POMS PRODUCTION AND OPERATIONS MANAGEMENT

PRODUCTION AND OPERATIONS MANAGEMENT POMS

Vol. 22, No. 5, September–October 2013, pp. 1287–1298
ISSN 1059-1478 | EISSN 1937-5956 | 13 | 2205 | 1287

DOI 10.1111/poms.12020

Broadband Network Management and the Net Neutrality Debate

Hong Guo
Department of Management, Mendoza College of Business, University of Notre Dame, Notre Dame, Indiana 46556-5646, USA
hguo@nd.edu

Hsing Kenneth Cheng, Subhajyoti Bandyopadhyay
Department of Information Systems and Operations Management, Warrington College of Business Administration, University of Florida, Gainesville, Florida 32611-7169, USA
kenny.cheng@warrington.ufl.edu, shubho.bandyopadhyay@warrington.ufl.edu

The debate of net neutrality and the potential regulation of net neutrality may fundamentally change the dynamics of data consumption and transmission through the Internet. The existing literature on economics of net neutrality focuses only on the supply side of the market, that is, a broadband service provider (BSP) may charge content providers for priority delivery of their content to consumers. In this article, we explore a complete spectrum of broadband network management options based on both the supply and demand sides of the market. We find that although the BSP always prefers the non-neutral network management options, it does not always discriminate both sides of the market. From the social planner's perspective, we find that some network management options maximize the social welfare under certain market conditions while other options reduce the social welfare. Using the terminology from a recent Federal Communications Commission report and order, we categorize the social welfare maximizing options as "reasonable network management" and the social welfare reducing options as "unreasonable discrimination." We also identify conditions under which the BSP's network management choices deviate from the social optimum. These conditions help establish the criteria under which the social planner might wish to regulate the BSP's actions.

Key words: net neutrality; broadband network management; traffic prioritization, public policy
History: Received: June 2011; Accepted: September 2012 by Cheryl Gaimon, after 2 revisions.

1. Introduction

The issue of net neutrality first received widespread media attention when some broadband service providers (BSPs) like Verizon, Comcast, and AT&T (among others) proposed to charge popular online websites for priority delivery of the latter's content to their residential and commercial customers (Helm 2006, Waldmeir 2006). The proposal encountered stiff resistance from those who were supposed to be charged, and thus erstwhile competitors like Google, Yahoo!, and Microsoft were soon lobbying before the US Congress to pass legislation that would prevent the broadband service providers from carrying out their proposed plan (Newmark and McCurry 2006), and thereby maintain what was termed the "neutrality" of the Internet.

Essentially, net neutrality is a broadband network management principle. The core issue of net neutrality is as follows: the original design of the Internet communication protocols abides by the principle that every data packet is treated equally, so that no data packet gets priority over another (leading to the coinage of the term net neutrality—i.e., the net is neutral in its handling of any data packet that passes through it). Thus, any network management practice that treats different data packets differently would be considered a violation of the net neutrality principle. In the spectrum of the net neutrality debate (Turner 2010), the violation of the net neutrality principle can occur by providing a preferential treatment of one content provider's packets over another's or by slowing down a group of users' traffic. Some BSPs have already been experimenting with the latter form of deviation from net neutrality. For example, Comcast, the largest provider for broadband Internet access service, was found to have slowed down network traffic that originated from the popular peer-to-peer (P2P) networks in 2007 (McCullagh 2007). After initially denying any such behavior, Comcast defended its actions by claiming that the traffic from the P2P networks, which was dominated by just a small fraction of the total number of users, was slowing down the network traffic for the rest of the users. The US Federal Communications Commission (FCC) later declared Comcast's actions to be illegal, but on April

6, 2010, the District of Columbia court of appeals ruled that the FCC had no authority to regulate how the BSPs manage traffic to their consumers. The FCC in turn responded a few months later on December 23, 2010, arguing that it indeed had the authority to "adopt open Internet rules"—thus providing further fuel to the net neutrality debate that is currently making the rounds in the US Congress.

Note that the net neutrality principle does not preclude product differentiations—in other words, charging two users different prices for two different bandwidths (which in turn determines the speed of accessing content) is not a violation of the net neutrality principle. However, if that pricing scheme is accompanied by a technological mechanism whereby the packets received by one user gets *priority* over the packets received by another user, such a situation would amount to the violation of the net neutrality principle. If two users do not share a common bandwidth resource, so that giving a higher quality of service to one user does not degrade the experience of the other user, then it is a case of product differentiation and not a violation of the net neutrality principle. If, on the other hand, the users share the same common bandwidth resource, so that giving a higher quality of service to one user degrades the experience of the other user, it is a violation of the net neutrality principle. In general, there are many ways by which a BSP can deviate from net neutrality, and we will discuss several such deviations in section 2.4—"Network Management Options."

