A Two-Sided Market Analysis of Provider Investment Incentives with an Application to the Net-Neutrality Issue

JOHN MUSACCHIO *

Technology and Information Management, University of California, Santa Cruz

GALINA SCHWARTZ

Department of Electrical Engineering and Computer Sciences, University of California, Berkeley

JEAN WALRAND

Department of Electrical Engineering and Computer Sciences, University of California, Berkeley

Abstract

We address whether local ISPs should be allowed to charge content providers, who derive advertising revenue, for the right to access end-users. We compare two-sided pricing where such charges are allowed to one-sided pricing where they are prohibited. By deriving provider equilibrium actions (prices and investments), we determine which regime is welfare-superior as a function of a few key parameters. We find that two-sided pricing is more favorable when the ratio between parameters characterizing advertising rates and end-user price sensitivity is either low or high.

1 Introduction

Today, an Internet service provider (ISP) charges both end-users who subscribe to that ISP for their last-mile Internet access as well as content providers that are directly connected to the ISP. However, an ISP generally does not charge content providers that are not directly attached to it for delivering content to end-users. One of the focal questions in the network neutrality policy debate is whether these current charging practices should continue and be mandated by law, or if ISPs ought to be allowed to charge all content providers that deliver content to the ISP's end-users. Indeed the current network neutrality debate began when the CEO of AT&T suggested that such charges be allowed (see Whitacre, 2005).

To address this question, we develop a two-sided market model of the interaction of ISPs, end-users, and content providers. The model is closely related to the existing two-sided markets literature as we detail later in this section. In our model, the ISPs play the “platform” role that intermediates the two sides: content providers and end-users. We model a “neutral”

* Contact author: John Musacchio, Technology and Information Management Program, Jack Baskin School of Engineering, University of California Santa Cruz, 1156 High Street, MS: SOE 3. E-mail: johnm@soe.ucsc.edu
This work was supported by the National Science Foundation under Grants ITR/ANIR 0331659 and NeTS-FIND 0627161 and 0626397. We sincerely thank the anonymous reviewers and the editors, Julian Wright and in particular Marius Schwartz, whose thoughtful and detailed suggestions for improving this paper are greatly appreciated.

network as a regime in which ISPs are allowed to charge only content providers that buy their Internet access from them. We assume that with such a charging structure, the ISPs compete on price to attract content providers to buy their access from them, driving these prices to the ISP's costs. By normalizing the price content providers pay ISPs to be net of ISP connection cost, we can model a neutral network as having a zero price for content providers. This could be viewed as a market with one-sided pricing where ISPs only charge end-users. Conversely, two-sided pricing (corresponding to a "non-neutral" network) allows all ISPs to charge all content providers, thus permitting the ISPs to extract revenues from both content providers and end-users. Our contribution is the development of a model that relates the level of end-user usage to the investment decisions of providers of both types under both network regimes. We use our model to compare the welfare of both regimes.

The question we address in this work is part of the larger debate on network neutrality, which includes diverse issues such as whether service differentiation should be allowed, or whether charges for content constitute an impingement of freedom of speech (see Odlyzko, 2008; and Chong, 2007). In 2006 there was a considerable divergence of opinions on the subject of net-neutrality. Indeed, the issue was intensely debated by law and policy makers, and the imposition of restrictive network regulations on Internet service providers (ISPs) in order to achieve network neutrality seemed likely. Recently, the situation has begun to change. In June 2007, the Federal Trade Commission (FTC) issued a report, forcefully stating the lack of FTC support for network neutrality regulatory restraints, and warning of "potentially adverse and unintended effects of regulation" (FTC, 2007, p.159). Similarly, on September 7, 2007 the Department of Justice issued comments "cautioning against premature regulation of the Internet," (DOJ, 2007). Thus by the fall of 2007, the imminent threat of new regulation has diminished somewhat, and a consensus favoring the current (or unregulated) network regime seemed to have emerged. Still, the debate about network neutrality is far from over. Many prominent members of Congress support new network neutrality legislation, including two leading candidates for president in the 2008 election. We do not attempt to address all of the questions in the network neutrality debate. We only study the issue of whether ISPs ought to be allowed to charge content providers for accessing the end-users.

Our model is based on the ideas of two-sided markets, and there is a large literature on the subject. For a survey of two-sided markets, see for example, Rochet and Tirole (2006) and Armstrong (2006). The two-sided market literature studies markets in which a platform provider needs to attract two types of participants, and the presence of more of one type makes the platform more valuable to the other type. Rochet and Tirole (2006) define the market as two-sided, when the volume of realized transactions depends not solely on the aggregate price level charged to the two parties, but also on how this aggregate is divided between them. Using the two-sided market parlance, the ISPs of our model provide the platform, while end-users are one type of participant and content providers are the other type. As will become clear when we describe the details of our model, the end-users "single-home" or connect to one ISP. In a two-sided network, the content providers are forced to "multi-home" or pay multiple ISPs for delivering their content (see for example Armstrong, 2006 on "competitive bottlenecks"). In contrast, the content providers in a network with one-sided pricing "single-home" or pay just one ISP for connectivity. However, a content provider that pays one ISP enjoys the benefits of having connectivity to all the ISPs, because all the ISPs are interconnected. This is in contrast to most two-sided market models where the participants of one platform do not benefit from the presence of

participants of another platform, that is, Microsoft Xbox users do not benefit from more game makers writing Nintendo Wii games. This is an important structural difference.

