Netnomics (2010) 11:149–179
DOI 10.1007/s11066-010-9052-3

Competitive resource sharing by Internet Service Providers

Gireesh Shrimali

Accepted: 18 April 2010 / Published online: 18 May 2010

Abstract We look at non-cooperative resource sharing (a generalization of paid peering) among Internet Service Providers (ISPs), where individually rational providers who not only compete for customers but also participate in resource sharing, in order to utilize underlying complementarities in cost structures. In particular, we are interested in the following question: would simple, easy-to-implement access pricing mechanisms guarantee *ex-ante participation* in resource sharing even by providers who, subsequent to deciding participation, engage in competition for customers, set access prices and make routing decisions? We first show that, in presence of linear access pricing, participation in the sharing arrangement is possible, but not guaranteed. We then show that a two-part tariff guarantees participation in the sharing agreement—this is not obvious given that resource sharing alters customer bases. We also show that our mechanism is robust to providers mis-reporting their types. Next, we show that, though both providers choose strictly positive customer bases, one of the them has no incentive to utilize the resources of the other and effectively acts as a resource supplier, whereas the other provider utilizes both resources. Finally, we show the robustness of our results to different cost structure and game forms, and provide some policy implications.

The author would like to thank Prof. Sunil Kumar at Stanford University for his help and support on this work.

Electronic supplementary material The online version of this article (doi:10.1007/s11066-010-9052-3) contains supplementary material, which is available to authorized users.

G. Shrimali (✉)
Indian School of Business, Hyderabad, 500032, India
e-mail: gireesh_shrimali@isb.edu

Our results have significant implications not only for policy design since they suggest that paid peering should be encouraged but also for design of realistic traffic engineering protocols.

Keywords Resource sharing · ISP peering · Game theory · Multi-stage Nash game · Backward induction

1 Introduction

1.1 Motivation

Consider a market with service providers that operate resources with sunk fixed costs and non-negligible variable costs. A key characteristic of this market is that the variable costs depend on the amount of traffic carried on a provider's own resource. Suppose that the customer demands result in loads that are infinitely divisible such that these loads can be executed on either of the providers' resources, and it were feasible for these providers to share their resources using simple, easy-to-implement, access pricing schemes, such as linear or affine pricing. Then, we are interested in whether "strategic complementarities" [42] exist among providers, i.e., whether rational providers would engage in resource sharing even in the presence of competition for customers (Fig. 1).

Our motivating example is Internet Service Provider (ISP) peering [28]. Today's Internet is a hierarchical collection of many distinct telecommunication networks (Fig. 2), where, due to the rapidly changing technology and the esoteric market for the equipment, most of the *fixed costs are sunk* [30]. Most of these networks are owned and operated by independent

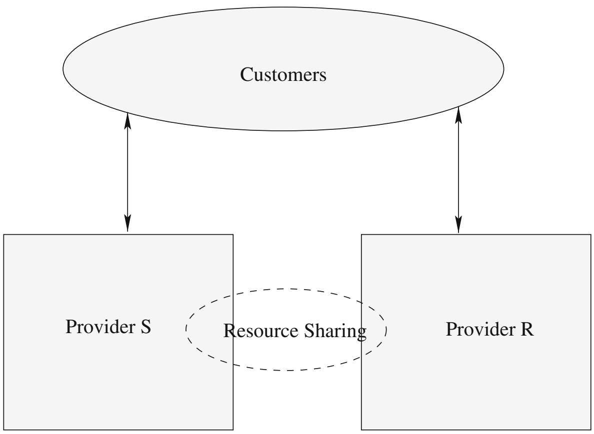

Fig. 1 Markets with competing providers, who could potentially share resources

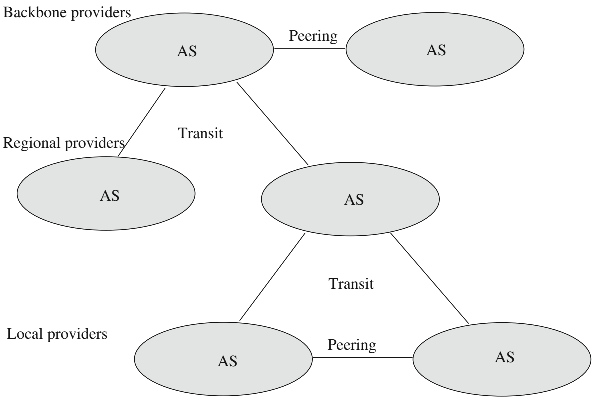

Fig. 2 Internet hierarchy and interconnection arrangements between ISPs

commercial entities, also referred to as Internet Service Providers (ISPs), which not only sell access to the Internet but also participate in providing end-to-end connectivity. This end-to-end connectivity is achieved through a series of bilateral contracts, called interconnection agreements that specify the following [28]. First, how the ISPs would physically connect to each other using either dedicated interconnection links or Internet exchange points (IXPs)—they may connect to each other at multiple locations. Second, given the choice of multiple interconnection points, how the ISPs would route traffic to each other over these interconnection points and, in addition, how they would carry traffic sent by other providers. Finally, how they would charge each other for carrying traffic.

Most interconnection relationships between providers may be classified under one of two categories [28]: “transit” and “peer” (Fig. 2). Transit is typically used by providers that are geographically separated. In a transit relationship, in a two level sub-hierarchy, a traffic-originating (or traffic-terminating) provider pays a transit provider to carry traffic destined to nodes outside (inside) the originator’s (terminator’s) local network. That is, a transit provider charges both the originator and terminator. On the other hand, peering relationships are typically between providers that have geographical overlap. For example, all backbone providers have peering arrangements between them on a pair-wise basis. As another example, two providers at a lower level peer with each other and eliminate the use of the transit provider for traffic destined to each other—they would, however, still use the transit provider for traffic destined to the rest of the Internet.

In today’s Internet, peering relationships are mostly “Bill-and-Keep” (BAK) [6]. In this arrangement, the providers do not charge each other for the traffic accepted on the peering links. This arrangement is also referred to as “Zero-Dollar” peering or “Sender-Keep-All” peering [14]. This worked well when traffic flows were light and symmetric [9]. The low network utilization allowed provides to accommodate peered traffic with little incremental congestion. Furthermore, traffic symmetry implied near-zero net volume between providers, making sophisticated network monitoring and accounting unnecessary. However, barter arrangements such as BAK expose peering participants to the opportunistic behavior of their peers. For example, ISPs predominantly use the nearest-exit or “hot-potato” routing [26], where outgoing traffic exits a provider’s network as quickly as possible. In this case, the lack of explicit control over usage of resources may result in the problem of “free-riding” [37], where the sender predominantly uses the peer’s resources for carrying it’s traffic. Since this is done in a self-interested manner, it may lead to suboptimal network utilization, data loss, and delay in the receiver’s network.

This free-riding may eventually prevent participation in the peering arrangement. Many established providers, such as Sprint and UUnet, have opted to remove their presence at public network access points (NAPs) and instead negotiate interconnection agreements at private exchange points [19]. While the exact terms are unknown, private interconnection agreements are based on several technical and operating criteria, such as network capacity, geographic coverage, number and dispersion of interconnection points, volume of traffic exchanged, traffic flow symmetry, and network management capability (for example, see the stated criteria by [41] and [22]). These range from settlement-free arrangements to paid contracts (see [43] for an introduction to paid peering) that have usage-based and congestion-based components.

Though the question of how peering can be established in an economically optimal manner has attracted considerable interest (for example, [5, 8, 11, 15, 18]), there have been very few studies [39] that address the optimal design of interconnection contracts using actual network engineering parameters, such as average traffic flows at various interconnection points [1]. We believe that the inclusion of operational aspects is essential in grounding the interconnection contracts in reality. Further, there is no work that examines interconnection between competing providers in presence of routing considerations. Therefore, the focus of our work is on paid peering among competing providers in presence of operational aspects.

1.2 Our work

In a traditional peering arrangement, providers would accept traffic destined to nodes within their respective networks only. For example (Fig. 1), ISP *R* would accept traffic from ISP *S* destined to nodes within ISP *R*’s network only. We extend this definition to include competitive resource sharing, where ISPs would accept traffic from peers irrespective of the final destination. That is, ISP *R* would accept traffic from ISP *S* destined to nodes in either network. This way ISP *R* would just act as a resource that ISP *S* could utilize. These ISPs, who otherwise compete for the same pool of customers, may want to pool their networks together and carry each others’ traffic if they can improve their profits due to improvement in variable costs of operating their respective networks (Fig. 1).[1] These *variable costs* are generally related to size of the customer base, which eventually manifests itself in the load on the network. For example, in objectives currently used by ISPs, these are related to either the maximum load on a link or the average delay within the network [7].

We are then interested in access pricing mechanisms that, in presence of strategic routing, guarantee non-trivial participation in peering by self-interested, asymmetric providers that also compete for customers. The asymmetry among the providers is a key feature of our work—it essentially means asymmetry in networks and traffic engineering practices, which result in asymmetry in cost experienced in the networks. An example of such asymmetry is networks with different capacities, which results in different average delay guarantees.

The common features of competitive resource sharing are the following. First, the resource sharing is at the level of the providers. Individual customers, though free to switch between providers, do not explicitly control the sharing arrangements. For example, though source routing [1] is gaining popularity in the Internet, most of the traffic is still routed by the service providers. Second, since it directly affects their revenues, the service providers care about the variable costs of operating their resources, defined appropriately. Third, there is no third party, such as the government, ensuring cooperation. So, each provider must assume that the other provider will act in its self-interest. Therefore, any contractual agreement for sharing must be self-enforcing, self-regulating, and budget-balanced. Fourth, at the operational level, once providers engage in resource sharing, the owner of the resource cannot explicitly regulate usage by the other providers,[2] and each provider would be able to load balance across all resources without any restrictions. However, the providers may implicitly control the load sent by the other providers by using access pricing.

