and Ozdaglar (2007b) generalize the preceding study to slightly more general topologies in which parallel paths with multiple links may replace the parallel links. They show that even this slight generalization can make the game arbitrarily inefficient when the efficiency is measured as mentioned.

Hayrapetyan et al. (2007) consider instances on parallel links in which the demand to be routed is elastic and decreases in a concave way as the cost for using the network increases. The social cost in that work is the sum of the players' profits plus a term that represents the utility gathered by the traffic that gets routed with a nice trade-off occurring between these two terms. For that game, they show that, in a network with parallel links and linear latencies, there is always a pure Nash equilibrium with the price of anarchy, that is, the measure for inefficiency, being bounded by a constant factor even when the latency functions are relaxed to be convex. For the case in which latencies have zero value under zero flow, they improve the constant.

Following that work, Ozdaglar (2008) studies the same model and manages to prove tight bounds on the efficiency of that game. Musacchio (2009) and Musacchio and Wu (2007) rederive and generalize those (upper) bounds for the case of series-parallel networks via a connection to electrical circuits; see

the survey by Ozdaglar and Srikant (2007) for further discussion. Johari et al. (2010) study an extension of network pricing games in which operators compete in prices and investments.

Harks et al. (2019) use a very similar policy to regulate competition between link operators. There, a regulator is able to set a unique price cap for all link operators. As it turns out, this restricts the regulator so that the induced network performance is not always optimal. For two-link parallel networks, this reduction in performance is characterized for different classes of latency function.

Our price competition model corresponds to Bertrand competition in a network setting (Dixon 2001, chapter 6). Under this setting Chawla et al. (2008) address questions regarding the price of anarchy and of stability with respect to two objectives: the social welfare of all the players (users and sellers) and the total profit obtained by all the sellers. Their work only considers capacity-based congestion (this corresponds to latency functions that are identically zero until capacity is reached and then jump to infinity), and no regulation is imposed on the game. Their results show a dependence of the price of anarchy or stability on the number of monopolistic links, namely the links whose removal disconnects an origin–destination pair. In contrast, our model does not suffer from monopolies: setting caps on prices prevents monopolistic links from charging arbitrarily large prices. Following the same model and focusing on the social welfare of the consumers as the objective, Chawla and Niu (2009) extend the results of Chawla et al. (2008).

An interesting and related model is raised by Anshelevich and Sekar (2015). They consider the edges of a network as goods, and each edge is owned by a different profit-maximizing seller. In the first stage of the game, sellers set prices for the use of their edges and have production costs depending on the level of use. In the second stage, the users of the network, that is, the buyers, choose origin-to-destination paths so that each of them maximizes the buyer's utility minus the payments to the sellers. The main difference with the model presented here is that users impose an externality on the sellers via the production costs and not on the other users as is the case in our model. This crucially affects the equilibrium pricing and the social welfare (e.g., for single commodity networks, if no monopolies are present, there is always an optimal equilibrium, which is not the case in our model).

In other related work, Papadimitriou and Valiant (2010) consider the case in which the routing is no longer selfish but is controlled by the edges of the network, and each edge either minimizes its average latency or announces a suitable price to its neighbors in order to maximize its profit. Caragiannis et al. (2017)

consider a model of buyers and sellers of a similar product, which, under some reformulation, can be seen as a variant of the parallel links model of Acemoglu and Ozdaglar (2007a) with heterogeneous buyers but constant latencies. Instead of minimizing traffic costs, maximizing the profit from tolls is considered in Kuiteing et al. (2016), Castelli et al. (2017), and Briest et al. (2012). There, a central authority or unique owner has control of all the toll-able edges, yet more importantly, the edge costs are constants rather than flow-dependent. Recently, Schmand et al. (2019) consider a two-stage game in which operators compete in investments so as to increase the bandwidth of a link to attract users.

The study of network congestion games in which a central operator is allowed to charge tolls in order to improve efficiency has a long history, starting with Beckmann et al. (1956). Cole et al. (2003) and Fleischer (2005) provide upper bounds on tolls that induce the optimal flow as an equilibrium, and Dial (1999) considers the objective of minimizing users' payments among optimal flow-inducing tolls. The study of network user games in which, for each link, a (potentially adversarially chosen) upper bound on the toll is present was first considered by Bonifaci et al. (2011) and later by Jelinek et al. (2014) and Fotakis et al. (2015). Results in these papers show that, when upper bounds are present, optimality cannot, in general, be achieved, yet on the positive side, algorithms are proposed under rather restrictive settings with provable guarantees regarding the efficiency of the network.

The transportation literature also addresses questions regarding the pricing of privately operated roads in transportation networks. A recent stream of papers focuses on the so-called build-operate-transfer scheme in which private investors build and operate roads at their own expense and, in return, receive the revenue from tolls charged within some years, after which the roads are transferred to the government. See, for example, Yang and Meng (2000), Yang et al. (2002), and Meng and Lu (2017).

2. Preliminaries

We study a network pricing game, in which non-atomic players, which we call *users*, selfishly minimize their cost (delays plus tolls) across a network; on top of this, each network link is operated by a different selfish agent that maximizes profit by charging tolls on users traversing the agent's link.

2.1. The Network Users' Game: Selfish Routing

Let $G = (V, E)$ be a network with V the set of nodes and E the set of directed edges/links of the network. We consider a multicommodity flow instance described by origin–destination node pairs $\{(o^k, d^k)\}_{k \in K}$ for a finite set of commodities K . For each commodity k ,

we assume that the total amount of traffic is dependent on the costs the traffic experiences; the higher the costs, the lower the traffic. We analyze the elastic traffic demand model as introduced by Beckmann et al. (1956) and later used for network pricing games by Hayrapetyan et al. (2007) and Ozdaglar (2008). We model elastic demand with a *utility function* $u^k : [0, r^k] \rightarrow \mathbb{R}_+$ for each $k \in K$, where $u^k(x)$ captures the reservation value for travel of the particle of the demand (i.e., the infinitesimally small user) at level x and $r^k \in \mathbb{R}_+$ is the maximum demand of commodity k . We assume that $u^k(\cdot)$ is nonincreasing and continuous for each $k \in K$ so that, in a sense, the users are ordered decreasingly according to their utility for traveling. Let $\mathbf{u} = (u^k)_{k \in K}$ be the vector of all utility functions. Define the aggregate utility function $U^k : [0, r^k] \rightarrow \mathbb{R}_+$ by $U^k(x) = \int_0^x u^k(y) dy$. By definition, this function is nondecreasing, concave, and continuously differentiable.

For each link $e \in E$, there is a latency function $\ell_e : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, that represents the delay experienced by users traversing this link as a function of the total flow on the link. We assume this function to be nondecreasing, convex, and smooth.

2.1.1. Paths and Flows. For each commodity $k \in K$, let \mathcal{P}^k denote the set of $o^k - d^k$ paths and let $\mathcal{P} = \cup_k \mathcal{P}^k$ be the union of all these paths. A *flow* for commodity k is a nonnegative vector $\mathbf{x}^k = (x_p^k)_{p \in \mathcal{P}^k}$ such that $\sum_{p \in \mathcal{P}^k} x_p^k \leq r^k$. For each commodity $k \in K$, let $r_x^k = \sum_{p \in \mathcal{P}^k} x_p^k$ be the amount of flow that is routed on the network by \mathbf{x}^k . A *flow* \mathbf{x} is a vector $(\mathbf{x}^k)_{k \in K}$, where each \mathbf{x}^k is a flow for commodity k . Let X denote the set of all flows. For a flow $\mathbf{x} \in X$ and $e \in E$, let $x_e^k = \sum_{p \in \mathcal{P}^k: e \in p} x_p^k$ be the amount of flow that \mathbf{x}^k routes on each link e and let $x_e = \sum_{k \in K} x_e^k$ be the amount of flow that \mathbf{x} routes on e . With a slight abuse of notation, we also denote $\mathbf{x} = (x_e)_{e \in E}$, the link-wise description of a flow.

2.1.2. Wardrop Equilibria and Optimal Flows. Given a flow $\mathbf{x} \in X$, the delay experienced on e is $\ell_e(x_e)$. In the case a toll $t_e \geq 0$ is charged for link usage, the combined cost of traversing e is $[\ell_e(x_e) + \alpha \cdot t_e]$, where $\alpha > 0$ represents the trade-off factor between delay and tolls and is assumed to be identical for all users. Without loss of generality (w.l.o.g.), we can assume α to be equal to one because we can always divide all tolls by α . A flow $\mathbf{x} \in X$ is a *Wardrop equilibrium* if all the routed traffic goes through shortest paths for the respective commodity, the utility for each traveling user is at least equal to the common shortest path cost of the user’s commodity, and the utility of each user not traveling is at most equal to the common shortest path cost of the user’s commodity. Formally, for every k , for every path $P \in \mathcal{P}^k$ with $x_P^k > 0$, and every path $P' \in \mathcal{P}^k$, $\sum_{e \in P} [\ell_e(x_e) + t_e] \leq \sum_{e \in P'} [\ell_e(x_e) + t_e]$ and

$\sum_{e \in P'} [\ell_e(x_e) + t_e] \leq u^k(r_x^k)$ with $r_x^k = r^k$ if $\sum_{e \in P} [\ell_e(x_e) + t_e] < u^k(r_x^k)$. Given such a flow, any $o^k - d^k$ path achieving the minimum end-to-end cost is called an *active path*, and any link that belongs to an active path is called an *active link*.

For any toll vector $\mathbf{t} \geq 0$, a Wardrop equilibrium exists; moreover, it minimizes the convex Wardrop potential $\Phi_{\mathbf{t}}(\mathbf{x})$ (Beckmann et al. 1956) and, thus, is given by

$$\mathbf{x}(\mathbf{t}) \triangleq \arg \max_{\mathbf{x} \in X} \left\{ \sum_{k \in K} U^k(r_x^k) - \sum_{e \in E} \int_0^{x_e} (\ell_e(y) + t_e) dy \right\}. \tag{1}$$

We restrict attention to instances in which the Wardrop equilibrium is unique for all $\mathbf{t} \geq 0$. In particular, this can be achieved by assuming that $\ell_e(x)$ is strictly increasing for all $e \in E$ or by assuming that there is no pair of paths with common endpoints and constant latency. In particular, $(x_e(\mathbf{t}))_{e \in E}$ is a well-defined function that is moreover continuous by Berge’s (1963) theorem.

Given a flow $\mathbf{x} \in X$ and toll vector \mathbf{t} , define the *total users’ cost* by $C(\mathbf{x}, \mathbf{t}) = \sum_{e \in E} [\ell_e(x_e) + t_e] x_e$ and the *users’ surplus* (or *consumer surplus*) by

$$CS(\mathbf{x}, \mathbf{t}, \mathbf{u}) = \sum_{k \in K} U^k(r_x^k) - C(\mathbf{x}, \mathbf{t}).$$

We define the *social welfare* by

$$SW(\mathbf{x}, \mathbf{u}) = \sum_{k \in K} U^k(r_x^k) - \sum_{e \in E} \ell_e(x_e) \cdot x_e.$$

Notice that the social welfare is simply the users’ surplus plus the toll operators’ profits, yet the tolls do not appear in the definition as they are transfers from users to toll operators.

An *optimal flow* $\mathbf{x}^*(\mathbf{u})$ is a flow that maximizes the social welfare with respect to (w.r.t.) \mathbf{u} . The classical result by Beckmann et al. (1956) shows that a flow $\mathbf{x}^*(\mathbf{u})$ is an optimal flow if and only if $\mathbf{x}^*(\mathbf{u}) = \mathbf{x}(\hat{\mathbf{t}}(\mathbf{u}))$, where $\hat{t}_e(\mathbf{u}) = x_e^*(\mathbf{u}) \ell'_e(x_e^*(\mathbf{u}))$: the toll vector $\hat{\mathbf{t}}$ is known as the marginal tolls. By our assumptions on the latency functions, optimal flows are unique. Any toll vector that induces the optimal flow is called optimal.

2.2. The Network Operators’ Game: Price Competition on Toll

In our model, additionally, every link $e \in E$ is operated by a different operator: these are the players of the price competition game. Each player e is allowed to charge a nonnegative toll t_e for the usage of the player’s link. Under the resulting toll vector \mathbf{t} , each link gets flow $x_e(\mathbf{t})$ according to the induced Wardrop equilibrium, and thus, the profit of player e is given by

$\pi_e(\mathbf{t}) \triangleq t_e x_e(\mathbf{t})$. We are interested in the equilibrium outcomes of this game.

2.2.1. Profit Functions. For each player $e \in E$, strategy is given by toll t_e , and profit is given by $\pi_e(t_e, \mathbf{t}_{-e}) = t_e x_e(t_e, \mathbf{t}_{-e})$, where we use the standard game-theoretic notation $\mathbf{t} = (t_e, \mathbf{t}_{-e})$.