So far, the growing literature on economics of net neutrality (see, e.g., Cheng et al. 2011, Choi and Kim 2010, Economides and Hermalin 2013, Economides and Tåg 2012, Guo et al. 2010, 2012, Hermalin and Katz 2007) has modeled the net neutrality debate from the supply side, whereby the BSP charges content providers for preferential delivery of their packets. However, as the aforementioned Comcast example shows, the BSP might also be interested in charging some of the consumers themselves who are willing to pay a fee to have their requested packets delivered with priority (or the BSP might de-prioritize the requested content in the absence of the priority fee). In other words, a data packet traveling from its origin to its destination can be made "non-neutral" by the BSP at various stages of its journey—either on the supply side or on the demand side, or both. Figure 1 shows a schematic of the different aspects of the net neutrality debate. It clearly brings out the role of the BSP as the gatekeeper who can manage its network traffic by prioritizing data packets based on either the content providers (at the left of the figure) or the consumers (at the right of the figure) or both and charge the content providers and/or the consumers accordingly.

Figure 1 Schematic Model

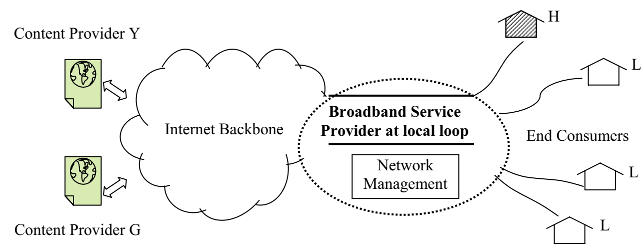

In this article, we propose a comprehensive framework of BSPs' network management options by exploring both the supply and demand sides of the market simultaneously. As the BSP has the freedom to make a packet non-neutral at either ends of the market, we consider various network management options for the BSP that moderate the traffic on both sides of the market. Specifically, by considering the supply side with competing content providers with different capabilities of generating advertising revenues from their consumers as well as the demand side of different classes of users with dynamic and differentiated data usage patterns, we formulate the BSP's decision problem of managing its network traffic to maximize profits.

This general analysis of network management with both content provider discrimination and user discrimination yields more refined results compared to the existing literature. We find that although the BSP always prefers the non-neutral network management options (which is consistent with the literature), it does not always discriminate both sides of the market. The literature on economics of net neutrality finds mixed results regarding the impact of the potential net neutrality regulation on social welfare. Some studies (Economides and Hermalin 2013, Economides and Tåg 2012) conclude that social welfare is higher under net neutrality while others (Cheng et al. 2011, Choi and Kim 2010, Hermalin and Katz 2007) find that content provider discrimination is likely to be welfare enhancing. We find that different network management options would have different impacts on social welfare. Depending on market conditions of both the supply side and the demand side, some deviations from net neutrality may result in higher social welfare. Looking at just the supply side gives an incomplete picture, especially when BSPs in real life have started experimenting with network management on both sides. Our research is therefore more reflective of the reality and shows the complicated and nuanced nature of the problem.

Although the BSP might prefer a certain network management option under certain market conditions, such a choice might be detrimental to the consumers or the society as a whole. Thus, from the social planner's perspective, we further assess each available network management option. We find that some network management options maximize the social welfare under certain market conditions while other options reduce the social welfare. Using the terminology from the FCC report and order (FCC 2010a), we then categorize the social welfare maximizing options as "reasonable network management" and the social welfare reducing options as "unreasonable discrimination." We also identify conditions under which the BSP's network management choices deviate from the social optimum. These conditions help establish the criteria under which the social planner might wish to regulate the BSP's actions to maximize the social welfare.

The remainder of the article is organized as follows: In section 2, we set up a stylized game-theoretical model that captures the incentives of the various players in a two-sided market. We then propose a framework of the BSP's network management options and analyze these options. Section 3 identifies the BSP's preferred choices of network management under different market conditions. Section 4 evaluates the network management options from the social planner's perspective. We further categorize these network management options into reasonable network management and unreasonable discrimination. Finally, we conclude with a summary of our findings and some directions for future research in section 5.

2. The Model

In this section, we model the system of digital content provision, data transmission, and data consumption as a two-sided market. There are three types of players: content providers generate digital content; broadband service providers deliver the content from content providers to end consumers; and end consumers receive and consume digital content. Among these players, broadband service providers act as an intermediary serving the two sides of the market, that is, both the content providers (the "supply" side) and the consumers (the "demand" side). We discuss the details of these three players in the following subsections.

2.1. Broadband Service Provider

Following the existing literature on economics of net neutrality (Cheng et al. 2011, Choi and Kim 2010, Hermalin and Katz 2007), we consider a monopolist BSP who provides Internet access to end consumers and delivers digital content from websites to end consumers. The BSP serves both the supply and demand sides of the market—specifically, two content providers Y and G on the supply side as well as two types of end consumers H and L on the demand side (the modeling details of the content providers and the consumers are described in subsequent subsections). The BSP's capacity is fixed and denoted by μ. A summary of all the notations used in the article is provided in Table 1.