Finally, the providers are individually rational. This means that the providers would participate in resource sharing, i.e., accept proposed access pricing schemes and choose load splits, only if they benefit from resource sharing. Note that, by benefiting from resource sharing, we mean that providers are weakly better off due to resource sharing. We emphasize here that the comparison relevant to individual rationality is between the following two cases. In one, the providers do not participate in resource sharing and the only

[1] Though this paper focuses on competitive resource sharing from the perspective of ISPs, such potential resource sharing among service providers is not limited to ISPs and may arise in other settings. For example, consider utility computing [31], where providers are owners of computing facilities and consumers that are owners of computing tasks. These providers may pool their computing facilities together by participating in a “grid”, which allows networked computers to be part of a virtual computer architecture, to improve the average load on their respective resources. As another example, consider any service industry, such as hospitals and call centers, multiple competing vendors providing a sharable service may pool their resources together to improve their respective variable costs.

[2] This may happen due to various reasons, including implementation difficulties, contractual forms, and traditional operational procedures.

decision is to select the customer bases. This is our baseline.[3] In the other, providers participate in resource sharing and the decisions outlined above are made in the prescribed sequence. Note that the customers bases chosen in these two cases are potentially different and, therefore, the participation in resource sharing is considerably more complex to analyze compared to the situation where the customer bases are (artificially) held fixed across the two cases.

In this context, a high level decision is whether to participate in resource sharing. The notion of "individual rationality" is appropriate for deciding participation in resource sharing.

Providers may make three additional decisions in these settings. The first set of decisions (i.e., two decisions), assuming that they are participating in resource sharing, is to select customer bases (which is equivalent to setting retail prices for the customers) and to price the usage of their resource by the other providers. The second decision, given the previous decisions, is what should their utilization of each resource be.

We analyze these decisions using non-cooperative game theory [10] in a simple two-provider model that represents individual networks by single links, and embeds network complexity in cost functions. We use a two-stage sequential Nash game in which the providers set the customer bases and access prices in the first stage and choose the resource usage in the second stage (we provide a justification for choosing this game form in Section 2.1). We note that the notion of what is called "subgame-perfect equilibrium" is appropriate for this pricing and usage game.

We first show the *existence of subgame-perfect equilibria* in our Nash game. Due to simultaneous moves by providers in each of the stages, the Nash game is a sequential game with imperfect information [10]. Characterization of equilibria in such games is difficult in general [13]. Therefore, we show existence of subgame-perfect equilibria indirectly. First, assuming that these equilibria exist, we identify special structural properties of the same. Then, using these properties, we show the existence of the subgame-perfect equilibria through first principles [10] in each of the two stages.

[3]Our baseline is referred to as not peering which is equivalent to *both* the providers not using peers' networks and carrying all their traffic themselves—this is also referred to as "cold-potato" routing [26], where the traffic leaves a provider's network as late as possible. We intentionally stay away from any comparison with transit for two reasons. First, it unnecessarily complicates the analysis, given that the focus is on when ISPs would choose to participate in resource sharing. For this reason, existing literature (e.g., [18]) also focuses on paid peering independently of transit. Second, transit can be viewed as a baseline [37]. Another way to think about this is that transit comes as an addition to baseline (i.e., cold-potato) in which case inclusion of transit increases provider cost compared to baseline (assuming the cost of setting up peering is zero—this can be justified on the basis of peering links being a fixed, sunk cost) and baseline itself would be preferred compared to transit. Then, again, a meaningful comparison from our perspective is between baseline and pair peering where traffic is routed in a strategic manner.

As mentioned above, in the process of demonstrating the existence of subgame-perfect equilibria, we establish the following structural properties of the same:

- We show that both providers select strictly positive customer bases—i.e., in presence of competitive resource sharing, no provider decides to stay out of the market.
- We also show that one of the providers acts as an *effective supplier*—a provider that has no incentive to send any load to the other, whereas the other provider splits its traffic over both the resources.
- We further show that the structural analysis used in this paper can be extended to different settings, and focus on two representative cases. In the first case, the providers compete for customers but, instead of caring about variable costs of operating their own networks, they care about customer disutility.[4] In the second case, providers serve distinct, fixed markets.[5]
- In addition, though we justify our game form in Section 2.2, we show that our structural results are insensitive to the game form chosen. We demonstrate this by looking at a three stage sequential Nash game in which providers choose customer demands in the first stage, set access prices in the second stage, and decide on load splits in the third stage.

However, through numerical examples, we demonstrate that participation in resource sharing is possible but not guaranteed for general cost forms. These numerical examples also show that the determination of effective supplier would depend not only on provider costs but also on the price sensitivity of customers—given marginal resource cost functions, the identity of the effective supplier would toggle as customers become more price sensitive.

To overcome these issues, we shift our attention to the special case of providers operating resources with linear (or affine) marginal costs. This case, though a simplification, is relevant for two reasons: first, it provides insight into the complex sequential Nash game; and second, affine marginal costs can be used to accurately represent the behavior of more complex forms close to any operating point [25]. We then obtain the following results.

- We start by showing that participation in resource sharing is possible but not guaranteed under linear access pricing. We also identify sufficient conditions that guarantee participation in resource sharing.
- More importantly, we show that the providers as a whole benefit from such interaction, and that a two part tariff would *ensure participation in resource sharing* even in the presence of competition for customers. This is

[4]For example, this would happen if ISPs competed for customers who are sensitive to end-to-end quality of service. We see this becoming a more relevant model in the future with the emergence of quality sensitive services (such as streaming media, voice over IP, and virtual private network provision).

[5]For example, this would happen if ISPs were serving distinct customer segments, such as ISP monopolies in access markets.

significant since, though it is somewhat intuitive that two-part tariffs would ensure participation if customer demands were (artificially) held fixed, it is not intuitively obvious that they would do so even when the customer demands themselves potentially change due to resource sharing.

- Further, we identify the effective supplier *ex-ante* as the provider with the lower marginal cost coefficient. This essentially translates to the intuitive ex-post condition that the effective supplier's marginal cost is lower than the other provider's marginal cost after customer bases are decided but before the flow splits are.
- Finally, we explore the sensitivity of our results to the case when there is uncertainty about the provider types, as manifested in marginal cost coefficients. We show that, when the identity of the effective supplier is known, the providers have no incentive to mis-report their types—i.e., participation in resource sharing can always be guaranteed.

To summarize, we show that simple, easy-to-implement access pricing mechanisms guarantee that competing providers benefit from cooperation implicit in resource sharing. This has *significant implications for policy makers* since our results suggest that resource sharing should be encouraged so that providers, who otherwise compete for customers, can all benefit from strategic complementarities in their cost structures.

While stylized models lead to somewhat stylized prescriptions, our results also provide useful insights for real settings. First, they suggest that simple, and therefore implementable, access pricing mechanisms may suffice to ensure participation in resource sharing even under the possibility of free riding. Second, for each route in a network one or the provider would act as an effective supplier, and some kind of settling mechanism (e.g., averaging over routes) across routes would suffice in addition to what we propose. Finally, as in [38], the actual access pricing mechanism may be implementable via an iterative procedure. This would involve probing by each peer of the other's network to estimate marginal costs followed by an adjustment of the access price.

1.3 Related work

The work on ISP peering can be divided into two broad categories: one that focuses on the interaction of various modes of Bill-and-Keep peering with transit (e.g., [2, 12, 28, 44]), and another that contrasts Bill-and-Keep peering with paid peering. Our work belongs to the second category, and we focus on related work for the same.

Various aspects of paid ISP peering in presence of competition for customers have been analyzed by [18, 33, 35], and [12]. Our work differs from this literature in three significant ways. First, these papers (e.g., [12, 18, 33, 35]) mostly focus on market structure in presence for competition for customers and *do not worry about strategic routing of the traffic once peering is decided upon*. That is, these papers take the routing of traffic as exogenous. In our work, routing is endogenous and strategic, although simple, and it provides significant insights. Second, most of this work (e.g., [18, 33, 35]) *takes access pricing as exogenous* in order to focus on competition for customers. The focus of our work is endogenous determination of access prices, similar to [12] where providers bargain over settlement. Third, much of this literature (e.g., [18, 33]) assumes that the providers are symmetric in their cost structures. This equal and constant marginal cost assumption results in fixed and symmetric access charges for the other provider's traffic. We analyze a more general case where the providers are not symmetric, similar to [12]. However, unlike [12], who do not have any notion of traffic dependent cost, we consider general traffic-dependent cost functions.

Tan et al. [39] and Johari [16] are similar to our work since they model peering between asymmetric providers as a parallel arrangement where the availability of a parallel resource may help reduce provider cost. In addition, they use endogenous access pricing in presence of strategic routing. However, they suffer from a major limitation in that they assume no competition for customers (i.e., they use fixed customer bases). Overall, we believe that ours is the first comprehensive work on ISP peering among asymmetric providers that combines strategic routing with endogenous access pricing in presence for competition for customers.

Finally, many papers (e.g., [20, 29]) are similar to our work since they model customer pricing in presence of competition for customers. In some cases, the modeling is even more intricate than ours. However, a major difference is that they do not model interconnection agreements, a critical feature of our work.

1.4 Organization

The organization of the paper is as follows. Section 2 presents the problem formulation and the two-stage sequential Nash game. Section 3 establishes the existence of Nash equilibria as well as structural results under general marginal costs. Section 4 establishes results on participation under linear marginal costs. Section 5 shows the robustness our structural results to other settings. In particular, Section 5.3 shows the insensitivity of our results to the chosen game form. We conclude the paper with Section 6 that suggests policy implications and future directions. In this section, we also address some practical implications of our work. We present auxiliary results as well as the proofs of all results in the Electronic Supplementary Material.