Figure 1: ISPs $\{T_1, \dots, T_6\}$, content providers $\{C_1, C_2\}$ and end-user u .

Other researchers have used the ideas of two-sided markets to study network neutrality. Hermalin and Katz (2007) model network neutrality as a restriction on the product space, and consider whether ISPs should be allowed to offer more than one grade of service. Hogendorn (2007) studies two-sided markets where intermediaries sit between “conduits” and content providers. In his context, net-neutrality means content has open access to conduits where an “open access” regime affords open access to the intermediaries. Weiser (2007) discusses policy issues related to two-sided markets.

The novelty of our model over other work in the two-sided market literature is our explicit modeling of platform investment choices. In the existing literature, the platform incurs the cost of serving the users, which usually is assumed linear in the number of users, but does not make an investment choice.

To introduce the notation and illustrate the arguments for and against the imposition of regulations that prohibit the ISPs from engaging in two-sided pricing, consider the network shown in Figure 1. The figure shows end-user u , ISPs $\{T_1, \dots, T_6\}$, and content providers $\{C_1, C_2\}$. In this network, T_1 and T_2 are transit ISPs (that is, the ISPs who operate portions of the Internet backbone), and $T_3 - T_6$ are local ISPs (i.e, residential ISPs who provide last-mile access for end-users). In the figure, the content providers are attached to a transit ISP whereas a typical end-user u is attached to a local ISP.

In our model, in a market with one-sided pricing, end-users and content providers pay only for their direct access. The transit ISPs charge the local ISPs for carrying their traffic. The transit ISPs typically enter in peering agreements under which they agree to carry each other’s traffic, usually free of charge. The transit ISPs charge the content providers for their attachment. We model one-sided pricing (“neutral” network), as a case where the local ISPs are prohibited from charging content providers not buying access directly from them. Thus, with one-sided pricing, content provider C_1 pays ISP T_1 for its access to the Internet but does not pay any of the other ISPs. In contrast, with two-sided pricing (“non-neutral” network) ISP T_3 is able to charge content provider C_1 for carrying its content to end-user u , even though C_1 is not directly attached to T_3 .

So far, we have distinguished between two classes of access providers: transit ISPs and local ISPs. The local ISPs provide the last-mile link to end-users (that is, residential consumers). The transit ISPs have no residential customers and serve only content providers, though there are notable examples of companies operating both transit and local networks. The core economic distinction between these ISP classes is that local ISPs are thought to

have more market power. This market power is a result of substantial scale economies in providing mass-market (residential) broadband access. These scale economies limit (today at least) the number of local ISPs in a specific residential area. In particular, today's end-user choice of local ISP is limited to one to three at most. To reflect this market power, in our model each local ISP has a monopoly over its user base.

In reverse, "first-mile" access for major content providers is supplied by transit ISPs, who are believed to be considerably more competitive, because individual content providers (i) generate a lot of traffic and therefore can economically support dedicated links from several suppliers, and (ii) can pick to locate themselves at the points where the network access is the cheapest (that is, to attach directly to the backbone).

Our main purpose of distinguishing transit and local ISPs was to help explain why under today's institutions, local ISPs do not charge content providers. First, content providers can buy access at their end (competitively) from transit ISPs; second, typically, transit ISPs do not pay for access to local ISPs (enjoy settlement-free "peering"); and third, local ISPs do not receive any other payments from content providers. Thus under current practices, the local ISPs do not charge content providers, except perhaps in the case where a company operates both transit and local networks. Even in such a case, one could argue that the ISP only charges competitively for general access rather than for the right to reach the ISP's users in particular. This is because, for instance, a content provider always has the option to switch to a non-integrated competitive transit provider. Thus, in our further analysis we abandon the distinction between local ISPs and transit ISPs, and focus on local ISPs only, which from now on we simply call ISPs.

The arguments for and against net-neutrality having to do with local ISPs charging content providers can be summarized as follows. See ACLU (2006), Farber (2007), Felten (2006), Lessig¹, Owen and Rosston (2003), Sidak (2006), and Yoo (2006) for more elaboration on these points.

Against Neutrality: This line of reasoning is usually expressed by ISPs. Say that C_1 is a source of video streams that require a large bandwidth. ISP T_3 may argue that to accommodate the traffic originating from C_1 , he must make substantial investments, which he may not be able to recover from the end-users who buy access directly from him. However, the content providers make additional advertising revenues when end-users consume new high-bandwidth services, which justify their investments. If the ISPs cannot get a share of these additional revenues, they will not invest as much to increase the network capacity. The situation causes poor network quality, which reduces end-user demand, which in turn leads to further reduction of incentives to invest for both provider types.

For Neutrality: This line of reasoning is usually expressed by content providers. If every ISP can charge the content provider C_i for the right to access its end-users, the ability of ISPs to charge content providers would increase dramatically. ISPs would then charge content providers more. In turn, this would reduce the investment incentives for content providers thus lowering content quality.

To sum up, both lines of reasoning argue that their preferred regime would generate higher overall welfare. Clearly, both sides cannot be right; a more detailed analysis is required to clarify the trade-off.

This paper explores how provider investments and revenues differ with network regime. We assume that the number of ISPs and content providers is fixed. That is, we do not

¹ "The Lessig Blog," <http://www.lessig.org/blog/archives/003391.shtml>

consider the longer term impact of regulations of the ISPs on provider entry incentives (that is, we assume the fixed structure of the network industry).

