

Competition in Parallel-Serial Networks

Daron Acemoglu and Asuman Ozdaglar, *Member, IEEE*

Abstract—We study the efficiency implications of competition among profit-maximizing service providers in communication networks. Service providers set prices for transmission of flows through their (sub)network. The central question is whether the presence of prices will help or hinder network performance. We investigate this question by considering the difference between users’ willingness to pay and delay costs as the efficiency metric. Previous work has demonstrated that in networks consisting of parallel links, efficiency losses from competition are bounded. Nevertheless, parallel-link networks are special, and in most networks, traffic has to simultaneously traverse links (or subnetworks) operated by independent service providers. The simplest network topology allowing for this feature is the parallel-serial structure, which we study in this paper. In contrast to existing results, we show that in the presence of serial links, the efficiency loss relative to the social optimum can be arbitrarily large. The reason for this degradation of performance is the *double marginalization* problem, whereby each serial provider charges high prices not taking into account the effect of this strategy on the profits of other providers along the same path. Nevertheless, when there are no delay costs without transmission (i.e., latencies at zero are equal to zero), irrespective of the number of serial and parallel providers, the efficiency of strong oligopoly equilibria can be bounded by 1/2, where strong oligopoly equilibria are equilibria in which each provider plays a strict best response and all of the traffic is transmitted. However, even with strong oligopoly equilibria, inefficiency can be arbitrarily large when the assumption of no delay costs without transmission is relaxed.

Index Terms—Pricing, competition, congestion externalities, Wardrop equilibrium, social optimum, oligopoly equilibrium, efficiency, price of anarchy.

I. INTRODUCTION

There has been growing interest in pricing as a method of allocating scarce network resources. The standard approach assumes that prices may be used as control parameters to achieve a system optimum in a decentralized manner (see, e.g., [1], [2], [3]). Nevertheless, in practice many prices are controlled by for-profit service providers, whose objective is, at least in part, to increase their revenues and profits. A central question is whether prices chosen to maximize service provider revenues will also play a useful role in the allocation of network resources across users.

Research to date suggests that profit-maximizing pricing may improve the allocation of resources in communication networks with self-interested users. Let the metric of efficiency be the difference between users’ willingness to pay and delay

costs in the decentralized equilibrium relative to that in the social optimum (which would be chosen by a network planner with full information and full control over users). In the absence of prices, the decentralized equilibrium can be highly inefficient (see [4]). In contrast, Acemoglu and Ozdaglar [5] show that with inelastic and homogeneous users, pricing by a monopolist controlling all links in a parallel-link network always achieves efficiency (i.e., the efficiency metric is equal to 1). Huang, Ozdaglar and Acemoglu [6] extend this result to a general network topology. More recently, Acemoglu and Ozdaglar [7] show that in a parallel-link network with inelastic and homogeneous users, the efficiency metric with an arbitrary number of competing network providers is always greater than or equal to 5/6.

Most communication networks cannot be represented by parallel-link topologies, however. A given source-destination pair will typically transmit through multiple interconnected subnetworks (or links), potentially operated by different service providers. Existing results on the parallel-link topology do not address how the cooperation and competition between service providers will impact on efficiency in such general networks.

In this paper, we take a step in this direction by considering the simplest network topology that allows for serial interconnection of multiple links/subnetworks, which is the parallel-serial topology. We focus on a single source-destination pair, with flows choosing one of multiple parallel paths. We allow each path to consist of multiple links/subnetworks operated by independent service providers. Our main results show that that the efficiency losses resulting from competition are considerably higher with this topology. The source of additional inefficiency is the presence of serial service providers and will thus be present in more general network topologies. This suggests that unregulated competition in general communication networks may have considerable costs in terms of the efficiency of resource allocation and certain types of regulation may be necessary to make sure that service provider competition does not lead to significant degradation of network performance.

In our model, an origin-destination pair is linked by multiple parallel paths, each potentially consisting of an arbitrary number of serial links. Congestion costs are captured by link-specific non-decreasing convex latency functions, denoted by $l_i(\cdot)$ for link i . Each link is owned by a different service provider. All users are inelastic and homogeneous.

This environment induces the following two-stage game: each service provider simultaneously sets the price for transmission of bandwidth on its link, denoted by p_i . Observing all the prices, in the second stage users route their information through the path with the lowest effective cost, where effective cost consists of the sum of prices and latencies of the links

Manuscript received May 22, 2006; revised January 15, 2007. This research was funded in part by the National Science Foundation CAREER award DMI-0545910. A preliminary version of this work was presented at INFOCOM, 2006.

D. Acemoglu is with the Department of Economics, Massachusetts Institute of Technology, Cambridge, MA, 02139 (e-mail: daron@mit.edu).

A. Ozdaglar is with the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA, 02139 (e-mail: asuman@mit.edu).

along a path [i.e., sum of $p_i + l_i(\cdot)$'s over the links comprising a path]. Our objective is to study the efficiency properties of the subgame perfect equilibria of this game.

The main novel aspect of this model compared to the parallel-link topology is the pricing decisions of different (serial) service providers along a single path. When a particular provider charges a higher price, it creates a negative externality on other providers along the same path, because this higher price reduces the transmission that all the providers along this path receive. This is the equivalent of the *double marginalization* problem in economic models with multiple monopolies and is the source of the significant degradation in the efficiency performance of the network.

In its most extreme form, the double marginalization problem leads to a type of “coordination failure”, whereby all providers, expecting others to charge high prices, also charge prohibitively high prices, effectively killing all data transmission on a given path. Such coordination failures can lead to arbitrarily low efficiency. This type of pathological behavior can happen in subgame perfect equilibria (what we refer to as oligopoly equilibria, OE), but we show that it cannot happen in strict subgame perfect equilibria, *strict OE*, which follows the notion of strict equilibrium introduced in Harsanyi [8]. In strict OE, each service provider must play a strict best response to the pricing strategies of other service providers. We show that this requirement is sufficient to rule out the pathological coordination failures mentioned above.

Nevertheless, we show that strict OE can also have arbitrarily large efficiency losses again owing to the double marginalization problem. Even in a strict OE, serial providers ignore the negative externality they create on other providers along the same path and charge too high prices, which can once again prevent any transmission on a particular path, even when such transmission is socially optimal.

Interestingly, however, these extreme inefficient outcomes occur when high prices on a particular path prevent the entire available traffic from being transmitted. To investigate implications of price competition when all traffic is transmitted, we define an even stronger notion of equilibrium, *strong OE*, as a strict OE in which all traffic is transmitted.¹ We show that when latency without any traffic is equal to zero [i.e., $l_i(0) = 0$], there is a tight bound of 1/2 on the efficiency of strong OE irrespective of the number of paths and service providers in the network. This bound is reached by simple examples. In strong OE, the double marginalization problem is still present, and this is the reason why the bound of 1/2 is lower than the 5/6 bound in our previous work, [7].

However, the assumption that $l_i(0) = 0$ is important for this result. We show that when this assumption is relaxed, the efficiency loss of strong OE relative to the social optimum can be again arbitrarily large.

These results shed doubt on the conjecture that unregulated competition among service providers might lead to prices approximating those that would be set as control parameters

by a centralized network operator. Instead, they show that competition among service providers with general network topologies can lead to significant degradation of network performance (Example 3 below can be part of any network topology and cause arbitrary efficiency losses). Nevertheless, it has to be borne in mind that the examples that have very poor performance relative to the social optimum are somewhat pathological, and this begs the question of whether better performance bounds could be obtained in more realistic topologies, which is an area left for future work.

Work related to our paper includes studies quantifying efficiency losses of selfish routing without prices (e.g., Koutsoupias and Papadimitriou [10], Roughgarden and Tardos [4], Correa, Schulz, and Stier-Moses [9], and Friedman [11]); of resource allocation by different market mechanisms (e.g., Johari and Tsitsiklis [12], Sanghavi and Hajek [13]); and of network design (e.g., Anshelevich et. al. [14]). Basar and Srikant [15] analyze monopoly pricing in a network context under specific assumptions on the utility and latency functions, while He and Walrand [16] study competition and cooperation among Internet service providers under specific demand models. Most closely related to the current paper are our previous work [7], where we study the existence and efficiency of oligopoly equilibria in parallel-link networks, as well as Hayrapetyan, Tardos, and Wexler [17] and Ozdaglar [18], who study pricing in a parallel-link network with elastic demand. No other paper has investigated price competition in the presence of serial providers or more general topologies.

The rest of the paper is organized as follows. Section II outlines the basic environment. It defines the concept of Wardrop equilibrium for the routing of flows given prices set by service providers. Section III defines the concept of equilibrium in the game among the service providers and establishes the existence of a pure strategy equilibrium with linear latency functions, and the existence of a mixed strategy equilibrium more generally. Section IV focuses on the efficiency analysis of oligopoly equilibrium and contains the main results of the paper. This section first shows that an oligopoly equilibrium can be arbitrarily inefficient because of the coordination failures resulting from double marginalization. It then introduces the concepts of strict and strong oligopoly equilibria, and provides a characterization of equilibrium prices in strict oligopoly equilibria. This section also establishes the existence of a unique strong oligopoly equilibrium with linear latencies and a sufficiently high reservation utility, and presents bounds on the inefficiency of strong oligopoly equilibria. It concludes by showing how even strong oligopoly equilibria can be arbitrarily inefficient when latencies at zero congestion are positive. Section V concludes, while the Appendices contain some of the proofs not provided in the text.

II. MODEL

We consider a network with I parallel paths that connect a single source-destination pair. Each path i consists of n_i links. Let $\mathcal{I} = \{1, \dots, I\}$ denote the set of paths and \mathcal{N}_i denote the set of links on path i . Let x_i denote the flow on path i , and $x = [x_1, \dots, x_I]$ denote the vector of path flows. Each link

¹Models of selfish routing without prices, e.g., [4] or [9], assume that all traffic is always transmitted. Our model incorporates a reservation utility for users, so that this is not necessarily the case. All traffic will be transmitted in equilibrium when this reservation utility is sufficiently large.