2 Problem formulation

2.1 The model

In this section we derive provider profit functions as implied by customer behavior. We denote the endogenously determined customer demands for provider S and R by x_S and x_R, respectively. We assume that these result in

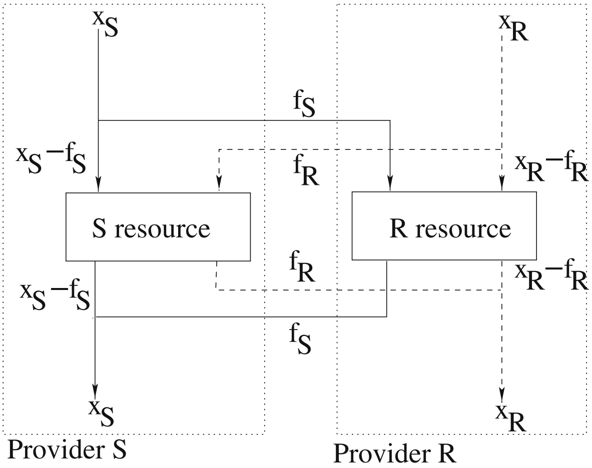

Fig. 3 The resource sharing model

equivalent loads x_S and x_R on provider S and R, respectively. et al. in ISP peering, the total amount of traffic sent by the customers of provider S and R would be proportional to x_S and x_R, respectively and, ignoring proportionality constants, x_S and x_R would represent loads.[6]

Under the resource sharing arrangement, a provider may choose to execute its load on either providers' resource through load balancing. We make the resource sharing arrangement more precise by looking at load balancing by provider S first (Fig. 3): it splits the load x_S such that the component $0 \leq f_S \leq x_S$ executes over provider R's resource whereas the component $(x_S - f_S)$ executes over its own resource. The components for provider R, i.e., $0 \leq f_R \leq x_R$ and $(x_R - f_R)$, can be described in a similar way. This load splitting results in provider S executing loads $(x_S - f_S)$ and f_R on its resource, and in provider R executing loads f_S and $(x_R - f_R)$ on its resource. Thus, the total loads on provider S and provider R's resources are $(x_S - f_S + f_R)$ and $(x_R + f_S - f_R)$, respectively. The load splitting arrangement may also be interpreted as providers S and R sending fractions $\frac{f_S}{x_S}$ and $\frac{f_R}{x_R}$, respectively, of the total loads to the other provider.

We now look at customer behavior. We assume that the market faces an elastic demand consisting of non-atomic (infinitesimal) rational customers. We further assume that customers are sensitive to the price charged by a provider. In this scenario, equilibrium customer demands $x_S \geq 0$ and $x_R \geq 0$ would be

[6]Note that we have made the implicit assumption that traffic loads are deterministic. It is well known that for Poisson traffic average (i.e., deterministic) loads may be used to calculate average disutilities, such as average delay [24]. Even when the traffic at the individual customer level is nowhere close to the "smooth" Poisson process [21], it goes through a lot of multiplexing before getting to the providers' networks. This multiplexing results in smoothing of the traffic characteristics and the resulting traffic looks like a Poisson stream [4]. Thus, given our focus on provider backbone networks, the use of deterministic loads is justified in our models.

realized in response to providers S and R setting linear customer prices $p_S \geq 0$ and $p_R \geq 0$ such that the prices are equal to each other and to the inverse demand [27]. Given a market size x_M and realized demands x_S and x_R, this can be formalized as

$$p_S = p_R = P_M\left(1 - \frac{x_S + x_R}{x_M}\right), \tag{1}$$

where $P_M(1 - \frac{x_S+x_R}{x_M})$ is the inverse demand seen by the market as a whole.[7]

Then, though each of the providers has a choice to compete over price or quantity, we assume that the providers compete using quantities, i.e., a la Cournot.[8] We are now ready to derive the profit functions for the providers. We assume that the providers price the load sent by the other provider as follows: provider S charges provider R an amount $a_S f_R$ and provider R charges provider S an amount $a_R f_S$, where $a_S, a_R \in \mathbb{R}_+$.[9] Then, since the providers operate resources with sunk fixed costs and non-negligible operating costs, the provider profits are simply given by the sum of the revenue generated from their own customers and the effective revenue from access pricing. These profits can be written as

$$\pi_S = P_M\left(1 - \frac{x_S + x_R}{x_M}\right) x_S - J_S\big(x_S, x_R, a_S, a_R, f_S, f_R\big) \tag{2a}$$

$$\pi_R = P_M\left(1 - \frac{x_S + x_R}{x_M}\right) x_R - J_R\big(x_S, x_R, a_S, a_R, f_S, f_R\big), \tag{2b}$$

where the *provider effective costs* are defined as

$$J_S\big(x_S, x_R, a_S, a_R, f_S, f_R\big) \triangleq \tilde{J}_S\big(x_S, x_R, f_S, f_R\big) - a_S f_R + a_R f_S \tag{3a}$$

$$J_R\big(x_S, x_R, a_S, a_R, f_S, f_R\big) \triangleq \tilde{J}_R\big(x_S, x_R, f_S, f_R\big) + a_S f_R - a_R f_S, \tag{3b}$$

and $\tilde{J}_S(x_S, x_R, f_S, f_R)$ as well as $\tilde{J}_R(x_S, x_R, f_S, f_R)$ are the provider variable costs—we define these as the *resource costs* (4). From (3) it is clear that *the*

---

[7]The idea here is that, if the prices were not equal, customers would switch to the provider with the lower price. In addition, if the prices were not equal to the inverse demand, the total demand would change until they are. That is, if the prices were below the inverse demand, the total demand would increase and, if the prices were above the inverse demand, the total demand would decrease. This downward sloping symmetric inverse demand (or price) function is what we would see in a market with elastic demand, with price at demand meeting or exceeding market capacity. Note that this inverse demand formulation implies $x_S + x_R \leq x_M$.

[8]First, it is well known that Cournot competition is appropriate in markets with rising marginal costs whereas Bertrand competition is appropriate in markets with flat marginal costs [40]. The marginal resource costs in (4) satisfy the rising marginal cost condition. Second, it is also well known that if providers pre-commit to quantity (or capacity) and then compete a la Bertrand, a Cournot outcome results [17]. In all our examples, since providers would typically be capacity constrained in the short run, they would set upfront limits on how much demand they would like to sign up for, and Cournot competition can be appealed to without loss of generality.

[9]As a realistic example, note that pricing based on flows (i.e., arrival rates) is typical in ISP interaction - transit pricing being one example [28].

providers' profits are constrained by the variable costs. That is, providers profit maximization would translate to variable cost minimization for every fixed x_S and x_R.

We assume that the *resource per unit cost*, defined as the cost to carry a unit load on a resource, depends only on the total loads on the resource. These resource per unit costs for provider S and R are defined as $c_S(x_S - f_S + f_R)$ and $c_R(x_R + f_S - f_R)$, respectively. These can also be written as $c_S(x_S - f_d)$ and $c_R(x_R + f_d)$, where $f_d = (f_S - f_R)$. We next define the *resource costs* $C(.)$ in terms of per-unit costs $c(.)$ as $C(x) = xc(x),\ x \geq 0$.[10] Then, we get

$$\tilde{J}_S(x_S, x_R, f_S, f_R) \triangleq C_S(x_S - f_d) \triangleq c_S(x_S - f_d)(x_S - f_d) \tag{4a}$$

$$\tilde{J}_R(x_S, x_R, f_S, f_R) \triangleq C_R(x_R + f_d) \triangleq c_R(x_R + f_d)(x_R + f_d). \tag{4b}$$

Two realistic examples of such per unit cost functions are as follows.

Example 1 The per unit delay in an $M/G/1$ queue is given by $c(x) = \frac{\theta^N}{(\theta^D - x)}$, where θ^N (> 0) is proportional to the variance of service times and θ^D ($> x \geq 0$) is the capacity of the resource.

Note that this form is in fact used by ISPs as a proxy for the network cost—they optimize the average delay through their networks as a first step towards determining weights for the commonly used shortest path algorithms [7].

Example 2 The drop probability in a finite buffer of size $\theta^P > 0$ is given by $c(x) = \theta^C x^{\theta^P}$.

2.2 The Nash game

We consider a two-stage Nash game of complete information in which providers pick the customer bases as well as access prices in the first stage and loads splits in the second stage.[11] The customer bases and access prices in the first stage are picked having committed to picking load splits in their self-interest in the second stage. The game is formulated below and is solved using the standard backward induction technique.

In the first stage, the providers simultaneously solve the following optimization problems: given (x_R, a_R), provider S solves

$$\max_{x_S \in [0, x_M],\ a_S \in \mathbb{R}_+} \pi_S^S \triangleq P_M \left(1 - \frac{x_S + x_R}{x_M}\right) x_S - J_S(x_S, x_R, a_S, a_R) \tag{5}$$

[10] We caution the reader to fully understand the distinction between resource costs $C(.)$ and per-unit costs $c(.)$ to avoid confusion later in the paper.

[11] We also provide a rationale for using this particular game at the end of this subsection.

and, given (x_S, a_S), provider R solves

$$\max_{x_R \in [0, x_M],\ a_R \in \mathbb{R}_+} \pi_R^S \triangleq P_M \left(1 - \frac{x_S + x_R}{x_M}\right) x_R - J_R(x_S, x_R, a_S, a_R), \tag{6}$$

where π_S^S and π_R^S are the profits of provider S and R, respectively, as defined in (2), and $J_S(x_S, x_R, a_S, a_R)$ and $J_R(x_S, x_R, a_S, a_R)$ are response functions from stage two, as defined in (7) and (8).

Then, in the second stage, for fixed x_S, x_R, a_S and a_R, the providers simultaneously solve the following optimization problems: given f_R, provider S solves

$$J_S(x_S, x_R, a_S, a_R) \triangleq \min_{f_S \in [0, x_S]} J_S(x_S, x_R, a_S, a_R, f_S, f_R) \tag{7}$$

and, given f_S, provider R solves

$$J_R(x_S, x_R, a_S, a_R) \triangleq \min_{f_R \in [0, x_R]} J_R(x_S, x_R, a_S, a_R, f_S, f_R), \tag{8}$$

where $J_S(x_S, x_R, a_S, a_R, f_S, f_R)$ and $J_R(x_S, x_R, a_S, a_R, f_S, f_R)$ are as defined in (3).

The individual rationality conditions that the provider profits be greater under resource sharing when compared to not sharing (i.e., $\pi_S^S \geq \pi_S^{NS}$ and $\pi_R^S \geq \pi_R^{NS}$) ensure participation in resource sharing. Here π_S^{NS} and π_R^{NS} are provider profits when the first stage is solved with the artificial restriction that $f_S = f_R = 0$ in the second stage, i.e., assuming that the second stage does not exist and access prices as well as load splits do not matter. We follow this convention throughout the paper: the quantities in presence of sharing are denoted with the superscript S and the quantities in absence of sharing are denoted with the superscript NS.