Under net neutrality, the BSP charges consumers a fixed fee F for Internet access service and does not charge content providers for content delivery. When net neutrality is not enforced, the BSP has various ways to deviate from net neutrality—discriminate

Table 1 List of Notations

Notation	Description
μ	Capacity of the BSP measured in packets per unit of time, which corresponds to the service rate of the server in a M/M/1 queueing system
F	A uniform fixed fee charged by the BSP to end consumers
F_H, F_L	Fixed fees charged to H-type and L-type consumers, respectively
p	Unit price per packet for preferential delivery of data packets
α	Percentage of H-type consumers
λ_H, λ_L	Rate of content requested from H-type and L-type consumers in packets per unit of time
V_H, V_L	The gross value of retrieving content for H-type and L-type consumers, respectively
t	Consumers' fit cost parameter
w_{HY}, w_{HG}, w_{LY}, w_{LG}	Expected delay per packet for consumer groups HY, HG, LY, and LG, respectively
d	Consumers' congestion cost parameter that converts the delay for consumers waiting for the content to arrive from the websites to a disutility due to congestion
U_{HY}, U_{HG}, U_{LY}, U_{LG}	Utility function for consumer groups HY, HG, LY, and LG, respectively
x_H, x_L	Marginal H-type or L-type consumer who is indifferent between content provider Y and content provider G0
π_{BSP}	BSP's profit
π_Y, π_G	Content providers' profit
r_Y, r_G	Revenue rate of content providers Y and G, respectively
I_Y, I_G	Indicator function that denotes whether content provider Y or G pays for the premium data delivery or not
q_1, q_2, q_3, q_4	Total traffic of priority classes 1, 2, 3, and 4, respectively
q_{HY}, q_{HG}, q_{LY}, q_{LG}	Total traffic of consumer groups HY, HG, LY, and LG, respectively
q_Y, q_G	Content providers' total traffic
SW	Social welfare

between different content providers, different types of consumers, or both. Figure 2 presents a framework of the network management options available to the BSP.

In the case of content provider discrimination, the BSP may provide preferential delivery for a particular provider if that provider pays a usage-based fee p. Based on the content providers' choices on whether to pay the preferential delivery fee or not, there are four possible outcomes. In the case of consumer discrimination, there are two ways that the BSP can manage traffic on the shared network resource for different types of consumers. As a result, the BSP has eight traffic management options considering both content provider discrimination and consumer discrimination (see Table 2). Note that some options involve content provider discrimination only; some options involve consumer discrimination only; other options involve discrimination on both sides. We will analyze these options in Subsection 2.4.

2.2. Consumers

We consider a unit mass of end consumers for broadband Internet access service. To capture the heterogeneous features of end consumers, we assume that there are two types of users: a fraction α of heavy (H) consumers and $1-\alpha$ fraction of light (L) consumers. Following prior literature (Cheng et al. 2011), consumers' requests for digital content (measured by data packets) follow a Poisson process. The requested rates of data packets by the two user types are given by λ_H and λ_L, respectively. We assume $\alpha < 1/2$, $\lambda_H > \lambda_L$, and $\alpha\lambda_H > (1-\alpha)\lambda_L$ to capture consumers' usage patterns.[1] In other words, there is a small proportion ($\alpha < 1/2$) of heavy users who request more data than light users ($\lambda_H > \lambda_L$). At the aggregate level, the small group of heavy users generates the majority of traffic on the network, that is, $\alpha\lambda_H > (1-\alpha)\lambda_L$. Finally, we assume that $\mu > \alpha\lambda_H + (1-\alpha)\lambda_L$ to ensure a stable queueing system. Consumers enjoy the received content with a gross valuation. We use V_H (or V_L) to denote heavy (or light) users' valuation for content.

Figure 2 Broadband Service Provider's Network Management Framework

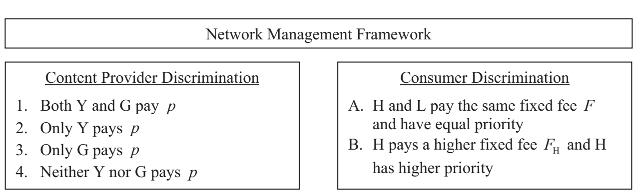

Combining the four outcomes for content provider discrimination and the two outcomes for consumer discrimination yields eight different traffic management options (see Table 2).