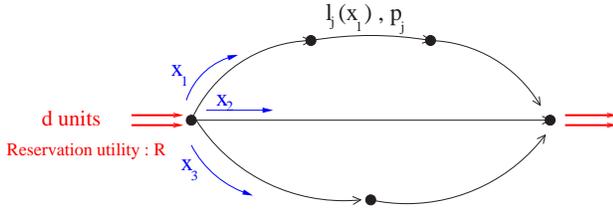


Fig. 1. A network with serial and parallel links.

in the network has a flow-dependent latency function $l_i(x_i)$, which measures the delay as a function of the total flow on link i (see Figure 1). We denote the price per unit flow (bandwidth) of link j by p_j . Let $p = [p_j]_{j \in \mathcal{N}_i, i \in \mathcal{I}}$ denote the vector of prices.

We are interested in the problem of routing d units of flow across the I paths. We assume that this is the aggregate flow of many “small” users and thus adopt the Wardrop’s principle (see [19]) in characterizing the flow distribution in the network; i.e., the flows are routed along paths with minimum effective cost, defined as the sum of the latencies and prices of the links along that path (see the definition below). Wardrop’s principle is used extensively in modelling traffic behavior in transportation networks ([20], [21], [22]) and communication networks ([4], [9]). We also assume that users have a *reservation utility* R and decide not to send their flow if the effective cost exceeds the reservation utility. This implies that user preferences can be represented by the piecewise linear aggregate utility function $u(\cdot)$ depicted in Figure 2. This assumption also implies that all users are “homogeneous” since they all have the same reservation utility, R .²

Definition 1 For a given price vector $p \geq 0$, a vector $x^{WE} \in \mathbb{R}_+^I$ is a *Wardrop equilibrium* (WE) if $\sum_{i \in \mathcal{I}} x_i^{WE} \leq d$ and for all i with $x_i^{WE} > 0$,

$$\begin{aligned} \sum_{j \in \mathcal{N}_i} (l_j(x_i^{WE}) + p_j) &= \min_{k \in \mathcal{I}} \left\{ \sum_{j \in \mathcal{N}_k} (l_j(x_k^{WE}) + p_j) \right\}, (1) \\ \sum_{j \in \mathcal{N}_i} (l_j(x_i^{WE}) + p_j) &\leq R, \end{aligned}$$

with $\sum_{i \in \mathcal{I}} x_i^{WE} = d$ if $\min_{k \in \mathcal{I}} \left\{ \sum_{j \in \mathcal{N}_k} l_j(x_k^{WE}) + p_j \right\} < R$. We denote the set of WE at a given p by $W(p)$.

We adopt the following assumption on the latency functions throughout the paper except in Section IV-F.

Assumption 1 For each $i \in \mathcal{I}$, the latency function $l_i : [0, \infty) \mapsto [0, \infty)$ is convex, continuously differentiable, non-decreasing, and satisfies $l_i(0) = 0$.

Proposition 1 (Existence and Continuity) Let Assumption 1 hold. For any price vector $p \geq 0$, the set of WE, $W(p)$, is nonempty. Moreover, the correspondence $W : \mathbb{R}_+^I \rightrightarrows \mathbb{R}_+^I$ is upper semicontinuous.

²We discuss potential issues in extending this work to users with elastic demands and heterogeneous quality-of-service requirements in the concluding section.

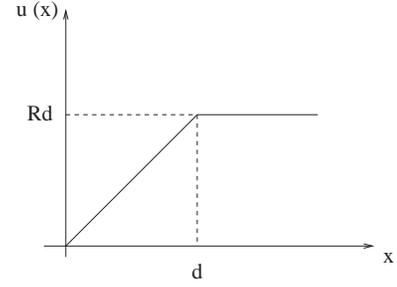


Fig. 2. Aggregate utility function.

Proof sketch: Given any $p \geq 0$, the proof is based on using Assumption 1 (in particular the nondecreasing assumption on the latency functions) to show that the set of WE is given by the set of optimal solutions of the following optimization problem

$$\begin{aligned} \max_{x \geq 0} \quad & \sum_{i \in \mathcal{I}} \left((R - \sum_{j \in \mathcal{N}_i} p_j) x_i - \int_0^{x_i} \sum_{j \in \mathcal{N}_i} l_j(z) dz \right) \quad (2) \\ \text{s.t.} \quad & \sum_{i \in \mathcal{I}} x_i \leq d. \end{aligned}$$

Q.E.D.

For a given price vector p , the WE need not be unique in general. Under further restrictions on the l_i , we obtain:

Proposition 2 (Uniqueness) Let Assumption 1 hold. Assume further that for all $i \in \mathcal{I}$, there exists some $j \in \mathcal{N}_i$, such that l_j is strictly increasing. Then, for any price vector $p \geq 0$, the set of WE, $W(p)$, is a singleton. Moreover, the function $W : \mathbb{R}_+^I \mapsto \mathbb{R}_+^I$ is continuous.

Proof: Under the given assumptions, for any $p \geq 0$, the objective function of problem (2) is strictly convex, and therefore has a unique optimal solution. This shows the uniqueness of the WE at a given p . Since the correspondence W is upper semicontinuous from Proposition 1 and single-valued, it is continuous. **Q.E.D.**

We next define the social problem and the social optimum, which is the routing (flow allocation) that would be chosen by a central network planner that has full control and information about the network.

Definition 2 A flow vector x^S is a *social optimum* if it is an optimal solution of the *social problem*

$$\begin{aligned} \text{maximize}_{x \geq 0} \quad & \sum_{i \in \mathcal{I}} \left(R - \sum_{j \in \mathcal{N}_i} l_j(x_i) \right) x_i \quad (3) \\ \text{subject to} \quad & \sum_{i \in \mathcal{I}} x_i \leq d. \end{aligned}$$

By Assumption 1, the social problem has a continuous objective function and a compact constraint set, guaranteeing the existence of a social optimum, x^S . Moreover, using the

optimality conditions for a convex program, we see that a vector $x^S \in \mathbb{R}_+^I$ is a social optimum if and only if $\sum_{i \in \mathcal{I}} x_i^S \leq d$ and there exists a $\lambda^S \geq 0$ such that $\lambda^S (\sum_{i=1}^I x_i^S - d) = 0$ and for each $i \in \mathcal{I}$,

$$\begin{aligned} R - \sum_{j \in \mathcal{N}_i} l_j(x_i^S) - x_i^S \sum_{j \in \mathcal{N}_i} l'_j(x_i^S) &\leq \lambda^S \quad \text{if } x_i^S = 0, \\ &= \lambda^S \quad \text{if } x_i^S > 0. \end{aligned} \quad (4)$$

For a given vector $x \in \mathbb{R}_+^I$, we define the value of the objective function in the social problem,

$$\mathbb{S}(x) = \sum_{i \in \mathcal{I}} \left(R - \sum_{j \in \mathcal{N}_i} l_j(x_i) \right) x_i, \quad (5)$$

as the *social surplus*, i.e., the difference between the users' willingness to pay and the total latency.

III. OLIGOPOLY PRICING AND EQUILIBRIUM

We assume that there are multiple service providers, each owning one of the links on the paths in the network. Each link may represent a more general subnetwork operated by independent service providers.³ Service provider j charges a price p_j per unit bandwidth on link $j \in \mathcal{N}_i$. Given the vector of prices of links owned by other service providers, $p_{-j} = [p_k]_{k \neq j}$, the profit of service provider j with $j \in \mathcal{N}_i$ is

$$\Pi_j(p_j, p_{-j}, x) = p_j x_i,$$

where $x \in W(p_j, p_{-j})$.

The objective of each service provider is to maximize profits. Because their profits depend on the prices set by other service providers, each service provider forms conjectures about the actions of other service providers, as well as the behavior of users, which they do according to the notion of subgame perfect Nash equilibrium. We refer to the game among service providers as the *price competition game*.

Definition 3 A vector $(p^{OE}, x^{OE}) \geq 0$ is a (pure strategy) *Oligopoly Equilibrium* (OE) if $x^{OE} \in W(p_j^{OE}, p_{-j}^{OE})$ and for all $i \in \mathcal{I}$, $j \in \mathcal{N}_i$, and $p_j \geq 0$, $x \in W(p_j, p_{-j}^{OE})$,

$$\Pi_j(p_j^{OE}, p_{-j}^{OE}, x^{OE}) \geq \Pi_j(p_j, p_{-j}^{OE}, x). \quad (6)$$

We refer to p^{OE} as the *OE price*.

The next proposition shows that for linear latency functions, there exists a pure strategy OE. The proof relies on the explicit characterization of the OE prices (see Proposition 4 below), and therefore is provided in Appendix B.

Proposition 3 Let Assumption 1 hold and assume that the latency functions are linear. Then the price competition game has a pure strategy OE.

The existence result cannot be generalized to general convex latency functions as shown in the following example.

³This is because users care about the overall latency of a subnetwork, not about patterns of routing flows within the subnetwork operated by a service provider. The possibility of service providers operating multiple paths is considered in [7] and is ruled out here to simplify the notation.

Example 1 Consider a two path network with one link on each path. Let the total flow be $d = 1$. Assume that the latency functions are given by

$$l_1(x) = 0, \quad l_2(x) = \begin{cases} 0 & \text{if } 0 \leq x \leq \delta \\ \frac{x-\delta}{\epsilon} & \text{if } x \geq \delta, \end{cases}$$

for some $\epsilon > 0$ and $\delta > 1/2$, with the convention that when $\epsilon = 0$, $l_2(x) = \infty$ for $x > \delta$. It can be easily verified that there exists no pure strategy OE for small ϵ (see [7] for details).

Nevertheless, it can be shown that the price competition game always has a mixed strategy OE (see the analysis for a parallel link network in [7]).

IV. EFFICIENCY ANALYSIS

A. Inefficiency of OE

In this section, we study the efficiency properties of OE, and strict and strong OE (defined below). We consider price competition games that have pure strategy OE or strict OE (this set includes, but is larger than, networks with linear latency functions covered by Proposition 3). Given a parallel-link network with I paths, n_i links on path i , and latency functions $\{l_j\}_{(j \in \mathcal{N}_i, i \in \mathcal{I})}$, let $\overline{OE}(\{l_j\})$ denote the set of flow allocations $x^{OE} = [x_i^{OE}]_{i \in \mathcal{I}}$ at an OE (or strict OE depending on the context).

