Interactions, Competition and Innovation in a Service-Oriented Internet: An Economic Model

Zhi-Li Zhang, Papak Nabipay & Andrew Odlyzko
University of Minnesota
zhzhang@cs.umn.edu, {nabi0003, odlyzko}@umn.edu

Roch Guerin
University of Pennsylvania
guerin@ee.upenn.edu

Abstract—This paper presents a new economic approach for studying competition and innovation in a complex and highly interactive system of network providers, users, and suppliers of digital goods and services (i.e., service providers). It employs Cournot and Bertrand games to model competition among service providers and network providers, respectively, and develops a novel unified model to capture the interaction and competition among these players in a "service-oriented" Internet. Incentives for service and network innovation are studied in this model.

I. INTRODUCTION

As the Internet evolves and diffuses through society, economic factors are often more important than technological ones in determining what services are offered, and how they are priced. But the economics of the Internet is far more mysterious than the underlying technology. That is true even at the level of network providers, where the basic connectivity is fairly well known, but the business relationships that give rise to observed connections, and the incentives that led to them, are mostly hidden from view. When we come to services offered over the Internet, the complexity increases, and our knowledge drops even further. And the literature in this area is still limited, in spite of the extensive interest that exists.

In this paper, we consider the Internet as a service delivery platform (i.e., we disregard its other roles, such as in providing connectivity among users, etc.), and study the relationships among the entities that provide transport (*network providers*, in our language), those that provide services (*service providers*), and *users*. We assume a certain industry structure (with network providers completely separate from service providers), and several rules imposed by regulators (such as some non-discrimination conditions, and possibly even some price regulation), and some other common market features (such as basic network connectivity being provided to users on a flat rate basis). This constrains the problem enough to provide opportunity to build models that are tractable and yet reflect what is observed in the marketplace, and are rich enough to show interesting dynamics.

We propose a simple economic model of the interactions and competition among service providers, network providers and users. Using this model, we explore how competition affects the network and service providers, and in particular how to maximize the incentives for innovation on the part

of network providers and service providers. Our main tools come from the standard economic literature on *Bertrand* and *Cournot* competition (see Section II for a brief overview of these concepts). While both types of competition are well known in economics, *one of the key novelties and contributions of our paper is combining these two different types of games in a single unified framework to capture the co-dependence or interaction between service and network providers.* With our assumptions, we model the competition between service providers using Cournot games, and the competition between network providers using Bertrand games.

The two types of competition (or games) are tied together in a *two-stage Stackelberg game* where service providers determine the optimal (equilibrium) amount of services each produces/offers to meet user demands, and network providers determine what optimal (equilibrium) prices to charge service providers for transporting the accompanying services. We are able to explicitly solve this unified Cournot-Bertrand model, and thereby study the effects of competition between service providers and competition between network providers on the overall equilibrium market demand/supply and prices. Furthermore, it enables us to investigate and quantify incentives for service and network providers to innovate and further spur the market demand for services.

Because of its simplifying assumptions, our economic model clearly does not capture the intricate and complex relationships and industry structures that exist in the real Internet, and represents only a *modest* attempt in analyzing and understanding these relationships in a formal economic setting. Nonetheless, our work provides a useful and *tractable* model to generate some *qualitative* insights into these relationships. To our best knowledge, our paper is perhaps the first attempt to explicitly model the interaction, competition and innovation among service providers, network providers and users in a "service-oriented" Internet. As more services–not only content (news, music, videos, etc.), but also software, computing and storage resources (as in *cloud computing*)–are being offered online, we believe that understanding the economic factors that affect the interacting, co-dependent and yet competing relationships among various players in this service-oriented Internet is of critical importance. As the research community contemplates "clean-slate" designs for future Internet architectures, models for assessing the economic viability of new network architectures are especially needed. We hope our work

can inspire more studies to follow.