Considering the consumers' heterogeneous demand patterns, the BSP may charge a uniform fixed fee (F) per unit time to all consumers or different fixed fees (F_H and F_L) per unit time to different types of consumers for Internet access. We assume the consumers are differentiated based on their preference for content. Following the classic Hotelling model (Hotelling 1929), both H-type and L-type consumers are uniformly distributed along the line segment of [0, 1] based on their content preference (Figure 3). Consumers incur a linear fit cost (with unit fit cost parameter t), which increases in the distance between their ideal content and their preferred provider. Consumer x_H (or x_L) represents the marginal H-type (or L-type) consumer who is indifferent between content provider Y and content provider G. As a result, the consumer market will be divided into four consumer groups: the H-type consumers located to the left of x_H will choose content provider Y (we call this consumer group HY) with market share αx_H; the H-type consumers located to the right of x_H will choose content provider G (we call this consumer group HG) with market share $\alpha(1-x_H)$. Similarly, consumer group LY consists of consumers located to the left of x_L with market share $(1-\alpha)x_L$ and consumer group LG consists of consumers located to the right of x_L with market share $(1-\alpha)(1-x_L)$. The two indifferent consumers (x_H and x_L) realize the lowest net utilities

Table 2 Broadband Service Provider's Network Management Options

CPs \ Consumers		Consumer Discrimination	
		A. Same fixed fee; H and L have equal priority	B. H pays a higher fixed fee; H has higher priority
Content Provider Discrimination	1. neither pay	Outcome 1A (Net Neutrality)	Outcome 1B
	2. Y pays	Outcome 2A	Outcome 2B
	3. G pays	Outcome 3A	Outcome 3B
	4. both pay	Outcome 4A	Outcome 4B

Options shaded with horizontal lines involve consumer discrimination only; options shaded with vertical lines involve content provider discrimination only; options shaded with grid involve discrimination on both sides.

Figure 3 Content Providers and their Share of Consumers

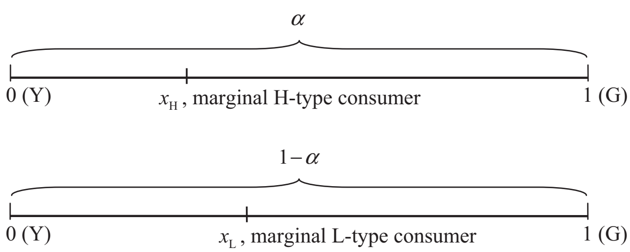

within their consumer type because they incur the highest fit costs and all consumers within each consumer group share the same congestion cost.

As the consumers are serviced by the BSP, which has a fixed network infrastructure capacity, congestion occurs in this shared network. Due to congestion, the consumers encounter a disutility while they wait for the requested data packets to be delivered. We use w_{ij} to denote the expected delay for a data packet from content provider j = Y or G to a type i = H or L consumer. Then, the aggregate delay for this type i consumer is $\lambda_i w_{ij}$ if he/she chooses content provider j. As a result, this type i consumer incurs a disutility of $d\lambda_i w_{ij}$, where d is the congestion cost parameter that captures the unit cost of delay for consumers waiting for the content to arrive from the websites.

For a type i = H or L consumer with content preference x, the utility functions of subscribing to content Y and G take the following form:

$$U_{iY}(x) = \begin{cases} V_i - tx - d\lambda_i w_{iY} - F, & \text{if the BSP charges a uniform fixed fee} \\ V_i - tx - d\lambda_i w_{iY} - F_i, & \text{if the BSP charges differential fixed fees} \end{cases}$$

$$U_{iG}(x) = \begin{cases} V_i - t(1-x) - d\lambda_i w_{iG} - F, & \text{if the BSP charges a uniform fixed fee} \\ V_i - t(1-x) - d\lambda_i w_{iG} - F_i, & \text{if the BSP charges differential fixed fees.} \end{cases}$$

Consumers request data from various websites and the requested data packets are transmitted through the BSP's network. As noted in the aforementioned literature, we assume an M/M/1 queue to model the data transmission service provided by the BSP under net neutrality. When net neutrality is enforced, that is, when no packet has priority over another, the time that a data packet spends in the system is $w_{ij} = \frac{1}{\mu - \alpha\lambda_H - (1-\alpha)\lambda_L}$ and the corresponding congestion cost is $d\lambda_i w_{ij} = \frac{d\lambda_i}{\mu - \alpha\lambda_H - (1-\alpha)\lambda_L}$, where i = H or L and j = Y or G.

In the absence of net neutrality, the BSP may prioritize data traffic based on content providers and/or user types. In this context, we note that the technology to discriminate packets and streamline Internet traffic has been available at minimal cost, and we therefore assume that there is no additional ongoing expense incurred by the BSP to implement a mechanism that enables preferential delivery of content. We use a priority queue to model the BSP's data transmission service with traffic prioritization. If all consumer groups receive the same priority for their traffic, then the congestion cost would remain the same as that under net neutrality. However, different traffic will be assigned different priorities under various network management options. We use priority classes to describe traffic with different priorities. Let q_i be the total traffic of priority class i. Priority classes are numbered in descending order, that is, traffic of class 1 has the highest priority. The expected delays for different priority classes are specified as follows (Hillier and Lieberman 1990, pp. 178–180). If there are two priority classes, then $w_1 = \frac{1}{\mu - q_1}$ and $w_2 = \frac{\mu}{(\mu - q_1)(\mu - q_1 - q_2)}$. If there are four priority classes, then $w_1 = \frac{1}{\mu - q_1}$, $w_2 = \frac{\mu}{(\mu - q_1)(\mu - q_1 - q_2)}$, $w_3 = \frac{\mu}{(\mu - q_1 - q_2)(\mu - q_1 - q_2 - q_3)}$, and $w_4 = \frac{\mu}{(\mu - q_1 - q_2 - q_3)(\mu - q_1 - q_2 - q_3 - q_4)}$. Different network management options involve different numbers of priority classes. All data packets receive the same priority in outcomes 1A/4A. There are two priority classes in outcomes 2A/3A/1B/4B and four priority classes in outcomes 2B/3B. Details of the waiting times can be found in Table 3.