We define the efficiency metric at some $x^{OE} \in \overline{OE}(\{l_j\})$ as

$$r_I(\{l_j\}, x^{OE}) = \frac{\mathbb{S}(x^{OE})}{\mathbb{S}(x^S)}, \quad (7)$$

where x^S is a social optimum given the latency functions $\{l_j\}$ [cf. Eq. (5)]. We adopt the convention that for the efficiency metric $0/0 = 1$. Following the literature on the ‘‘price of anarchy,’’ (see [10]), we are interested in the worst performance of an oligopoly equilibrium, so we look for a lower bound on

$$\inf_{\{l_j\}} \inf_{x^{OE} \in \overline{OE}(\{l_j\})} r_I(\{l_j\}, x^{OE}).$$

We first show that the performance of an OE can be arbitrarily bad.

Example 2 Consider a two path network, which has 3 links on path 1 with identically 0 latency functions and one link on path 2 with latency function $l(x_2) = kx_2$, where $k \geq 1/2$. Let the total flow be $d = 1$ and the reservation utility be $R = 1$.

The unique social optimum for this example is $x^S = (1, 0)$. Now consider the following strategy combination. Each of the three service providers on path 1, denoted by $i = 1, 2$, and 3, charge price $p_1^i = 1$, while the service provider on path 2 charges $p_2 = 1/2$. It can be verified that there is no deviation that is profitable for any of the service providers. First, consider the serial providers on path 1; given the prices of two of the serial service providers, there will always be zero traffic on path 1, so the remaining service provider is playing a best response (since any price for this provider would lead to zero profits). Moreover, it can be verified that these strategies are not weakly dominated, since if $i = 1, 2$ were to play $p_1^i = 0$ and the provider on path 2 were to set a high enough p_2 , $i = 3$

would choose to play $p_1^3 = 1$. (This also establishes that the OE will be trembling hand perfect, see [23], pp. 351-356 for a definition). Finally, let us consider the provider on path 2. Given the strategies of the serial providers on path 1 and a fixed $k > 0$, it can be verified that the optimal strategy of this provider is to set $p_2 = 1/2$. The resulting equilibrium flow allocation is

$$x^{OE} = \left[0, \frac{1}{2k}\right],$$

which involves routing all the admitted traffic on path 2 (though not all of the traffic is necessarily admitted). Therefore, the efficiency metric for this example is

$$r_2(\{l_j\}, x^{OE}) = \frac{\sum_{i=1}^2 x_i^{OE} - l_2(x_2^{OE})x_2^{OE}}{1} = \frac{1/4k}{1},$$

which goes to 0 as $k \rightarrow \infty$.

Example 2 establishes that pure strategy OE with the parallel-serial link topology can be arbitrarily inefficient. This result is at some level pathological, however. The reason for the arbitrary inefficiency of the pure strategy OE is the “unreasonably” high prices charged by the service providers along path 1. It is a best response (even weakly undominated) strategy for them to do so, because other providers also charge unreasonably high prices, so there is no transmission on this path and they suffer no adverse consequences from charging unreasonable prices.

We may expect this pathological situation not to arise in practice for a number of reasons. First, firms may not coordinate on such an equilibrium (especially when other equilibria exist). In this case, for example, we may expect them to realize that if they all reduced their prices, they would all make higher profits and would still be playing equilibrium actions. Second and relatedly, we may expect providers on a path to form a “coalition” and coordinate their pricing decisions (at least within some bounds). The coalition-proof Nash equilibrium concept of Bernheim, Peleg and Whinston [24] naturally leads to such an outcome and can be shown to rule out the situation in Example 2. The application of this concept to the current setup is further discussed in Section V. Finally, an application of the stronger concept of equilibrium introduced by Harsanyi, the *strict equilibrium* (see [8], or [23], pp. 11-12), will be shown to rule out the pattern in Example 2. In the current paper, we pursue this last approach and strengthen the concept of equilibrium to *strict oligopoly equilibrium*.

B. Strict OE and Price Characterization

Harsanyi’s concept of strict equilibrium requires each player’s best response to be unique. Recall that the standard Nash equilibrium and our OE concept only require each player, in particular each service provider, to play a weak best response. We now strengthen this condition.

Definition 4 A vector $(p^{OE}, x^{OE}) \geq 0$ is a *strict OE* (Oligopoly Equilibrium) if $x^{OE} \in W(p_j^{OE}, p_{-j}^{OE})$ and for

all $i \in \mathcal{I}$, $j \in \mathcal{N}_i$, and for all $p_j \geq 0$, $p_j \neq p_j^{OE}$, $\forall x \in W(p_j, p_{-j}^{OE})$,

$$\Pi_j(p_j^{OE}, p_{-j}^{OE}, x^{OE}) > \Pi_j(p_j, p_{-j}^{OE}, x). \quad (8)$$

We refer to p^{OE} as the *strict OE price*.

In the remainder of this paper, we focus on strict OE and we use the notation $\overrightarrow{OE}(\{l_j\})$ to denote the set of flow allocations $x^{OE} = [x_i^{OE}]_{i \in \mathcal{I}}$ at a strict OE for a network with latency functions $\{l_j\}_{(j \in \mathcal{N}_i, i \in \mathcal{I})}$.

The difference between Definitions 3 and 4 is obvious. The latter requires service providers to play a strict best response, while the former does not. Notice that in both equilibria, we have not changed the behavior of the users given by the WE (as in Definition 1). Notice also that we have removed the qualifier “pure strategy,” since as is well known, strict equilibria always have to be pure strategy (because mixed strategy equilibria, by definition, involve players being indifferent among the strategies over which they are mixing). Therefore, there are situations in which a mixed strategy OE exists, but strict OE does not. Moreover, it can be verified that there are also situations in which a pure strategy OE exists, but a strict OE does not.

We do not view this as a serious shortcoming, however, since, as Example 1 above showed, even pure strategy OE do not always exist. Moreover, Proposition 5 below establishes that when latency functions are linear, a unique strong OE (which is a stronger version of strict OE) exists.

We next provide an explicit characterization of the strict OE prices, which will be essential for the subsequent efficiency analysis. The following lemma establishes that all path flows are positive at a strict OE.

Lemma 1 Let (p^{OE}, x^{OE}) be a strict OE. Let Assumption 1 hold. Then $p_j^{OE} x_i^{OE} > 0$ for all $i \in \mathcal{I}$ and $j \in \mathcal{N}_i$.

Proof: Assume to arrive at a contradiction that $p_j^{OE} x_i^{OE} = 0$ for some $i \in \mathcal{I}$ and $j \in \mathcal{N}_i$. Then, at any price \bar{p}_j with $\bar{p}_j > p_j^{OE}$, we have

$$\Pi_j(\bar{p}_j, p_{-j}^{OE}, x) = \Pi_j(p_j^{OE}, p_{-j}^{OE}, x^{OE}),$$

contradicting the definition of the strict OE (cf. Definition 8).

Q.E.D.

As shown in Example 2, the result of the preceding lemma does not extend to non-strict OE prices, i.e., there may be OE in which some of the providers make zero profit while others are making positive profits. We have shown in [7] that for parallel-link topology, if at any OE one of the providers makes positive profit, all of the providers make positive profits (see [7], Lemma 4). Example 2 shows that this result no longer holds for non-strict OE for the parallel-serial topology. Lemma 1, on the other hand, ensures that it holds for strict OE and allows us to write the optimization problems for each provider in terms of equality and inequality constraints. We can then use the first order optimality conditions to obtain an explicit characterization of the strict OE prices (see Appendix A for the proof).

Proposition 4 Let (p^{OE}, x^{OE}) be a strict OE. Let Assumption 1 hold. Then, for all $i \in \mathcal{I}$, $j \in \mathcal{N}_i$, we have

$$(a) \quad p_j^{OE} \geq x_i^{OE} \sum_{k \in \mathcal{N}_i} l'_k(x_i^{OE}).$$

$$(b) \quad p_j^{OE} = \begin{cases} x_i^{OE} \sum_{k \in \mathcal{N}_i} l'_k(x_i^{OE}), & \text{if } l'_k(x_s^{OE}) = 0, \\ & \text{for some } s \neq i, \text{ for all } k \in \mathcal{N}_s, \\ \min \left\{ \frac{1}{n_i} \left[R - \sum_{k \in \mathcal{N}_i} l_k(x_i^{OE}) \right], \right. \\ \left. x_i^{OE} \left[\sum_{k \in \mathcal{N}_i} l'_k(x_i^{OE}) + \frac{1}{\sum_{s \neq i} \sum_{k \in \mathcal{N}_s} l'_k(x_s^{OE})} \right] \right\}, & \text{otherwise.} \end{cases} \quad (9)$$

In particular, for two links, when the minimum effective cost is less than R , for $i = 1, 2$, $j \in \mathcal{N}_i$, the strict OE prices are given by

$$p_j^{OE} = x_i^{OE} \left[\sum_{k \in \mathcal{N}_1} l'_k(x_1^{OE}) + \sum_{k \in \mathcal{N}_2} l'_k(x_2^{OE}) \right].$$

The price characterization in Proposition 4 is a generalization of the price characterization in [7], and as in that paper, it will be useful in providing bounds on the inefficiency of price competition. However, the next example shows that even with strict OE, efficiency losses can be arbitrarily large.

C. Inefficiency of Strict OE

Example 3 Consider a one path network, which has n links with identical latency functions $l(x) = x/n$. Let the total flow be $d = 1$ and the reservation utility be $R = 1$.

For any n , the unique social optimum for this example is $x^S = 1/2$, with a corresponding social surplus $\mathbb{S}(x^S) = 1/4$. Using the price characterization given in Proposition 4 and the definition of a WE, it follows that there exists a unique strict OE, in which all providers charge the price $p^{OE} = 1/(n+1)$, and the equilibrium flow is $x^{OE} = 1/(n+1)$. The efficiency metric for this example is therefore

$$r_1(\{l_j\}, x^{OE}) = \frac{\left(1 - \frac{1}{n+1}\right) \frac{1}{n+1}}{\left(1 - \frac{1}{2}\right) \frac{1}{2}} = \frac{4n}{(n+1)^2},$$

which goes to 0 as $n \rightarrow \infty$.