In section 2, we propose an economic model that relates the investments and prices to revenue for content providers and ISPs. We analyze two-sided pricing (non-neutral) in section 3.1 and we analyze one-sided pricing (neutral) in section 3.2. Section 4 is devoted to

Figure 2: The direction of payments in the model. The dotted lines indicate payments made only with two-sided pricing (“non-neutral”).

a comparison of the two regimes of the network. In section 5, we summarize our findings. The technical details of the analysis can be found in our working paper, see Musacchio *et al* (2009).

2 Model

Figure 2 illustrates our setting. In the model, there are M content providers and N ISPs. Each ISP T_n is attached to end-users U_n ($n = 1, 2, \dots, N$) and charges them p_n per click. The ISP T_n has a monopoly over its end-user base U_n . Thus, the end-users are divided between the ISPs, with each ISP having $1/N$ of the entire market. This assumption reflects the market power of local ISPs. Each ISP T_n charges each content provider C_m an amount equal to q_n per click. Content provider C_m invests c_m and ISP T_n invests t_n .

Recall from the Introduction that we measure q net of content providers’ access payment (for network attachment), which is set at marginal cost due to the ISPs’ competition. Accordingly, we measure the content provider per-user charges to advertisers (which we denote as a) net of the content provider’s access payment.

We characterize usage of end-users U_n by the number of “clicks” B_n they make. Since Internet advertising is most often priced per click, clicks are a natural metric for expressing advertising revenue. It is a less natural metric for expressing ISP revenue from end-users, because ISPs do not charge users per click but rather base their charges on bits. However, it is convenient to use only one metric and argue that one could approximate one metric from knowledge of the other using an appropriate scaling factor. The rate B_n of clicks of end-users U_n depends on the price p_n but also on the quality of the network, which in-turn is determined by provider investments, as follows

$$B_n = \left\{ \frac{1}{N^{1-w}} (c_1^v + \dots + c_M^v) \left[(1 - \rho) t_n^w + \frac{\rho}{N} (t_1^w + \dots + t_N^w) \right] \right\} e^{-p_n/\theta} \quad (1)$$

where $\rho \in (0,1)$, $\theta > 0$, and $v, w \geq 0$ with $v + w < 1$. For a given network quality (the expression in the curly brackets) the rate of clicks exponentially decreases with price p_n .

The term $c_1^v + \dots + c_M^v$ is the value of the content providers' investments as seen by a typical end-user. This expression is concave in the investments of the individual providers, and the interpretation is that each content provider adds value to the network. Also note that the structure of the expression is such that end-users value a network in which content is produced by numerous content providers higher than a network in which the content is provided by a single provider with the same cumulative investment. Our end-users' preference for content variety is similar to that of the classical monopolistic competition model by Dixit and Stiglitz (1977). The term in square brackets reflects the value of the ISP's investments for end-users. Clearly, users U_n value the investment made by their ISP, but they may also value the investments made by other ISPs. For instance, a user of one ISP might derive more value by having a better connection with users of another ISP. In our model, the parameter ρ captures this spill-over effect. When $\rho = 1$, end-users U_n value investments of all ISPs equally while when $\rho = 0$, they value only the investment of their ISP. When $\rho \in (0,1)$ end-users U_n value the investment of his ISP T_n more than investments of other ISPs $k \neq n$. The term ρ reflects the direct network effect of ISPs on end-users (not between them and content providers). This effect captures a typical network externality (see Thijssen, 2004 for a discussion of investment spill-over effects). The factor $1/N^{1-w}$ is a convenient normalization. It reflects the division of the end-user pool among N providers and it is justified as follows. Suppose there were no spill-over and each ISP were to invest t/N . The total rate of clicks should be independent of N . In our model, the total click rate is proportional to $(1/N^{1-w})(N(t/N)^w)$, which is indeed independent of N .

At this point, it is worth discussing why we have included certain features in the model. We chose to include N ISPs in the model, rather than just one, because there is a potential of free-riding in the two-sided pricing (non-neutral) case, which gives multiple ISPs the power to charge each content provider. This is because an ISP can increase his price to content providers and enjoy the additional revenue this increase causes, while the downside of inducing the content provider to invest less has to be borne by all of the ISPs. The magnitude of this effect will likely increase with N . We have also allowed for an arbitrary number of content providers, and included the spill-over term ρ , to study whether these features have a strong effect.

Returning to the model specification, the rate R_{mn} of clicks from end-users U_n to C_m is given by

$$R_{mn} = \frac{c_m^v}{c_1^v + \dots + c_M^v} B_n. \quad (2)$$

Thus, the total rate of clicks for content provider C_m is given by

$$D_m = \sum_n R_{mn}. \quad (3)$$

For further intuition behind the expressions (1) and (3) for click rates, consider a symmetric case, with all providers investing equally ($t_n = t/N$ and $c_m = c/M$) and $\rho = 1$. Then, (1) and (3) reduce to

$$B_n(\text{symmetric case}) = \left\{ \frac{1}{N} M^{1-v} c^v t^w \right\} e^{-p/\theta}, \quad D_m(\text{symmetric case}) = \left\{ \frac{1}{M^v} c^v t^w \right\} e^{-p/\theta}$$

respectively. Thus for both provider types, the click rates exponentially decrease with end-user price, increase with provider investments c and t , and exhibit decreasing return to investments. The only difference in the content provider and ISP click rate expressions are the effects of M and N . Recall that we included a normalization factor in (1) so that the same total investment split evenly across all ISPs should result in the same total click rate, regardless of N . On the other hand, we assumed that a user values variety in content, so that the same total investment split amongst multiple content providers is more valued than if it were split amongst fewer content providers. That is why B_n increases in M for fixed c , but D_m does not increase with N .