Related Work. Network economics has been a very active area of research. While there are many papers on network pricing and other economic issues of networks (see, e.g., the proceedings of *NetEcon* workshops or *Review of Network Economics* journal), relatively few studies have focused on economic aspects of network architectures. Examples of the latter include recent studies of network neutrality (see, e.g., [1]–[3]), evaluating network architectures via the theory of real options [4], and modeling technology adoption dynamics [5].

II. PROBLEM SETTING

A. Problem Setting and Basic Assumptions

We consider the Internet as a service delivery platform, and assume three *separate* types of entities: *service providers*, *network providers* and *users*. Fig. 1 schematically depicts the relations among the three entities. Clearly, users are the key drivers in the relations among the three. Users pay a fixed monthly fee to their network providers for basic connectivity, which enables them to access various application services. While users pay for basic connectivity through a flat rate, we assume that access to services such as music or video, or cloud computing, is *not* free (and so, in particular, is not paid for through advertising). Instead, users pay for those services to the service providers. The payment might be per unit of service, e.g., $0.99 per song, or might be on a monthly subscription basis for a newspaper or music site. But in some rough sense users' fees to service providers are proportional to those users' volumes of consumption of those services. Service providers derive their revenues from user fees, and have to pay for the costs of creating their offerings, as well as for their transport to the users. Network providers derive their revenues from the flat monthly fees of their users and the usage-sensitive transport charges from service providers (see eq. (4)), but have to pay for their infrastructure.
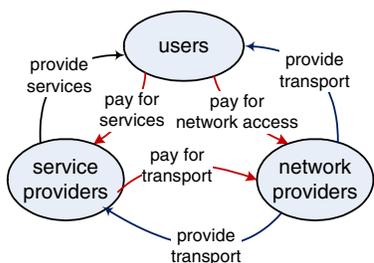

Fig. 1. Relations between network providers, service providers and users.

In our model we assume that there are only two service providers who offer the same type of services (in the economic parlance, they are *substitutable goods*) and compete between themselves for user accesses. There are also only two network providers, each connecting their respective users to the two service providers and competing between themselves for transporting services to users. For simplicity, we assume that each user stays with the same network provider throughout (e.g., due to geographical locations and local incumbent monopoly).

Fig. 2 depicts the flow of commodities (namely, services and the associated traffic) and prices. As in the current Internet, the action of any one of the actors, say in improving the transmission infrastructure, or in changing a price, affects all other actors, and leads to interesting dynamics.

B. A Quick Economic Primer and Model Justification

In economics, *demand (function/curve)* is defined as a function between the price of a commodity (*services* in our model) and the quantity of the commodity (the amount of service in our model) that consumers/users are willing and able to purchase at the given price [6]. Likewise, *supply (function/curve)* is defined as a function between the price of a commodity and the quantity of the commodity that a firm (a service provider in our case) is able to produce at the given price. Demand curves are used in economic models to estimate behaviors in *competitive* markets. At the *equilibrium price*, demand and supply are equal, and the *equilibrium quantity* (the amount of service) will be produced and consumed at the equilibrium price.

We use two different types of standard economic games in our model: *Cournot competition* between service providers, and *Bertrand competition* between network providers. In economics [7], *Cournot competition/game* is a model of competition in which firms (service providers in our case) *compete on quantity* (i.e., the volume of services) they produce–which they decide independently and at the same time– based on their cost of production, represented by the *marginal cost*, or the cost for producing one additional unit of service. In a Cournot game, *the market price is set at a level such that demand equals the total quantity produced by all firms.* In other words, the service price is determined by the supply and demand for services in the market. In contrast, in a *Bertrand competition/game* firms (network providers in our case) *compete on price*–with each one choosing its price independently and at the same time–and supply the quantities demanded at those prices.