This example establishes that even with strict OE, which rules out the pathological coordination failures discussed above, efficiency losses can be arbitrarily large. The reason for this is again the double marginalization problem, which increases the cost of transmission so much that there is no transmission in equilibrium along certain paths (e.g., along the single path in the example as $n \rightarrow \infty$). It is also evidence that the structure depicted in Example 3 can be part of any

general network topology, and thus establishes that strict OE can be arbitrarily inefficient in general networks.

Despite this simplicity and potential generality, the behavior in Example 3 is still somewhat pathological, since it relies on the double marginalization problem reducing the transmission. This may be thought to be unlikely particularly in networks where the reservation utility, R , of users is high enough. This leads us to define an even stronger notion of equilibrium, *strong OE*.⁴

Definition 5 A vector $(p^{OE}, x^{OE}) \geq 0$ is a *strong OE* (Oligopoly Equilibrium) if it is a strict OE, and $\sum_{i \in \mathcal{I}} x_i^{OE} = d$. In this case, we refer to p^{OE} as the *strong OE price* and denote the set of strong OE flow allocations in a network with latency functions by $\{l_j\}$ by $\overrightarrow{OE}^d(\{l_j\})$.

The only difference between Definition 4 and Definition 5 is that in the latter we require all of the potential flow, d , to be transmitted. This will be the case when the reservation utility, R , of users is large enough. The following proposition establishes that this is indeed the case and in fact when R is large enough there exists a unique strong OE. The proof of the proposition is provided in Appendix C.

Proposition 5 Let Assumption 1 hold. Assume further that the latency functions are linear, i.e., $l_j(x) = a_j x$ for all $j \in \mathcal{N}_i$, and all $i \in \mathcal{I}$. Define the set

$$\mathcal{I}_0 = \left\{ i \in \mathcal{I} \mid \sum_{j \in \mathcal{N}_i} a_j = 0 \right\}.$$

Let I_0 denote the cardinality of set \mathcal{I}_0 and assume that $I_0 \leq 1$. Then, there exists some \bar{R} , such that for all reservation utilities $R \geq \bar{R}$, the price competition game has a unique strong OE.

Note that the assumption $I_0 \leq 1$ in the preceding proposition cannot be dispensed with, i.e., without this assumption, we can have situations in which there exists a pure strategy OE, but no strict OE. To see this, consider a two-link parallel network where both latency functions are identically equal to zero [i.e., $l_i(x) = 0$, for $i = 1, 2$ and for all x]. In this case, there exists a unique pure strategy OE, identical to the standard Bertrand equilibrium, where both service providers charge $p_1 = p_2 = 0$. It can be verified, however, that this is not a strict best response for either of them, thus a strict OE does not exist (there are no other pure strategy OE and mixed strategy OE cannot be strict).

D. Efficiency of Strong OE with Two Paths

We now characterize the efficiency properties of strong OE. We start with a two path network, with n_i links on path $i = 1, 2$, where each link is owned by a different provider. First, consider the following example, which illustrates that even

⁴Note that this notion is unrelated to Aumann's notion of "strong equilibrium," which requires a Nash equilibrium to have the property that no coalition of players should be able to jointly deviate, taking the actions of all other players as given, and increase the payoffs to all the members of the coalition (see [25], [24]). The notion of coalition-proof Nash equilibrium discussed in Section V is a weaker version of Aumann's strong equilibrium.

with strong OE the efficiency loss can be worse than that in parallel link networks (which was shown to be bounded below by 5/6 in [7]).

Example 4 Consider a two path network, which has n links on path 1 with identically 0 latency functions and one link on path 2 with latency function $l(x_2) = x_2/2$. Let the total flow be $d = 1$ and the reservation utility be $R = 1$.

The unique social optimum for this example is $x^S = (1, 0)$. Using Proposition 4 and the definition of a WE, OE flows x^{OE} must satisfy

$$\begin{aligned} & \sum_{j \in \mathcal{N}_1} l_j(x_1^{OE}) + x_1^{OE} \left[\sum_{j \in \mathcal{N}_1} l'_j(x_1^{OE}) + \sum_{j \in \mathcal{N}_2} l'_j(x_2^{OE}) \right] \\ &= \sum_{j \in \mathcal{N}_2} l_j(x_2^{OE}) + x_2^{OE} \left[\sum_{j \in \mathcal{N}_1} l'_j(x_1^{OE}) + \sum_{j \in \mathcal{N}_2} l'_j(x_2^{OE}) \right]. \end{aligned}$$

Substituting for the latency functions and solving the above together with $x_1^{OE} + x_2^{OE} = 1$ shows that unique strong OE involves

$$x^{OE} = \left(\frac{2}{n+2}, \frac{n}{n+2} \right),$$

which goes to $(0, 1)$ as $n \rightarrow \infty$. The social surplus at the social optimum is 1, while the social surplus at the strong OE goes to 1/2 as $n \rightarrow \infty$.

We next present two lemmas, which will be useful in providing a bound on the efficiency metric for strong OE. Note that these lemmas are valid for all OE. The first lemma is straightforward and allows us to assume without loss of generality that $R \sum_{i=1}^I x_i^S - \sum_{i=1}^I l_i(x_i^S) x_i^S > 0$ in the subsequent analysis.

Lemma 2 Given a set of latency functions $\{l_j\}_{j \in \mathcal{N}_i}$, $i \in \mathcal{I}$, assume that

$$\sum_{i \in \mathcal{I}} \left(\sum_{j \in \mathcal{N}_i} l_j(x_i^S) \right) x_i^S = R \sum_{i \in \mathcal{I}} x_i^S,$$

for some social optimum x_S . Then every $x^{OE} \in \overrightarrow{OE}(\{l_j\})$ is a social optimum, implying that $r_I(\{l_j\}, x^{OE}) = 1$.

The following lemma provides a relation between the total flow admitted at an OE and at a social optimum.

Lemma 3 For a set of latency functions $\{l_j\}_{(j \in \mathcal{N}_i, i \in \mathcal{I})}$, let Assumption 1 hold. Let (p^{OE}, x^{OE}) be an OE and x^S be a social optimum. Then

$$\sum_{i \in \mathcal{I}} x_i^{OE} \leq \sum_{i \in \mathcal{I}} x_i^S.$$

Proof: Assume to arrive at a contradiction that $\sum_{i \in \mathcal{I}} x_i^{OE} > \sum_{i \in \mathcal{I}} x_i^S$. This implies that $x_i^{OE} > x_i^S$ for some i . Hence,

$$l_j(x_i^{OE}) \geq l_j(x_i^S), \quad \forall j \in \mathcal{N}_i.$$

We also have $l_j(x_i^{OE}) > l_j(x_i^S)$ for some $j \in \mathcal{N}_i$. [Otherwise, we would have $l_j(x_i^S) = l'_j(x_i^S) = 0$ for all $j \in \mathcal{N}_i$, which yields a contradiction by the optimality conditions (4) and the

fact that $\sum_{i \in \mathcal{I}} x_i^S < d$.] Using the definition of the WE and the optimality conditions (4), we obtain

$$R - \sum_{j \in \mathcal{N}_i} \left(l_j(x_i^{OE}) - p_j^{OE} \right) \geq R - \sum_{j \in \mathcal{N}_i} \left(l_j(x_i^S) - x_i^S l'_j(x_i^S) \right).$$

Combining the preceding with $l_j(x_i^{OE}) \geq l_j(x_i^S)$ for all $j \in \mathcal{N}_i$, with strict inequality for some j , and

$$p_j^{OE} \geq x_i^{OE} l'_j(x_i^{OE}) \geq x_i^S l'_j(x_i^S),$$

[using Proposition 4(a) and the fact that $x l'(x)$ is nondecreasing, cf. Assumption 1], we obtain a contradiction. **Q.E.D.**

The next theorem provides a tight lower bound on $r_2(\{l_j\}, x^{OE})$ [cf. (7)] for a strong OE. In the following, we assume without loss of generality that $d = 1$.

Theorem 1 Consider a two path network, with n_i links on path $i = 1, 2$, where each link is owned by a different provider, and link $j \in \mathcal{N}_i$ has a latency function l_j . Suppose that Assumption 1 holds and the price competition game has a strong OE. Then

$$r_2(\{l_j\}, x^{OE}) \geq \frac{1}{2}, \quad \forall x^{OE} \in \overrightarrow{OE}^d(\{l_j\}). \quad (10)$$

Moreover, the bound is tight, i.e., there exists $\{l_j\}$ and $x^{OE} \in \overrightarrow{OE}^d(\{l_j\})$ that attains the lower bound in (10).