At this point, a note on the choice of the game structure is in order. First, let us ignore competition for customers for a moment and focus on resource sharing through access pricing. It is intuitive to envisage the interaction between providers as a secondary market for resources, and the choice of setting access prices before resource usage decisions is reasonable since providers (as users) would like to see prices on resources before making usage decisions. Therefore, we find it appropriate to use a sequential game form in which access prices are chosen first and load splits second.[12]

[12] We may also point out that possibilities other than this sequential form include the following: a form in which access prices and load splits are chosen simultaneously; a sequential form in which load splits are chosen first and prices second; and an alternating offer bargaining game. For the first two forms, it is fairly straightforward to show that $f_S = f_R = 0$ always: that is, there is no resource sharing. This is not intuitive since the providers should be able to improve on their profits, if possible, as allowed by our sequential form. The last form—the bargaining game—though appealing, is not only complicated but also does not result in crisp predictions, as provided by our simple sequential form.

Having established a reasonable sequence of access pricing and load splits decisions, we next look into the timing of competition for customers. We look at the simplest form where the providers decide all price variables simultaneously. That is, we look at a two stage game in which providers decide on customer demands and set access prices in the first stage and decide on load splits in the second stage. In our opinion, not only it is the simplest realistic model of the interaction between service providers but it also ensures participation in resource sharing by competing providers.

We recognize that there may still be some objections to this choice of game form. We show that our main results are robust to the choice of game form provided the sequence of access price and load split decisions is maintained. We demonstrate this in Section 5.3 where we analyze a three stage game in which customer demands are realized in the first stage, access prices are set in the second stage, and load splits are realized in the third stage.

3 General marginal costs: existence of Nash equilibria

In this section, we look at the existence of Nash equilibria of the Nash game (5)–(8). Recall that this Nash game is independent of the decision to participate in competitive resource sharing. We then examine participation in competitive resource sharing based on the outcome of the Nash game.

3.1 Preliminaries

First, we assume that

Assumption 1 *The per unit cost functions $c(x) \geq 0,\ x \geq 0$ are strictly increasing, convex and three times continuously differentiable. That is, the resource cost functions $C(x),\ x \geq 0$ are increasing, strictly convex and three times continuously differentiable, with $C'''(x) \geq 0$.*

Note that the linear (and affine) per unit costs in Section 4 as well as in Examples 1 and 2 satisfy Assumption 1.

Second, we make the simplifying assumption that the resource marginal costs are equal at zero market realizations. That is,

Assumption 2 *$C'_S(0) = C'_R(0)$, or equivalently, $c_S(0) = c_R(0)$.*

For example, for the per unit cost functions specified in Example 1, this is equivalent to $\frac{\theta_S^N}{\theta_S^D} = \frac{\theta_R^N}{\theta_R^D}$. Similarly, for the per unit cost functions specified in Example 2, this is always true since $\theta^{p_S}\theta^{p_R} > 0$. In general, since we have assumed $c(.)$ to be continuous, from the Stone-Weierstrass theorem [25], this requires the constant terms in the polynomial approximations to be equal. Assumption 2 is limiting in the sense that it restricts the class of allowable per unit cost functions. It is not necessary in order to get our results, however, it greatly simplifies our analysis.

Third, we assume that, at zero market realization, the inverse demand is strictly greater than both the provider marginal costs.

Assumption 3 $P_M > \max(C'_S(0), C'_R(0))$.

For example, for the per unit cost functions specified in Example 1, this is equivalent to $P_M > \max(\frac{\theta_S^N}{\theta_S^D}, \frac{\theta_R^N}{\theta_R^D})$. Similarly, for the per unit cost functions specified in Example 2, this is equivalent to $P_M > 0$ for $\theta_S^p \theta_R^p > 0$. This assumption is critical since, as a baseline, it ensures that the market exists even if resource sharing were not an option and providers were engaging only in competition for customers.

3.2 Results

We start by showing the existence of subgame-perfect equilibria for the Nash game (5)–(8), as follows. We also derive some useful structural properties which emphasize that resource sharing result in non-trivial outcomes.

Theorem 1 *Subgame-perfect equilibria exist in the Nash game* (5)–(8)*. Further* $x_S x_R > 0$*,* $f_S f_R = 0$*, and* $f_S + f_R > 0$ *in any subgame-perfect equilibrium. In addition, in any subgame-perfect equilibrium,* $C'_S(x_S) \neq C'_R(x_R)$*.*

This result is instructive in many ways. First, it demonstrates that the phenomenon of effective supplier behavior holds. That is, it shows that in these markets it is never the case that both providers send strictly positive flows to each other. Second, it shows not only that both the providers would select strictly positive customer bases ($x_S x_R > 0$) in presence of resource sharing but also that they would collectively exchange *strictly* positive flows ($f_S + f_R > 0$) and that boundary conditions, such as the effective supplier opting out of the market (e.g., $x_S > 0,\ x_R = 0,\ f_S > 0,\ f_R = 0$), are ruled out. That is, it shows not only that under resource sharing providers would first choose customer demands and then choose *non-trivial* flow splits such that the loads carried on each resource are different from the customer demands themselves but also that providers who were choosing strictly positive demands without resource sharing may still choose to do so even though resource sharing may cause them to potentially select demands different from ones chosen if resource sharing were not an option.

However, Theorem 1 does not guarantee participation in resource sharing. That is, it does not show that both providers benefit from resource sharing. The following example looks at this issue.[13]

[13] We delegate the robustness results on existence of Nash equilibria and game forms to Section 5.

Example 3 When $c_S(x) = \frac{x}{2}$, $c_R(x) = \frac{x^2}{3}$ (Fig. 4) and $P_M = 6$, $x_M = 4$, provider S is the effective supplier, with $\pi_S^S = 2.4049 < 2.4155 = \pi_S^{NS}$ as well as $\pi_R^S = 2.5635 > 2.5568 = \pi_R^{NS}$. However, with the same per unit cost functions, with $P_M = 6$, $x_M = 0.4$, provider R is the effective supplier, with $\pi_S^S = 0.2585 > 0.2541 = \pi_S^{NS}$ as well as $\pi_R^S = 0.2653 < 0.2712 = \pi_R^{NS}$. In both cases, the non effective supplier benefits from resource sharing whereas the effective supplier does not.

For general cost forms, Example 3 hints towards the following difficulties. First, though participation in resource sharing is possible, it is not guaranteed. That is, it is possible that the effective supplier does not benefit from resource sharing. Second, the identification of the effective supplier is not straightforward. That is, with the same per unit cost functions, the identity of the effective supplier changes from one provider to the other as the consumers' price sensitivity changes.

Given these difficulties with general cost forms, and the fact that a comprehensive analysis proves hard, we simplify our analysis, and assume that the provider's marginal cost forms are restricted to linear (or affine) functions. This case, though simple, is relevant for two reasons: first, it provides insight into the complex sequential Nash game; and second, linear (or affine) marginal costs can be used to accurately represent the behavior of more complex forms close to any operating point [25].

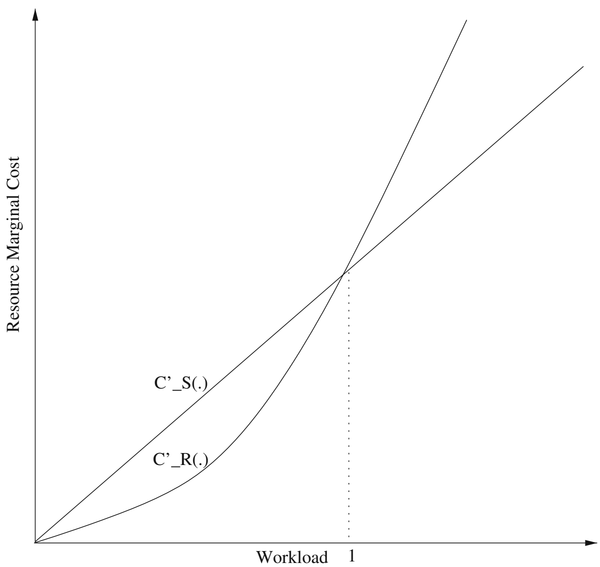

Fig. 4 Marginal resource cost vs. workload (Example 3)

4 Linear marginal costs: participation in resource sharing

This section establishes two main results. First, it establishes conditions under which resource sharing is guaranteed under linear access pricing (Section 4.1). Second, it examines non-linear access pricing schemes under which resource sharing is always guaranteed (Section 4.2). We also show that these results are robust under mis-reporting of cost functions by the providers (Section 4.3).

4.1 Participation is possible under linear access pricing

In this subsection we look at conditions under which participation in resource sharing is possible given that providers subsequently engage in the Nash game (5)–(8). We start with the specific case of linear marginal costs and establish (a negative result) that participation in resource sharing is impossible. However, we also show that participation is possible in the more general case of affine marginal costs. We then derive conditions under which such participation is guaranteed.

We first look at the case of linear marginal costs (Fig. 5). That is, we assume that the provider's resource costs are given by $C_S(x) = \theta_S x^2$, $C_R(x) = \theta_R x^2$ for ISP S and R, respectively, which gives the marginal costs as $C'_S(x) = 2\theta_S x$, $C'_R(x) = 2\theta_R x$. Then it can be verified that it is not possible for both providers to be better off by resource sharing.