In our model we use a Cournot game to characterize the competition between two service providers, as we assume they produce *substitutable* services (e.g., music or video downloads), and the price of these services is essentially determined by competition and market supply and demand. In other words, due to competition, different service providers have limited ability to choose their own prices, and in equilibrium, they converge to the same price and consequently determine the equilibrium amount of services each produces or offers. On the other hand, we use the Bertrand game to characterize the competition between two network providers, as they do *not* directly produce services, but instead transport the full amount of services produced by service providers to their respective users. They *can* indirectly influence the amount of services produced by service providers, thus the amount (of traffic) they have to transport, by varying the prices they charge them. Hence, network providers compete on the network transport prices/fees they charge the service providers to influence or attract the amount of services they transport.
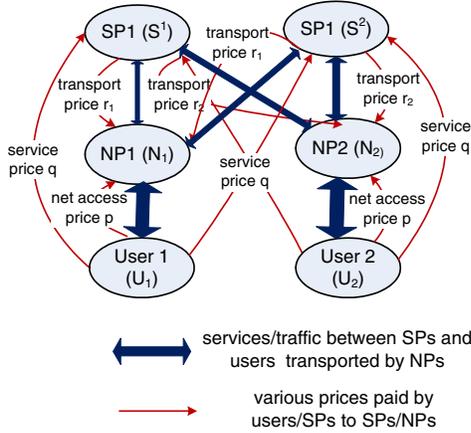
Fig. 2. Model illustration: flows of prices and services/traffic.

III. FORMAL MODEL AND ITS SOLUTION

A. The Model

As depicted in Fig. 2, there are two network providers[1], denoted by N_i, $i = 1, 2$, and two service providers, S^j, $j = 1, 2$. The user (base), U_i, of network provider N_i is assumed fixed. Each user U_i, $i = 1, 2$, pays a price of p (e.g., a monthly network access charge) to access the network, and buys services from either one of the two service providers, paying a price of q per unit of service. Network provider N_i charges service provider S^j a price of r_i to transport one unit of service (or rather the associated traffic) between S^j and U_i. Note that while the transport prices r_1 and r_2 charged by the network providers may be different, they do not charge different prices to different service providers. (Hence a form of network neutrality is assumed in our model.) Note also that in our model both q and r_i represent usage-sensitive (i.e., "volume-based") pricing, while p is a flat-rate price.

The overall user demand for services is a function of both the network access charge p and the service price q. Hence when p or q increases, user demands for services decrease. Let x_i^j, $i, j = 1, 2$, denote the amount of service produced by service provider S^j and consumed by user U_i. The relation between the service price q, the network access charge p and the service quantities ("supplies"), x_i^j, is assumed to be given by the following (inverse) demand function:

$$q = 1 - \beta p - \sum_{j=1}^{2} \sum_{i=1}^{2} \frac{1}{\gamma_i^j} x_i^j. \tag{1}$$

We first note that eq. (1) implicitly assumes that the service price is normalized to be within the range [0,1]; hence service quantities and other parameters are appropriately normalized. The parameter $\beta(\geq 0)$ in eq. (1) captures the effect of the network access charge p on the overall user demand, viz., the

[1]Throughout this paper we use the convention that *subscripts* denote *network* providers, whereas *superscripts* denote *service* providers. For conciseness, we will also use the singular term "user" to collectively refer to all users (the *user base*) of one network provider, and treat the collection of them as if they were a single "user."

"user demand sensitivity" to p. The parameters $\gamma_i^j (\geq 0)$, $i, j = 1, 2$, capture the combined "abilities" of service provider S^j and network provider N_i to offer and transport services to meet the user demand. The role of γ_i^j's can be intuitively understood as follows: $\gamma_i^j = -\partial x_i^j / \partial q$, so that γ_i^j represents the proportional "market share" of an increase in user demand that service provider S^j and network provider N_i can jointly capture by increasing the service supply (by ∂x_i^j amount) to user U_i of network provider N_i when the service price q drops by ∂q. Hence the larger γ_i^j is, the better S^j (and N_i)'s ability to increase its supply and gain market share.