Leaving the symmetric case behind and returning to the general model, we assume that content providers charge a fixed amount a per click to the advertisers. Each content provider's objective is to maximize its profit which is equal to revenues from end-user clicks net of investment costs. Thus

$$\Pi_{Cm} = \sum_{n=1}^N (a - q_n) R_{mn} - \beta c_m \quad (4)$$

where the term $\beta > 1$ is the outside option (alternative use of funds c_m).

ISP T_n profit is

$$\Pi_{Tn} = (p_n + q_n) B_n - \alpha t_n \quad (5)$$

where $\alpha > 1$ is the outside option of the ISP. We assume providers of each type are identical and we will focus on finding symmetric equilibria for both one- and two-sided pricing.

3 The analysis of one- and two-sided pricing

To compare one-sided and two-sided pricing (neutral and non-neutral networks), we make the following assumptions.

- One-sided pricing (neutral network): In stage 1, each T_n simultaneously chooses (t_n, p_n) . The price q_n charged to content providers is constrained to be 0. In stage 2, each C_m chooses c_m .
- Two-sided pricing (non-neutral network): In stage 1, each T_n simultaneously chooses (t_n, p_n, q_n) . In stage 2, each C_m chooses c_m .

As we discussed in the Introduction, in one-sided pricing, the price q_n charged to content providers is zero by the following argument. Since content providers need only pay the ISP with which they are directly connected, the resulting price competition forces prices down to ISP costs. Thus the prices $\{q_n\}$ are zero after one normalizes them to be net of ISP cost. In contrast, in a network with two-sided pricing (non-neutral), all ISPs have the ability to charge each content provider. This permits ISPs to exercise their market power on content providers and allows them to extract a non-zero price.

In both cases, we assume that content providers observe ISP investments, and can subsequently adjust their investments based on the ISPs' choices. We justify this assumption by the difference in time and scale of the required initial investments. The investments of ISP tend to be longer-term investments in infrastructure, such as deploying networks of fibre-optic cable. Conversely, the investments of content-providers tend to be shorter-term and more ongoing in nature, such as development of content, making ongoing improvements to a search algorithm, or adding/replacing servers in a server farm.

3.1 Two-sided pricing

In a network with two-sided pricing (non-neutral network), each ISP chooses (t_n, p_n, q_n) and each content provider chooses c_m . To analyze this situation, we study how C_m chooses the optimal c_m for a given set of (t_n, p_n, q_n) . We then substitute that value of c_m in the expression for Π_T and we optimize for (t_n, p_n, q_n) .

The best choice for c_m given (t, p, q) maximizes

$$\begin{aligned}\Pi_{Cm} &= aD_m - \sum_n q_n R_{mn} - \beta c_m \\ &= N^{w-1} c_m^v \left[\sum_n (a - q_n) \left((1 - \rho) t_n^w + \frac{\rho}{N} (t_1^w + \dots + t_N^w) \right) e^{-p_n/\theta} \right] - \beta c_m.\end{aligned}\quad (6)$$

As a result of the cancelation between the denominator of R_{mn} (see (2)) and the expression for the click rate B_n (see (1)), the dependency of content provider m 's revenue on the investments of other content providers disappears. Thus the expression for profit Π_{Cm} is independent of other content provider investments c_j , $j \neq m$. Therefore, each content provider's optimization is unaffected by the simultaneously made (but correctly anticipated in equilibrium) investment decisions of the other content providers. Assuming that the term in square brackets is positive, we find that

$$\beta c_m^{1-v} = v N^{w-1} \left[\sum_k (a - q_k) \left((1 - \rho) t_k^w + \frac{\rho}{N} (t_1^w + \dots + t_N^w) \right) e^{-p_k/\theta} \right] =: \beta c^{1-v}. \quad (7)$$

For that value of c_m , we find that

$$\Pi_{Tn} = MN^{w-1}(q_n + p_n)F_n e^{-p_n/\theta} \left(\frac{v}{N\beta}\right)^{v/(1-v)} \left[\sum_k (a - q_k) e^{-p_k/\theta} F_k \right]^{v/(1-v)} - \alpha t_n \quad (8)$$

where

$$F_n = (1 - \rho)t_n^w + \frac{\rho}{N}(t_1^w + \dots + t_N^w) = \phi t_n^w + \frac{\rho}{N} \sum_{k \neq n} t_k^w \quad (9)$$

with

$$\phi = 1 - \rho + \frac{\rho}{N} < 1, \text{ if } N > 1. \quad (10)$$

The ISP T_n chooses investment and prices (t_n, p_n, q_n) that maximize his profit given by equation (8). The simultaneous decisions of each of the ISPs affect each other, therefore in order to find a Nash equilibrium we need to identify a point where the best response functions intersect. Writing that the three corresponding partial derivatives of (8) are equal to zero, and then finding the symmetric intersection point of the best response functions, we find the following solution (see our working paper, Musacchio et al. (2008)):