Proof: The proof follows a number of steps:

Step 1: We are interested in finding a lower bound for the problem

$$\inf_{\{l_j\}} \inf_{x^{OE} \in \overrightarrow{OE}^d(\{l_j\})} r_2(\{l_j\}, x^{OE}). \quad (11)$$

Given $\{l_j\}$, let $x^{OE} \in \overrightarrow{OE}(\{l_j\})$ and let x^S be a social optimum. By Lemma 3 and the fact that $x^{OE} \in \overrightarrow{OE}^d(\{l_j\})$ (i.e., it is a strong OE), we have

$$\sum_{i=1}^2 x_i^{OE} = \sum_{i=1}^2 x_i^S = 1.$$

This implies that there exists some i such that $x_i^{OE} < x_i^S$. Since the problem is symmetric, we can restrict ourselves to $\{l_j\}$ for which $x_1^{OE} < x_1^S$. We claim that

$$\inf_{\{l_i\} \in \mathcal{L}_2} \inf_{x^{OE} \in \overrightarrow{OE}^d(\{l_i\})} r_2(\{l_i\}, x^{OE}) \geq r_2^{OE}, \quad (12)$$

where problem (E) is given by

$$r_2^{OE} = \min_{\substack{l_{i,j}^S, (l_{i,j}^S)' \geq 0 \\ l_{i,j}, l'_{i,j} \geq 0 \\ y_i^S, y_i^{OE} \geq 0}} \frac{R - y_1^{OE} \left(\sum_{j \in \mathcal{N}_1} l_{1,j} \right) - y_2^{OE} \left(\sum_{j \in \mathcal{N}_2} l_{2,j} \right)}{R - y_1^S \left(\sum_{j \in \mathcal{N}_1} l_{1,j}^S \right) - y_2^S \left(\sum_{j \in \mathcal{N}_2} l_{2,j}^S \right)}$$

subject to

$$l_{i,j}^S \leq y_i^S (l_{i,j}^S)', \quad i = 1, 2, j \in \mathcal{N}_i, \quad (13)$$

$$\left(\sum_{j \in \mathcal{N}_2} l_{2,j}^S \right) + y_2^S \left(\sum_{j \in \mathcal{N}_2} (l_{2,j}^S)' \right) = \left(\sum_{j \in \mathcal{N}_1} l_{1,j}^S \right) + y_1^S \left(\sum_{j \in \mathcal{N}_1} (l_{1,j}^S)' \right), \quad (14)$$

$$\begin{aligned} \left(\sum_{j \in \mathcal{N}_1} l_{1,j}^S \right) + y_1^S \left(\sum_{j \in \mathcal{N}_1} (l_{1,j}^S)' \right) &\leq R, \\ \sum_{i=1}^2 y_i^S &= 1, \\ l_{1,j} + l'_{1,j} (y_1^S - y_1^{OE}) &\leq l_{1,j}^S, \quad \forall j \in \mathcal{N}_1, \quad (15) \\ l_{i,j} &\leq y_i^{OE} l'_{i,j}, \quad i = 1, 2, j \in \mathcal{N}_i, \quad (16) \\ \sum_{i=1}^2 y_i^{OE} &= 1, \\ &+ \text{ Strict OE Constraints.} \end{aligned}$$

Problem (E) can be viewed as a finite dimensional problem that captures the equilibrium and social optimum characteristics of the infinite dimensional problem given in (11). This implies that instead of optimizing over the entire function l_j for some $j \in \mathcal{N}_i$, $i \in \mathcal{I}$, we optimize over the possible values of $l_j(\cdot)$ and $l'_j(\cdot)$ at the equilibrium and the social optimum, which we denote by $l_{i,j}, l'_{i,j}, l_{i,j}^S, (l_{i,j}^S)'$. The constraints of the problem guarantee that these values satisfy the necessary optimality conditions for a social optimum and a strict OE (which are the same as the conditions for a strong OE). In particular, conditions (13) and (16) capture the convexity assumption on $l_j(\cdot)$ by relating the values $l_{i,j}, l'_{i,j}$ and $l_{i,j}^S, (l_{i,j}^S)'$ [note that the assumption $l_j(0) = 0$ is essential here]. Condition (14) is the optimality condition for the social optimum. Condition (15) uses the nondecreasing and the convexity assumption on the latency functions; since we are focusing on $\{l_j(\cdot)\}$ such that $x_1^{OE} \leq x_1^S$, we must have

$$l_{1,j} + l'_{1,j} (y_1^S - y_1^{OE}) \leq l_{1,j}^S,$$

for all $j \in \mathcal{N}_1$. Finally, the last set of constraints are the necessary conditions for a pure strategy OE. In particular, for a two path network, using Proposition 4, the Strict OE Constraints are given by

$$\begin{aligned} &n_1 y_1^{OE} \left[\sum_{j \in \mathcal{N}_1} l'_{1,j} + \sum_{j \in \mathcal{N}_2} l'_{2,j} \right] + \sum_{j \in \mathcal{N}_1} l_{1,j} \\ &= n_2 y_2^{OE} \left[\sum_{j \in \mathcal{N}_1} l'_{1,j} + \sum_{j \in \mathcal{N}_2} l'_{2,j} \right] + \sum_{j \in \mathcal{N}_2} l_{2,j}, \end{aligned}$$

[and therefore n_1 and n_2 are also decision variables in problem (E)]. Note that given any feasible solution of problem (11), we have a feasible solution for problem (E) with the same objective function value. Therefore, the optimum value of problem (E) is indeed a lower bound on the optimum value of problem (11).

Step 2: Consider the following change of variables for problem (E)

$$\begin{aligned} l_1^S &= \sum_{j \in \mathcal{N}_1} l_{1,j}^S, & l_2^S &= \sum_{j \in \mathcal{N}_2} l_{2,j}^S \\ l_1 &= \sum_{j \in \mathcal{N}_1} l_{1,j}, & l_2 &= \sum_{j \in \mathcal{N}_2} l_{2,j}, \\ (l_1^S)' &= \sum_{j \in \mathcal{N}_1} (l_{1,j}^S)', & (l_2^S)' &= \sum_{j \in \mathcal{N}_2} (l_{2,j}^S)', \end{aligned}$$

$$l'_1 = \sum_{j \in \mathcal{N}_1} l'_{1,j}, \quad l'_2 = \sum_{j \in \mathcal{N}_2} l'_{2,j},$$

and rewrite problem (E) as

$$r_2^{OE} = \text{minimize}_{\substack{l_i^S, (l_i^S)' \geq 0 \\ l_i, l'_i \geq 0 \\ y_i^S, y_i^{OE} \geq 0}} \frac{R - l_1 y_1^{OE} - l_2 y_2^{OE}}{R - l_1^S y_1^S - l_2^S y_2^S} \quad (E')$$

subject to

$$\begin{aligned} l_i^S &\leq y_i^S (l_i^S)', \quad i = 1, 2, \\ l_2^S + y_2^S (l_2^S)' &= l_1^S + y_1^S (l_1^S)', \\ l_1^S + y_1^S (l_1^S)' &\leq R, \end{aligned}$$

$$\sum_{i=1}^2 y_i^S = 1,$$

$$l_1 + l'_1 (y_1^S - y_1^{OE}) \leq l_1^S,$$

$$l_i \leq y_i^{OE} l'_i, \quad i = 1, 2,$$

$$\sum_{i=1}^2 y_i^{OE} = 1,$$

+ Strict OE Constraints.

Note that this problem has a very similar structure to the finite-dimensional problem considered in the proof of Theorem 1 of [7] for parallel-link networks. Let $(\bar{l}_i^S, (\bar{l}_i^S)', \bar{l}_i, \bar{l}'_i, \bar{y}_i^S, \bar{y}_i^{OE})$ denote the optimal solution of problem (E'). We have shown in [7] that $\bar{l}_i^S = 0$ for $i = 1, 2$.

Step 3: Using $\bar{l}_i^S = 0$ for $i = 1, 2$, and $\bar{l}_1 = 0, \bar{l}'_1 = 0$, we see that

$$r_2^{OE} = \min_{\substack{l_2, l'_2 \\ y_1^{OE}, y_2^{OE} \geq 0 \\ n_1, n_2 \geq 1}} 1 - \frac{l_2 y_2^{OE}}{R} \quad (17)$$

$$\begin{aligned} \text{subject to} \quad &l_2 \leq y_2^{OE} l'_2, \\ &l_2 + n_2 y_2^{OE} l'_2 = n_1 y_1^{OE} l'_2, \\ &n_1 y_1^{OE} l'_2 \leq R. \\ &\sum_{i=1}^2 y_i^{OE} \leq 1. \end{aligned}$$

Next, using the transformation $m_1 = n_1 y_1^{OE}$ and $m_2 = n_2 y_2^{OE}$ to write:

$$r_2^{OE} = \min_{\substack{l_2, l'_2 \\ y_1^{OE}, y_2^{OE} \geq 0 \\ m_1, m_2 \geq 0}} 1 - \frac{l_2 y_2^{OE}}{R} \quad (18)$$

$$\begin{aligned} \text{subject to} \quad &l_2 \leq y_2^{OE} l'_2, \\ &l_2 + m_2 l'_2 = m_1 l'_2, \\ &m_1 l'_2 \leq R. \\ &\sum_{i=1}^2 y_i^{OE} = 1, \end{aligned}$$

though we also have to ensure that the solution to this program ensures that n_1 and n_2 are integers.

Now it can be verified that $(\bar{l}_2, \bar{l}_2, \bar{y}_1^{OE}, \bar{y}_2^{OE}, \bar{m}_1, \bar{m}_2) = (\frac{R}{2}, \frac{R}{2}, 0, 1, 2, 1)$ is an optimal solution to the program (18), and moreover, it satisfies $n_1, n_2 \geq 1$, thus it is also a solution to (17). The corresponding optimum value is $r_2^{OE} = 1/2$. By (12), this implies that

$$\inf_{\{l_j\}} \inf_{x^{OE} \in \overrightarrow{OE}(\{l_j\})} r_2(\{l_j\}, x^{OE}) \geq \frac{1}{2}.$$

Finally, Example 4 shows that this bound is tight, i.e.,

$$\min_{\{l_j\}} \min_{x^{OE} \in \overrightarrow{OE}(\{l_j\})} r_2(\{l_j\}, x^{OE}) = \frac{1}{2}.$$

Q.E.D.

This theorem shows that there exists a tight bound of $1/2$ for strong OE under the assumption of zero latency without any congestion [$l_i(0) = 0$]. In contrast to the case in Example 3, strong OE ensures that all of the traffic is transmitted in equilibrium, which is the key to the existence of a bound on the inefficiency of equilibrium.

The bound with strong OE is nonetheless worse than the efficiency bound in the parallel-link topology considered in [7]. This is again because of the double marginalization problem: each provider along path 1 has a greater incentive to increase its price (relative to the benchmark where all these links are owned by the same provider), because it does not internalize the reduction in the profits of the other link owners along the same path. Consequently, in Example 4, there are higher prices along path 1, and this induces greater fraction of users to choose path 2, increasing inefficiency. To see the role of serial links more clearly, consider a modified version of Example 4, where all n links along path 1 are owned by the same service provider. This would make the example equivalent to a parallel-link topology. In this case the unique strict OE flows are given by $x_1^{OE} = 2/3$ and $x_2^{OE} = 1/3$, and this example reaches the $5/6$ bound of [7] rather than $1/2$ bound of Example 4.

E. Efficiency of Strong OE with Multiple Paths

We next consider a more general network consisting of I paths, with n_i links on path i , where each link is owned by a different provider. The following example illustrates the efficiency properties of a strong OE in an I path network.