2.3. Content Providers

We consider two horizontally differentiated content providers (Y and G) competing for end consumers. Content providers adopt an advertisement-assisted revenue model. Let q_{HY}, q_{HG}, q_{LY}, and q_{LG} be the total traffic of consumer groups HY, HG, LY, and LG, respectively. Then $q_{HY} = \alpha x_H \lambda_H$, $q_{HG} = \alpha(1 - x_H)\lambda_H$, $q_{LY} = (1-\alpha)x_L\lambda_L$, and $q_{LG} = (1-\alpha)(1-x_L)\lambda_L$. If we use q_Y (or q_G) to denote the total traffic for content provider Y (or G), then $q_Y = q_{HY} + q_{LY}$ and $q_G = q_{HG} + q_{LG}$. Thus, content providers' profits can be specified as $\pi_j = (r_j - I_j p)q_j$, j = Y or G, where r_j is the revenue rate of content provider j and indicator function I_j denotes whether content provider j pays for the premium data delivery or not. Without loss of generality, we assume $r_G \geq r_Y$ to capture the asymmetric nature of firms' efficiency. In other words, content provider G is either more efficient than or equally efficient as content provider Y in terms of generating revenue from their consumer base. Details of content providers' profit functions can be found in Table 4.

Table 3 Waiting Times

Option	Waiting times
1A/4A	$w_{HY\_1A} = w_{HG\_1A} = w_{LY\_1A} = w_{LG\_1A} = w_{HY\_4A} = w_{HG\_4A} = w_{LY\_4A} = w_{LG\_4A} = \dfrac{1}{\mu - (q_{HY\_1A} + q_{HG\_1A} + q_{LY\_1A} + q_{LG\_1A})}$
1B/4B	$w_{HY\_1B} = w_{HG\_1B} = w_{HY\_4B} = w_{HG\_4B} = \dfrac{1}{\mu - (q_{HY\_1B} + q_{HG\_1B})}$
	$w_{LY\_1B} = w_{LG\_1B} = w_{LY\_4B} = w_{LG\_4B} = \dfrac{\mu}{[\mu - (q_{HY\_1B} + q_{HG\_1B})][\mu - (q_{HY\_1B} + q_{HG\_1B} + q_{LY\_1B} + q_{LG\_1B})]}$
2A	$w_{HY\_2A} = w_{LY\_2A} = \dfrac{1}{\mu - (q_{HY\_2A} + q_{LY\_2A})}$
	$w_{HG\_2A} = w_{LG\_2A} = \dfrac{\mu}{[\mu - (q_{HY\_2A} + q_{LY\_2A})][\mu - (q_{HY\_2A} + q_{HG\_2A} + q_{LY\_2A} + q_{LG\_2A})]}$
2B	$w_{HY\_2B} = \dfrac{1}{\mu - q_{HY\_2B}}$
	$w_{LY\_2B} = \dfrac{\mu}{(\mu - q_{HY\_2B})[\mu - (q_{HY\_2B} + q_{LY\_2B})]}$
	$w_{HG\_2B} = \dfrac{\mu}{[\mu - (q_{HY\_2B} + q_{LY\_2B})][\mu - (q_{HY\_2B} + q_{LY\_2B} + q_{HG\_2B})]}$
	$w_{LG\_2B} = \dfrac{\mu}{[\mu - (q_{HY\_2B} + q_{LY\_2B} + q_{HG\_2B})][\mu - (q_{HY\_2B} + q_{HG\_2B} + q_{LY\_2B} + q_{LG\_2B})]}$
	$w_{HG\_3A} = w_{LG\_3A} = \dfrac{1}{\mu - (q_{HG\_3A} + q_{LG\_3A})}$
3A	$w_{HY\_3A} = w_{LY\_3A} = \dfrac{\mu}{[\mu - (q_{HG\_3A} + q_{LG\_3A})][\mu - (q_{HY\_3A} + q_{HG\_3A} + q_{LY\_3A} + q_{LG\_3A})]}$
3B	$w_{HG\_3B} = \dfrac{1}{\mu - q_{HG\_3B}}$
	$w_{LG\_3B} = \dfrac{\mu}{(\mu - q_{HG\_3B})[\mu - (q_{HG\_3B} + q_{LG\_3B})]}$
	$w_{HY\_3B} = \dfrac{\mu}{[\mu - (q_{HG\_3B} + q_{LG\_3B})][\mu - (q_{HG\_3B} + q_{LG\_3B} + q_{HY\_3B})]}$
	$w_{LY\_3B} = \dfrac{\mu}{[\mu - (q_{HG\_3B} + q_{LG\_3B} + q_{HY\_3B})][\mu - (q_{HY\_3B} + q_{HG\_3B} + q_{LY\_3B} + q_{LG\_3B})]}$