$$p_n = p = \theta - a; \quad (11)$$

$$q_n = q = a - \theta \frac{v}{N(1-v) + v}; \quad (12)$$

$$t_n = t \text{ with } (Nt)^{1-v-w} = x^{1-v} y^v e^{-(\theta-a)/\theta}; \quad (13)$$

$$c_m = c \text{ with } c^{1-v-w} = x^w y^{1-w} e^{-(\theta-a)/\theta}; \quad (14)$$

$$\Pi_{cm}^{1-v-w} = \Pi_c^{1-v-w} = \left(\frac{\theta v(1-v)}{N(1-v) + v} \right)^{1-v-w} x^w y^v e^{-(\theta-a)/\theta}; \quad (15)$$

$$\Pi_{Tn}^{1-v-w} = \Pi_T^{1-v-w} = \left(\frac{M\theta(N(1-v) - wN\phi(1-v) - vw)}{N(N(1-v) + v)} \right)^{1-v-w} x^w y^v e^{-(\theta-a)/\theta}; \quad (16)$$

$$\frac{\Pi_c}{c} = \frac{\beta(1-v)}{v}; \quad (17)$$

$$\frac{\Pi_T}{t} = \frac{\alpha}{w} \left[\frac{N(1-v)}{N\phi(1-v) + v} - w \right]; \quad (18)$$

$$B^{1-v-w} = M^{1-v-w} x^w y^v e^{-(\theta-a)/\theta} \quad (19)$$

where $B = \sum_n B_n = \sum_m D_m$ is the total click rate and

$$x = \frac{M\theta w N\phi(1-v) + v}{\alpha} \text{ and } y = \frac{\theta}{\beta} \frac{v^2}{N(1-v) + v}. \quad (20)$$

3.2 One-sided pricing

A network with one-sided pricing (neutral network) is similar to one with two-sided pricing (non-neutral), except that $q_n = 0$ as we argued in section 2 for $n = 1, \dots, N$. The best choice of c given $\{q_n = 0, p_n, t_n\}$ is such that

$$\beta c_m^{1-v} = vN^{-1} \left[\sum_k a((1-\rho)t_k^w + \frac{\rho}{N}(t_1^w + \dots + t_N^w))e^{-p_k/\theta} \right] =: \beta c^{1-v}.$$

For that value of c_m , we find that

$$\Pi_{Tn} = MN^{-1}p_n F_n e^{-p_n/\theta} \left(\frac{v}{\beta} \right)^{v/(1-v)} \left[\sum_k a e^{-p_k/\theta} F_k \right]^{v/(1-v)} - \alpha t_n \quad (21)$$

where

$$F_n = \phi t_n^w + \frac{\rho}{N} \sum_{k \neq n} t_k^w.$$

The ISP T_n chooses investment and price (t_n, p_n) that maximize the above expression. We find a symmetric Nash equilibrium by writing that the two corresponding partial derivatives of (21) with respect to a single ISP actions are zero, and that the other ISPs make the same actions, and solving all of the resulting equations. This analysis leads to the following solutions (see our working paper, Musacchio et al. (2008)):

$$p_n = p_0 := \frac{\theta N(1-v)}{N(1-v) + v}; \quad (22)$$

$$q_m = 0; \quad (23)$$

$$t_n = t_0 \text{ where } (Nt_0)^{1-v-w} = x^{1-v} y_0^v e^{-p_0/\theta}; \quad (24)$$

$$c_m = c_0 \text{ where } c_0^{1-v-w} = x^w y_0^{1-w} e^{-p_0/\theta}; \quad (25)$$

$$\Pi_{cm}^{1-v-w} = \Pi_{c0}^{1-v-w} := (a(1-v))^{1-v-w} x^w y_0^v e^{-p_0/\theta}; \quad (26)$$

$$\Pi_{Tn}^{1-v-w} = \Pi_{T0}^{1-v-w} := \left(\frac{M\theta(N(1-v) - wN\phi(1-v) - wv)}{N(N(1-v) + v)} \right)^{1-v-w} x^w y_0^v e^{-p_0/\theta}; \quad (27)$$

$$\frac{\Pi_{c0}}{c_0} = \frac{\beta(1-v)}{v}; \quad (28)$$

$$\frac{\Pi_{T0}}{t_0} = \frac{\alpha}{w} \left[\frac{N(1-v)}{N\phi(1-v) + v} - w \right]; \quad (29)$$

$$B_0^{1-v-w} = M^{1-v-w} x^w y_0^v e^{-p_0/\theta} \quad (30)$$

where B_0 is the total click rate, x is given in (20), and

$$y_0 := \frac{av}{\beta}. \quad (31)$$

4 Comparison

In this section, we compare the Nash equilibria of the two regimes. In section 4.1, we derive expressions for the welfare of end-users, and the ratio of social welfare with one- vs. two-sided pricing. In section 4.2, we demonstrate that the return on investments is the same in both regimes. In section 4.3, we compare the profits and social welfare of the two regimes for a range of parameters.