Example 5 Consider an I path network, which has n links on path 1 with identically 0 latency functions and one link on each of the paths $2, \dots, I$ with the same latency function $l(x) = x(I-1)/2$. Let the total flow be $d = 1$ and the reservation utility be $R = 1$.

Clearly, the unique social optimum for this example is $x^S = [1, 0, \dots, 0]$. Using Proposition 4 and the definition of a WE, it can be seen that the flow allocation at the unique strict (strong) OE is

$$x^{OE} = \left[\frac{2/n}{1 + 2/n}, \frac{1}{(I-1)(1 + 2/n)}, \dots, \frac{1}{(I-1)(1 + 2/n)} \right].$$

Hence the efficiency metric for this example is

$$r_I(\{l_j\}, x^{OE}) = 1 - \frac{1}{2} \left(\frac{1}{1 + 2/n} \right)^2,$$

which goes to $1/2$ as $n \rightarrow \infty$.

The next theorem generalizes Theorem 1. The proof is similar to that of Theorem 1 and is omitted.

Theorem 2 Consider a general I path network, with n_i links on path $i \in \mathcal{I}$, where each link is owned by a different provider, and link j , $j \in \mathcal{N}_i$, has a latency function l_j . Suppose that Assumption 1 holds and the price competition game has a strong OE. Then

$$r_I(\{l_j\}, x^{OE}) \geq \frac{1}{2}, \quad \forall x^{OE} \in \overrightarrow{OE}^d(\{l_j\}). \quad (19)$$

Moreover, the bound is tight, i.e., there exists $\{l_j\}$ and $x^{OE} \in \overrightarrow{OE}^d(\{l_j\})$ that attains the lower bound in (19).

An important implication of this theorem and of Example 5 is that the bound of $1/2$ is tight even with an arbitrarily large number of paths. Naturally, such a tight bound could be obtained trivially when all paths except two have arbitrarily high latencies. Nevertheless, Example 5 shows that this bound is reached with positive flows on all paths for arbitrarily large networks. This result is interesting in part because models where a large number of oligopolists compete often converge to a competitive and efficient equilibrium, and yet this example shows this not to be the case in our model. It is important to note, however, that such a convergence result would apply as $n \rightarrow \infty$ when a given network is replicated n times. Instead, in examples where the metric of efficiency remains at $1/2$, the network in question is not a n -replication of another network.

F. Positive Latency at 0 Congestion

Interestingly, the bound on the efficiency loss of strong OE does not generalize once we relax the assumption that $l_i(0) = 0$.

Example 6 Consider a two path network, which has n links on path 1 with identically 0 latency functions and one link on path 2 with latency function $l(x_2) = \epsilon x_2 + b$ for some scalars $\epsilon > 0$ and $b > 0$. Again the unique social optimum is $\bar{x}^S = (1, 0)$. The flows at the unique strict (strong) OE are given by

$$\bar{x}^{OE} = \left(\frac{2\epsilon + b}{\epsilon(n+2)}, \frac{\epsilon n - b}{\epsilon(n+2)} \right).$$

Let $\epsilon = b/\sqrt{n}$. Then, as $b \rightarrow 1$ and $n \rightarrow \infty$, we have that $\bar{x}^{OE} \rightarrow (0, 1)$, and the efficiency metric $r_2(\{l_j\}, x^{OE}) \rightarrow 0$.

This example shows that the efficiency loss could be arbitrarily high even at a strong OE for a network that involves parallel and serial links if the assumption $l_i(0) = 0$ is relaxed. This establishes:

Proposition 6 In the presence of positive latency at zero congestion, strong OE with the parallel-serial topology can be arbitrarily inefficient.

It is useful to note that in the same example with the parallel-link topology (i.e., all n links along path 1 owned by the same provider), we would have

$$x^{OE} = \begin{cases} \left(\frac{b+2\epsilon}{3\epsilon}, \frac{\epsilon-b}{3\epsilon} \right), & \text{if } \epsilon \geq b, \\ (1, 0), & \text{otherwise.} \end{cases}$$

Consequently, $b \rightarrow 1$ and $\epsilon \rightarrow 0$, we have that $x^{OE} \rightarrow (1, 0)$, and $r_2(\{l_j\}, x^{OE}) \rightarrow 1$. Therefore, the highly inefficient equilibrium is a result of the parallel-serial topology, not of the assumption that there is positive latency at 0 congestion. In fact, [7] shows that with parallel topology, but positive latency at 0 congestion, there is again a tight bound of $2\sqrt{2} - 2$ on efficiency, which is quite close to, but slightly lower than $5/6$.

V. CONCLUSIONS

In this paper, we presented an analysis of price competition in communication networks with congestion. The focus has been the efficiency implications of price competition in networks with the serial-parallel topology.

Our major result is that contrary to the case of parallel-link topology studied in [7], the parallel-serial topology leads to significant efficiency losses relative to the social optimal. In particular, OE can now be arbitrarily inefficient. This is partly due to an extreme (pathological) form of double marginalization, whereby all serial providers on a particular path charge prohibitively high prices expecting others on that path to do so as well.

We showed that the concept of strict OE, which requires all service providers to play strict best responses, removes this pathological behavior, but the efficiency loss of strict OE is also unbounded because of the related double marginalization problem. In particular, the total cost of transmission on a path consisting of many serial providers can be prohibitively high that most of the users do not transmit in equilibrium, even though transmission of all the traffic is socially optimal.

Yet, when users value transmission sufficiently, we may expect them to transmit even with high costs. Motivated by this, we defined a stronger notion of equilibrium, strong OE, which is a strict OE with all of the traffic transmitted in equilibrium. For strong OE, we showed that as long as there is zero latency at zero congestion, there is a tight bound of $1/2$ on the inefficiency resulting from price competition.

Once the zero latency at zero congestion assumption is removed, however, there is no such tight bound even with strong OE, and the equilibrium can once again be arbitrarily inefficient.

The results in this paper add to the growing literature on the impact of game-theoretic interactions between service providers and users in communication networks. A number of questions in this area require further analysis:

- 1) All the examples of arbitrarily high inefficiency presented in this paper are under extreme configurations. Therefore, we suspect that these worst-case results may not be informative about the extent of degradation of performance in real-world network structures. An open area for further study is the quantification of inefficiency arising from price competition in “average” or “typical”

networks. The methods used by Friedman in his analysis of genericity of inefficiency of selfish routing are likely to be useful in this context as well (see [11]).

- 2) Our results suggest that competition between service providers can have significant costs in more general topologies, as long as these include serial service providers, causing the double marginalization problem. In fact, Example 3, which shows the possibility of arbitrarily large inefficiencies, can be part of any general network topology. Nevertheless, a characterization of the structure of networks that would lead to worst-case scenarios is an area for future research.
- 3) The most important simplifying feature of our analysis is the assumption that users are “homogeneous” in the sense that the same reservation utility, R , applies to all users. It is possible to conduct a similar analysis with elastic and heterogeneous users (or traffic) as in [17], [18], assuming that service providers are restricted to charge uniform prices to all users. Perhaps a more attractive alternative is to allow non-linear pricing and price differentiation, whereby service providers may charge different prices for different qualities of service and different delay guarantees (and let users sort between different types of services or contracts). This is an important research area for understanding equilibria in communication networks, where users often have heterogeneous quality of service requirements.
- 4) All of our efficiency bounds concern pure strategy equilibrium. Possible bounds for mixed strategy oligopoly equilibria is another open research question.
- 5) As discussed above, another interesting equilibrium notion to consider in models of competition in parallel-serial or more general topologies would be the coalition-proof Nash equilibrium concept of Bernheim, Peleg and Whinston, [24]. It can be shown that the pure strategy oligopoly equilibrium in Example 2 where all serial service providers charge prohibitively high prices is not coalition proof. In particular, in that game, the coalition consisting of all the serial providers along path 1 would have a “self-enforcing” deviation that would increase the payoff to each of them (this would be simply to reduce their prices simultaneously to allow positive transmission through path 1). The concept of coalition-proof Nash equilibrium may be attractive in the context of competition in general communication networks, since we may expect self-enforcing agreements between providers that are both cooperating and competing to emerge. The problem with coalition-proof Nash equilibria, however, is that such equilibria often fail to exist. Despite this potential shortcoming, an investigation of the structure and efficiency of coalition-proof Nash equilibria in communication networks with general topologies would be an interesting area for further study.

VI. APPENDIX A: PROOF OF PROPOSITION 4

Since (p^{OE}, x^{OE}) is a strict OE, it follows by Lemma 1 that $p_j^{OE} x_i^{OE} > 0$ for all $i \in \mathcal{I}$ and $j \in \mathcal{N}_i$. Using the definition

of a Wardrop equilibrium (cf. Definition 1), we have that for all $i \in \mathcal{I}$ and $j \in \mathcal{N}_i$, (p_j^{OE}, x^{OE}) is an optimal solution of the problem

$$\max_{(p_j, x) \geq 0} p_j x_i \quad (20)$$

subject to

$$\begin{aligned} p_j + \sum_{\substack{k \in \mathcal{N}_i \\ k \neq j}} p_k^{OE} + \sum_{k \in \mathcal{N}_i} l_k(x_i) &= \sum_{k \in \mathcal{N}_s} p_k^{OE} + l_k(x_s), \\ &\forall s \neq i, \\ p_j + \sum_{\substack{k \in \mathcal{N}_i \\ k \neq j}} p_k^{OE} + \sum_{k \in \mathcal{N}_i} l_k(x_i) &\leq R, \\ \sum_{s \in \mathcal{I}} x_s &\leq d. \end{aligned} \quad (21)$$

Using the first order optimality conditions of the preceding problem, we obtain

$$p_j^{OE} = x_i^{OE} \sum_{k \in \mathcal{N}_i} l'_k(x_i^{OE}) - \theta, \quad (22)$$

where

$$\theta = \begin{cases} 0, & \text{if } l'_k(x_s^{OE}) = 0 \text{ for some } s \neq i, \\ & \text{for all } k \in \mathcal{N}_s, \\ \frac{-x_i^{OE} - \xi}{\sum_{s \neq i} \frac{1}{\sum_{k \in \mathcal{N}_s} l'_k(x_s^{OE})}}, & \text{otherwise,} \end{cases}$$

and $\xi \geq 0$ is the Lagrange multiplier associated with constraint (21). Since $\theta \leq 0$, Eq. (22) yields part (a) of the proposition.