2.2.2. Regulated Network Pricing Game and Nash Equilibria. The regulated network pricing game we consider is the following. A central planner may choose a cap vector $\bar{\mathbf{t}} \geq 0$ for tolls, and each player wants to maximize the player’s own profit under this constraint. We study two solution concepts: (pure) Nash equilibrium and strong (pure) Nash equilibrium. Tolls \mathbf{t} are a Nash equilibrium for the network pricing game if, for every $e \in E$, t_e is the best response of player e to \mathbf{t}_{-e} ; that is, we have $t_e \in \mathcal{BR}_e(\mathbf{t}_{-e})$, where the best response mapping $\mathcal{BR}_e(\mathbf{t}_{-e})$ is defined as

$$\mathcal{BR}_e(\mathbf{t}_{-e}) \triangleq \arg \max \{ \pi_e(s_e, \mathbf{t}_{-e}) : 0 \leq s_e \leq \bar{t}_e \}.$$

Tolls \mathbf{t} are a strong Nash equilibrium if there is no coalition that jointly decides for a deviation so that all operators in the coalition increase their profit (Aumann 1959). Formally, tolls \mathbf{t} are a strong Nash equilibrium if, for any set $E^\Delta \subseteq E$, there exists no \mathbf{t}' with $\{e \in E : t'_e \neq t_e\} \subseteq E^\Delta$ such that $\pi_e(\mathbf{t}') > \pi_e(\mathbf{t})$ for all $e \in E^\Delta$. Notice that any strong Nash equilibrium is a Nash equilibrium.

Definition 1 (Cap Equilibrium and Great Tolls). Given an instance of the profit maximization game, we say that a nonnegative vector $\bar{\mathbf{t}} = (\bar{t}_e)_{e \in E}$ is

- a. A (strong) cap equilibrium if, when restricting the strategy space for every player e to tolls $s_e \in [0, \bar{t}_e]$, then $(s_e)_{e \in E} = (\bar{t}_e)_{e \in E}$ is the unique (strong) Nash equilibrium.
- b. A great set of tolls if it is optimal (i.e., induces the optimal flow) and a cap equilibrium.

3. Regulated Network Pricing Game

In this section, we study the regulated profit-maximization game and some of its structural properties. We prove that the marginal tolls, when used as caps, are always a Nash equilibrium for the profit-maximization game, thus resolving the issues of equilibrium existence, uniqueness, and inefficiency raised in the literature. Our main result, presented in Theorem 1, strengthens this by showing that a potentially larger set of optimal tolls, when used as caps, leads to a unique Nash equilibrium, which is, furthermore, robust to coalitions.

We start the section by proving two monotonicity properties of the Wardrop flow as a function of tolls. These properties are related to the lower level game.

Lemma 1. Let $\mathbf{t}, \mathbf{t}' \geq 0$ be two toll vectors such that $\mathbf{t} \leq \mathbf{t}'$ and $E^< \triangleq \{e \in E : t_e < t'_e\}$ is nonempty. Then, there exist $e_1, e_2 \in E^<$ such that

- i. $x_{e_1}(\mathbf{t}') \leq x_{e_1}(\mathbf{t})$.
- ii. $[x_{e_2}(\mathbf{t}') - x_{e_2}(\mathbf{t})][\ell_{e_2}(x_{e_2}(\mathbf{t}')) + t'_{e_2} - \ell_{e_2}(x_{e_2}(\mathbf{t})) - t_{e_2}] \leq 0$.

Proof. To prove (i), we compare flows $\mathbf{x}(\mathbf{t})$ and $\mathbf{x}(\mathbf{t}')$ with respect to the Wardrop potentials. By the optimality of the Wardrop flow on its respective potential, we get the following inequalities:

$$\begin{aligned} & \sum_{k \in K} U^k(r_{x(\mathbf{t})}^k) - \sum_{e \in E} \int_0^{x_e(\mathbf{t})} (\ell_e(x) + t_e) dx \\ & \geq \sum_{k \in K} U^k(r_{x(\mathbf{t}')}^k) - \sum_{e \in E} \int_0^{x_e(\mathbf{t}')} (\ell_e(x) + t_e) dx \\ & \sum_{k \in K} U^k(r_{x(\mathbf{t}')}^k) - \sum_{e \in E} \int_0^{x_e(\mathbf{t}')} (\ell_e(x) + t'_e) dx \\ & \geq \sum_{k \in K} U^k(r_{x(\mathbf{t})}^k) - \sum_{e \in E} \int_0^{x_e(\mathbf{t}')} (\ell_e(x) + t'_e) dx. \end{aligned}$$

Combining these inequalities, we get

$$\sum_{e \in E^<} [t'_e - t_e][x_e(\mathbf{t}) - x_e(\mathbf{t}')] \geq 0,$$

and thus, because $E^<$ is nonempty, there must exist an $e_1 \in E^<$ such that $x_{e_1}(\mathbf{t}') \leq x_{e_1}(\mathbf{t})$, proving (i).

Let us now prove (ii). Notice that, because $\mathbf{x}(\mathbf{t}) + (\mathbf{x}(\mathbf{t}') - \mathbf{x}(\mathbf{t})) = \mathbf{x}(\mathbf{t}')$ is a feasible flow, then $\mathbf{x}(\mathbf{t}') - \mathbf{x}(\mathbf{t})$ is a feasible direction for $\mathbf{x}(\mathbf{t})$ in the space of feasible flows. By the first order optimality conditions of the Wardrop potential,

$$\begin{aligned} & \sum_{k \in K} \sum_{p \in \mathcal{P}^k} \left[U^k(r_{x(\mathbf{t})}^k) - \sum_{e \in P} [\ell_e(x_e(\mathbf{t})) + t_e] \right] \\ & \quad \times [x_p^k(\mathbf{t}') - x_p^k(\mathbf{t})] \leq 0. \end{aligned}$$

Analogously, $\mathbf{x}(\mathbf{t}) - \mathbf{x}(\mathbf{t}')$ is a feasible direction for $\mathbf{x}(\mathbf{t}')$, and thus,

$$\begin{aligned} & \sum_{k \in K} \sum_{p \in \mathcal{P}^k} \left[U^k(r_{x(\mathbf{t}')}^k) - \sum_{e \in P} [\ell_e(x_e(\mathbf{t}')) + t'_e] \right] \\ & \quad \times [x_p^k(\mathbf{t}) - x_p^k(\mathbf{t}')] \leq 0. \end{aligned}$$

Adding up these inequalities and using that, for any two flows \mathbf{y}, \mathbf{z} , we have

$$\sum_{k \in K} \sum_{p \in \mathcal{P}^k} \left[y_p^k \cdot \sum_{e \in P} (\ell_e(z_e) + t_e) \right] = \sum_{e \in E} [y_e \cdot (\ell_e(z_e) + t_e)],$$

we obtain

$$\begin{aligned} & \sum_{k \in K} \left[r_{x(t')}^k - r_{x(t)}^k \right] \left[u^k \left(r_{x(t')}^k \right) - u^k \left(r_{x(t)}^k \right) \right] \\ & + \sum_{e \in E^=} [x_e(t') - x_e(t)] [\ell_e(x_e(t')) - \ell_e(x_e(t))] \\ & + \sum_{e \in E^<} [x_e(t') - x_e(t)] [\ell_e(x_e(t')) + t'_e \\ & - \ell_e(x_e(t)) - t_e] \leq 0, \end{aligned}$$

where $E^= \triangleq \{e \in E : t_e = t'_e\}$. Observe now that the first and second summation terms are nonnegative as u^k is decreasing for all k and ℓ_e is noncreasing for all e . Thus, $\sum_{e \in E^<} [x_e(t') - x_e(t)] [\ell_e(x_e(t')) + t'_e - \ell_e(x_e(t)) - t_e] \leq 0$, implying that there exists $e_2 \in E^<$ such that $[x_{e_2}(t') - x_{e_2}(t)] [\ell_{e_2}(x_{e_2}(t')) + t'_{e_2} - \ell_{e_2}(x_{e_2}(t)) - t_{e_2}] \leq 0$, proving (ii). \square

The following result gives an intriguing inequality satisfied by any profit-maximizing toll. It says that in a (local) maximum, the toll is at least as high as the induced marginal costs of the users. The interpretation of this result is that firms have an incentive to increase prices above what is socially desired because of the congestion effects of the users. This property is a consequence of the first order optimality conditions, in combination with the monotonicity properties stated. We later see that this lemma is crucial for our main result.

Lemma 2. *Let $t \geq 0$ be a toll vector. If t_e is a local optimum for the profit-maximization problem (that is, for objective $\pi_e(\cdot, t_{-e})$), then $x_e(t) \cdot \ell'_e(x_e(t)) \leq t_e$.*

Proof. Notice first that, in the case $x_e(t) = 0$, the result obviously holds, and thus, we may restrict ourselves to the case $x_e(t) > 0$. By continuity of the induced Wardrop flow, $t_e = 0$ is not a local maximizer. Also by continuity and t_e being a local maximizer, there exists a $\delta > 0$ so that $0 < x_e(s_e, t_{-e}) < x_e(t)$ for any $s_e \in (t_e, t_e + \delta)$.

Now, because t_e is a local maximizer, we can use the first-order optimality conditions, $D^+[\pi_e(t_e, t_{-e})] \leq 0$, where $D^+[f(x)] \triangleq \limsup_{h \rightarrow 0^+} \frac{f(x+h) - f(x)}{h}$ is the upper right Dini derivative. Using the linearity of Dini derivatives, we get

$$x_e(t) + t_e D^+[x_e(t)] = D^+[\pi_e(t_e, t_{-e})] \leq 0. \quad (2)$$

On the other hand, because $x_e(s_e, t_{-e}) < x_e(t)$ for any $s_e \in (t_e, t_e + \delta)$, by Lemma 1(ii), we have that $[\ell_e(x_e(s_e, t_{-e})) + s_e - \ell_e(x_e(t)) - t_e] / [x_e(s_e, t_{-e}) - x_e(t)] \leq 0$. After rearranging terms, we get

$$\frac{\ell_e(x_e(s_e, t_{-e})) - \ell_e(x_e(t))}{x_e(s_e, t_{-e}) - x_e(t)} \leq -\frac{s_e - t_e}{x_e(s_e, t_{-e}) - x_e(t)}.$$

Taking $\limsup_{s_e \rightarrow t_e^+}$ in the expression, we get that the left-hand side converges to $\ell'_e(x_e(t))$, whereas the

right-hand side converges to $-(D^+[x_e(t)])^{-1}$. Because $D^+[x_e(t)] < 0$, we conclude that $\ell'_e(x_e(t)) \leq -1/D^+[x_e(t)]$. This, in combination with (2), gives

$$x_e(t) \ell'_e(x_e(t)) \leq -\frac{x_e(t)}{D^+[x_e(t)]} \leq t_e,$$

which proves the result. \square

Our main result shows a strong consequence of the preceding lemma. All optimal tolls upper bounded by the marginal tolls, when used as caps, lead to a unique (strong) Nash equilibrium and, thus, are great tolls.

Theorem 1. *Let \bar{t} be an optimal toll vector with $\bar{t} \leq \hat{t}$, where $\hat{t}_e(u) = x_e^*(u) \ell'_e(x_e^*(u))$. Then, \bar{t} is a (strong) cap equilibrium.*

Remark 1. Our proof shows a property that is stronger than a strong Nash equilibrium. In fact, what we prove is that, for any $E^\Delta \subseteq E$, there exists no t' with $\{e \in E : t'_e \neq t_e\} \subseteq E^\Delta$ such that $\pi_e(t') \geq \pi_e(t)$ for all $e \in E^\Delta$ with at least one strict inequality. This property clearly implies the classical notion of strong Nash equilibrium stated in Section 2, but we believe it might be of independent interest. The same stronger property is proved in Theorems 4 and 6.