The two service providers engage in a Cournot competition. Given a service price q as in eq. (1), S^1 and S^2 compete to determine the "optimal" amount of services, x_i^j, $i = 1, 2$, to offer users of both networks so as to maximize their respective profits. Let s^j denote the *marginal cost* of S^j (for producing one unit of service). The profit generated by S^j for supplying an amount x_i^j of services to user U_i is given by

$$\Pi_i^{S^j} := (q - r_i - s^j) x_i^j, \tag{2}$$

where r_i is the transport price that N_i charges for transporting the traffic associated with *one unit of service*. Let $X^j := x_1^j + x_2^j$ denote the total amount of service produced by S^j. The total profit of S^j is then given by

$$\Pi^{S^j} := \Pi_1^{S^j} + \Pi_2^{S^j} = (q - r_1) x_1^j + (q - r_2) x_2^j - s^j X^j. \tag{3}$$

In contrast, network providers engage in a *Bertrand competition*. For $i = 1, 2$, define $Y_i := x_i^1 + x_i^2$, the total amount of services produced by S^1 and S^2 and consumed by user U_i of N_i. Thus Y_i is the amount of services (or the accompanying traffic) that network provider N_i must transport between its user and the service providers. Network provider N_i does not "produce" Y_i; it only indirectly controls it by adjusting the price r_i it charges the two service providers. Let n_i denote the *marginal cost* of N_i (for transporting the traffic associated with one unit of service). The total profit of N_i is thus

$$\Pi^{N_i} := (r_i - n_i) Y_i + p, \ i = 1, 2, \tag{4}$$

where the first term in the right hand side represents the net profit for transporting service traffic, and the second term is the price paid by the user for network access. In a Bertrand game, the two network providers, N_i, $i = 1, 2$, compete by determining the "optimal" price r_i each charges the service providers to maximize their respective profit.

Combining the Cournot competition between service providers, and the Bertrand competition between network providers, produces a *Stackelberg* game consisting of a Cournot (sub-)game and a Bertrand (sub-)game. The parameters and variables used in the model are summarized in Table I. In the next subsection, we show how the model can be solved.

B. Solving the Model

We solve the model using a two-stage procedure. In the first stage, given the network transport prices r_i's, the service providers S^1 and S^2 compete in a Cournot game. The Nash equilibrium state of the Cournot game yields the "optimal"

TABLE I
SUMMARY OF MODEL PARAMETERS AND VARIABLES

List of Parameters		
Parameters	β	user demand sensitivity to p
	γ_i^j	(joint) "market share" of S^j & N_i
	$\gamma^j = \gamma_1^j + \gamma_2^j$	total "market share" of S^j
	$\gamma_i = \gamma_i^1 + \gamma_i^2$	total "market share" of N_i
List of Variables		
Quantities	x_i^j	service of S^j consumed by U_i
	$X^j = x_1^j + x_2^j$	total service produced by S^j
	$Y_i = x_i^1 + x_i^2$	total service transported by N_i
Prices	q	user service price
	r_i	network transport price (by N_i)
	p	user network access charge

equilibrium service quantities, x_i^j, $i, j = 1, 2$, and service price q, all as functions of p, r_i's. In the second stage, the network providers N_1 and N_2 compete in a Bertrand game to determine the network transport prices, r_i's, so as to maximize their respective profit. Substituting r_i's into the (optimal) x_i^j's obtained in the Cournot game, we obtain the final optimal equilibrium service quantities, x_i^j's, in terms of p (and the model parameters). Using this two-stage procedure, the solution to the Stackelberg game with two-stage Cournot-Bertrand sub-games is stated in the following theorem. A more detailed derivation of this theorem can be found in [8], where numerical examples are also provided to illustrate the results.