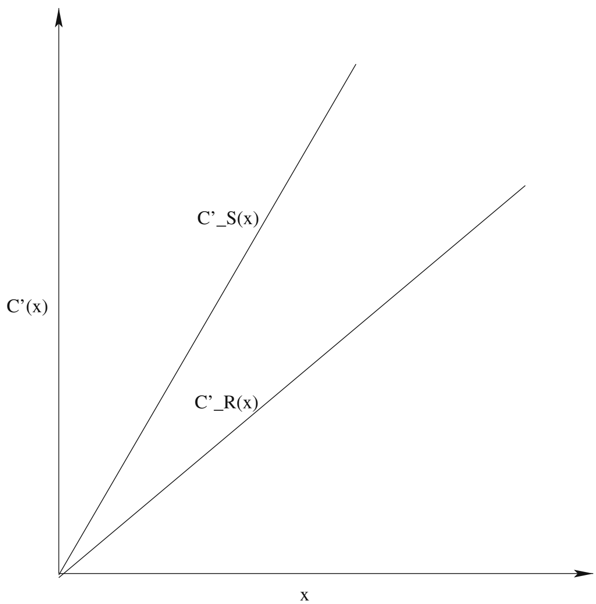

Fig. 5 Provider marginal costs: the linear case

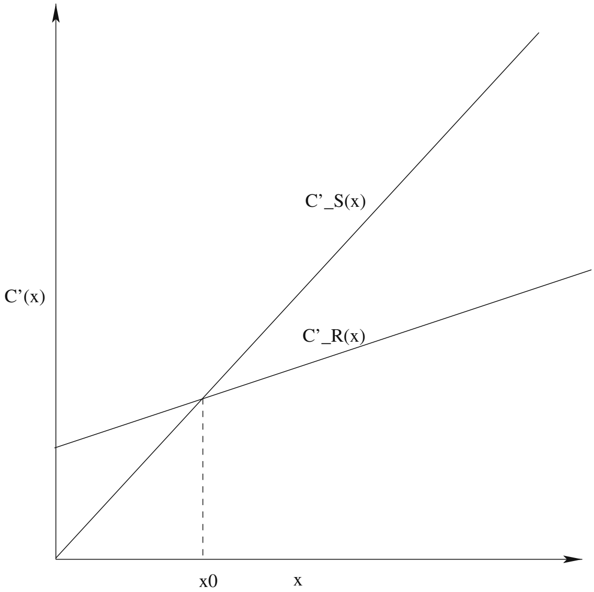

Fig. 6 Provider marginal costs: the affine case

Theorem 2 *When* $C_S(x) = \theta_S x^2$, $C_R(x) = \theta_R x^2$, $\theta_S > 0$, $\theta_R > 0$, $p_M > 0$, $x_M > 0$, *subgame-perfect equilibria exist in the Nash game* (5)–(8)*, with* $x_S x_R > 0$, $f_S + f_R > 0$, $f_S f_R = 0$. *However,* $(\pi_S^S - \pi_S^{NS})(\pi_R^S - \pi_R^{NS}) < 0$. *Further, if* $\theta_S > \theta_R > 0$ *then* $f_S^S > 0 = f_R^S$, $\pi_S^S - \pi_S^{NS} > 0$, $\pi_R^S - \pi_R^{NS} < 0$, *and* $(\pi_S^S + \pi_R^S) - (\pi_S^{NS} + \pi_R^{NS}) > 0$.

That is, though the society as well as the non-effective supplier are always better off by resource sharing, the effective supplier is not, and participation in resource sharing is not guaranteed. Thus, even though the combined surplus of the providers increases due to resource sharing, both providers are not able to benefit from the interaction, with the non-effective supplier free riding on the effective supplier. This free riding, even in presence of access prices, is worrisome because it calls into question whether providers that compete for customers would ever participate in a peering arrangement.

The obvious question is whether the same, somewhat negative, result extends to the case of affine marginal costs. To answer this, we look at affine marginal costs that would cross at x_0 (Fig. 6). That is,[14] $C_S(x) = \theta_S x^2$ and $C_R(x) = \theta_R x^2 + 2(\theta_S - \theta_R)x_0 x$, which gives $C'_S(x) = 2\theta_S x$ and $C'_R(x) = 2\theta_R(x - x_0) + 2\theta_S x_0$. When sharing resources, depending on the value of the parameters, there can be two Nash equilibria—one on each side of x_0. We show that, for the Nash equilibrium to the left of x_0, it is indeed possible for both providers to participate in resource sharing, and derive sufficient conditions

[14]This formulation assumes that $\theta_S > \theta_R$, which ensures that $C'_S(x) \geq 0$, $C'_R(x) \geq 0$, $x \geq 0$.

for the same. However, for the Nash equilibrium to the right of x_0, we show that providers will not participate in resource sharing. This is consistent with Theorem 2 since the graph to the right of x_0 is the identical to the case of linear marginal costs.

Theorem 3 *When* $C_S(x) = \theta_S x^2$ *and* $C_R(x) = \theta_R x^2 + 2(\theta_S - \theta_R)x_0 x$, $P_M > 0$, $x_M > 0$, $\theta_S > 0$, $\theta_R > 0$, *subgame-perfect equilibria exist in the Nash game* (5)–(8), *with* $x_S x_R > 0$, $f_S + f_R > 0$, $f_S f_R = 0$. *In addition, assume* $\theta_S > \theta_R > 0$.

- *If* $x_0^* > x_0 > 0$, *then* $x_S^S > x_0$, $x_R^S > x_0$, $f_S^S > 0 = f_R^S$; *and* $(\pi_S^S - \pi_S^{NS})(\pi_R^S - \pi_R^{NS}) < 0$. *Further,* $\pi_S^S - \pi_S^{NS} > 0$ *and* $\pi_R^S - \pi_R^{NS} < 0$.
- *If* $x_0^* < x_0 < \psi_2$, *then* $0 < x_S^S < x_0$, $0 < x_R^S < x_0$, $f_R^S = 0 < f_S^S < x_S^S$. *Further, if* $x_0^* < \psi_1 < x_0$, *then* $\pi_S^S - \pi_S^{NS} > 0$ *and* $\pi_R^S - \pi_R^{NS} > 0$.

With

$$x_0^* = \frac{P_M x_M}{3P_M + 2\theta_S x_M}$$

$$\psi_1 = \frac{P_M x_M (3P_M^3(5\theta_S + \theta_R) + 2P_M^2(8\theta_S^2 + 19\theta_S\theta_R + 3\theta_R^2)x_M + 8P_M\theta_S\theta_R(4\theta_S + 3\theta_R)x_M^2 + 16\theta_S^2\theta_R^2 x_M^3)}{(\theta_S - \theta_R)(27P_M^4 + 54P_M^3(\theta_S + \theta_R)x_M + 4P_M^2(8\theta_S^2 + 29\theta_S\theta_R + 8\theta_R^2)x_M^2 + 64P_M\theta_S\theta_R(\theta_S + \theta_R)x_M^3 + 32\theta_S^2\theta_R^2 x_M^4)}$$

$$\psi_2 = \frac{P_M x_M (P_M(2\theta_S + \theta_R) + 2\theta_S\theta_R x_M)}{(\theta_S - \theta_R)(3P_M^2 + 4P_M(\theta_S + \theta_R)x_M + 4\theta_S\theta_R x_M^2)}.$$

Note that the relationship between x_0 and x_0^* determines whether the equilibrium demands fall to the left or to the right of x_0. The case $x_0^* > x_0$ is similar to Theorem 2 and does not require further explanation. However, the case $x_0^* < x_0$ requires more explanation. In this case, the sufficient condition for non-trivial participation in resource sharing is given by $x_0^* < \psi_1 < x_0 < \psi_2$. The lower bound on x_0 (i.e., $\psi_1 < x_0$) is what guarantees participation in resource sharing—if $\psi_1 > x_0$ then participation in resource sharing is not guaranteed. The upper bound on x_0 (i.e., $x_0 < \psi_2$) ensures that all equilibrium demands stay strictly positive and that the non effective supplier does not send all its traffic to the effective supplier. To summarize the sufficient condition for participation in resource sharing, when the equilibrium demands lie to the left of the crossover point, and the marginal cost crossover point lies within an interval, providers will participate in resource sharing by choosing strictly positive demands, and the non-effective supplier will use both resources.

This result shows that, in theory, it is possible to characterize the conditions under which both providers would participate in resource sharing. This potential self selection of resource sharing strongly hints at underlying complementarities in an otherwise competitive environment. The sufficiency conditions depend on parameters defining the Cournot market as well as provider marginal cost functions. One potential avenue to proceed would be to tease out these relationships further.

However, as we have seen, participation in resource sharing is not guaranteed, even when the providers benefit as a whole from resource sharing (Theorem 2). The potential failure of resource sharing under linear access pricing indicates that, due the lack of flexibility available in the linear access pricing scheme, the effective supplier is unable to extract enough surplus to compensate for the potential loss in profit due to loss in revenue as well as increase in costs. Then, as a mechanism designer it is of greater interest to us whether a simple, and easy to implement, access pricing mechanism would guarantee participation independently of the parameters of the problem. This is what we focus on next.

Given the similarities between linear and affine marginal costs we focus on linear marginal costs in the remainder of this section, unless specified otherwise.

4.2 Non-linear access pricing guarantees participation

We now show that, if non-linear access pricing schemes were allowed, the effective supplier would be able to deliver mutually beneficial solutions under resource sharing. For example, one such scheme is a two-part tariff where suppliers would charge a lump sum amount b in addition to the per unit access charge a. This implies that the first stage of the Nash game would be modified to: given (x_R, a_R, b_R), provider S solves

$$\max_{x_S\in[0,x_M],\ a_S\in\mathbb{R}_+,\ b_S\in\mathbb{R}_+} \pi_S^{GS} \triangleq P_M\left(1-\frac{x_S+x_R}{x_M}\right)x_S - J_S(x_S,x_R,a_S,a_R,b_S,b_R) \tag{9}$$

and, given (x_S, a_S, b_S), provider R solves

$$\max_{x_R\in[0,x_M],\ a_R\in\mathbb{R}_+,\ b_R\in\mathbb{R}_+} \pi_R^{GS} \triangleq P_M\left(1-\frac{x_S+x_R}{x_M}\right)x_R - J_R(x_S,x_R,a_S,a_R,b_S,b_R), \tag{10}$$

where $J_S(x_S, x_R, a_S, a_R, b_S, b_R)$ and $J_R(x_S, x_R, a_S, a_R, b_S, b_R)$ are response functions from stage two, as defined in (11) and (12). Similarly, the second stage would be modified to: for fixed $(x_S, x_R, a_S, a_R, b_S, b_R, f_R)$, provider S solves

$$J_S(x_S,x_R,a_S,a_R,b_S,b_R) \triangleq \min_{f_S\in[0,x_S]} J_S(x_S,x_R,a_S,a_R,b_S,b_R,f_S,f_R) \tag{11}$$

and, given $(x_S, x_R, a_S, a_R, b_S, b_R, f_S)$, provider R solves

$$J_R(x_S,x_R,a_S,a_R,b_S,b_R) \triangleq \min_{f_R\in[0,x_R]} J_R(x_S,x_R,a_S,a_R,b_S,b_R,f_S,f_R), \tag{12}$$

where, modifying (3), we get

$$J_S(x_S,x_R,a_S,a_R,b_S,b_R,f_S,f_R) \triangleq \tilde{J}_S(x_S,x_R,f_S,f_R) - (a_S f_R + b_S) + (a_R f_S + b_R) \tag{13a}$$

$$J_R(x_S,x_R,a_S,a_R,b_S,b_R,f_S,f_R) \triangleq \tilde{J}_R(x_S,x_R,f_S,f_R) + (a_S f_R + b_S) - (a_R f_S + b_R), \tag{13b}$$

and $\tilde{J}_S(x_S, x_R, f_S, f_R)$ and $\tilde{J}_R(x_S, x_R, f_S, f_R)$ are the resource costs for provider S and R, respectively, as defined in (4).