Proof. First, we prove that \bar{t} is a strong Nash equilibrium. By way of contradiction, let E^Δ be a set of links and $t \leq \bar{t}$ with $e \in E^\Delta \triangleq \{e \in E : t_e < \bar{t}_e\} \subseteq E^\Delta$ such that $\pi_e(t) \geq \pi_e(\bar{t})$ for all $e \in E^\Delta$ with at least one strict inequality. Then, for all $e \in E^<$, we have $t_e x_e(t) \geq \bar{t}_e x_e(\bar{t})$, which implies $x_e(t) > x_e(\bar{t})$ and

$$t_e \geq \frac{\bar{t}_e x_e(\bar{t})}{x_e(t)}. \quad (3)$$

By Lemma 1(ii), there exists $e \in E^<$ such that $[x_e(\bar{t}) - x_e(t)] [\ell_e(x_e(\bar{t})) + \bar{t}_e - \ell_e(x_e(t)) - t_e] \leq 0$, which, in combination with $x_e(t) > x_e(\bar{t})$, gives

$$\ell_e(x_e(\bar{t})) + \bar{t}_e \geq \ell_e(x_e(t)) + t_e. \quad (4)$$

Let $e \in E^<$ be a link satisfying (4). We have the following inequalities:

$$\begin{aligned} \ell'_e(x_e(\bar{t})) (x_e(t) - x_e(\bar{t})) & \leq \ell_e(x_e(t)) - \ell_e(x_e(\bar{t})) \\ & \leq \bar{t}_e - t_e \leq \bar{t}_e \left(1 - \frac{x_e(\bar{t})}{x_e(t)} \right), \end{aligned} \quad (5)$$

where the first inequality follows from the convexity of ℓ_e , the second from (4), and the third from (3). Because $\bar{t}_e \leq \hat{t}_e = \ell'_e(x_e(\bar{t})) x_e(\bar{t})$ as \bar{t} is optimal, we obtain from (5) that

$$\ell'_e(x_e(\bar{t})) (x_e(t) - x_e(\bar{t})) \leq \ell'_e(x_e(\bar{t})) x_e(\bar{t}) \left(\frac{x_e(t) - x_e(\bar{t})}{x_e(t)} \right).$$

Because $\ell'_e(x_e(\bar{t})) > 0$ and $x_e(t) > x_e(\bar{t})$, we conclude that $x_e(t) \leq x_e(\bar{t})$, a contradiction.

Now, we show that there is a unique Nash equilibrium. By way of contradiction, suppose there exists another Nash equilibrium $t \leq \bar{t}$ (with at least one strict inequality) for the profit-maximization game with caps \bar{t} . Because all players are playing their best response, we can use Lemma 2, concluding that $x_e(t)\ell'_e(x_e(t)) \leq t_e$ for all $e \in E^<$ (which is nonempty by assumption), which, in turn, gives

$$x_e(t)\ell'_e(x_e(t)) \leq t_e < \bar{t}_e \leq \hat{t}_e = x_e^* \ell'_e(x_e^*),$$

concluding that $x_e(t) < x_e^* = x_e(\bar{t})$ for all $e \in E^<$. But, from Lemma 1(i), there exists $e_1 \in E^<$ such that $x_{e_1}(t) \geq x_{e_1}(\bar{t})$, a contradiction. \square

Remark 2. If, for an optimal flow, we have that, for every commodity, there are some users that travel and some users that do not travel, then marginal tolls are the unique optimal toll vector and, thus, the unique toll vector that satisfies the condition of Theorem 1. Otherwise, there might be multiple optimal toll vectors. We refer to Appendix A for a related discussion on the linear program MPT, whose feasible set is all optimal toll vectors, and its solution minimizes the users' surplus at equilibrium.

4. Maximizing Consumer Surplus


We now consider the question of how efficient great tolls are with respect to consumer surplus. The main result in this section, Theorem 2, only applies to single-source, single-sink networks, and thus, we restrict ourselves to these instances and omit any dependence on $k \in K$.

In what follows, we compare the maximum consumer surplus of an optimal toll vector upper bounded by the marginal tolls (this is a great set of tolls by Theorem 1) to the maximum consumer surplus of any optimal toll vector. Among all optimal toll vectors, let t^M be the one that maximizes consumer surplus, and let t^B be the one that maximizes consumer surplus restricted to being upper bounded by the marginal tolls. Both these toll vectors can be calculated by means of a linear program, that is, linear programs Below Marginal Toll (BMT) and Minimum Payment Toll (MPT) in Appendix A for t^B and t^M , respectively, and in general, they may differ. The following example shows that the consumer surplus under t^B can be arbitrarily lower than the consumer surplus under t^M .

Example 1. Consider the Braess network of Figure 2 with r units of flow to be routed and $u(x) = 2r$ for $x \in [0, r]$.

An optimal flow in this case is to split the demand into half on the upper and half on the lower path; in

Figure 2. (Color online) The Network of Example 1



particular, no flow traverses the middle uv link. This way, marginal tolls are given by $(\hat{t}_{ou}, \hat{t}_{uv}, \hat{t}_{vd}) = (r/2, 0, r/2)$, and notice that this is the only feasible toll vector that satisfies the conditions of Theorem 1, achieving a value of $2r^2 - 2r^2 = 0$. On the other hand, it is easy to see that the consumer surplus-maximizing optimal toll vector just needs to assign a sufficiently large value for t_{uv} (more precisely, $t_{uv} \geq r/2$), and the rest of the tolls can be zero; therefore, the optimal consumer surplus is $2r^2 - 3r^2/2 = r^2/2$.

The main result of this section is a comparison in the spirit of the bicriteria bound of Roughgarden and Tardos (2002): by how much should we decrease the utility function in order to induce the same level of consumer surplus when comparing t^B to t^M ? A consequence of this result is that the consumers' surplus under t^B is at least as much as the optimal surplus if each of them had half the valuation.

For particular values of the parameter $\beta \in [1, \infty)$ in the next theorem, we refer to Theorems A.1 and A.2. For the proof of Theorem 2 we need the following lemma. The proof of Lemma 3 can be found in Appendix A.

Lemma 3. Let \tilde{u} and u be utility functions so that $\tilde{u}(x) = \alpha \cdot u(x)$, where $0 < \alpha \leq 1$. Then,

- i. $r_{x^*(u)} \geq r_{x^*(\tilde{u})}$.
- ii. $CS(x^*(u), t^B, u) \geq CS(x^*(\tilde{u}), t^B, \tilde{u})$.
- iii. $CS(x^*(u), t^M, u) \geq CS(x^*(\tilde{u}), t^M, \tilde{u})$.

Theorem 2. Let $C(x^*(u), t^B) \leq \beta \cdot C(x^*(u), t^M)$ for some $\beta \in [1, \infty)$. Then,

$$CS(x^*(u), t^B, u) \geq CS(x^*(\tilde{u}), t^M, \tilde{u}),$$

where $\tilde{u}(x) = \beta/(2\beta - 1) \cdot u(x)$.

Proof. By the optimal flow characterization of Beckmann et al. (1956), we have that $r_{x^*(u)} < r$ for all u with $C(x^*(u), \hat{t}) > r \cdot u(r)$ and $r_{x^*(u)} = r$ for all u with $C(x^*(u), \hat{t}) \leq r \cdot u(r)$.

We consider the following three cases: (1) $u(r) < C(x^*(u), \hat{t})/r$, (2) $C(x^*(u), \hat{t})/r \leq u(r) < (2\beta - 1)/\beta \cdot C(x^*(u), \hat{t})/r$, and (3) $u(r) \geq (2\beta - 1)/\beta \cdot C(x^*(u), \hat{t})/r$.

Case 1. Assume that $u(r) < C(x^*(u), \hat{t})/r$. Then, $r_{x^*(u)} < r$, and thus,

$$\begin{aligned} CS(x^*(u), t^B, u) &= CS(x^*(u), t^M, u) \\ &\geq CS(x^*(\tilde{u}), t^M, \tilde{u}), \end{aligned}$$

where the equality follows because $CS(x^*(u), t, u) = U(r_{x^*(u)}) - r_{x^*(u)} \cdot u(r_{x^*(u)})$ for all optimal t and the inequality by Lemma 3.

Case 2. Assume that $C(x^*(u), \hat{t})/r \leq u(r) < (2\beta - 1)/\beta \cdot C(x^*(u), \hat{t})/r$. Then, $r_{x^*(u)} = r$, and because $\tilde{u}(r) = \beta/(2\beta - 1) \cdot u(r) < C(x^*(u), \hat{t})/r$, we have $r_{x^*(\tilde{u})} < r$. Thus,

$$\begin{aligned} CS(x^*(u), t^B, u) &\geq CS(x^*(\tilde{u}), t^B, \tilde{u}) \\ &= CS(x^*(\tilde{u}), t^M, \tilde{u}), \end{aligned}$$

where the inequality follows by Lemma 3 and the equality because $CS(x^*(u), t, u) = U(r_{x^*(u)}) - r_{x^*(u)} \cdot u(r_{x^*(u)})$ for all optimal t .

Case 3. Assume that $u(r) \geq (2\beta - 1)/\beta \cdot C(x^*(u), \hat{t})/r$. Then, $r_{x^*(u)} = r$, and because $\tilde{u}(r) = \beta/(2\beta - 1) \cdot u(r) \geq C(x^*(u), \hat{t})/r$, we have $r_{x^*(\tilde{u})} = r$. Thus,

$$\begin{aligned} &C(x^*(u), t^B) - C(x^*(u), t^M) \\ &\leq \left(1 - \frac{1}{\beta}\right) \cdot C(x^*(u), t^B) \\ &\leq \left(1 - \frac{1}{\beta}\right) \cdot C(x^*(u), \hat{t}) \\ &\leq r \cdot \frac{\beta - 1}{2\beta - 1} \cdot u(r) - \int_0^r u(x)dx + \int_0^r u(x)dx \\ &\leq - \int_0^r \frac{\beta}{2\beta - 1} \cdot u(x)dx + \int_0^r u(x)dx \\ &= - \int_0^r \tilde{u}(x)dx + \int_0^r u(x)dx = U(r) - \tilde{U}(r), \end{aligned}$$


where the first inequality follows from $C(x^*(u), t^B) \leq \beta \cdot C(x^*(u), t^M)$, which is true by assumption; the second from $C(x^*(u), t^B) \leq C(x^*(u), \hat{t})$; the third from $\beta/(2\beta - 1) \cdot u(r) \geq C(x^*(u), \hat{t})/r$; and the fourth from $r \cdot u(r) \leq \int_0^r u(x)dx$ and the definition of $\tilde{u}(x)$. Rearranging terms yields

$$\begin{aligned} CS(x^*(u), t^B, u) &\geq \tilde{U}(r) - C(x^*(u), t^M) \\ &= \tilde{U}(r) - C(x^*(\tilde{u}), t^M) \\ &= CS(x^*(\tilde{u}), t^M, \tilde{u}), \end{aligned}$$

as needed. \square

Remark 3. Theorem 2 cannot be extended to multi-commodity networks. Consider, for example, the network of Figure 3. Commodities $o - d_1$, $o - d_2$ have a maximum demand of one. If $u^1(x) = 8$ for $x \in [0, 1]$ and $u^2(x) = 4$ for $x \in [0, 1]$, then the optimal flow is 3/4 for

Figure 3. (Color online) The Network of Remark 3



commodity 1 and one for commodity 2. The unique optimal toll vector is $(3/4, 7/4)$, inducing a consumer surplus of 1/2.

If $u^1(x) = 4$ for $x \in [0, 1]$ and $u^2(x) = 2$ for $x \in [0, 1]$, then the optimal flow is zero for commodity 1 and one for commodity 2. The unique optimal toll vector is $(0, 0)$, inducing a consumer surplus of one.

5. Extensions

In this last section, we consider two different extensions of our model. First, we study an extension in which operators face operating costs. The main result states that any optimal toll vector upper bounded by the marginal tolls is a great set of tolls with a unique equilibrium. Second, we investigate the setting in which operators own multiple links. We show that marginal tolls induce a unique Nash equilibrium as long as the links owned by each operator are complementary.

5.1. When Links Have Operating Costs

5.1.1. Extra Preliminaries. In this case, each operator $e \in E$ is allowed to charge a nonnegative toll t_e for its usage while additionally facing operating costs as a function of the amount of flow on the link. We assume that this (cost) function $c_e : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is nondecreasing, convex, and smooth. Convexity of operating costs reflects that, with higher traffic, there is a higher chance for accidents and correspondingly higher costs for clean up and maintenance for which the toll operator is responsible. Under the resulting toll vector t , each link gets flow $x_e(t)$ according to the induced Wardrop equilibrium, and thus, the profit of player e is given by $\pi_e(t) \triangleq t_e x_e(t) - c_e(x_e(t))$.

5.1.1.1. Profit Functions. For each player $e \in E$, the player's strategy is given by toll $t_e \geq 0$, and the player's profit is given by $\pi_e(t_e, t_{-e}) = t_e x_e(t_e, t_{-e}) - c_e(x_e(t))$.

5.1.1.2. Optimal Flows. Given a flow x , the social welfare is defined by

$$SW(x, u) = \sum_{k \in K} U^k(r_x^k) - \sum_{e \in E} \ell_e(x_e) \cdot x_e - \sum_{e \in E} c_e(x_e).$$

An optimal flow $x^*(u)$ is a flow that maximizes the social welfare w.r.t. u . The vector of marginal tolls \hat{t} , now defined as $\hat{t}_e = x_e^*(u) \ell'_e(x_e^*(u)) + c'_e(x_e^*(u))$, induces the optimal flow, that is, $x(\hat{t}) = x^*(u)$.