4.1 User welfare and social welfare

Before proceeding we define the following notation.

$$\pi := \frac{v}{N(1-v) + v} \quad \text{and} \quad \delta := \frac{a}{\theta}. \quad (32)$$

In order to compute end-user welfare, we use the total click rate to proxy the aggregate user demand. This enables us to calculate consumer surplus and use it to measure end-user welfare. We compute the consumer surplus by taking the integral of the demand function from the equilibrium price to infinity. This integral is taken with the investment levels of content providers and ISPs fixed. We find

$$W_U(\text{two-sided}) = M\theta x^{w/(1-v-w)} y^{v/(1-v-w)} e^{-\frac{\theta-a}{\theta(1-v-w)}}.$$

The expression for the one-sided case is the same, but with y exchanged for y_0 and $\theta - a$ in the exponent exchanged with p_0 . The ratio of the social welfare with one- vs. two-sided pricing has the form

$$\begin{aligned} & \frac{W_U(\text{one-sided}) + N\Pi_T(\text{one-sided}) + M\Pi_C(\text{one-sided})}{W_U(\text{two-sided}) + N\Pi_T(\text{two-sided}) + M\Pi_C(\text{two-sided})} \\ &= \frac{1 + \delta(1-v) + (\pi/v)(N(1-v) - wN\phi(1-v) - wv)}{1 + \pi(1-v) + (\pi/v)(N(1-v) - wN\phi(1-v) - wv)} [(\delta/\pi)^v e^{\pi-\delta}]^{1/(1-w-v)}. \end{aligned}$$

4.2 Return on investment

We use the definition $\delta = a/\theta$ and equations (11), (12), (22), and (23) to relate prices in the one- and two-sided pricing regimes. This permits us to formulate the following proposition.

Proposition 1: The price p end-users pay in the two-sided pricing case is given by

$$p = \theta(1 - \delta) = \theta - a.$$

Also, we note that

$$p + q = p_0 = \theta(1 - \pi).$$

Moreover,

$$\frac{\Pi_C}{c} = \frac{\Pi_{C0}}{c_0} \quad \text{and} \quad \frac{\Pi_T}{t} = \frac{\Pi_{T0}}{t_0}.$$

Thus, from Proposition 1, the total revenue per click of the ISPs is the same in both regimes and so are the rate of return on investments of the content and ISPs.

Despite the fact that the rate of return on investments are the same in both regimes, as we will see in the next subsection, the size of those investments and resulting profits may be quite different across regimes.

4.3 Comparative statics

Dividing the expressions for one-sided pricing by the corresponding expressions for the two-sided pricing case, we define ratios such as

$$r(\Pi_C) := \left(\frac{\Pi_C(\text{one-sided})}{\Pi_C(\text{two-sided})} \right)^{1-v-w}$$

where $\Pi_C(\text{two-sided})$ is the profit per content provider in the two-sided case as expressed in (15) and $\Pi_C(\text{one-sided})$ is the profit per content provider with one-sided pricing (26). We define $r(c)$, $r(t)$, and $r(\Pi_T)$ similarly. We find

$$r(\Pi_T) = r(t) = r(B) = \left(\frac{\delta}{\pi} \right)^v e^{\pi-\delta}; \quad (33)$$

$$r(\Pi_C) = r(c) = \left(\frac{\delta}{\pi} \right)^{1-w} e^{\pi-\delta}. \quad (34)$$

Figure 3 shows the ratios of profits with one- vs. two-sided pricing for both content providers and ISPs. Figure 4 shows the same ratios, but for different values of v and w . The figures show that for small or large values of a/θ , the ratio of advertising revenue per click to the constant characterizing price sensitivity of end-users, two-sided pricing is preferable to both content providers and ISPs. (Here we say “preferable” in that the profits are larger, though we have seen that the rate of return on investments are the same.) For mid range values of a/θ , one-sided pricing is preferable to both, though the values of a/θ where the transition between one-sided being preferable to two-sided are not exactly the same for content providers and ISPs. Furthermore, as N , the number of ISPs increases, the range of a/θ values for which one-sided pricing is superior increases, while also the degree by which it is superior (in terms of revenues to content providers and ISPs) increases.

These results can be explained by the following reasoning. When a/θ is large, the content providers’ revenues from advertising are relatively high, and the ISPs’ revenue from end-users are relatively low. Because of this, the ISPs incentives to invest are suboptimal (too low relative to the socially optimal ones), unless they can extract some of the content providers’ advertising revenue by charging the content providers. Thus in the one-sided pricing case, the ISPs under invest, making the rewards for them as well as content providers less than it could have been with two-sided pricing.

It is important to note that when a/θ is larger than 1, the price p charged end-users becomes negative in the two-sided case as can be seen from (11). If end-users were actually paid to click, intuition suggests that they would click an unbounded amount and therefore our exponential model of demand (1) would not be valid in this region of a/θ . However, one could interpret price p to be net any variable costs v to end users – similar to how we

define q .² With this interpretation, the price p could be negative while the actual prices users see positive, so long as $|p| < v$. We therefore show numerical results in our plots for a/θ as large as 1.2.

Figure 3: The ratios of profits ($v = 0.5, w = 0.25$) for different values of N .

Figure 4: The ratios of profits ($v = 0.25, w = 0.5$) for different values of N .

When a/θ is very small, the content providers' advertising revenue is relatively low, and the ISP's end-user revenue is relatively high. In order to get the content providers to invest adequately, the ISPs need to pay the content providers. That is why for small enough a/θ the price q is negative (see Figure 5), representing a per click payment from the ISPs to the content providers.