Since b's are constants, the subgame-perfect equilibrium of this Nash game would involve the same first-order conditions as the Nash game (5)–(8). However, now the effective supplier would set b such that the non-effective supplier is indifferent between sharing and non-sharing. That is, if provider R is the effective supplier,

$$b_R = \pi_S^S - \pi_S^{NS},$$

where π_S^S and π_S^{NS} are as defined for the Nash game (5)–(8). Since the society is always better off (Theorem 2), this guarantees that the effective supplier would be better off due to sharing since

$$(\pi_R^S + b) - \pi_R^{NS} = \pi_S^S + \pi_R^S - (\pi_S^{NS} + \pi_R^{NS}) > 0,$$

and participation in resource sharing is always guaranteed. This can be summarized in the following result.

Theorem 4 *When* $C_S(x) = \theta_S x^2$, $C_R(x) = \theta_R x^2$, $\theta_S > 0$, $\theta_R > 0$, $p_M > 0$, $x_M > 0$, *subgame-perfect equilibria exist in the Nash game* (9)–(12), *with* $x_S x_R > 0$, $f_S + f_R > 0$, $f_S f_R = 0$. *Further,* $\pi_S^{GS} - \pi_S^{NS} = 0$ *and* $\pi_R^{GS} - \pi_R^{NS} > 0$.

4.3 Participation is guaranteed even when providers mis-report types

This is significant since Theorem 4 indicates that participation in resource sharing can be guaranteed even among competing providers. But, one challenge still remains. We have assumed so far that the providers are honest about reporting their types (θs). But, do they really have the incentive to report honestly? This can again be analyzed in the game theoretic framework we have developed. Under mis-reporting of θs, we would have the following two-stage sequential Nash game: in the first stage, the providers would report their types, choose customer demands, and set access prices; and in the second stage decide on flow splits.

In this case, (9) and (10) can be modified to: for fixed $(\hat{\theta}_R, x_R, a_R, b_R)$ provider S solves

$$\max_{\hat{\theta}_S \in \mathbb{R}_+,\ x_S \in [0, x_M],\ a_S \in \mathbb{R}_+,\ b_S \in \mathbb{R}_+} \pi_S^{TS} \triangleq P_M \left(1 - \frac{x_S + x_R}{x_M}\right) x_S - J_S(\hat{\theta}_S, \hat{\theta}_R, x_S, x_R, a_S, a_R, b_S, b_R) \quad (14)$$

and, for fixed $(\hat{\theta}_S, x_S, a_S, b_S)$ provider R solves

$$\max_{\hat{\theta}_R \in \mathbb{R}_+,\ x_R \in [0, x_M],\ a_R \in \mathbb{R}_+,\ b_S \in \mathbb{R}_+} \pi_R^{TS} \triangleq P_M \left(1 - \frac{x_S + x_R}{x_M}\right) x_R - J_R(\hat{\theta}_S, \hat{\theta}_R, x_S, x_R, a_S, a_R, b_S, b_R), \quad (15)$$

where $J_S(\hat{\theta}_S, \hat{\theta}_R, x_S, x_R, a_S, a_R, b_S, b_R)$ and $J_R(\hat{\theta}_S, \hat{\theta}_R, x_S, x_R, a_S, a_R, b_S, b_R)$ are response functions from stage two, as defined in (16) and (17). Similarly, the second stage would be modified to: for fixed $(\hat{\theta}_S, \hat{\theta}_R, x_S, x_R, a_S, a_R, b_S, b_R, f_R)$, provider S solves

$$J_S(\hat{\theta}_S, \hat{\theta}_R, x_S, x_R, a_S, a_R, b_S, b_R) \triangleq \min_{f_S \in [0, x_S]} J_S(\hat{\theta}_S, \hat{\theta}_R, x_S, x_R, a_S, a_R, b_S, b_R, f_S, f_R) \tag{16}$$

and, given $(\hat{\theta}_S, \hat{\theta}_R, x_S, x_R, a_S, a_R, b_S, b_R, f_S)$, provider R solves

$$J_R(\hat{\theta}_S, \hat{\theta}_R, x_S, x_R, a_S, a_R, b_S, b_R) \triangleq \min_{f_R \in [0, x_R]} J_R(\hat{\theta}_S, \hat{\theta}_R, x_S, x_R, a_S, a_R, b_S, b_R, f_S, f_R), \tag{17}$$

where, modifying (3), we get

$$J_S(\hat{\theta}_S, \hat{\theta}_R, x_S, x_R, a_S, a_R, b_S, b_R, f_S, f_R) \triangleq \tilde{J}_S(\hat{\theta}_S, \hat{\theta}_R, x_S, x_R, f_S, f_R) - (a_S f_R + b_S) + (a_R f_S + b_R) \tag{18a}$$

$$J_R(\hat{\theta}_S, \hat{\theta}_R, x_S, x_R, a_S, a_R, b_S, b_R, f_S, f_R) \triangleq \tilde{J}_R(\hat{\theta}_S, \hat{\theta}_R, x_S, x_R, f_S, f_R) + (a_S f_R + b_S) - (a_R f_S + b_R), \tag{18b}$$

and $\tilde{J}_S(\hat{\theta}_S, \hat{\theta}_R, x_S, x_R, f_S, f_R)$ and $\tilde{J}_R(\hat{\theta}_S, \hat{\theta}_R, x_S, x_R, f_S, f_R)$ are the resource costs for provider S and R, respectively, as defined in (4).

Then we get the following result, given providers' beliefs about their types.

Theorem 5 *When $C_S(x) = \theta_S x^2$, $C_R(x) = \theta_R x^2$, $p_M > 0$, $x_M > 0$, subgame-perfect equilibria exist in the Nash game (14)–(17), with $x_S x_R > 0$, $f_S + f_R > 0$, $f_S f_R = 0$. Further, if $\theta_S \in [\theta_S^{\min}, \theta_S^{\max}]$, $\theta_R \in [\theta_R^{\min}, \theta_R^{\max}]$, and $\theta_S^{\min} > \theta_R^{\max}$, then $\hat{\theta}_S = \theta_S$ and $\hat{\theta}_R = \theta_R$ in subgame-perfect equilibrium.*

That is, when the providers types have non-overlapping supports, providers report their types honestly, and participation in resource sharing is guaranteed under non-linear access pricing. Note that in this case both providers believe that the types belong to non-overlapping intervals. This ensures that providers can not credibly signal a type outside these intervals. We acknowledge that this leaves out the case when the beliefs may overlap. This turns out to be much harder to analyze, and we leave this to future work, assuming for the time being that providers can do proper due diligence to narrow the supports down enough to ensure that they do not overlap.

This finishes our treatment of the special case where provider marginal costs are linear. For the remainder of our paper we switch back to general cost forms.

5 Robustness of structural results

In this section, we demonstrate that our results are not limited to the particular setting described in Section 1, and are applicable in different settings. In particular, we present two settings where our results would apply.

5.1 When customers care about disutility

The first setting is similar to our original setting except that though providers operate networks with negligible operating costs, customers (and hence providers) care about disutility and choose providers based on prices as well as expected disutility. For example, as mentioned in Section 1, this disutility may be related to end-to-end quality of service (QOS) guarantees required by services such as voice over IP (VoIP) and video on demand.

We assume that the load splitting described in Section 2.1 occurs at the level of every customer. For example, in ISP peering, this would happen if every packet originating from provider S is sent on provider S and R's networks with probabilities $1 - \frac{f_S}{x_S}$ and $\frac{f_S}{x_S}$, respectively, and every packet originating from provider R is sent on provider S and R's networks with probabilities $\frac{f_R}{x_R}$ and $1 - \frac{f_R}{x_R}$, respectively. We consider the cost of carrying load on each resource and define the expected *per unit disutility* for provider S and R's customers as[15]

$$\hat{J}_S(x_S, x_R, f_S, f_R) \triangleq c_S(x_S - f_S + f_R)(1 - \frac{f_S}{x_S}) + c_R(x_R + f_S - f_R)\frac{f_S}{x_S} \quad (19a)$$

$$\hat{J}_R(x_S, x_R, f_S, f_R) \triangleq c_S(x_S - f_S + f_R)\frac{f_R}{x_R} + c_R(x_R + f_S - f_R)(1 - \frac{f_R}{x_R}), \quad (19b)$$

respectively, where each per unit disutility is the weighted sum of the per unit costs (i.e., $c_S(.)$ and $c_R(.)$) on the two resources - each term in the sum is given by the product of the resource per unit cost and the fraction of traffic sent on the resource. This formulation assumes the following. First, we assume that the customers see the same per unit costs on a particular resource. For example, all customers see the per unit cost on resource S as $c_S(.)$. Second, we assume that customers know the per unit costs on both resources, in particular on the resource of the provider they are not connected to. This will be possible, for example, if the customers actively probe the performance of both resources, as indicated in [34]. Finally, we assume that the cost to split and recombine load across the resources is zero. For example, in ISP peering, the delay through the (complex) networks would be much higher than the delay over the peering links.