The timing of the game (Figure 4) is as follows: In stage 1, the BSP announces the network management option, which includes traffic prioritization options and pricing schemes for content providers and consumers. In stage 2, two content providers choose service levels, that is, either pay for the premium service or not pay. In stage 3, both types of consumers choose either content provider Y or content provider G.

2.4. Network Management Options

In this subsection, we analyze all the potential network management options. As the BSP considers both the supply and demand sides, each network management option involves two components—content provider discrimination and consumer discrimination. On the supply side (content provider discrimination), content providers may pay the BSP a usage-based fee for preferential delivery of their data packets. Based on whether each of the two content providers pays for the preferential delivery service or not, there are four potential outcomes—outcomes 1, 2, 3, and 4. On the demand side (consumer discrimination), traffic for different types of consumers may be managed in two ways: the two classes of users pay the same fee and receive the same priority (outcome A) or one class

Table 4 Profits for the BSP and Content Providers

Option	The BSP's profit	Content providers' profits
1A	$\pi_{BSP\_1A} = F_{1A}$	$\pi_{Y\_1A} = r_Y q_{Y\_1A}$ $\pi_{G\_1A} = r_G q_{G\_1A}$
1B	$\pi_{BSP\_1B} = \alpha F_{H\_1B} + (1-\alpha)F_{L\_1B}$	$\pi_{Y\_1B} = r_Y q_{Y\_1B}$ $\pi_{G\_1B} = r_G q_{G\_1B}$
2A	$\pi_{BSP\_2A} = F_{2A} + p_{2A} q_{Y\_2A}$	$\pi_{Y\_2A} = (r_Y - p_{2A}) q_{Y\_2A}$ $\pi_{G\_2A} = r_G q_{G\_2A}$
2B	$\pi_{BSP\_2B} = \alpha F_{H\_2B} + (1-\alpha)F_{L\_2B} + p_{2B} q_{Y\_2B}$	$\pi_{Y\_2B} = (r_Y - p_{2B}) q_{Y\_2B}$ $\pi_{G\_2B} = r_G q_{G\_2B}$
3A	$\pi_{BSP\_3A} = F_{3A} + p_{3A} q_{G\_3A}$	$\pi_{Y\_3A} = r_Y q_{Y\_3A}$ $\pi_{G\_3A} = (r_G - p_{3A}) q_{G\_3A}$
3B	$\pi_{BSP\_3B} = \alpha F_{H\_3B} + (1-\alpha)F_{L\_3B} + p_{3B} q_{G\_3B}$	$\pi_{Y\_3B} = r_Y q_{Y\_3B}$ $\pi_{G\_3B} = (r_G - p_{3B}) q_{G\_3B}$
4A	$\pi_{BSP\_4A} = F_{4A} + p_{4A}(q_{Y\_4A} + q_{G\_4A})$	$\pi_{Y\_4A} = (r_Y - p_{4A}) q_{Y\_4A}$ $\pi_{G\_4A} = (r_G - p_{4A}) q_{G\_4A}$
4B	$\pi_{BSP\_4B} = \alpha F_{H\_4B} + (1-\alpha)F_{L\_4B} + p_{4B}(q_{Y\_4B} + q_{G\_4B})$	$\pi_{Y\_4B} = (r_Y - p_{4B}) q_{Y\_4B}$ $\pi_{G\_4B} = (r_G - p_{4B}) q_{G\_4B}$

Figure 4 Sequence of Events in the Game

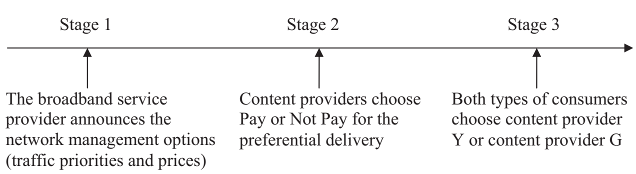

gets higher priority with respect to the other class by paying a higher fee[2] (outcome B). These network prioritization options coupled with corresponding pricing mechanisms result in eight potential network management options for the BSP.