Theorem 1 (Solution of the Model): For a fixed network access charge p, the competitions between service providers and network providers, modeled as a Stackelberg game with two-stage Cournot-Bertrand sub-games, yield the *equilibrium state* $\langle \{x_i^j, i = 1, 2, j = 1, 2\}, \tilde{q}, \tilde{r}_1, \tilde{r}_2 \rangle$, where

$$\tilde{x}_i^j = \frac{\gamma_i^j}{5}\left[\frac{3}{8}(1 - \beta p) - \frac{3}{4}(s^1 + s^2) - \frac{7}{16}(n_1 + n_2)\right.$$
$$\left. + \frac{7s^1}{16}\left(\frac{\gamma_1^1}{\gamma_1} + \frac{\gamma_2^1}{\gamma_2}\right) + \frac{7s^2}{16}\left(\frac{\gamma_1^2}{\gamma_1} + \frac{\gamma_2^2}{\gamma_2}\right)\right] + \frac{\gamma_i^j}{5}(w_{-i} - s^j), \quad (5)$$

here $w_{-i} := \frac{1}{4}[n_{-i} - \frac{1}{\gamma_{-i}}(s^1\gamma_{-i}^1 + s^2\gamma_{-i}^2)]$, and for a given player i, the notation $-i$ indicates the other player (network provider); furthermore,

$$\tilde{r}_i = \frac{1}{8}(1 - \beta p) - \frac{1}{4}(s^i + s^j) + \frac{1}{16}(3n_i - n_{-i})$$
$$- \frac{s^i}{16}\left(\frac{3\gamma_i^1}{\gamma_i} - \frac{\gamma_{-i}^1}{\gamma_{-i}}\right) - \frac{s^j}{16}\left(\frac{3\gamma_i^2}{\gamma_i} - \frac{\gamma_{-i}^2}{\gamma_{-i}}\right). \quad (6)$$

and

$$\tilde{q} = \frac{3}{10}(1 - \beta p) + \frac{1}{5}(s^1 + s^2) + \frac{1}{20}(n_1 + n_2)$$
$$- \frac{s^1}{20}\left(\frac{\gamma_1^1}{\gamma_1} + \frac{\gamma_2^1}{\gamma_2}\right) - \frac{s^2}{20}\left(\frac{\gamma_1^2}{\gamma_1} + \frac{\gamma_2^2}{\gamma_2}\right). \quad (7)$$

∎

As a result of the Cournot/Bertrand competitions between service/network providers, the amount of service, $\tilde{x}_i^j$'s, offered at equilibrium by S^j to user U_i of network provider N_i is (roughly) in the proportion $\frac{\gamma_i^j}{\gamma}$ of the total market. Thus the "market share" of service provider S^j in network N_i is (roughly) γ_i^j, and its total "market share" in both networks is (roughly) $\gamma^j := \gamma_1^j + \gamma_2^j$. Similarly, network provider N_i

carries (roughly) a fraction γ_i/γ of the total traffic generated by the service market, with a market share of (roughly) γ_i.

Effect of Network Access Price p and Its Optimal Choice. In our model, we assume that the price p is determined by a policy maker with the objective of maximizing *social welfare*. Using Theorem 1, we now determine the optimal price p^* that maximizes the total social welfare. Following standard economic principles, we define the *social welfare function*, denoted by Π^{SW}, as follows:

$$\Pi^{SW}(p) := \Pi^U(p) + \sum_{i=1}^{2} \Pi^{N_i}(p) + \sum_{j=1}^{2} \Pi^{S^j}(p), \quad (8)$$

where $\Pi^U(p)$ is the *consumer (i.e., user) surplus* when the network access price is p. Since the total market demand, $\tilde{X}(p) := \sum \tilde{x}_i^j$, as a function of p is given by eq. (5), the consumer surplus is therefore given by

$$\Pi^U := \int_p^{p_{max}} \tilde{X}(p) = \frac{1}{2}(p_{max} - p) \cdot \tilde{X}(p). \quad (9)$$

The social welfare is a convex function. In particular, its maximum is attained when either $p = 0$ or $p = p_{max}$. When $p = p_{max}$, there is no user demand for services. The network providers only profit from network access charges paid by the users. This scenario obviously only makes sense if there is intrinsic value in the network besides the services offered by the service providers. Otherwise, users would have no incentive to use the network and pay a network access charge; the network provider would then lose users and thus receive zero profit in the end. On the other hand, when $p = 0$, the market demand is maximized, and therefore the consumer surplus is also maximized. Further, both the profits of the service providers and the network providers are non-zero.