5.1.2. Results. By the next theorem, even with operating costs, the existence of great tolls remains guaranteed. The proofs of the two theorems are along the lines of the proof of Theorem 1 (their analogue in the basic model) and can be found in Appendix B.

Theorem 3. Let \bar{t} be an optimal toll vector with $\bar{t} \leq \hat{t}$. Then, \bar{t} is a cap equilibrium.

Remark 4. Theorem 3 does not guarantee that each provider earns a positive profit. If the cap is too low or the operating costs are too high, the profit of a provider might be negative. Notice that it is easy to check whether a given cap induces a negative profit as tolls and flows are known in the equilibrium.

Theorem 4. Let $\bar{t} \leq \hat{t}$ be an optimal toll vector. If $\pi_e(\bar{t}) \geq 0$ for all $e \in E$ and operating costs are affine, that is, $c_e(x) = a_e \cdot x_e + b_e$ with $a_e \geq 0$ for all $e \in E$, then \bar{t} is a strong cap equilibrium.

Remark 5. If operating costs are linear, that is, $c_e(x) = a_e \cdot x_e$ with $a_e \geq 0$ for all $e \in E$, we have, by definition of \hat{t} , that $\pi_e(\hat{t}) \geq 0$ for all $e \in E$. Thus, marginal tolls are a strong cap equilibrium if operating costs are linear.

5.2. Allowing Multiple Links per Operator

5.2.1. Extra Preliminaries. In the case of multiple links per operator, the network users’ game remains the same, but we have to redefine some notions for the network operators’ game. For every player $i \in [n]$, where n is the number of players, player i owns a subset $E_i \subseteq E$ and is allowed to charge a nonnegative toll t_e to every $e \in E_i$. We assume that E_i and E_j are disjoint for all $i, j \in [n]$ with $i \neq j$. Under the resulting toll vector t , each link gets flow $x_e(t)$ according to the induced Wardrop equilibrium, and the profit from link e is $\pi_e(t) := t_e x_e(t)$. The profit of player i is given by $\pi_i(t) := \sum_{e \in E_i} \pi_e(t)$.

5.2.1.1. Profit Functions. For each player $i \in [n]$, the player’s strategy is given by toll vector $t_i = (t_e)_{e \in E_i}$, and the player’s profit, as a function of the player’s strategy, is given by $\pi_i(t_i, t_{-i}) = \sum_{e \in E_i} \pi_e(t_i, t_{-i}) = \sum_{e \in E_i} t_e x_e(t_i, t_{-i})$. We also need the profit function of a link e when tolls on other links are fixed according to some vector t_{-e} , defined as $\pi_e(t_e, t_{-e})$. Whenever t_{-e} is clear from context, we simply write $x_e(t_e)$ and $\pi_e(t_e)$. Note that all these profit functions are continuous because, for every toll vector t , the flow function $x(t)$ is continuous.

The following definition is important for the analysis when operators own multiple links.

Definition 2. A set of links $E' \subseteq E$ is called complementary if, for all $e \in E'$, all $0 \leq t_e \leq t'_e$, and all $t_{-e} \geq 0$, it is $x_{e'}(t_e, t_{-e}) \leq x_{e'}(t'_e, t_{-e})$ for all $e' \in E'$ with $e' \neq e$.

5.2.1.2. Series-Parallel Graphs. A directed $o - d$ multigraph is *series parallel* if it consists of a single link (o, d) or from two series-parallel graphs with terminals (o_1, d_1) and (o_2, d_2) composed either in series or in parallel. In a *series composition*, d_1 is identified with o_2 , o_1 becomes o , and d_2 becomes d . In a *parallel composition*, o_1 is identified with o_2 and becomes o , and d_1 is identified with d_2 and becomes d . Any series-parallel graph has a decomposition tree that reveals all the “building blocks” of the graph; that is, every parent is a series or a parallel combination of its children.

The following proposition shows that a set of links whose pairs cannot belong to the same path in a series-parallel graph is complementary. The proof lies in Appendix B.

Proposition 1. Let G be series parallel and $E' \subseteq E$. If, for all $e, e' \in E'$ and all paths $P : e \in P$, we have $e' \notin P$, then E' is complementary.

Note that this proposition implies that, in parallel-link networks, any set $E' \subseteq E$ is complementary.

The following example shows that two links of different paths need not be complementary.


Example 2. Consider the graph of Figure 4 with two units of flow to be routed from o to d and assume that the utilities are sufficiently large so that all users travel. If $t_1 = 0$ and $t_2 = 0$, then one unit of flow is routed on the zigzag path, and one unit of flow is routed on the direct path from o to d . If $t_1 = 1$ and $t_2 = 0$, then $3/5$ units of flow are routed on the upper and lower paths, and $4/5$ units of flow are routed on the direct path from o to d . Hence, by increasing t_1 , the flow on the direct path is reduced.

5.2.2. Results. The existence of great tolls is guaranteed if all players own complementary links.

Theorem 5. Let \mathcal{G} be an instance in which each player owns a set of complementary links and let $\bar{t} \leq \hat{t}$ be an optimal toll vector. Then, \bar{t} is a cap equilibrium.

The proof can be found in Appendix B and is quite technical. The main idea of the proof is that an operator

Figure 4. (Color online) Noncomplementary Links



that operates a single link has an incentive to use the upper bound as a toll, and this incentive remains if, by doing so, the operator only gains more flow on the operator’s other links and, thus, only gains more profit.

For series-parallel graphs, we get the following corollary. For its proof, one simply has to combine Proposition 1, which yields that players own complementary links, with Theorem 5.


Corollary 1. *Let \mathcal{G} be an instance in which each player owns a set of links that in pairs cannot belong to the same path and let $\bar{t} \leq \hat{t}$ be an optimal toll vector. Then, \bar{t} is a cap equilibrium.*

Theorem 5 is tight in the sense that, for instances in which some player operates noncomplementary links, there exist optimal tolls upper bounded by the marginal ones and such that, when they are set as caps, they are not Nash equilibria. An example follows. Note that the network is series parallel, which also makes Corollary 1 tight.

Example 3. Consider the series-parallel graph of Figure 5 with one unit of flow to be routed, and assume that the utilities are sufficiently large so that all users travel. Let player 1 operate links e_1 and e_3 and the other links be operated by other players, one link per player. It is easy to check that links e_1 and e_3 are noncomplementary; for example, increasing the toll on e_1 may decrease the flow on e_3 . In this instance, the optimal flow x^* routes $x_1^* = 1/6$ units through e_1 , $x_2^* = 1/6$ units through e_2 , $x_3^* = 1/3$ units through e_3 , and $x_4^* = 2/3$ units through e_4 . Thus, the vector of marginal tolls is $\hat{t} = (1/6, 1/6, 1/3, 0)$. If all players play the marginal tolls, then player 1 has profit equal to $(1/6)^2 + (1/3)^2 = 5/36$. But playing marginal tolls is not a best response for player 1 as player 1 can play $t_1 = 1/8 < \hat{t}_1$ and $t_3 = 1/3 = \hat{t}_3$, getting a profit of $121/864 > 5/36$ (the equilibrium flow under $t = (1/8, 1/6, 1/3, 0)$ is $x(t) = (7/36, 11/72, 25/72, 47/72)$). Thus, when \hat{t} are set as caps, they are not Nash equilibria.

It is worth noting that, on instances in which players operate complementary links, \bar{t} is not a strong Nash equilibrium (despite being a Nash equilibrium). Example 4 has such an instance.

Figure 5. (Color online) The Single-Commodity Series-Parallel Network Used in Example 3



Example 4. Consider the series-parallel graph of Figure 6 with two units of flow to be routed, and assume that utilities are high enough so that all users travel. Let player 1 operate links e_1 and e_7 , player 2 operate links e_3 and e_5 , and the other links be operated by other players, one link per player. It is easy to check that links e_1 and e_7 are complementary, and the same holds for links e_3 and e_5 . Note that Figure 6 is essentially two copies of Figure 5 connected in parallel. Based on that, one can derive that the optimal flow x^* is such that $(x_1^*, x_2^*, x_3^*, x_4^*) = (1/6, 1/6, 1/3, 0)$ and $(x_5^*, x_6^*, x_7^*, x_8^*) = (1/6, 1/6, 1/3, 0)$, and the marginal tolls for links e_1 through e_8 are $(\hat{t}_1, \hat{t}_2, \hat{t}_3, \hat{t}_4) = (1/6, 1/6, 1/3, 0)$ and $(\hat{t}_5, \hat{t}_6, \hat{t}_7, \hat{t}_8) = (1/6, 1/6, 1/3, 0)$. If all players play the marginal tolls, then players 1 and 2 both have profit equal to $(1/6)^2 + (1/3)^2 = 5/36$. But players 1 and 3 can form a coalition and together deviate and play $t_1 = 1/8 < \hat{t}_1$ and $t_3 = 1/3 = \hat{t}_3$ and $t_5 = 1/8 < \hat{t}_5$ and $t_7 = 1/3 = \hat{t}_7$, getting this way a profit of $121/864 > 5/36$ each. Thus, when \hat{t} are set as caps, they are not a strong Nash equilibrium.


Yet we have the following theorem. Its proof can be found in Appendix B.

Theorem 6. *Let \mathcal{G} be an instance in which all links are complementary links, and let $\bar{t} \leq \hat{t}$ be an optimal toll vector. Then, \bar{t} is a strong cap equilibrium.*

6. Conclusion

In this work, we propose a simple regulation policy for network owners in network pricing games. Network owners have an incentive to set prices that are higher than their marginal externality, yielding inefficiencies for the network users. Moreover, Nash equilibria in the network owners’ market need not even exist. We show that, if a central authority regulates the market by setting the appropriate upper bounds on the prices, all of these issues get resolved. In particular, the well-known marginal tolls as caps have the property that the network owners’ market has a unique (strong) Nash equilibrium and induces a socially optimal flow. In general, every toll vector that

Figure 6. (Color online) The Single-Commodity Series-Parallel Network Used in Example 4



is upper bounded by marginal tolls and induces the optimal flow results as the unique Nash equilibrium when being imposed as caps.

Given that optimal social welfare can be achieved by imposing great tolls as caps, we asked the question of how to fairly distribute the social welfare among the network owners and users. We study the efficiency of great tolls that are upper bounded by marginal tolls from the perspective of the network users. Another approach could be to study the efficiency of great tolls from the perspective of the network owners. In order to do so, a complete characterization of great tolls would be needed.

The main result, that marginal tolls are great tolls, applies to the setting in which each link owners owns only one link in the network. The result extends to the setting in which link owners own complementary links (i.e., links for which an increase in toll value may only increase the flow on the other owned links). In general, the result that marginal tolls are great tolls breaks down if network owners own non-complementary links, for example, multiple links on the same path. This implies that market regulation works well in markets with a high degree of competition, whereas it is less effective in markets with less competition.

Acknowledgments

A preliminary version of this paper was accepted for publication in the *Proceedings of the 19th ACM Conference on Economics and Computation*. Part of the work was done while a subset of the authors were at the Simons Institute for the Theory of Computing in Berkeley, California.

Appendix A. Missing Proofs of Section 4

We start with the proof of Lemma 3 that we used to prove Theorem 2.

Proof of Lemma 3. We first note that, by the definition of the total users' cost $C(x(t), t)$, for an equilibrium $x(t)$ under tolls t , the common users' path cost equals $C(x(t), t)/r_x$.

To prove (i), we can assume that $r_{x^*(u)} < r$ as otherwise the result follows trivially because $r \geq r_{x^*(\tilde{u})}$. By recalling that the optimal flow is a Wardrop equilibrium w.r.t. the marginal latency functions, $r_{x^*(u)} < r$ implies that the common users' path cost at that equilibrium is equal to $u(r_{x^*(u)})$, that is, $C(x^*(u), \hat{t}(u))/r_{x^*(u)} = u(r_{x^*(u)})$.