We assume that the market faces an elastic demand consisting of non-atomic (infinitesimal) rational customers. We further assume that customers

[15] We use the terms cost and disutility interchangeably.

are sensitive to the price charged by a provider as well as to the expected per unit disutility experienced as a customer of the same provider. That is, a customer would choose the provider where it faces a lower *effective price*, defined as the sum of customer price charged by the provider and expected per unit disutility (19) faced as a customer of the same provider. In this scenario, equilibrium customer demands $x_S \geq 0$ and $x_R \geq 0$ would be realized at linear customer prices $p_S \geq 0$ and $p_R \geq 0$ such that the effective prices are equal to each other and to the inverse demand, which is the (equal) customer price the providers would charge in a market in absence of disutility effects. Given a market size x_M and realized demands x_S and x_R, this can be formalized as

$$p_S + \hat{J}_S(x_S, x_R, f_S, f_R) = p_R + \hat{J}_R(x_S, x_R, f_S, f_R) = P_M(1 - \frac{x_S + x_R}{x_M}), \quad (20)$$

where $\hat{J}_S(x_S, x_R, f_S, f_R)$ and $\hat{J}_R(x_S, x_R, f_S, f_R)$ are the expected per unit disutility costs seen by the customers of providers S and R, respectively, as defined in (19), and $P_M(1 - \frac{x_S+x_R}{x_M})$ is the inverse demand seen by the market as a whole. The idea here is that, if the effective prices were not equal, customers would switch to the provider with the lower effective price, thus equalizing effective prices through increasing expected per unit disutility. In addition, if the effective prices were not equal to the inverse demand, the total demand would change until they are. That is, if the effective prices were below the inverse demand, the total demand would increase and, if the effective prices were above the inverse demand, the total demand would decrease.

Using p_S and p_R as functions of x_S and x_R (20), these profits can be written as

$$\pi_S = p_S x_S + a_S f_R - a_R f_S = P_M(1 - \frac{x_S + x_R}{x_M})x_S - J_S(x_S, x_R, a_S, a_R, f_S, f_R)$$

$$\pi_R = p_R x_S - a_S f_R + a_R f_S = P_M(1 - \frac{x_S + x_R}{x_M})x_R - J_R(x_S, x_R, a_S, a_R, f_S, f_R),$$

which is the same as (2), and $J_S(x_S, x_R, a_S, a_R, f_S, f_R)$ and $J_R(x_S, x_R, a_S, a_R, f_S, f_R)$ are as defined in (3), except that now we define the *providers' costs* as their customers' aggregate disutility as follows

$$\begin{aligned}\tilde{J}_S(x_S, x_R, f_S, f_R) &\triangleq x_S \hat{J}_S(x_S, x_R, f_S, f_R) \\ &= c_S(x_S - f_d)(x_S - f_S) + c_R(x_R + f_d) f_S \quad (22a) \\ \tilde{J}_R(x_S, x_R, f_S, f_R) &\triangleq x_R \hat{J}_R(x_S, x_R, f_S, f_R) \\ &= c_S(x_S - f_d) f_R + c_R(x_R + f_d)(x_R - f_R). \quad (22b)\end{aligned}$$

It is clear that *the providers' profits are constrained by the aggregate disutilities*. That is, providers profit maximization would translate to provider effective cost minimization for every fixed x_S and x_R.

We then get the following result that parallels Theorem 1. The proof of this result is similar to the one for Theorem 1, and is therefore omitted.

Theorem 6 *In the Nash game* (5)–(8), *with provider costs as defined in* (22), $x_S x_R > 0$, $f_S f_R = 0$, *and* $f_S + f_R > 0$ *in any subgame-perfect equilibrium. In addition, in any subgame-perfect equilibrium,* $C'_S(x_S) \neq C'_R(x_R)$.

This result not only assumes existence of subgame-perfect equilibria but also does not ensure participation in resource sharing. However, it is possible to do both for the special case of linear (or affine) marginal cost functions, in a manner similar to Section 4.

5.2 When providers serve distinct, inelastic markets

The second case is also similar to our original setting except that providers do not compete for customers—i.e., we have a setting where dedicated providers serve distinct, inelastic markets, with fixed x_S, x_R. For example, as mentioned in Section 1, this would happen if ISPs were serving distinct customer segments, such as ISP monopolies in access markets.

In this case, (1) would be modified to $p_S = p_R = p$, where p is what the customers would be willing to pay. In addition, (2) would be modified to

$$\pi_S = p_S x_S + a_S f_R - a_R f_S = p x_S - J_S(x_S, x_R, a_S, a_R, f_S, f_R)$$

$$\pi_R = p_R x_S - a_S f_R + a_R f_S = p x_R - J_R(x_S, x_R, a_S, a_R, f_S, f_R),$$

where $J_S(x_S, x_R, a_S, a_R, f_S, f_R)$ and $J_R(x_S, x_R, a_S, a_R, f_S, f_R)$ are as defined in (3). Now, since x_S, x_R, p are fixed, the providers' profit maximization problems are simply cost minimization problems, and analysis for the two-stage Nash game (5)–(8) applies. That is, all the results listed under the "Competition under Fixed Customer Bases" subsection in the Electronic Supplementary Material. In particular subgame-perfect equilibria always exist and participation in resource sharing is always guaranteed.[16]

5.3 An alternate sequential Nash game

Though we have justified the choice of the two stage Nash game (5)–(8) in Section 2.2, we now show that our results are robust with respect to the chosen game form. As a representative example, we consider a three-stage Nash game of complete information in which providers pick the customer bases in the first stage, access prices in the second stage, and load splits in the third stage. The access prices in the second stage are picked having committed to picking load splits in their self-interest in the third stage. Similarly, the customer bases in the first stage are picked having committed to picking the access prices and

[16] This work appeared in [36].

load splits in the second and third stages, respectively. The game is formulated below and is solved using the standard backward induction technique.

In the first stage, the providers simultaneously solve the following optimization problems: given x_R, provider S solves

$$\max_{x_S\in[0,x_M]} \pi_S^S \triangleq P_M\left(1-\frac{x_S+x_R}{x_M}\right)x_S - J_S(x_S,x_R) \tag{24}$$

and, given x_S, provider R solves

$$\max_{x_R\in[0,x_M]} \pi_R^S \triangleq P_M\left(1-\frac{x_S+x_R}{x_M}\right)x_R - J_R(x_S,x_R), \tag{25}$$

where π_S and π_R are the profits of provider S and R, respectively, as defined in (2), and $J_S(x_S, x_R)$ and $J_R(x_S, x_R)$ are response functions from stage two, as defined below in (26) and (27).

Then, in the second stage, for fixed x_S and x_R, the providers simultaneously solve the following optimization problems: given a_R, provider S solves

$$J_S(x_S,x_R) \triangleq \min_{a_S\in\mathbb{R}_+} J_S(x_S,x_R,a_S,a_R) \tag{26}$$

and, given a_S, provider R solves

$$J_R(x_S,x_R) \triangleq \min_{a_R\in\mathbb{R}_+} J_R(x_S,x_R,a_S,a_R), \tag{27}$$

where $J_S(x_S, x_R, a_S, a_R)$ and $J_S(x_S, x_R, a_S, a_R)$ are the response functions from stage one, as defined below in (7) and (8). Note that, due to fixed x_S and x_R, the providers' focus would shift from profit maximization to cost minimization in the last two stages of the game.

In the third stage of the game, for fixed x_S, x_R, a_S and a_R, providers solve the following optimization problems simultaneously: given f_R, provider S solves

$$J_S(x_S,x_R,a_S,a_R) = \min_{0\le f_S\le x_S} J_S(x_S,x_R,a_S,a_R,f_S,f_R), \tag{28}$$

and, given f_S, provider R solves

$$J_R(x_S,x_R,a_S,a_R) = \min_{0\le f_R\le x_R} J_R(x_S,x_R,a_S,a_R,f_S,f_R), \tag{29}$$

where $J_S(x_S, x_R, a_S, a_R, f_S, f_R)$ and $J_R(x_S, x_R, a_S, a_R, f_S, f_R)$ are as defined in (3) and $\tilde{J}_S(x_S, x_R, f_S, f_R)$ and $\tilde{J}_R(x_S, x_R, f_S, f_R)$ are the resource costs for provider S and R, respectively, as defined in (4). Note that (28) and (29) are the same as (7) and (8).

We then get the following result that parallels Theorem 1.

Theorem 7 *Assume that the following conditions hold.*[17]

- P_M, x_M, $C_S(.)$, *and* $C_R(.)$ *do not satisfy* $P_M(1 - \frac{3x}{x_M}) = C'_S(x) = C'_R(x)$.
- $C_S(.)$ *and* $C_R(.)$ *are not identical.*

Then, in the Nash game (24)–(29), $x_S x_R > 0$, $f_S f_R = 0$, *and* $f_S + f_R > 0$ *in any subgame-perfect equilibrium. In addition, in any subgame-perfect equilibrium,* $C'_S(x_S) \neq C'_R(x_R)$.

This result not only assumes existence of subgame-perfect equilibria but also does not ensure participation in resource sharing. However, it is possible to do both for the special case of linear (or affine) marginal cost functions, in a manner similar to Section 4.

6 Conclusion

6.1 Summary of results

This paper looks at strategic resource sharing by rational providers. We believe that this is the first attempt at combining competition for customers with resource sharing in presence of strategic routing and general costs functions.

We have introduced a sequential Nash game of imperfect information and have presented a constructive approach to derive many properties of the market. Most importantly, we have shown that simple pricing schemes can guarantee participation in resource sharing. Since ISP peering is always on a pairwise basis [28], our analysis from the perspective of number of ISPs is in the most general form.

6.2 Policy implications

We have shown that simple, easy-to-implement access pricing mechanisms guarantee that competing providers benefit from cooperation implicit in resource sharing. This resource sharing goes beyond the typical peering arrangements and also suggests that multi-hop resource sharing could even replace transit.