IV. INCENTIVES FOR INNOVATIONS

In this section, we consider the following questions: i) under what conditions would either a service provider or network provider have incentives for *service innovation* or *network upgrade*; and ii) how do the interactions and competition among them affect such incentives? To separate the effects of innovation by service and network providers, here we refer to innovation by a service provider as service innovation, and that by a network provider as network upgrade.

We define *service innovation* as an investment by a service provider which will result in an increase in the quality of the supply or a decrease in the market (service) price, thereby expanding its share of the user demand for services. Likewise, we define *network upgrade* as an investment by a network provider which will result in an increase in the quality and/or capability of the network infrastructure or a decrease in the market (transport) price, thereby expanding its share of the market demand for service/traffic transport. In both cases, any innovation or upgrade will increase the overall market demand.

A. Incentives for Service Innovation

We indirectly model the investment made by a service provider for service innovation by accounting for its *effect* in our model. Consider service provider S^j. Innovation by S^j would lead to an increase in the overall market demand for services, and therefore result in a decrease in the service price q. From the inverse demand function of eq. (1), $\partial x_i^j/(-\partial q) = \gamma_i^j$, and $\partial X^j/(-\partial q) = \partial x_1^j/(-\partial q) + \partial x_2^j/(-\partial q) = \gamma_1^j + \gamma_2^j = \gamma^j$. We assume that innovation by S^j would increase its competitiveness, and thus its corresponding "market share" of the overall user demand. Hence we model the *effect* of innovation by S^j by a positive increase in γ^j, and assume that service innovation by S^j has the same effect on its service offerings in both network providers. Namely, after the service innovation by S^j, we have for some $\eta^j > 0$ $\gamma_i^j := \gamma_i^j + \eta^j$, $i = 1, 2$, and $\gamma^j := \gamma^j + 2\eta^j$. On the other hand, service innovation by S^j may also affect its marginal cost s^j, namely from s^j to $s^j + \sigma^j$. The effect of service innovation on the marginal cost can be either positive (i.e., $\sigma^j > 0$), or negative (i.e., $\sigma^j \leq 0$), or neutral. All three options are accounted for in the model, which captures the effect of service innovation by S^j using two parameters (η^j, σ^j), where $\eta^j > 0$.

Let $\Pi^{j,before}$ denote the overall profit of service provider S^j before service innovation, and $\Pi^{j,after}$ its profit after service innovation. In order for S^j to have an incentive to innovate, its profit after service innovation should be larger than before service innovation. In other words, we must have

$$I^j(\eta^j, \sigma^j) := \Pi^{j,after} - \Pi^{j,before} > 0. \qquad (10)$$

It can be shown that $I(\eta^j, \sigma^j)$ is an increasing function in η^j and a decreasing function in σ^j. Hence if $\sigma^j < 0$, i.e., service innovation by S^j decreases its marginal cost, then for any $\eta^j > 0$, $I(\eta^j, \sigma^j) > 0$. Therefore when service innovation leads to a decrease in its marginal cost, it always pays for S^j to innovate. On the other hand, when $\sigma^j > 0$, service innovation by S^j would lead to an increase in its marginal cost. In this case, we can show that provided that $s^1 + s^2 + n_1 + n_2 << 1 - \beta p$, there exists a constant $c^j > 0$ such that if $\sigma^j < c^j \frac{\eta^j}{\gamma^j + \eta^j}(1 - \beta p)$, then $I(\eta^j, \sigma^j) > 0$. Hence in this case, it is only when the increase in its marginal cost is upper bounded by an appropriate gain in its market share that service innovation by S^j would then pay off. Otherwise, there is no incentive for S^j to innovate.