To reach a contradiction, let $r_{x^*(u)} < r_{x^*(\tilde{u})}$. Hall (1978) and Lin et al. (2004) show that the common path cost of a Wardrop equilibrium is nondecreasing in r , and additionally, the common path cost is strictly increasing in r when the latencies are strictly increasing. Because an optimal flow is a Wardrop equilibrium w.r.t. the marginal latency functions and the latency functions are strictly increasing, we have $C(x^*(u), \hat{t}(u))/r_{x^*(u)} < C(x^*(\tilde{u}), \hat{t}(\tilde{u}))/r_{x^*(\tilde{u})}$. Additionally, by the definition of equilibrium, $C(x^*(\tilde{u}), \hat{t}(\tilde{u}))/r_{x^*(\tilde{u})} \leq \tilde{u}(r_{x^*(\tilde{u})})$. Putting it

all together and using the hypothesis for $u(x)$ and $\tilde{u}(x)$ for the first inequality, we get

$$\begin{aligned} \tilde{u}(r_{x^*(u)}) &\leq u(r_{x^*(u)}) = C(x^*(u), \hat{t}(u))/r_{x^*(u)} \\ &< C(x^*(\tilde{u}), \hat{t}(\tilde{u}))/r_{x^*(\tilde{u})} \\ &\leq \tilde{u}(r_{x^*(\tilde{u})}), \end{aligned}$$

a contradiction because $\tilde{u}(x)$ is nonincreasing.

To prove (ii) and (iii), we consider the following three cases: (1) $r_{x^*(u)} < r$ and $r_{x^*(\tilde{u})} < r$, (2) $r_{x^*(u)} = r$ and $r_{x^*(\tilde{u})} < r$, and (3) $r_{x^*(u)} = r$ and $r_{x^*(\tilde{u})} = r$. We only give the proof of (ii); the proof of (iii) follows analogously, simply by changing t^B with t^M .

Case A.1. Assume that $r_{x^*(u)} < r$ and $r_{x^*(\tilde{u})} < r$. Then, by definition of a Wardrop equilibrium, $CS(x^*(u), t^B, u) = U(r_{x^*(u)}) - r_{x^*(u)} \cdot u(r_{x^*(u)})$ and $CS(x^*(\tilde{u}), t^B, \tilde{u}) = \tilde{U}(r_{x^*(\tilde{u})}) - r_{x^*(\tilde{u})} \cdot \tilde{u}(r_{x^*(\tilde{u})})$. Thus, we have

$$\begin{aligned} CS(x^*(u), t^B, u) &= \int_{x=0}^{r_{x^*(u)}} [u(x) - u(r_{x^*(u)})] dx \\ &= \int_{x=0}^{r_{x^*(u)}} \frac{1}{\alpha} [\tilde{u}(x) - \tilde{u}(r_{x^*(u)})] dx \\ &\geq \int_{x=0}^{r_{x^*(\tilde{u})}} [\tilde{u}(x) - \tilde{u}(r_{x^*(\tilde{u})})] dx \\ &= CS(x^*(\tilde{u}), t^B, \tilde{u}), \end{aligned}$$

where, for the second equality, we use the definition of $\tilde{u}(x)$, and for the inequality, we use $\alpha \in (0, 1]$, $r_{x^*(u)} \geq r_{x^*(\tilde{u})}$, and $\tilde{u}(r_{x^*(u)}) \leq \tilde{u}(r_{x^*(\tilde{u})})$.

Case A.2. Assume that $r_{x^*(u)} = r$ and $r_{x^*(\tilde{u})} < r$. Then, by definition of a Wardrop equilibrium, $CS(x^*(\tilde{u}), t^B, \tilde{u}) = \tilde{U}(r_{x^*(\tilde{u})}) - r_{x^*(\tilde{u})} \cdot \tilde{u}(r_{x^*(\tilde{u})})$. Thus, we have

$$\begin{aligned} CS(x^*(u), t^B, u) &= U(r) - C(x^*(u), t^B) \\ &\geq \int_{x=0}^r [u(x) - u(r)] dx \\ &= \int_{x=0}^r \frac{1}{\alpha} [\tilde{u}(x) - \tilde{u}(r)] dx \\ &\geq \int_{x=0}^{r_{x^*(\tilde{u})}} [\tilde{u}(x) - \tilde{u}(r_{x^*(\tilde{u})})] dx \\ &= CS(x^*(\tilde{u}), t^B, \tilde{u}), \end{aligned}$$

where the first inequality follows from $C(x^*(u), t^B) \leq C(x^*(u), \hat{t}(u)) \leq r \cdot u(r)$ because of all the demand being routed under $x^*(u)$ and the constraint $v_t - v_s \leq u(r)$ in (BMT2), and the second inequality follows from $\alpha \in (0, 1]$, $r > r_{x^*(\tilde{u})}$, and $\tilde{u}(r) \leq \tilde{u}(r_{x^*(\tilde{u})})$.

Case A.3. Assume that $r_{x^*(u)} = r$ and $r_{x^*(\tilde{u})} = r$. Then, we have

$$\begin{aligned} CS(x^*(u), t^B, u) &= U(r) - C(x^*(u), t^B) \\ &\geq \tilde{U}(r) - C(x^*(\tilde{u}), t^B) \\ &= CS(x^*(\tilde{u}), t^B, \tilde{u}), \end{aligned}$$

where the inequality follows from $u(x) \geq \tilde{u}(x)$ for all $x \in [0, r]$, which implies $U(r) \geq \tilde{U}(r)$, and $x^*(u) = x^*(\tilde{u}) = x^*$ because all the flow is routed by the optimal solution under either utility function. \square

Theorem A.1, Corollary A.1, and Theorem A.2, presented next, can be used to get an upper bound on β in Theorem 2. The family of Braess graphs (see, e.g., Roughgarden 2006, Kleer and Schäfer 2016) can be used to show that these bounds are tight.

Theorem A.1 considers an approximation bound based on properties of the latency functions. The result can be interpreted as follows: if the sensitivity of links to changes in the flows is bounded by a factor, then the performance of the proposed solution is also bounded by that factor (plus one) times the value of the optimal solution.

Theorem A.1. *Suppose all latency functions ℓ in the profit-maximization game satisfy $\sup_{x \geq 0} \frac{x\ell'(x)}{\ell(x)} \leq \gamma$. Then, $C(x^*(u), t^B) \leq \beta \cdot C(x^*(u), t^M)$, where $\beta = \gamma + 1$.*

Proof. By the assumption on latency functions,

$$\begin{aligned} C(x^*(u), t^B) &\leq \sum_e [\ell_e(x_e^*(u)) + x_e^*(u)\ell'_e(x_e^*(u))] \\ &\quad \times x_e^*(u) \leq \sum_e (1 + \gamma)\ell_e(x_e^*(u)) \\ &\quad \times x_e^*(u) \leq (\gamma + 1) C(x^*(u), t^M), \end{aligned}$$

proving the result. \square

The following result is a direct corollary of Theorem A.1.

Corollary A.1. *For polynomial latency functions of degree at most d and nonnegative coefficients, $C(x^*(u), t^B) \leq (d + 1)C(x^*(u), t^M)$.*

For Theorem A.2, we need Lemma A.1, but first we define programs BMT and MPT that compute tolls t^B and t^M , respectively. Both programs are in variables (t, v) and define a potential v^k for each commodity k in such a way that any flow-carrying path is indeed a shortest path. The following linear program encodes all possible great set of tolls upper bounded by marginal:

$$\text{(BMT)} \quad \left\{ \begin{array}{ll} \max \sum_{k \in K} U^k(r_{x^*}^k) - \sum_{e \in E} [\ell_e(x_e^*(u)) + t_e] x_e^*(u) & \\ v_u^k - v_v^k + t_e = -\ell_e(x_e^*(u)) & \forall k, e = (u, v): x_e^{*k} > 0 \\ v_u^k - v_v^k + t_e \geq -\ell_e(x_e^*(u)) & \forall k, e = (u, v): x_e^{*k} = 0 \\ v_{d^k}^k - v_{o^k}^k = u^k(r_{x^*}^k) & \forall k: 0 < r_{x^*}^k < r^k \\ v_{d^k}^k - v_{o^k}^k \leq u^k(r^k) & \forall k: r_{x^*}^k = r^k \\ v_{d^k}^k - v_{o^k}^k \geq u^k(0) & \forall k: r_{x^*}^k = 0 \\ t_e \leq \hat{t}_e & \forall e \in E \\ t_e \geq 0 & \forall e \in E. \end{array} \right.$$

The first two sets of constraints ensure that the flow is routed on the shortest paths for the respective commodity. The third, fourth, and fifth sets of constraints ensure that the utilities of the users traveling (or not) have the correct relation with the cost of the flow-carrying paths of the respective commodity.

With a similar linear program, by dropping the constraints that upper bound the tolls, we encode all possible optimal tolls:

$$\text{(MPT)} \quad \left\{ \begin{array}{ll} \max \sum_{k \in K} U^k(r_{x^*}^k) - \sum_{e \in E} [\ell_e(x_e^*(u)) + t_e] x_e^*(u) & \\ v_u^k - v_v^k + t_e = -\ell_e(x_e^*(u)) & \forall k, e = (u, v): x_e^{*k} > 0 \\ v_u^k - v_v^k + t_e \geq -\ell_e(x_e^*(u)) & \forall k, e = (u, v): x_e^{*k} = 0 \\ v_{d^k}^k - v_{o^k}^k = u^k(r_{x^*}^k) & \forall k: 0 < r_{x^*}^k < r^k \\ v_{d^k}^k - v_{o^k}^k \leq u^k(r^k) & \forall k: r_{x^*}^k = r^k \\ v_{d^k}^k - v_{o^k}^k \geq u^k(0) & \forall k: r_{x^*}^k = 0 \\ t_e \geq 0 & \forall e \in E. \end{array} \right.$$

Because Theorem A.2 considers single-commodity instances, and great tolls other than the marginal tolls exist only when all the demand is routed, we may simplify BMT to BMT', which considers a single commodity of unit demand (this assumption is w.l.o.g. as latency functions can always be adjusted appropriately) in which all the flow is routed (so we omit the dependence on u) and minimizes the users' cost (because the aggregate utility of the optimal flow is still fixed):

$$\text{(BMT')} \quad \left\{ \begin{array}{ll} \min \sum_{e \in E} [\ell_e(x_e^*) + t_e] x_e^* & \\ v_u - v_v + t_e = -\ell_e(x_e^*) & \forall e = (u, v): x_e^* > 0 \\ v_u - v_v + t_e \geq -\ell_e(x_e^*) & \forall e = (u, v): x_e^* = 0 \\ t_e \leq \hat{t}_e & \forall e \in E \\ t_e \geq 0 & \forall e \in E. \end{array} \right.$$

Lemma A.1 is a structural lemma that allows us to upper bound the value of (BMT'). This is naturally important in order to derive an approximation bound. We first need some definitions. Consider a directed network G and the undirected network G^u that comes from G if we drop the directions of its links. Any path in G^u is called an undirected path in G . For an undirected path P , the links that are traversed in their actual direction are called *forward* links, denoted by P^+ , and the ones traversed in their reversed direction are called *backward* links, denoted by P^- . Finally, an undirected path has J alternations if, when traversing it, there are exactly J times at which a forward link is followed by a backward link. For more on such alternating paths, see, for example, Lin et al. (2011), Nikolova and Stier-Moses (2015), and Kleer and Schäfer (2016).