In each of these eight options, the BSP's objective is to maximize its profit π_{BSP} subject to the participation constraints of the two types of consumers. The BSP's profit consists of two parts—Internet access fee from consumers and preferential delivery fee from content providers if at least one content provider chooses to pay the fee. Detailed profit functions for the BSP under different network management options can be found in Table 4. As mentioned earlier, the two indifferent consumers (x_H and x_L) have the lowest utility within their type. Thus, the consumers' participation constraints can be simplified to ensuring that these two marginal consumers have nonnegative utilities. Furthermore, for a particular outcome to be a Nash equilibrium, the content providers should not have any incentive to deviate from that equilibrium. For example, consider outcome 1A (where neither content provider pays for priority delivery), given G's choice of not paying the priority fee, Y's profit when Y does not pay (which corresponds to outcome 1A) has to be no less than Y's profit when it does pay (which corresponds to outcome 2A). Similarly, given Y's choice of not paying the priority fee, G will prefer to also not pay the fee (outcome 1A) over the outcome where it pays the fee (that is outcome 3A). Therefore, the BSP's profit maximization problem in every outcome will include two additional constraints that ensure that Y and G will prefer this outcome over the other outcome if they deviate from the current choice. The general form of the BSP's profit maximization problem can be specified as:

$$\begin{aligned} &\max_{F,p} \pi_{BSP} \\ \text{s.t.}\quad & U_{HY}(x_H) = U_{HG}(x_H) \geq 0 \\ & U_{LY}(x_L) = U_{LG}(x_L) \geq 0 \\ & \pi_Y(I_Y, I_G) \geq \pi_Y(1 - I_Y, I_G) \\ & \pi_G(I_Y, I_G) \geq \pi_G(I_Y, 1 - I_G). \end{aligned}$$

Depending on whether content provider Y or G pays the premium delivery fee to the BSP or not, that is, the values of I_Y and I_G, each outcome can be considered a realization of the above formulation (with Outcome 1A corresponding to the case of net neutrality).

The details of consumers' utility functions are as specified in subsection 2.2, and the details of the profits of the BSP and content providers can be found in Table 4. All derivations of solving the BSP's profit maximization problems with different network management options are provided in the Online Appendix. We find that it is never possible that only the less efficient content provider Y pays the preferential delivery fee while the more efficient competitor G does not. Therefore, outcomes 2A and 2B are not feasible and the BSP has six feasible network management options at its disposal.

In the next section, we compare the six feasible network management options and explore the conditions under which the BSP might choose certain options in equilibrium.

3. BSP's Network Management Choices

In this section, we compare the BSP's profits under the six feasible network management options and derive conditions for the BSP's preferred network management choices. In other words, we solve for the parameter regions and their corresponding outcomes that yield the highest profit for the BSP. Our findings are summarized in Proposition 1.

PROPOSITION 1. (BSP'S NETWORK MANAGEMENT CHOICES). *Outcomes 3A, 3B, 4A, and 4B are the only possible equilibria. Figure 5 visualizes the BSP's network management choices.*

PROOF. See the Online Appendix for detailed proof and mathematical expressions of the separating conditions of the different equilibria.

From the BSP's point of view, user discrimination only (outcome 1B) dominates net neutrality (1A) while discriminating both sides (4B) dominates user discrimination only (1B). Thus, we know that outcomes 3A, 3B, 4A, 4B are the only possible equilibria since outcomes 1A and 1B are dominated while outcomes 2A and 2B are infeasible. We further derive conditions for these equilibria and summarize them in the proof of Proposition 1.

Figure 5 Broadband Service Provider's Network Management Choices

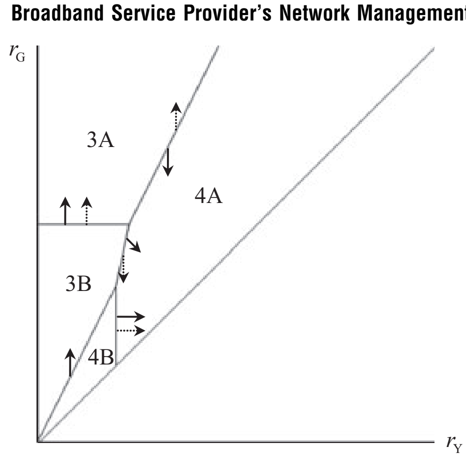

Notes: Figures 5–7 are derived based on parameter values $\alpha = 0.05$, $\mu = 2$, $\lambda_L = 1$, $\lambda_H = 12.67$, $d = 1$, $t = 65$, and $V_H = V_L$. Figures with other parameter values have qualitatively identical structures.[3] The solid arrows represent the changing directions of the separating lines when the BSP's capacity μ increases (utilization rate decreases). The dotted arrows represent the changing directions of the separating lines when $V_H - V_L$ increases.