B. Incentives for Network Upgrade

Likewise, we model the *effect* of network upgrade by network provider N_i using two parameters, (η_i, μ_i), where $\eta_i > 0$ reflects an increase in N_i's market share as a result of its upgrade, i.e, $\gamma_i^1 := \gamma_i^1 + \eta_i$, $\gamma_i^2 := \gamma_i^2 + \eta_i$ and thus $\gamma_i := \gamma_i + 2\eta_i$; whereas μ_i reflects the resulting change in its marginal cost, $n_i := n_i + \mu_i$, which can again be either positive, negative or neutral. The *incentive function* I_i for N_i is given by

$$I_i(\eta_i, \mu_i) := \Pi_{i,after} - \Pi_{i,before}, \qquad (11)$$

where $\Pi_{i,before}$ and $\Pi_{i,after}$ denote the profit of N_i before and after network upgrade, respectively. N_i has incentive to upgrade its network if and only if $I_i(\eta_i, \mu_i) > 0$. Again we can show that when network upgrade leads to a decrease in its marginal cost, i.e., $\mu_i < 0$, it always pays for N_i to upgrade its network. On the other hand, when $\mu_i > 0$, it is only when the increase in its marginal cost is upper bounded by an appropriate gain in its market share that network upgrade by N_i would pay off. Otherwise, N_i has no incentive to upgrade.

V. CONCLUSIONS

In this paper we have developed a simple economic model to study the interactions and competition among service providers, network providers and users. The novelty of the model lies in combining Cournot and Bertrand competition to capture the co-dependent, interacting and competitive relationships among service providers and network providers. Using this model, we explored how these relationships can affect incentives for innovation on the part of network providers and service providers. Our work represents a modest first attempt in characterizing and modeling the intricate and complex interactions that exist between various actors in a "service-oriented" Internet. We plan to further expand this model to include more realistic inter-connection structures among network providers and service providers, and thereby more faithfully capture the complex interactive and competitive relationships among them. We would also like to incorporate competition among network providers for users, and study the dynamics of evolving relationships among these actors. Conducting empirical studies of these relations based on real data is also a research area of great interest to us.

REFERENCES

[1] N. Economides and J. Tag, "Net neutrality on the Internet: A two-sided market analysis," 2007, sSRN eLibrary, http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1019121.

[2] J. Musacchio, G. Schwartz, and J. Walrand, "A two-sided market analysis of provider investment incentives with an application to the net-neutrality issue," *Review of Network Economics*, vol. 8, no. 1, pp. 22–39, March 2009.

[3] A. Odlyzko, "Network neutrality, search neutrality, and the never-ending conflict between efficiency and fairness in markets," *Review of Network Economics*, vol. 8, no. 1, pp. 40–60, March 2009.

[4] M. Gaynor and S. Bradner, "A real options metric to evaluate network, protocol, and service architecture," *ACM Computer Communication Review*, vol. 34, no. 5, October 2004.

[5] Y. Jin, S. Sen, R. Guerin, K. Hosanagar, and Z.-L. Zhang, "Dynamics of competition between incumbent and emerging network technologies," in *Proc. ACM SIGCOMM Workshop on the Economics of Networks, Systems, and Computation (NetEcon'08)*, 2008.

[6] P. Krugman and R. Wells, *Microeconomics*. New York: Worth Publishers, 2005.

[7] J. Tirole, *The Theory of Industrial Organization*. Cambridge, MA: MIT Press, 1988.

[8] Z.-L. Zhang, P. Nabipay, A. Odlyzko, and R. Guerin, "Interactions, competition and innovation in a service-oriented Internet: An economic model," 2009, technical Report Version, University of Minnesota, http://www-users.cs.umn.edu/~zhzhang/publications.html.