Our results have significant implications for policy makers since our they suggest that resource sharing (i.e., generalized peering) should be encouraged as a policy so that providers, who otherwise compete for customers, can benefit from strategic complementarities in their cost structures. The end result is an increase in social surplus.

[17] This first assumption is reasonable since it rules out the case that three arbitrarily chosen functions cross at the same point. The second assumption is what makes our analysis interesting since resource sharing makes sense in presence of providers with asymmetric cost functions.

One way to support resource sharing would be to encourage the formation of Internet exchanges (IX). These are physical infrastructures (essentially, interconnected network switches) that allow different ISPs to exchange Internet traffic without the need for costly and time-consuming bilateral connections.

6.3 Future work

Though we have shown participation in resource sharing for linear marginal costs, as a first step, we need to show that simple pricing schemes can guarantee participation in resource sharing even in presence of general cost functions.

Further, a shortcoming of this paper is that the analyzed model is stylized. Although this model is quite informative as a first step, we would like to extend this model to include the following features to make it more realistic.

6.3.1 Network structure

This model reduces an ISP network to a single path, and embeds network complexity in per-unit cost functions, as in previous research in this area that (e.g., [16, 35]). Given that sum of convex functions are still convex [3], the resource per unit costs are a good first order approximation for realistic networks costs [7]. However, they may not fully absorb the complexities of the underlying complex networks, with actual aggregate costs being determined by network topologies, source and destination profiles, and routing and other traffic engineering practices. To be more realistic, the model needs to incorporate some network structure.

Our model is a first step in this direction, and can be used as a basic building block. For example, if the per-unit cost functions were constants then our analysis would trivially be applicable on a per-path basis, and the identity of the effective supplier may change depending on the path in focus. Another possible extension would be to model the dynamics of the provider interaction, using our basic building block—in this analysis, the idea of the effective supplier may change based not only on the path in question but also on the time-of-day. Then, a suitable averaging scheme may need to be applied to figure out the final settlement between the providers.

6.3.2 Access prices

We have discussed the existence of simple pricing schemes that guarantee participation in resource sharing. However, the question of how these prices may be determined in a realistic setting has not been answered.

One way to do this would be to use network probing [34]. For example, if providers used intra-network average delay as the variable costs, providers could send probe packets to approximate the actual delays incurred in the other networks. These probes can then be judicially controlled to calculate the marginal delays (e.g., through successive approximations [32]) that can in turn be used to set access prices.

6.3.3 Multiple providers

A competitive market may have more than two providers, and resource sharing arrangements among multiple providers may be more complicated. A more comprehensive analysis would include resource sharing among multiple providers, and potentially investigate what happens to the structural properties as the number of providers becomes large.

6.3.4 Heterogeneous cost structures

In the case where providers care about aggregate customer disutility (Section 5.1), it is quite possible that providers do not see the same cost structures on the same resource. This may happen, e.g., if a resource owner treats its own traffic differently from the traffic of the other provider. This may also happen in the case of ISP peering, where the resources are complex networks, with actual aggregate costs being determined by network topologies, source and destination profiles, and routing and other traffic engineering practices. Incorporating different cost structures on the same resource would be one way to capture these details.

6.3.5 Uncertainty

Finally, we have studied a model without uncertainty or private information. In reality, there is considerable uncertainty about not only network structure but also traffic flows. In addition, the characteristics of a provider's network is its private information, which may only be implicitly derived. Thus, using random network structures and traffic flows, and employing the idea of Bayesian Nash equilibrium would be a logical next step.

References

1. Anderson, D., Snoeren, A., & H. Balakrishnan (2003). Best path vs. multi path overlay routing. In *Proceedings of the 3rd ACM SIGCOMM conference on Internet measurement* (pp. 91–100).
2. Baake, P., & Wichman, T. (1999). On the economics of Internet peering. *Netnomics, 1*, 89–105.
3. Boyd, S., & Vanderberghe, L. (2004). *Convex optimization*. Cambridge University Press.
4. Cao, J., & Ramanan, K. (2002). A Poisson limit for buffer overflow probabilities. In *Proceedings of IEEE INFOCOM* (pp. 994–1003).
5. Cremer, J., Rey, P., & Tirole, J. (2000). Connectivity in the commercial Internet. *Journal of Industrial Economics, 48*, 433–472.
6. DeGraba, P. (2000). Bill and keep at the central office as the efficient interconnection regime. In *OPP working paper series* (Vol. 33).
7. Fortz, B., & Thorup, M. (2000). Internet traffic engineering by optimizing OSPF weights. In *Proceedings of IEEE INFOCOM* (pp. 519–528).
8. Foros, O., & Hansen, B. (2001). Competition and compatibility among Internet Service Providers. *Information Economic and Policy, 13*, 411–425.
9. Frieden, R. (1998). Without public peer: The potential regulatory and universal service consequences of Internet balkanization. *Virginia Journal of Law and Technology, 3*, 1522–1687.

10. Fudenberg, D., & Tirole, J. (2000). *Game theory*. MIT Press.
11. Giovannetti, E. (2002). Interconnection, differentiation, and bottlenecks in the Internet. *Information Economics and Policy, 14*(3), 385–404.
12. Jahn, E., & Prufer, J. (2006). Interconnection and competition among asymmetric networks in the network backbone market. TILEC discussion paper no. 2006-033. Available at http://arno.uvt.nl/show.cgi?fid=56729.
13. Harris, C., Reny, P., & Robson, A. (1995). The existence of subgame-perfect equilibrium in continuous games with almost perfect information: A case for public randomization. *Econometrica, 63*(3), 507–544.
14. Huston, G. (1999). *ISP survival guide: Strategies for running a competitive ISP*. New York: Wiley.
15. Jew, B., & Nichols, R. (1999). Internet connectivity: Open competition in the face of commercial expansion. In *Pacific telecommunication conference*. Honolulu, HI.
16. Johari, R. (2003). Routing and peering in a competitive Internet. MIT LIDS Publication 2570, January 2003. Available at http://www.stanford.edu/~rjohari/.
17. Kreps, J. M., & Scheinkman, J. A. (1983). Quantity precommitment and Bertrand competition yield Cournot outcomes. *The Bell Journal of Economics, 14*(2), 326–337.
18. Laffont, J. J., Marcus, S., Rey, P., & Tirole, J. (2003). Internet interconnection and the off-net-cost pricing principle. *RAND Journal of Economics, 34*(2), 370–390.
19. Lambert, P. (1997). *Sprint, UUNet to cut ties to Internet exchange*.
20. Le Cadre, H., Bouhtou, M., & Tuffin, B. (2009). Consumers' preference modeling to price bundle offers in the telecommunications industry: A game with competition among operators. *Netnomics, 10*, 171–208.
21. Leland, W. E., Taqqu, M. S., Willinger, W., & Wilson, D. V. (1994). On the self-similar nature of ethernet traffic. *IEEE Transactions on Networking, 2*(1), 1–15.
22. Level 3 (2002). *Interconnection policy*. Available at http://www.level3.com/us/products/crossroads/policy.
23. Luenberger, D. (1984). *Linear and nonlinear programming*. Springer.
24. Kleinrock, L. (1974). *Queueing systems*. Wiley.
25. Marsdenm, J. E., & Hoffman, M. J. (1993). *Elementary classical analysis*. W. H. Freeman and Co.
26. Marcus, J. S. (1999). *Designing wide area networks and internetworks: A practical guide*. Addison Wesley.
27. Mendelson, H., & Hwang, S. (1990). Optimal incentive compatible pricing for the M/M/1 queue. *Operations Research, 38*(5), 870–883.
28. Norton, W. B. (2001). *Internet service providers and peering*. Available at http://www.equinix.com.
29. Nucciarelli, A., Gastaldi, M., & Levialdi, N. G. (2009). Asymmetric competition and collection rates differentials: Determinants of prices in international telephone service markets. *Netnomics*. doi:10.1007/s11066-009-9042-5.
30. Pindyck, R. (2003). Mandatory unbundling and irreversible investment in telecom networks. Working paper 4452-3. MIT Sloan School of Management.
31. Rappa, M. A. (2004). The utility business model and the future of computing services. *IBM System Journal, 43*(1), 32–42.
32. Robbins, H., & Monro, S. (1951). A stochastic approximation method. *Annals of Mathematical Statistics, 22*(3), 400–407.
33. Shakkottai, S., & Srikant, R. (2005). Economics of network pricing with multiple ISPs. In *Proceedings of IEEE Infocom* (Vol. 1, pp. 184–194).
34. Shih, M. F., & Hero, A. (2001). Unicast inference of network link delay distributions from edge measurements. In *Proceedings of IEEE ICASSP* (Vol. 6, pp. 3421–3424).
35. Shneorson, S., & Mendelson, H. (2003). *Internet pricing, capacity and pricing*. Stanford Graduate School of Business working paper.
36. Shrimali, G., & Kumar, S. (2006). Paid peering among Internet Service Providers. In *Proceedings of IEEE workshop on game theory for networks*. Pisa, Italy.
37. Shrimali, G., & Kumar, S. (2008). Bill-and-keep peering. *Telecommunications Policy, 32*(1), 19–32.

38. Shrimali, G., Mutapcic, A., & Akella, A. (2007). Cooperative interdomain traffic engineering using Nash bargaining and dual decomposition. In *Proceedings of IEEE Infocom*, Anchorage, AK, pp. 330–338.
39. Tan, Y., Chiang, R., & Mookerjee, V. (2006). An economic analysis of interconnection agreements between Internet backbone providers. *Operations Research, 54*, 776–788.
40. Tirole, J. (1988). *The theory of industrial organization*. Cambridge: The MIT Press.
41. Verizon Communications (2005). *Verizon policy for settlement-free interconnection with Internet networks*.
42. Vives, X. (1990). Nash equilibrium with strategic complementarities. *Journal of Mathematical Economics, 19*(3), 305–321.
43. Voxel (2009). Paid peering: Everything you wanted to know (but were afraid to ask). Available at http://www.voxel.net/blog/2009/01/paid-peering/.
44. Weiss, M. B., & Shin, S. J. (2004). Internet interconnection economic model and its analysis: Peering and settlement. *Netnomics, 6*(1), 43–57.