Lemma A.1. *There exists an undirected $o - d$ path P such that the first and last link of P belong in P^+ , all $e \in P^+$ are flow carrying, and*

$$C(x^*, t^B) = \sum_{e \in P^+} \ell_e(x_e^*) - \sum_{e \in P^-} [\ell_e(x_e^*) + \hat{t}_e].$$

Proof. We start by taking the dual of (BMT'), at which variables \mathbf{y}' are associated with shortest path constraints and variables \mathbf{z} with the upper bounds on tolls.

$$(DBMT') \quad \left\{ \begin{array}{l} \max \quad \sum_{e \in E} [-y'_e \ell_e(x_e^*) - z_e \hat{t}_e + \ell_e(x_e^*) x_e^*] \\ \sum_{e \in \delta^+(u)} y'_e - \sum_{e \in \delta^-(u)} y'_e = 0 \quad \forall u \in V \\ y'_e - z_e \leq x_e^* \quad \forall e \in E \\ y'_e \geq 0 \quad \forall e : x_e^* = 0 \\ z_e \geq 0 \quad \forall e \in E. \end{array} \right.$$

Let now $y_e \triangleq y'_e - x_e^*$, and notice that the nonnegativity constraints for y'_e , where $x_e^* = 0$, can be rewritten as $y_e \geq 0$. Let us observe that $-\mathbf{y}$ is a unit demand undirected flow (i.e., a flow without sign constraints). By our unit demand assumption,

$$\begin{aligned} \sum_{e \in \delta^+(o)} (-y_e) - \sum_{e \in \delta^-(o)} (-y_e) &= \sum_{e \in \delta^+(o)} x_e^* - \sum_{e \in \delta^-(o)} x_e^* \\ &= 1 \\ \sum_{e \in \delta^+(d)} (-y_e) - \sum_{e \in \delta^-(d)} (-y_e) &= \sum_{e \in \delta^+(d)} x_e^* - \sum_{e \in \delta^-(d)} x_e^* \\ &= -1 \\ \sum_{e \in \delta^+(u)} (-y_e) - \sum_{e \in \delta^-(u)} (-y_e) &= \sum_{e \in \delta^+(u)} x_e^* - \sum_{e \in \delta^-(u)} x_e^* \\ &= 0, \end{aligned}$$

where, in the last equation, $u \neq o, d$. This way, we can reformulate (DBMT') as

$$(DBMT) \quad \left\{ \begin{array}{l} -\min \quad \sum_{e \in E} [y_e \ell_e(x_e^*) + z_e \hat{t}_e] \\ -\mathbf{y} \text{ undirected unit flow} \\ z_e - y_e \geq 0 \quad \forall e \in E \\ y_e \geq 0 \quad \forall e : x_e^* = 0 \\ z_e \geq 0 \quad \forall e \in E. \end{array} \right.$$

Now, we make the following observation: given any \mathbf{y} satisfying these constraints, there is a unique best choice of $\mathbf{z} = \mathbf{z}(\mathbf{y})$; namely, if $y_e \geq 0$, then $z_e = y_e$, and if $y_e < 0$, then $z_e = 0$. Furthermore, observe that $\mathbf{z}(\mathbf{y})$ is piece-wise affine: as long as no \mathbf{y} nonnegativity constraint becomes active, \mathbf{z} changes linearly as a function of \mathbf{y} . Now, let $(\mathbf{y}, \mathbf{z}(\mathbf{y}))$ be an optimal solution for (DBMT): we show that we can choose this optimal solution in such a way that it does not support any cycles. Suppose there exists an undirected cycle $C = C^+ \cup C^-$ (where C^+ and C^- are defined so that the links in C^+ traverse the cycle in the opposite direction from the links in C^-) with variables $y_e \neq 0$ for all $e \in C$. Given $\varepsilon \in \mathbb{R}$, consider the perturbation \mathbf{y}^ε such that, for all $e \in C^+$, $y_e^\varepsilon = y_e + \varepsilon$, and for all $e \in C^-$, $y_e^\varepsilon = y_e - \varepsilon$. Notice that, up to the point at which, for the first time, some y_e becomes zero, the perturbations $(\mathbf{y}^\varepsilon, \mathbf{z}(\mathbf{y}^\varepsilon))$ are feasible for (DBMT), and the objective function changes linearly with ε . Because neither of these perturbations can improve the objective, the objective has to be constant for these perturbations. This way, we choose either a positive or negative ε until one of the y variables reaches zero. The resulting

perturbation $(\mathbf{y}^\varepsilon, \mathbf{z}(\mathbf{y}^\varepsilon))$ is, thus, optimal, and it does not contain C in its support. We can continue this procedure until all cycles are eliminated.

As a conclusion, there exists an optimal solution $(\mathbf{y}, \mathbf{z}(\mathbf{y}))$ whose support is an undirected simple $o-d$ path that we call P . By following this path and using the conservation constraints, we have that each $e \in P$ satisfies that either $y_e = -1$ (and, thus, $z_e = 0$) or $y_e = 1$ (and, thus, $z_e = 1$). Letting $P^+ = \{e \in P : y_e = -1\}$ and $P^- = \{e \in P : y_e = 1\}$, we have that the optimal value of (DBMT) equals

$$\sum_{e \in P^+} \ell_e(x_e^*) - \sum_{e \in P^-} [\ell_e(x_e^*) + \hat{t}_e];$$

furthermore, because of the nonnegativity constraints for \mathbf{y} , all links $e \in P^+$ are necessarily flow carrying.

Finally, we prove that the first and last links of P are forward, which, additionally, by flow conservation, implies that links in P^+ are forward and links in P^- are backward. We just prove it for the first link as the argument for the last is analogous. Let e_1 be the first link of P : by flow conservation, $y_{e_1} = -1$, and thus, $x_{e_1}^* > 0$. Because x^* is optimal, then it is acyclic, so e_1 has to be forward. This completes the proof. \square

We make the following assumption on the instance: there exists a $J \geq 0$ such that any simple $o-d$ undirected path has at most J alternations. The smallest constant J satisfying this condition is called the *alternation number*, and our approximation bound only depends on this number.

Theorem A.2. Consider a single commodity and unit demand instance of the network pricing game whose underlying network has alternation number J . We have, $C(x^*, \mathbf{t}^B) \leq (J+1) \cdot C(x^*, \mathbf{t}^M)$.

Proof. By Lemma A.1, we have $C(x^*, \mathbf{t}^B) = \sum_{e \in P_+} \ell_e(x_e^*) - \sum_{e \in P_-} [\ell_e(x_e^*) + \hat{t}_e]$. Because the alternation number of G is J , we can decompose P into at most J segments of consecutive forward and backward links, $P = A_1 - B_1 - A_2 - B_2 - \dots - B_j - A_{j+1}$, from which $C(x^*, \mathbf{t}^B) \leq \sum_{j=1}^{J+1} [\sum_{e \in A_j} \ell_e(x_e^*)]$. Now, let $(\mathbf{t}^*, \mathbf{v}^*)$ be an optimal solution for (MPT); because all links $e \in A_j$, $j = 1, \dots, J+1$ are flow carrying, for each $j = 1, \dots, J+1$, we have $\sum_{e \in A_j} \ell_e(x_e^*) \leq v_d^* - v_o^*$. Combining the inequalities, we obtain $C(x^*, \mathbf{t}^B) \leq \sum_{j=1}^{J+1} [\sum_{e \in A_j} \ell_e(x_e^*)] \leq (J+1) [v_d^* - v_o^*] = (J+1) \cdot C(x^*, \mathbf{t}^M)$. \square

Appendix B. Missing Proofs of Section 5

B.1. Missing Proofs from Section 5.1

In order to prove Theorem 3, we first restate the profit-maximization property that yields the main result for the model with operating costs.

Lemma B.1. Let $\mathbf{t} \geq 0$ be a toll vector and $e \in E$ with $x_e(\mathbf{t}) > 0$. If t_e is a local optimum for the profit-maximization problem (i.e., for objective $\pi_e(\cdot, \mathbf{t}_{-e})$), then $x_e(\mathbf{t}) \cdot \ell'_e(x_e(\mathbf{t})) + c'_e(x_e(\mathbf{t})) \leq t_e$.

Proof. In order to use the first-order optimality conditions, we need to ensure that the flow on e neither suddenly drops to zero nor remains constant. By continuity, there exists an interval $[t_e, t_e + \delta']$, where $x_e(\cdot) > 0$ with $\delta' > 0$; furthermore, by local optimality, we may choose $0 < \delta < \delta'$ so that t_e is a profit-maximizing toll on the interval. This, in particular, implies that $0 < x_e(s_e) < x_e(\mathbf{t})$ for any $s_e \in (t_e, t_e + \delta)$.

Now, because t_e is a local maximizer, we can use the first-order optimality conditions, $D^+[\pi_e(t_e)] \leq 0$. Using the linearity of Dini derivatives, we get

$$\begin{aligned} x_e(t) + t_e D^+[x_e(t)] - c'_e(x_e(t)) \cdot D^+[x_e(t)] \\ = D^+[\pi_e(t_e)] \leq 0. \end{aligned} \tag{B.1}$$

On the other hand, because $x_e(s_e) < x_e(t)$ for any $s_e \in (t_e, t_e + \delta)$, by Lemma 1(ii), we have that $[\ell_e(x_e(s_e)) + s_e - \ell_e(x_e(t)) - t_e]/[x_e(s_e) - x_e(t)] \leq 0$. After rearranging terms, we get

$$\frac{\ell_e(x_e(s_e)) - \ell_e(x_e(t))}{x_e(s_e) - x_e(t)} \leq -\frac{s_e - t_e}{x_e(s_e) - x_e(t)}.$$

Taking $\limsup_{s_e \rightarrow t_e^+}$ in the expression, we get that the left-hand side converges to $\ell'_e(x_e(t))$, whereas the right-hand side converges to $-(D^+[x_e(t)])^{-1}$. Because $D^+[x_e(t)] < 0$, we conclude that $\ell'_e(x_e(t)) \leq -\frac{1}{D^+[x_e(t)]}$. This, in combination with (B.1), gives

$$\begin{aligned} x_e(t)\ell'_e(x_e(t)) + c'_e(x_e(t)) &\leq -\frac{x_e(t)}{D^+[x_e(t)]} + c'_e(x_e(t)) \\ &\leq t_e, \end{aligned}$$

which proves the result. \square

Proof of Theorem 3. First, we prove that \bar{t} is a Nash equilibrium. By way of contradiction, suppose that $t_e < \bar{t}_e$ maximizes profit given $t_f = \bar{t}_f$ for all $f \neq e$. Because $\bar{t}_e > 0$, we must have $x_e(t) > 0$, and thus, by Lemma 1(i), $x_e(t_e, \bar{t}_{-e}) > 0$. By Lemma B.1, we have $t_e \geq x_e(t_e, \bar{t}_{-e})\ell'_e(x_e(t_e, \bar{t}_{-e})) + c'_e(t_e, \bar{t}_{-e})$, implying

$$\begin{aligned} x_e(t_e, \bar{t}_{-e})\ell'_e(x_e(t_e, \bar{t}_{-e})) + c'_e(t_e, \bar{t}_{-e}) \\ \leq t_e < \bar{t}_e \leq \hat{t}_e = x_e^*\ell'_e(x_e^*) + c'_e(x_e^*), \end{aligned}$$

and thus, $x_e(t_e, \bar{t}_{-e}) < x_e^*$, contradicting Lemma 1(i).

Now that we have proven existence, we show that there is a unique Nash equilibrium. By way of contradiction, suppose there exists another Nash equilibrium $t \neq \bar{t}$ for the profit-maximization game with caps \bar{t} and let $E^< = \{e \in E : t_e < \bar{t}_e\}$. Observe that, for a deviation to be possible $\hat{t}_e \geq \bar{t}_e > 0$ for all $e \in E^<$, and also there is at least one player with $t_e < \bar{t}_e$ and $x_e(t) > 0$. If not, then, for all players $e \in E^<$, we have $x_e(t) = 0$, but this contradicts Lemma 1(i) as there should exist an $e_1 \in E^<$ such that $x_{e_1}(t) \geq x_{e_1}(\bar{t})$ and $x_{e_1}(\hat{t}) = x_{e_1}(\bar{t}) > 0$ because $\hat{t}_{e_1} > 0$.

So we can assume that there are some players with $t_e < \bar{t}_e$ and $x_e(t) > 0$. Because all players are playing their best response, we can use Lemma B.1, concluding that $x_e(t)\ell'_e(x_e(t)) + c'_e(x_e(t)) \leq t_e$ for all $e \in E^<$ with $x_e(t) > 0$, which, in turn, gives

$$\begin{aligned} x_e(t)\ell'_e(x_e(t)) + c'_e(x_e(t)) &\leq t_e < \bar{t}_e \leq \hat{t}_e \\ &= x_e^*\ell'_e(x_e^*) + c'_e(x_e^*), \end{aligned}$$

concluding that $x_e(t) < x_e^* = x_e(\bar{t})$ for all $e \in E^<$ with $x_e(t) > 0$. But, from Lemma 1(i), there exists $e_1 \in E^<$ such that $x_{e_1}(t) \geq x_{e_1}(\bar{t})$ and, by $\hat{t}_{e_1} > 0$, $x_{e_1}(\hat{t}) = x_{e_1}(\bar{t}) > 0$, a contradiction. \square

Proof of Theorem 4. By Theorem 3, it is sufficient to prove that \bar{t} is a strong Nash equilibrium. By way of contradiction,

let E^Δ be a set of links and $t \leq \bar{t}$ with $\{e \in E : t_e \neq \bar{t}_e\} \subseteq E^\Delta$ such that $\pi_e(t) \geq \pi_e(\bar{t})$ for all $e \in E^\Delta$ with at least one strict inequality. Then, for all $e \in E^< \subseteq E^\Delta$, we have $(t_e - a_e) \cdot x_e(t) - b_e \geq (t_e - a_e) \cdot x_e(\bar{t}) - b_e \geq 0$, where the second inequality holds by assumption, which implies $x_e(t) > x_e(\bar{t})$, and

$$t_e - a_e \geq \frac{(t_e - a_e) \cdot x_e(\bar{t})}{x_e(t)}. \tag{B.2}$$

By Lemma 1(ii), there exists $e \in E^<$ such that $[x_e(\bar{t}) - x_e(t)] \times [\ell_e(x_e(\bar{t})) + \bar{t}_e - \ell_e(x_e(t)) - t_e] \leq 0$, which, in combination with $x_e(t) > x_e(\bar{t})$, gives

$$\ell_e(x_e(\bar{t})) + \bar{t}_e \geq \ell_e(x_e(t)) + t_e. \tag{B.3}$$

Let $e \in E^<$ be a link satisfying (B.3). We have the following inequalities:

$$\begin{aligned} \ell'_e(x_e(\bar{t}))(x_e(t) - x_e(\bar{t})) &\leq \ell_e(x_e(t)) - \ell_e(x_e(\bar{t})) \\ &\leq \bar{t}_e - t_e \leq (\bar{t}_e - a_e) \\ &\quad \cdot \left(1 - \frac{x_e(\bar{t})}{x_e(t)}\right), \end{aligned} \tag{B.4}$$

where the first inequality follows from the convexity of ℓ_e , the second from (B.3), and the third from (B.2). Because $\bar{t}_e \leq \hat{t}_e = \ell'_e(x_e(\bar{t}))x_e(\bar{t}) + a_e$ as \bar{t} is optimal, we obtain from (B.4) that

$$\begin{aligned} \ell'_e(x_e(\bar{t}))(x_e(t) - x_e(\bar{t})) &\leq \ell'_e(x_e(\bar{t}))x_e(\bar{t}) \\ &\quad \times \left(\frac{x_e(t) - x_e(\bar{t})}{x_e(t)}\right). \end{aligned}$$

Because $\ell'_e(x_e(\bar{t})) > 0$ and $x_e(t) > x_e(\bar{t})$, we conclude that $x_e(t) \leq x_e(\bar{t})$, a contradiction.