Interestingly, even in the two-sided pricing case, our content providers obviously get some share of the surplus generated jointly by them and the ISPs. This is in contrast to the multi-homing case of Armstrong (2006) (which is roughly analogous to our two-sided pricing case), where the surplus is fully extracted from group-2 agents (content providers in our case) – the entire surplus is shared between the platform and the end-users. We do not have this extreme result, because in our model content providers invest *after* ISPs announce

² If we suppose such a normalization, it may be more natural to express end-user demand in terms of the price $p+v$ actually seen by users. However, supposing demand was exponential in $p+v$ rather than exponential in p would only change expression (1) for click rate by a constant factor. This constant factor disappears in any expression comparing the relative revenues or welfares of the two regimes.

Figure 5: The price q (that the ISPs charge the content providers).

prices. Thus, ISP commitment to the declared prices permits content providers to retain a positive fraction of the surplus.

Finally, when a/θ is in the intermediate range, in between the two extremes, both content providers and ISPs have adequate incentive to invest. However, another effect comes into play – ISP free riding becomes an important factor when N is large. As N increases in the two-sided pricing case there are more ISPs that levy a charge against each content provider. As the price ISPs charge content providers increases, it becomes less attractive for content providers to invest. Thus, an ISP choosing the price to charge content providers is balancing the positive effect of earning more revenue per click from content providers versus the negative effect of having fewer clicks because the content providers have reduced their investment. But each ISP sees the entire gain of raising its price, but the loss is borne by all N ISPs. Consequently, the ISPs overcharge the content providers in Nash equilibrium, and the degree of this overcharging increases with N . This is analogous to the tragedy of the commons where people overexploit a public resource. Another perhaps more direct analogy is the “castles on the Rhine effect” where each castle owner is incentivized to increase transit tolls on passing traffic excessively by ignoring the fact that the resulting reduction in traffic harms not only him, but also other castle owners. When all castles do the same, the traffic on the Rhine decreases (Kay, 1990). The extent of this negative externality, and hence the degree of over-charging, increases with N .

Figure 6 shows a three dimensional plot of the ratio of social welfare for one- vs. two-sided pricing. The plot shows how the ratio changes for different N and a/θ . The second panel of Figure 6 is a simplified version of the first panel. It depicts only the boundaries in the parameter space where one-sided pricing is preferable to two-sided and vice versa.

It is also worthwhile to note that spill-over parameter ρ and the number of content providers M do not appear in the expression for the ratio of content provider revenue between the two regimes (34) nor do they appear in the ratio of revenues for ISPs (33). This is in spite of the fact that ρ and M do appear in the expressions for both the one-sided and two-sided pricing equilibria. This suggests that the spill-over effect and number of content providers have little or no effect on the comparative welfare of the two regimes.

Figure 6: Left: The log of social welfare ratio (one-sided to two-sided). Right: Regions of social welfare superiority (one-sided vs. two-sided).

5 Conclusions

Our analysis seeks to address whether content providers should pay ISP(s) only for the right to access the network at their end, or whether other ISPs (that is, the ISPs other than the one through which a specific content provider is attached to the network) may also charge for the “right” to reach end-users. We call the former “one-sided pricing”, and the latter “two-sided pricing”. We suggest that the imposition of regulation that precludes the ISPs from two-sided pricing corresponds to a “neutral” network, while the practice of two-sided pricing corresponds to a “non-neutral” network.

We study how each pricing regime affects investment incentives of transit and content providers. We show that parameters such as advertising rate, end-user price sensitivity, and the number of ISPs influence whether one- or two-sided pricing achieves a higher social welfare. From our results, when the ratio of advertising rates to the constant characterizing price sensitivity is an extreme value, either large or small, two-sided pricing is preferable. If the ratio of advertising rates to the constant characterizing price sensitivity is not extreme, then an effect like the “castles on the Rhine effect” becomes more important. ISPs in a two-sided pricing regime have the potential to over charge content providers, and this effect becomes stronger as the number of ISPs increases.

In our comparison of one- and two-sided pricing we assume that the ISPs choose their strategy first and that the content providers follow. We justify this assumption by the difference in the time and scale of the required initial investments. While we feel this is a valid assumption, it is worth noting that if we had modelled the game such that the content providers and ISPs reverse their order of play, the game would not have a meaningful equilibrium. (In such a reversed model with two-sided pricing, the ISPs would charge an unbounded price to content providers after they had committed to their content investments. This is because at this point the content providers would have no action left to take in the game. This of course would lead to content providers investing zero in the first stage of the game.)

We make other key modeling choices, such as that the ISPs have local monopolies over their users. We believe that if we had studied a model where each ISP is a duopolist, which better models the degree of choice most end-users have today, our results would have been qualitatively similar. However, there are a number of competing effects such a model may

introduce. First, such a scenario would reduce the market power of ISPs over end-users, thus reducing the revenue the ISPs could extract from them. Since two-sided pricing provides ISPs with another source of revenue to justify their investments, this effect would tend to increase the parameter region for which two-sided pricing is social welfare superior. Second, a duopolist competing on product quality invests more than a monopolist, so this would tend to increase efficiency of one-sided pricing. Third, if the model were changed from having N to $2N$ ISPs, then the free riding or “castles on the Rhine” effect would grow, tending to reduce the welfare of the two-sided pricing case. The net effect of all these individual effects would of course depend on the detailed specifications of such a model.

Another assumption we make is that all ISPs have identical payoff functions. This would not be the case if the ISPs served different population sizes or different population densities for instance. Our assumption makes the analysis much easier, but we feel that if this assumption were relaxed, the results would be qualitatively similar though closed form solutions might not exist. We feel the same way about our assumption that content providers have identical payoff functions.