If $\min_{s \in \mathcal{I}} \left\{ \sum_{k \in \mathcal{N}_s} p_k^{OE} + l_k(x_s^{OE}) \right\} = R$, then, in view of the symmetry of the optimization problems of each provider on a serial link, it follows that

$$p_k^{OE} = \frac{1}{n_i} \left[R - \sum_{k \in \mathcal{N}_i} l_k(x_i^{OE}) \right], \quad \forall k \in \mathcal{N}_i, \forall i \in \mathcal{I}. \quad (23)$$

Assume next that $\min_{s \in \mathcal{I}} \left\{ \sum_{k \in \mathcal{N}_s} p_k^{OE} + l_k(x_s^{OE}) \right\} < R$. This implies that the constraint in Eq. (21) is slack, and therefore $\xi = 0$. Combining Eq. (22) with the relation in (23) yields the desired result. **Q.E.D.**

VII. APPENDIX B: PROOF OF PROPOSITION 3

Let $l_j(x) = a_j x$ for some $a_j \geq 0$. Define the set

$$\mathcal{I}_0 = \left\{ i \in \mathcal{I} \mid \sum_{j \in \mathcal{N}_i} a_j = 0 \right\},$$

(or equivalently, \mathcal{I}_0 is the set of $i \in \mathcal{I}$ such that $a_j = 0$ for all $j \in \mathcal{N}_i$). Let I_0 denote the cardinality of set \mathcal{I}_0 . There are two cases to consider:

Case 1: $I_0 \geq 2$. In this case, it is evident that a vector (p^{OE}, x^{OE}) with $p_j^{OE} = 0$ for all $i \in \mathcal{I}_0$, $j \in \mathcal{N}_i$ and $x^{OE} \in W(p^{OE})$ is an OE.

Case 2: $I_0 \leq 1$. In this case, for some $j \in \mathcal{N}_i$, let $B_j(p_{-j}^{OE})$ be the set of p_j^{OE} such that

$$(p_j^{OE}, x^{OE}) \in \arg \max_{\substack{p_j \geq 0 \\ x \in W(p_j, p_{-j}^{OE})}} p_j x_i. \quad (24)$$

Let $B(p^{OE}) = [B_j(p_{-j}^{OE})]$. By Berge's Theorem of the Maximum (see [26]), it follows that $B(p^{OE})$ is an upper semi-continuous correspondence. We next show that it is convex-valued.

Lemma 4 For all $i \in \mathcal{I}$, $j \in \mathcal{N}_i$, and $p_{-j}^{OE} \geq 0$, the set $B_j(p_{-j}^{OE})$ is a convex set.

Proof: Let $p_j \in B_j(p_{-j}^{OE})$ and $\bar{p}_j \in B_j(p_{-j}^{OE})$. Consider $x \in W(p_j, p_{-j}^{OE})$ and $\bar{x} \in W(\bar{p}_j, p_{-j}^{OE})$ such that (p_j, x) and (\bar{p}_j, \bar{x}) are optimal solutions of problem (24). If $p_j x_i = \bar{p}_j \bar{x}_i = 0$, then $\gamma p_j + (1 - \gamma) \bar{p}_j \in B_j(p_{-j}^{OE})$ for all $\gamma \in [0, 1]$, establishing convexity.

Assume next that $p_j x_i = \bar{p}_j \bar{x}_i > 0$. Assume to arrive at a contradiction that

$$p_j > \bar{p}_j, \quad (25)$$

which implies that $x_i < \bar{x}_i$. There are two cases to consider:

- $p_j + \sum_{\substack{k \in \mathcal{N}_i \\ k \neq j}} p_k^{OE} + \sum_{k \in \mathcal{N}_i} a_k x_i < \bar{p}_j + \sum_{\substack{k \in \mathcal{N}_i \\ k \neq j}} p_k^{OE} + \sum_{k \in \mathcal{N}_i} a_k \bar{x}_i$:

Since $p_j + \sum_{\substack{k \in \mathcal{N}_i \\ k \neq j}} p_k^{OE} + \sum_{k \in \mathcal{N}_i} a_k x_i < R$, it follows by the definition of a Wardrop equilibrium that $\sum_{s \in \mathcal{I}} x_s = d$. Moreover, in view of the relation between the effective costs, it can be seen that $x_s \leq \bar{x}_s$, for all $s \neq i$, which combined with $x_i < \bar{x}_i$, implies that $\sum_{s \in \mathcal{I}} x_s < \sum_{s \in \mathcal{I}} \bar{x}_s$, yielding a contradiction.

- $p_j + \sum_{\substack{k \in \mathcal{N}_i \\ k \neq j}} p_k^{OE} + \sum_{k \in \mathcal{N}_i} a_k x_i \geq \bar{p}_j + \sum_{\substack{k \in \mathcal{N}_i \\ k \neq j}} p_k^{OE} + \sum_{k \in \mathcal{N}_i} a_k \bar{x}_i$:

If both effective costs are equal to R , i.e.,

$$p_j + \sum_{\substack{k \in \mathcal{N}_i \\ k \neq j}} p_k^{OE} + \sum_{k \in \mathcal{N}_i} a_k x_i = \bar{p}_j + \sum_{\substack{k \in \mathcal{N}_i \\ k \neq j}} p_k^{OE} + \sum_{k \in \mathcal{N}_i} a_k \bar{x}_i = R,$$

then the optimization problem of provider j [cf. Problem 20] can be shown to have a strictly concave objective function over polyhedral constraints, thus implying that $p_j = \bar{p}_j$.

Assume next that $p_j + \sum_{\substack{k \in \mathcal{N}_i \\ k \neq j}} p_k^{OE} + \sum_{k \in \mathcal{N}_i} a_k x_i \leq R$, and $\bar{p}_j + \sum_{\substack{k \in \mathcal{N}_i \\ k \neq j}} p_k^{OE} + \sum_{k \in \mathcal{N}_i} a_k \bar{x}_i < R$. Define the sets

$$\mathcal{L} = \left\{ s \in \mathcal{I} \mid p_j + \sum_{\substack{k \in \mathcal{N}_i \\ k \neq j}} p_k^{OE} + \sum_{k \in \mathcal{N}_i} a_k x_i < \sum_{k \in \mathcal{N}_s} p_k^{OE} \right\},$$

and

$$\bar{\mathcal{L}} = \left\{ s \in \mathcal{I} \mid \bar{p}_j + \sum_{\substack{k \in \mathcal{N}_i \\ k \neq j}} p_k^{OE} + \sum_{k \in \mathcal{N}_i} a_k \bar{x}_i < \sum_{k \in \mathcal{N}_s} p_k^{OE} \right\}.$$

Following the line of proof of Proposition 4 (see Appendix A), we can show that

$$\begin{aligned} p_j &\leq x_i \left[\sum_{k \in \mathcal{N}_i} a_k + \frac{1}{\sum_{\substack{s \neq i \\ s \notin \mathcal{L}}} \frac{1}{\sum_{k \in \mathcal{N}_s} a_k}} \right], \\ \bar{p}_j &= \bar{x}_i \left[\sum_{k \in \mathcal{N}_i} a_k + \frac{1}{\sum_{\substack{s \neq i \\ s \notin \bar{\mathcal{L}}} \frac{1}{\sum_{k \in \mathcal{N}_s} a_k}} \right]. \end{aligned}$$

Moreover, in view of the relation between the effective costs, it can be seen that $\mathcal{L} \subset \bar{\mathcal{L}}$. Since $x_i < \bar{x}_i$, the preceding implies that $p_j < \bar{p}_j$, yielding a contradiction [cf. Eq. (25)].

Q.E.D.

Next, in view of the fact that $B(p^{OE})$ is upper semicontinuous, convex-valued and maps a compact set into a compact set, we can use Kakutani's fixed point theorem (e.g., [26]) to assert the existence of a p^{OE} such that $B(p^{OE}) = p^{OE}$. To complete the proof, it remains to show that there exists $x^{OE} \in W(p^{OE})$ such that (6) holds.

If $\mathcal{I}_0 = \emptyset$, we have by Proposition 2 that $W(p^{OE})$ is a singleton, and therefore (6) holds and $(p^{OE}, W(p^{OE}))$ is an OE.

Assume finally that $\mathcal{I}_0 = 1$, and that without loss of generality $1 \in \mathcal{I}_0$. We show that for all $\bar{x}, \tilde{x} \in W(p^{OE})$, we have $\bar{x}_i = \tilde{x}_i$, for all $i \neq 1$. Let

$$EC(x, p^{OE}) = \min_{k \in \mathcal{I}} \left\{ \sum_{j \in \mathcal{N}_k} l_j(x_k) + p_j^{OE} \right\}.$$

If at least one of the following inequalities

$$EC(\tilde{x}, p^{OE}) < R, \quad \text{or} \quad EC(\bar{x}, p^{OE}) < R$$

holds, then one can show that $\sum_{i=1}^I \tilde{x}_i = \sum_{i=1}^I \bar{x}_i = d$. Substituting $x_1 = d - \sum_{i \in \mathcal{I}, i \neq 1} x_i$ in problem (2), we see that the objective function of problem (2) is strictly convex in $x_{-1} = [x_i]_{i \neq 1}$, thus showing that $\tilde{x} = \bar{x}$. If both $EC(\tilde{x}, p^{OE}) = R$ and $EC(\bar{x}, p^{OE}) = R$, then

$$\sum_{j \in \mathcal{N}_i} l_j(\bar{x}_i) = \sum_{j \in \mathcal{N}_i} l_j(\tilde{x}_i),$$

which, by the assumption that l_j is strictly increasing for some $j \in \mathcal{N}_i$, implies that $\bar{x}_i = \tilde{x}_i$ for all $i \neq 1$, establishing our claim.