B.2. Missing Proofs from Section 5.2

Proof of Proposition 1. We first show that, for series-parallel graphs under any fixed demand r and for any link e , setting $t'_e > t_e$ only increases the common travel cost at equilibrium. The proof is by induction on the decomposition of the series-parallel graph. In the base case of a single link, this is true by Lemma 1.

Let G be a series composition of series-parallel graphs G_1 and G_2 , and w.l.o.g., let e belong to G_1 . Under toll t_e and demand r , let the travel costs through G_1 and G_2 be \mathcal{L}_1 and \mathcal{L}_2 , respectively. By induction, under toll t'_e , the travel cost through G_1 only increases to $\mathcal{L}'_1 \geq \mathcal{L}_1$, and the travel cost through G_2 remains the same. Thus, the common cost at equilibrium only increases.

For the other case, let now G be a parallel composition of series-parallel graphs G_1 and G_2 , and w.l.o.g., let e belong in G_1 . Under toll t_e and demand r , let the common travel cost through G_1 and G_2 be \mathcal{L} , and let the traffic routed through G_1 and G_2 be r_1 and r_2 , respectively. By induction, under toll t'_e , if we let r_1 units go through G_1 , the travel cost through G_1 only increases to some $\mathcal{L}' \geq \mathcal{L}$. Thus, at equilibrium $r'_1 \leq r_1$ are routed through G_1 and $r'_2 \geq r_2$ units are routed through G_2 in order to equalize the travel costs through G_1 and G_2 , which implies that the common cost at equilibrium only increases.

Given this, we go on to prove that, for elastic demand for any link e , setting $t'_e > t_e$ only increases the common travel

cost at equilibrium and only decreases the demand being routed. Let r be the demand routed at equilibrium under toll t_e and \mathcal{L} be the common cost at equilibrium. If r units are routed when t'_e is set, then, by the preceding, the travel cost only increases to some $\mathcal{L}' \geq \mathcal{L}$. Yet this may not be the equilibrium because some users may incur travel costs greater than their utility. By removing flow (while in equilibrium), we reduce the travel cost and increase the minimum utility among the users that travel until these two get equal, at which point we have an equilibrium with the traffic demand routed equal to some $r' \leq r$ and the travel cost equal to some $\mathcal{L}'' : \mathcal{L} \leq \mathcal{L}'' \leq \mathcal{L}'$, proving the claim.

Now, we go back to prove the proposition. Let E' be a set of links such that, for all $e, e' \in E'$ and all paths $P : e \in P$, we have $e' \notin P$. We prove that E' is complementary. Consider an arbitrary edge $e \in E'$. In order for e not to be in the same path with some other arbitrary edge $e' \in E'$, it should be that, in the decomposition of G , e lies in some graph G_1 , e' lies in some graph G_2 , and G_1 and G_2 are connected in parallel. To prove that $x_{e'}(t_e, t_{-e}) \leq x_{e'}(t'_{e'}, t_{-e})$, it suffices to show that G_1 gets at most the same flow and G_2 gets at least the same flow at equilibrium under t'_e compared with under t_e because, if this is the case, all of G_2 's edges may only gain flow.

We prove this by induction on the decomposition of the series-parallel graph starting from the full graph. For the base case, let G_1 and G_2 's composition happen the first time that a parallel composition between a graph containing e_1 and some other graph occurs (some series compositions may have occurred earlier). By what we have proven earlier, we know that, under t'_e , only less flow may go through G_1 and G_2 , yet their common travel costs only increase. This implies that the flow through G_2 only increases because its tolls remain unchanged, implying, at the same time, that the flow through G_1 only decreases.

For the induction, consider all compositions of a graph containing e_1 and some other graph that happens before the composition of G_1 and G_2 , and let the induction hypothesis hold for them. Using the induction hypothesis for the last time such a composition occurs, we get that the graph that contains e_1 gets at most the same flow under t'_e and has at least the same cost. This further implies the flow through G_1 and G_2 only decreases, and the cost only increases under t'_e . Similar to earlier, the latter implies the flow through G_2 only increases because its tolls remain unchanged, implying, at the same time, that the flow through G_1 only decreases, completing the induction and the proof. \square

Proof of Theorem 5. First, we prove that, if \mathcal{G} admits a Nash equilibrium, then it must be \bar{t} . To reach a contradiction, assume there exists a Nash equilibrium $t \neq \bar{t}$. Because \bar{t} are set as caps, $t \neq \bar{t}$ implies that $t \leq \bar{t}$, and the set $E^< = \{e \in E : t_e < \bar{t}_e\}$ is nonempty. By Lemma 1(i), there exists an $e_1 \in E^<$ such that $x_{e_1}(t) \geq x_{e_1}(\bar{t})$. Let player k be the player operating e_1 . We show that t is not a Nash equilibrium by showing that player k is not on the player's best response.

It is $x_{e_1}(\bar{t}) > 0$ or else $0 \leq t_{e_1} \leq \bar{t}_{e_1} \leq \hat{t}_{e_1} = x_{e_1}(\bar{t}) \ell'_{e_1}(x_{e_1}(\bar{t})) = 0$, implying $t_{e_1} = \bar{t}_{e_1}$, contradicting $e_1 \in E^<$. This further implies that $t_{e_1} > 0$; otherwise, player k is not on the player's best response because, by continuity, the player can increase the toll to $t_{e_1} = \delta$ for some small enough $\delta > 0$ so that e_1 still gets

positive flow, and thus, the player gains strictly more profit: positive profit from e_1 (instead of zero) and at least the same profit from the player's other links because they do not lose flow by the complementarity condition.

Consider the profit function $\pi_{e_1}(\cdot)$ of link e_1 when all other tolls are kept fixed. Toll t_{e_1} cannot be a local optimum of $\pi_{e_1}(\cdot)$ because, by Lemma 2 (recall $t_{e_1} > 0$) and $x \ell'(x)$ being increasing, it would be $t_{e_1} \geq x_{e_1}(t) \cdot \ell'_{e_1}(x_{e_1}(t)) \geq x_{e_1}(\bar{t}) \cdot \ell'_{e_1}(x_{e_1}(\bar{t})) = \hat{t}_{e_1} \geq \bar{t}_{e_1}$, contradicting $e_1 \in E^<$. Additionally, any local optimum $\pi_{e_1}(\cdot)$ must be above t_{e_1} as otherwise, if there was a local optimum $t'_{e_1} < t_{e_1}$, then, by Lemma 1(i), it would be $x_{e_1}(t'_{e_1}) \geq x_{e_1}(t_{e_1})$, and by a similar reasoning as before, it would be $t'_{e_1} \geq \bar{t}_{e_1}$, a contradiction.

By this, any local optimum of $\pi_{e_1}(\cdot)$ is strictly above t_{e_1} , and consequently, there exists a sufficiently small $\delta > 0$ such that increasing the toll of link e_1 from t_{e_1} to $t_{e_1} + \delta$, strictly increases the profit made by link e_1 . Now consider the deviation for player k in which the player increases the toll of e_1 by δ . By the previous discussion, the player strictly gains more profit from e_1 , but also the player does not lose profit from the other links the player operates because, by the complementarity condition, each of the player's links may only gain flow. Thus, player k is not choosing a best response.

It remains to prove that \bar{t} is a Nash equilibrium. To reach a contradiction, assume \bar{t} is not a Nash equilibrium, and let k be a player that, by deviating from \bar{t} in some of the player's links, strictly gains more profit, that is, for the profit of player k , it is $\pi_k(\bar{t}) < \pi_k(t^1)$, where t^1 is the resulting vector after player k 's deviation. W.l.o.g., we may assume that, for all links for which we have deviation from \bar{t} , that is, $t^1_e < \bar{t}_e$, it is $t^1_e > 0$ because, if a link has $t^1_e = 0$ and $x_e(t^1) = 0$, then setting $t^1_e > 0$ does not change the flow or the profits, and thus, we still have an improving deviation for player k although if a link has $t^1_e = 0$ and $x_e(t^1) > 0$, then player k is still on an improving deviation if the player sets $t^1_e = \delta$ for some small enough $\delta > 0$. The player is still at an improving deviation because, by that small increase on the toll, e_1 still has positive flow, and thus, the player gains strictly more profit: positive profit from e_1 (instead of zero) and at least the same profit from the player's other links because they do not lose flow by the complementarity condition.

The underlying idea for the proof is the following. Under t^1 , because of Lemmas 1(i) and 2, there exists a link for which player k can increase its toll and gain more profit from it and, at the same time, not lose profit from other links. The resulting toll vector gets closer to the cap vector \bar{t} and has strictly more profit, and as long as it does not get equal to \bar{t} , we may repeat the procedure and get even closer to \bar{t} with even bigger profits. Formally, though, we do it a bit differently. Starting from the deviation t^1 , we create a sequence of toll vectors t^i that all correspond to deviations of player k and are such that, for any $i \geq 2$, player k 's profit under t^i is only higher than the profit under t^{i-1} (and t^1 and \bar{t}). Then, we show that this sequence converges to \bar{t} , getting a contradiction, because the preceding implies $\pi_k(\bar{t}) < \pi_k(t^1) \leq \dots \leq \pi_k(t^i) \leq \dots \leq \pi_k(\bar{t})$.

Let E_k be the set of links that player k operates, and consider any arbitrary cyclic order of the links of E_k . Also, let e^1 be an arbitrary link of E_k . To inductively create t^i from t^{i-1} for $i \geq 2$ let e^i be the link that follows e^{i-1} in the cyclic order of E_k . Consider the profit function $\pi_{e^i}(\cdot)$ of link e^i when the tolls on

other links are fixed according to \mathbf{t}^{i-1} . Let t_e^{i*} be the smallest local optimum of $\pi_{e^i}(\cdot)$ and set $t_e^i = \min\{\max\{t_e^{i-1}, t_e^{i*}\}, \bar{t}_e\}$. For any other link e of E^k ($e \neq e^i$), set $t_e^i := t_e^{i-1}$. Note first that, by construction, $\pi_k(\mathbf{t}^i) \geq \pi_k(\mathbf{t}^{i-1})$ because the only toll that possibly changes is that of e^i , and it only increases toward the first maximizer of $\pi_{e^i}(\cdot)$, implying that the profit from e^i only increases, and by the complementarity condition, the flow and profit on any other link of E_k only increases. Additionally, note that, for all i , it is $t^i \leq \bar{\mathbf{t}}$, and for all $i \geq 2$, it is $t^{i-1} \leq t^i$ because, at every step, the only link getting its toll changed is e^i , and it gets a toll in between t_e^{i-1} and \bar{t}_e . Thus, \mathbf{t}^i is an increasing sequence of toll vectors that additionally is upper bounded by $\bar{\mathbf{t}}$, and consequently, it must converge to a toll vector, say \mathbf{t}^0 , that is, $\mathbf{t}^i \rightarrow \mathbf{t}^0$. Last, note that, for all i , $t^i \leq \bar{\mathbf{t}}$.