Our two-sided pricing model assumes that in-bound traffic to a local ISP could be identified as originating at a particular content provider, in order for the ISP to levy the appropriate charge to the content provider. This assumption would not strictly hold if content providers had some way of reaching end-users of an ISP without paying this ISP for end-user traffic. For instance, if there were a second ISP that enjoyed settlement free peering with the first ISP, the content provider could route its traffic through the second ISP and thus avoid the first ISP charge for end-user access. This strategy may be facilitated by the fact that the end-users of both ISPs send traffic to each other, and perhaps the traffic from the content provider could be masked in some way to look like traffic originating from the second ISP's end-users. However, the communication protocols of the Internet require that packets be labeled with the origin (IP) address. It seems unlikely today that a large content provider could have the origin addresses of its traffic falsified in a way that would both prevent ISPs from being able to charge the content provider while still enabling end-users to send traffic in the reverse direction back to the content provider. However, it is certainly possible that technology would be developed to enable such a strategy in the future, especially if there were an economic incentive for developing it.

There are several assumptions we make that might be viewed as limitations of our model. First, we fix the number of network providers, independent of the network regime. This assumption is realistic in the short-run, but in the long-run the number of providers entering the network industry is likely to differ with regime. Second, we do not consider heterogeneity in the providers nor in the end-users. Third, we assume full commitment to the declared prices, that is the ISPs cannot later adjust the prices that they declared initially. We also have not modeled the price content providers charge advertisers as a decision variable, but we have modeled the price ISPs charge end-users. Though we feel that the present model correctly captures focal features of the underlying network environment, in future work we plan to extend our model to permit endogenous industry structure, that is, we will study how network regime affects the number of content providers and ISPs.

6 References

- ACLU (2006) "Internet Freedom and Innovation at Risk: Why Congress Must Restore Strong Net Neutrality Protection," ACLU Press Release, September 9 2006. <http://www.aclu.org/freespeech/internet/26829res20060922.html>
- Armstrong, M. (2006) "Competition in Two Sided Markets," *RAND Journal of Economics*, 37: 668-691.
- Chong, R. (2007) "The 31 Flavors of the Net Neutrality Debate: Beware The Trojan Horse," *Advanced Communications Law and Policy Institute, Scholarship Series*, New York Law School, December 2007.
- Dixit, A. and J. Stiglitz (1977) "Monopolistic Competition and Optimum Product Diversity," *American Economic Review*, 67: 297-308.
- Department of Justice (2007) "Comments on Network Neutrality in Federal Communications Commission Proceeding," Press Release, September 7, 2007. http://www.usdoj.gov/opa/pr/2007/September/07_at_682.html
- Farber, D. and M. Katz (2007) "Hold Off On Net Neutrality," *Washington Post*, Jan. 22.
- Felten, E. (2006) "The Nuts and Bolts of Network Neutrality," Center for Information Technology Policy, Princeton, July 6, 2006. <http://itpolicy.princeton.edu/pub/neutrality.pdf>
- Federal Trade Commission (2007) "Broadband Connectivity Competition Policy," Report, June 2007.
- Hermalin, B. and M. Katz (2007) "The Economics of Product Line Restrictions With an Application to the Network Neutrality Controversy," *Information Economics and Policy*, 19: 215-48.
- Hogendorn, C. (2007) "Broadband Internet: Net Neutrality Versus Open Access," *International Economics and Economic Policy*, 4: 185-208.
- Kay, J. A. (1990) "Tax Policy: A Survey," *Economic Journal*, 100: 18-75.
- Musacchio, J., G. Schwartz and J. Walrand (2009) "A Two-Sided Market Analysis of Provider Investment Incentives With an Application to the Net-Neutrality Issue," Technical Report UCSC-SOE-09-08, School of Engineering, University of California, Santa Cruz, March 2009. <http://www.soe.ucsc.edu/johnm/Publications/neutrality.pdf>
- Odlyzko, A. (2008) "Network neutrality, Search Neutrality, and the Never-Ending Conflict Between Efficiency and Fairness in Markets," *Review of Network Economics*, 8: 40-60.
- Owen, B. and G. Rosston (2003) "Local Broadband Access: *Primum Non Nocere* or *Primum Processi*? A Property Rights Approach," Stanford Law and Economics Olin Working Paper No. 263.
- Rochet, J.-C. and J. Tirole (2006) "Two-Sided Markets: A Progress Report," *RAND Journal of Economics*, 37: 645-667.
- Sidak, G. (2006) "A Consumer-Welfare Approach to Network Neutrality Regulation of the Internet," *Journal of Competition Law & Economics*, 2: 349-474.

Thijssen, J. (2004) *Investment Under Uncertainty, Coalition Spillovers and Market Evolution in a Game Theoretic Perspective*. Series: Theory and Decision Library C, 35. Springer: New York.

Weiser, P. (2007) "Report from the Center for the New West Putting Network Neutrality in Perspective," Center for the New West Discussion Paper, January 11, 2007. <http://www.centerfornewwest.org/pdf/TelecomSummary.pdf>

Whitacre, E. (2005) Interview by R. O. Crockett "At SBC, It's All About 'Scale and Scope'," *Business Week*, November 7.

Yoo, C. S. (2006) "Network Neutrality and the Economics of Congestion," *Georgetown Law Journal*, 94: 1847-1908.