For some $x \in W(p^{OE})$, consider the vector $x^{OE} = (d - \sum_{i \neq 1} x_i, x_{-1})$. Since x_{-1} is uniquely defined and x_1 is chosen such that the providers on link 1 have no incentive to deviate, it follows that (p^{OE}, x^{OE}) is an OE. **Q.E.D.**

VIII. APPENDIX C: PROOF OF PROPOSITION 5

For brevity, we provide a sketch of this proof. The proof relies on the price characterization provided in Appendix A. Consider the following system of linear equations:

$$\begin{aligned} x_i & \left[2 \sum_{k \in \mathcal{N}_i} a_k + \frac{1}{\sum_{s \neq i} \frac{1}{\sum_{k \in \mathcal{N}_s} a_k}} \right] \\ & = x_h \left[2 \sum_{k \in \mathcal{N}_h} a_k + \frac{1}{\sum_{s \neq h} \frac{1}{\sum_{k \in \mathcal{N}_s} a_k}} \right], \quad \forall h \in \mathcal{I}, h \neq i, \end{aligned} \quad (26)$$

and

$$\sum_{i \in \mathcal{I}} x_i = d. \quad (27)$$

It is straightforward to see that, under the assumption $\mathcal{I}_0 \leq 1$, the preceding set of equations has a unique solution, which

we denote by \bar{x} , that satisfies $\bar{x}_i > 0$ for all i . Consider some $j \in \mathcal{N}_i, i \in \mathcal{I}$. For all $l \in \mathcal{N}_h, h \in \mathcal{I}$ with $l \neq j$, define

$$\bar{p}_l = \bar{x}_h \left[\sum_{k \in \mathcal{N}_h} a_k + \frac{1}{\sum_{s \neq h} \frac{1}{\sum_{k \in \mathcal{N}_s} a_k}} \right].$$

Consider the optimization problem

$$\max_{(p_j, x) \geq 0} p_j x_i \quad (28)$$

subject to

$$\begin{aligned} p_j + \sum_{\substack{k \in \mathcal{N}_i \\ k \neq j}} \bar{p}_k + \sum_{k \in \mathcal{N}_i} l_k(x_i) & = \sum_{k \in \mathcal{N}_s} \bar{p}_k + l_k(x_s), \quad \forall s \neq i, \\ \sum_{s \in \mathcal{I}} x_s & \leq d. \end{aligned}$$

An argument analogous to that in the proof of Proposition 4 in Appendix A immediately establishes that the vector (\bar{p}_j, \bar{x}) such that

$$\bar{p}_j = \bar{x}_i \left[\sum_{k \in \mathcal{N}_i} a_k + \frac{1}{\sum_{s \neq i} \frac{1}{\sum_{k \in \mathcal{N}_s} a_k}} \right]$$

is the unique optimal solution of problem (28). It also follows from Proposition 4 that for R sufficiently large, (\bar{p}, \bar{x}) satisfies the necessary conditions for a strict OE and $\sum_i \bar{x}_i = d$. It therefore follows that there exists some $\bar{R} < \infty$, such that for all $R \geq \bar{R}$, (\bar{p}, \bar{x}) is a strong OE.

To see that this is the unique strong OE for $R \geq \bar{R}$, note that any strong OE is also a strict OE, and Lemma 1 establishes that (p, x) with $p_j x_i = 0$ for some j and i cannot be a strict OE. Moreover, for $R \geq \bar{R}$ problem (28) has a unique solution, given by the unique solution to Equations (26) and (27), which is the unique candidate for a strong OE, thus completing the proof. **Q.E.D.**

ACKNOWLEDGMENTS

The authors would like to thank Attila Ambrus, Muhamet Yildiz and various seminar participants for very useful comments and suggestions.

REFERENCES

- [1] F. Kelly, A. Maulloo, and D. Tan, "Rate control for communication networks: shadow prices, proportional fairness, and stability," *Journal of the Operational Research Society*, vol. 49, pp. 237–252, 1998.
- [2] S. Low and D. Lapsley, "Optimization flow control, I: Basic algorithm and convergence," *IEEE/ACM Transactions on Networking*, vol. 7, no. 6, pp. 861–874, 1999.
- [3] H. Yaiche, R. Mazumdar, and C. Rosenberg, "A game theoretic framework for bandwidth allocation and pricing in broadband networks," *IEEE/ACM Transactions on Networking*, vol. 9, no. 5, 2000.
- [4] T. Roughgarden and É. Tardos, "How bad is selfish routing?" *Journal of the ACM*, vol. 49, no. 2, pp. 236–259, 2002.
- [5] D. Acemoglu and A. Ozdaglar, "Flow control, routing, and performance from service provider viewpoint," LIDS, Massachusetts Institute of Technology, Cambridge, MA, Tech. Rep. WP-1696, 2004.
- [6] X. Huang, A. Ozdaglar, and D. Acemoglu, "Efficiency and Braess' paradox under pricing in general networks," *IEEE Journal on Selected Areas in Communication, special issue on Price-Based Access Control and Economics*, vol. 24, no. 5, pp. 977–991, 2006.
- [7] D. Acemoglu and A. Ozdaglar, "Competition and efficiency in congested markets," *Mathematics of Operations Research*, 2007.

- [8] J. Harsanyi, "Games with randomly distributed payoffs: a new rationale for mixed strategy equilibrium points," *International Journal of Game Theory*, vol. 1, pp. 1–23, 1973.
- [9] J. Correa, A. Schulz, and N. Stier-Moses, "Selfish routing in capacitated networks," *Mathematics of Operations Research*, vol. 29, no. 4, pp. 961–976, 2002.
- [10] E. Koutsoupias and C. Papadimitriou, "Worst-case equilibria," in *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science*, 1999, pp. 404–413.
- [11] E. Friedman, "A generic analysis of selfish routing," in *Proceedings of IEEE CDC*, 2004.
- [12] R. Johari and J. Tsitsiklis, "Network resource allocation and a congestion game," *Mathematics of Operations Research*, vol. 29, no. 3, pp. 407–435, 2004.
- [13] S. Sanghavi and B. Hajek, "Optimal allocation of a divisible good to strategic buyers," in *Proceedings of IEEE CDC*, 2004.
- [14] E. Anshelevich, A. Dasgupta, J. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden, "The price of stability for network design with selfish agents," in *Proceedings of IEEE Symposium on the Foundations of Computer Science*, 2004, pp. 295–304.
- [15] T. Basar and R. Srikant, "Revenue-maximizing pricing and capacity expansion in a many-users regime," in *Proceedings of IEEE INFOCOM*, 2002.
- [16] L. He and J. Walrand, "Pricing Internet services with multiple providers," in *Proceedings of Allerton Conference*, 2003.
- [17] A. Hayrapetyan, E. Tardos, and T. Wexler, "A network pricing game for selfish traffic," in *Proceedings of Annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing*, 2005.
- [18] A. Ozdaglar, "Price competition with elastic traffic," *Networks*, 2007.
- [19] J. Wardrop, "Some theoretical aspects of road traffic research," in *Proceedings of the Institute of Civil Engineers*, 1952, pp. 325–378.
- [20] M. Beckmann, C. McGuire, and C. Winsten, *Studies in the Economics of Transportation*. Yale University Press, 1956.
- [21] S. Dafermos and F. Sparrow, "The traffic assignment problem for a general network," *Journal of Research of the National Bureau of Standards-B. Mathematical Sciences*, vol. 73, no. 2, pp. 91–118, 1969.
- [22] M. Patriksson, *The Traffic Assignment Problem: Models and Methods*, Netherlands, 1994.
- [23] D. Fudenberg and J. Tirole, *Game Theory*. Cambridge: MIT Press, 1991.
- [24] B. Bernheim, B. Peleg, and M. Whinston, "Coalition-proof Nash equilibria I. concepts," *Journal of Economic Theory*, vol. 42, no. 1, pp. 1–12, 1987.
- [25] R. Aumann, "Acceptable points in general cooperative n-person games," in *Contributions to the theory of games, IV*. Princeton: Princeton University Press, 1959.
- [26] C. Berge, *Topological Spaces*. Dover reprint, 1963.

PLACE
PHOTO
HERE

Daron Acemoglu is Charles P. Kindleberger Professor of Applied Economics in the Department of Economics at the Massachusetts Institute of Technology and a member of the Economic Growth program of the Canadian Institute of Advanced Research. He is also affiliated with the National Bureau Economic Research, Center for Economic Performance, and Center for Economic Policy Research, and is a fellow of the American Academy of Arts and Sciences, the Econometric Society, the European Economic Association and the Society of Labor Economists.

Daron Acemoglu has received a BA in Economics at the University of York, 1989, M.Sc. in Mathematical Economics and Econometrics at the London School of Economics, 1990, and Ph.D. in Economics at the London School of Economics in 1992. Since 1993, he has held the academic positions of Lecturer at the London School of Economics, and Assistant Professor, Pentti Kouri Associate Professor and Professor of Economics at MIT. He has received numerous awards and fellowships, including the award for best paper published in the Economic Journal in 1996 for his paper "Consumer Confidence and Rational Expectations: Are Agents' Beliefs Consistent With the Theory?", the inaugural T. W. Shultz prize at the University of Chicago in 2004, and the inaugural Sherwin Rosen award for outstanding contribution to labor economics in 2004.

He was awarded the John Bates Clark Medal in 2005, given every two years to the best economist in the United States under the age of 40 by the American Economic Association. Daron Acemoglu is the co-editor of the Review of Economics and Statistics (until July 2007), NBER Macroannual, and Econometrica (from July 2007).

PLACE
PHOTO
HERE

Asuman Ozdaglar (M '95) received the B.S. degree in Electrical Engineering from the Middle East Technical University in 1996, and the S.M. and the Ph.D. degrees in Electrical Engineering and Computer Science from the Massachusetts Institute of Technology (MIT) in 1998 and 2003, respectively. Since 2003, she has been with the Department of Electrical Engineering and Computer Science at MIT, where she is currently the Class of 43 Career Development Assistant Professor. She is affiliated with the Laboratory for Information and Decision

Systems, the Operations Research Center, and the Computation for Design and Optimization Program at MIT. Her research interests include game theory, optimization theory, with emphasis on nonlinear programming and convex analysis, and communication networks. She is the recipient of the Graduate Student Council Teaching award, and the NSF Career award.