We prove that $\mathbf{t}^0 = \bar{\mathbf{t}}$. To reach a contradiction, assume otherwise and let $E_k^< = \{e \in E_k : t_e^0 < \bar{t}_e\}$ be nonempty. By Lemma 1(i), there exists a link $f \in E_k^<$ such that $x_f(\mathbf{t}^0) \geq x_f(\bar{\mathbf{t}})$. Consider the profit function $\pi_f(\cdot)$ of link f when all other tolls are kept fixed (according to \mathbf{t}^0). Toll t_f^0 cannot be a local optimum of $\pi_f(\cdot)$ because in that case, by Lemma 2 (we can apply Lemma 2 because $t_f^0 \geq \bar{t}_f$, and for all e with $t_e^1 < \bar{t}_e$, it is $t_e^1 > 0$, implying $t_f^1 > 0$), it would be $t_f^0 \geq x_f(t_f^0) \cdot \ell'_f(x_f(t_f^0)) = x_f(\mathbf{t}^0) \cdot \ell'_f(x_f(\mathbf{t}^0)) \geq x_f(\bar{\mathbf{t}}) \cdot \ell'_f(x_f(\bar{\mathbf{t}})) = \hat{t}_f \geq \bar{t}_f$, contradicting $f \in E_k^<$. Additionally, any local optimum $\pi_f(\cdot)$ must be above t_f as, otherwise, if there is a local optimum $t'_f < t_f^0$, then, by Lemma 1(i) it would be $x_f(t'_f) \geq x_f(t_f^0)$, and by a similar reasoning, it would be $t'_f \geq \bar{t}_f$, a contradiction. Thus, any local optimum for $\pi_f(\cdot)$ is strictly above t_f^0 , say at $t_f^{0*} = t_f^0 + \delta$, for $\delta > 0$.

On the other hand, consider all those times when, while creating the sequence \mathbf{t}^i , link f was a candidate for getting its toll changed (recall the cyclic order when choosing which tolls to update). Intuitively, because $\mathbf{t}^i \rightarrow \mathbf{t}^0$, there is some high enough j such that the toll vectors \mathbf{t}^0 and \mathbf{t}^j are so close to each other that the corresponding profit functions of link f almost coincide. More formally, for the δ defined earlier, because $\mathbf{t}^i \rightarrow \mathbf{t}^0$, there must be some j for which link f is a candidate for getting its toll changed and is such that the smallest local optimum of the profit function of f under toll vector \mathbf{t}^j , say t_f^{j*} , is such that $|t_f^{0*} - t_f^{j*}| \leq \delta/2$. The latter implies that $t_f^{j*} > t_f^0$ because, by the preceding, $t_f^{0*} = t_f^0 + \delta$. But this readily gives a contradiction as, for that j , based on the construction of the sequence (namely $t_f^j = \min\{\max\{t_f^{j-1}, t_f^{j*}\}, \bar{t}_f\}$), link f would get either t_f^{j*} or \bar{t}_f as its toll, which are both strictly greater than t_f^0 (recall, $t_f^{j-1} \leq t_f^0$ by construction).

Thus, indeed $\mathbf{t}^0 = \bar{\mathbf{t}}$, and so $\mathbf{t}^i \rightarrow \bar{\mathbf{t}}$. Recall that the profit function $\pi_k(\cdot)$ of player k is a continuous function of tolls. This yields $\pi_k(\mathbf{t}^i) \rightarrow \pi_k(\bar{\mathbf{t}})$, which, combined with that for all i , $\pi_k(\mathbf{t}^i) \geq \pi_k(\mathbf{t}^1) > \pi_k(\bar{\mathbf{t}})$, gives the desired contradiction: $\pi_k(\bar{\mathbf{t}}) > \pi_k(\bar{\mathbf{t}})$. Consequently, every player is on the player's best response when playing according to $\bar{\mathbf{t}}$. \square

Proof of Theorem 6. We need only prove that $\bar{\mathbf{t}}$ is a strong Nash equilibrium. The proof is similar to that of Theorem 5. Assume there is a coalition of players that deviates from $\bar{\mathbf{t}}$ on some of their links so that no player loses profit and some player, say player k , strictly gains profit. Because all links are complementary, we may intuitively think that all deviated links belong to player k . We may get a contradiction similar to the proof of Theorem 5 because we may get a sequence of deviations $\mathbf{t}^i \rightarrow \bar{\mathbf{t}}$ such that, for every $i \geq 2$, no deviating

player gets worse profit in \mathbf{t}^i compared with \mathbf{t}^{i-1} , and for player k , it is $\pi_k(\mathbf{t}^i) \geq \pi_k(\mathbf{t}^{i-1}) \geq \pi_k(\mathbf{t}^1) > \pi_k(\bar{\mathbf{t}})$, yielding $\pi_k(\bar{\mathbf{t}}) > \pi_k(\bar{\mathbf{t}})$. \square

References

- Acemoglu D, Ozdaglar A (2007a) Competition and efficiency in congested markets. *Math. Oper. Res.* 32(1):1–31.
- Acemoglu D, Ozdaglar A (2007b) Competition in parallel-serial networks. *IEEE J. Selected Areas Comm.* 25(6):1180–1192.
- Anshelevich E, Sekar S (2015) Price competition in networked markets: How do monopolies impact social welfare. *Proc. 11th Internat. Conf. Web Internet Econom.* (Springer, Berlin, Heidelberg), 16–30.
- Aumann R (1959) Acceptable points in general cooperative n -person games. Luce RD, Tucker AW, eds. *Contributions to the Theory of Games*, vol. 4 (Princeton University Press, Princeton, NJ), 287–324.
- Beckmann M, McGuire C, Winsten C (1956) *Studies in the Economics of Transportation* (Yale University Press, New Haven, CT).
- Berge C (1963) *Topological Spaces: Including a Treatment of Multi-valued Functions, Vector Spaces, and Convexity*, Dover Books on Mathematics (Dover Publications, Mineola).
- Bonifaci V, Salek M, Schäfer G (2011) Efficiency of restricted tolls in non-atomic network routing games. *Proc. 4th Internat. Sympos. Algorithmic Game Theory* (Springer, Berlin, Heidelberg), 302–313.
- Briest P, Hoefer M, Krysta P (2012) Stackelberg network pricing games. *Algorithmica* 62(3):733–753.
- Caragiannis I, Chatzigeorgiou X, Kanellopoulos P, Krimpas G, Protopapas N, Voudouris A (2017) Efficiency and complexity of price competition among single-product vendors. *Artificial Intelligence* 248:9–25.
- Castelli L, Labbé M, Violin A (2017) Network pricing problem with unit toll. *Networks* 69(1):83–93.
- Chawla S, Niu F (2009) The price of anarchy in Bertrand games. *Proc. 10th ACM Conf. Electronic Commerce* (ACM, New York), 305–314.
- Chawla S, Niu F, Roughgarden T (2008) Bertrand competition in networks. *Proc. First Internat. Sympos. Algorithmic Game Theory* (Springer, Berlin, Heidelberg), 70–82.
- Cole R, Dodis Y, Roughgarden T (2003) Pricing network edges for heterogeneous selfish users. *Proc. 35th ACM Sympos. Theory Comput.* (ACM, New York), vol. 3, 444–467.
- Correa J, Stier-Moses N (2011) Wardrop equilibria. *Wiley Encyclopedia of Operations Research and Management Science* (Wiley, Hoboken, NJ).
- Dial R (1999) Minimal-revenue congestion pricing part i: A fast algorithm for the single-origin case. *Transportation Res. Part B: Methodological* 33(3):189–202.
- Dixon H (2001) *Surfing Economics: Essays for the Inquiring Economist* (Palgrave Macmillan, New York).
- Fleischer L (2005) Linear tolls suffice: New bounds and algorithms for tolls in single source networks. *Theoretical Comput. Sci.* 348:217–225.
- Fotakis D, Kalimeris D, Lianes T (2015) Improving selfish routing for risk-averse players. *Proc. 11th Internat. Conf. Web Internet Econom.* (Springer, Berlin, Heidelberg), 328–342.
- Gonzalez A (2016) Controversias por reducciones en la demanda en las concesiones de carreteras en Chile. *Revista de Derecho Económico* 76:61–97.
- Hall MA (1978) Properties of the equilibrium state in transportation networks. *Transportation Sci.* 12(3):208–216.
- Harks T, Schröder M, Vermeulen D (2019) Toll caps in privatized road networks. *Eur. J. Oper. Res.* 276(3):947–956.
- Hayrapetyan A, Tardos E, Wexler T (2007) A network pricing game for selfish traffic. *Distributed Comput.* 19(4):255–266.
- Jelinek T, Klaas M, Schäfer G (2014) Computing optimal tolls with arc restrictions and heterogeneous players. *Proc. 31th Internat. Sympos. Theoretical Aspects Comput. Sci.* (Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik).

- Johari R, Weintraub G, Van Roy B (2010) Investment and market structure in industries with congestion. *Oper. Res.* 58(5): 1303–1317.
- Kleer P, Schäfer G (2016) The impact of worst-case deviations in non-atomic network routing games. *Proc. Ninth Internat. Sympos. Algorithmic Game Theory* (Springer, Berlin, Heidelberg), 129–140.
- Kuiteing A, Marcotte P, Savard G (2016) Network pricing of congestion-free networks: The elastic and linear demand case. *Transportation Sci.* 51(3):791–806.
- Lin H, Roughgarden T, Tardos E (2004) A stronger bound on Braess's paradox. *Proc. 15th Annual ACM-SIAM Sympos. Discrete Algorithms* (SIAM, Philadelphia), 333–334.
- Lin H, Roughgarden T, Tardos E, Walkover A (2011) Stronger bounds on Braess's paradox and the maximum latency of selfish routing. *SIAM J. Discrete Math.* 25(4):1667–1686.
- Meng Q, Lu Z (2017) Quantitative analyses of highway franchising under build-operate-transfer scheme: Critical review and future research directions. *Transportation Res. Part B: Methodological* 102:105–123.
- Musacchio J (2009) The price of anarchy in parallel-serial competition with elastic demand. Technical Report No. UCSC-SOE-09-20, Baskin School of Engineering, University of California, Santa Cruz.
- Musacchio J, Wu S (2007) The price of anarchy in a network pricing game. *Proc. Allerton Conf.* (ACM, New York).
- Nikolova E, Stier-Moses N (2015) The burden of risk aversion in mean-risk selfish routing. *Proc. 16th ACM Conf. Econom. Comput.* (ACM, New York), 489–506.
- Ozdaglar A (2008) Price competition with elastic traffic. *Networks* 52(3):141–155.
- Ozdaglar A, Srikant R (2007) Incentives and pricing in communication networks. *Algorithmic Game Theory* (Cambridge University Press, New York), 571–591.
- Papadimitriou C, Valiant G (2010) A new look at selfish routing. *Proc. Internat. Comput. Sympos.*, 178–187.
- Roughgarden T (2005) *Selfish Routing and the Price of Anarchy* (MIT Press, Cambridge, MA).
- Roughgarden T (2006) On the severity of Braess's paradox: Designing networks for selfish users is hard. *J. Comput. System Sci.* 72(5): 922–953.
- Roughgarden T, Tardos E (2002) How bad is selfish routing? *J. ACM* 49(2):236–259.
- Schmand D, Schröder M, Skopalik A (2019) Network investment games with Wardrop followers. Baier C, Chatzigiannakis I, Flocchini P, Leonardi S, eds. *46th Internat. Colloquium Automata, Languages, Programming*, vol. 132, 151:1–151:14.
- Yang H, Meng Q (2000) Highway pricing and capacity choice in a road network under a build-operate-transfer scheme. *Transportation Res. Part A: Policy Practice* 24(3):207–222.
- Yang H, Tang W, Cheung W, Meng Q (2002) Profitability and welfare gain of private toll roads in a network with heterogeneous users. *Transportation Res. Part A: Policy Practice* 36(6):537–554.
-
- José Correa** is a professor in the department of industrial engineering at Universidad de Chile. His research interests are in game theory and mechanism design mostly from an algorithmic viewpoint. Recently he has mostly worked in data-driven pricing when users are forward-looking and network games, particularly in dynamic environments.
- Cristóbal Guzmán** is an assistant professor in the Institute for Mathematical and Computational Engineering at Pontificia Universidad Católica de Chile. His research interests include large-scale convex optimization, machine learning, and algorithmic game theory.
- Thanasis Lianeas** is a (nontenure track) lecturer and postdoctoral researcher at the School of Electrical and Computer Engineering of the National Technical University of Athens. His research interests lie in the fields of algorithmic game theory, online learning, and learning dynamics.
- Evdokia Nikolova** is an assistant professor in electrical and computer engineering at the University of Texas at Austin. Her current research is on developing algorithms for risk mitigation in networks with applications in energy, transportation, and healthcare. She is a recipient of a Google Research Faculty Award and a National Science Foundation CAREER award.
- Marc Schröder** is an assistant professor in the School of Business and Economics at Maastricht University. His research interests are various aspects of game theory with a current focus on algorithmic game theory and network games.