As shown in Figure 5, the BSP's preferences depend on the market conditions, specifically, the absolute and relative magnitudes of the revenue generation abilities of the content providers, r_G and r_Y. When the two content providers do not differ significantly in terms of their revenue generation rates, the BSP prefers either outcome 4A or 4B. In other words, regardless of what happens on the demand side of the market, the BSP prefers to charge both content providers a preferential delivery fee, which is low enough so that both Y and G decide to pay. This is tantamount to a Prisoner's dilemma for the two content providers—both would be better off by not paying the fee, but end up paying the fee nonetheless in order to ensure that its packets are not degraded with respect to the competitor's packets. However, when r_G is much higher than r_Y, the BSP would prefer either outcome 3A or 3B. In other words, the BSP would charge a high preferential delivery fee that only the more efficient content provider G could afford to pay. In return, G's packets get prioritized and therefore G ends up with a higher market share, and it is this larger surplus of G that the monopolist ISP can effectively extract under these conditions.

When the absolute magnitude of revenue generation rate (r_G or r_Y) is large, the BSP prefers outcome A (3A or 4A). In other words, the BSP prefers not to discriminate among users. However, when the absolute magnitude of revenue generation rate (r_G or r_Y) is small, the BSP prefers outcome B (3B or 4B), that is, the BSP now prefers to discriminate both content providers and consumers. The separating lines between the BSP's choice of A and B is essentially the relative profitability from the high-type consumers vs. that from the content providers. When the absolute magnitude of revenue generation rate (r_G or r_Y) is high enough, the BSP extracts sufficient rent from the content providers to offset the reduced profit from consumers under outcome A as compared with B; then the BSP prefers A. Consumers' valuation for content (V_H and V_L) also has impacts on the BSP's network management preferences. As indicated by the dotted arrows in Figure 5, when H-type consumers value the received content much more than L-type consumers, that is, $V_H - V_L$ increases, outcome B becomes the equilibrium for a wider range of parameter values.

Proposition 1 shows that the BSP does have the incentive to adopt discriminatory network management. In her recent critically acclaimed book *Internet Architecture and Innovation*, Barbara van Schewick (2010) mentions: "Whether network providers have an incentive to use the discriminatory functionality available to them is highly relevant to this debate. Proponents of regulation base their calls for regulatory intervention on the threat of discriminatory or exclusionary behavior by broadband network operators. According to them, regulation is needed to mitigate that threat. If, however, a network operator has no incentive to discriminate against independent portals, content, or applications, regulation is not necessary." (p. 219) Our analysis in the next section supports the need for regulation under certain market conditions. However, instead of limiting our discussion to comparing net neutrality to any discriminatory behavior, we examine the various network management options that are socially optimal and the conditions under which they are optimal: in other words, we use social welfare as the yardstick for evaluating the network management options and deciding on the need for regulation. Our social welfare analysis indicates that under certain market conditions, some types of discrimination actually yield the optimal level of social welfare. However, the BSP's profit-maximizing network management choice might be at odds with the socially optimal network management option, and regulatory intervention is necessary only when the BSP deviates from the socially optimal behavior.

4. Social Welfare Analysis: Reasonable Network Management or Unreasonable Discrimination?

The choice of the social planner with regards to network management options might be at odds with that of the BSP, since social welfare comprises not only the BSP's profit but also the profits of the content providers' profits as well as the consumers' surplus. Note that since the consumers' payments for the broadband services and content providers' payments for premium service are essentially internal transfers within the system as far as the calculation of the social welfare is concerned, the only measurable effect of the consumers on the social welfare comes from their valuations and the disutility that they attribute toward the congestion while the only measurable effect of the content providers on the social welfare comes from their revenue generation capability. We summarize the different outcomes that maximize social welfare (and the conditions under which they are the socially optimal options) in Proposition 2. Finally, we compare the results of Propositions 1 and 2 to find out the parameter values where the BSP's choice of network management deviates from the socially optimal choice, and these results are summarized in Proposition 3.

PROPOSITION 2. (SOCIAL WELFARE MAXIMIZING NETWORK MANAGEMENT OPTIONS). *Outcomes 1A, 1B, 4A, and 4B yield the same level of social welfare. Figure 6 visualizes the social welfare maximizing network management choices.*

PROOF. See the Online Appendix for detailed proof and mathematical expressions of the separating conditions of the different regions.

When content providers are very similar in their revenue generation capabilities, the social planner equally prefers outcomes 1A, 1B, 4A, and 4B, all of which yield the same social welfare. When one content provider has much higher revenue generation capability than the other, the social welfare maximizing outcome is 3A. At intermediate level of difference in revenue generation capabilities, the social welfare maximizing outcome is 3B. Proposition 3 summarizes the conditions where the BSP's network management preference differs from that of the social planner.

PROPOSITION 3. (THE BSP'S DEVIATION FROM THE SOCIAL OPTIMUM—"REASONABLE NETWORK MANAGEMENT" VS. "UNREASONABLE DISCRIMINATION"). *The BSP's network management preference sometimes differs from the social planner's preference. Figure 7 visualizes the BSP's deviation from the social optimum.*

PROOF. See the Online Appendix for detailed proof and mathematical expressions of the separating conditions of the different regions.

Figure 6 Social Welfare Maximizing Network Management Options

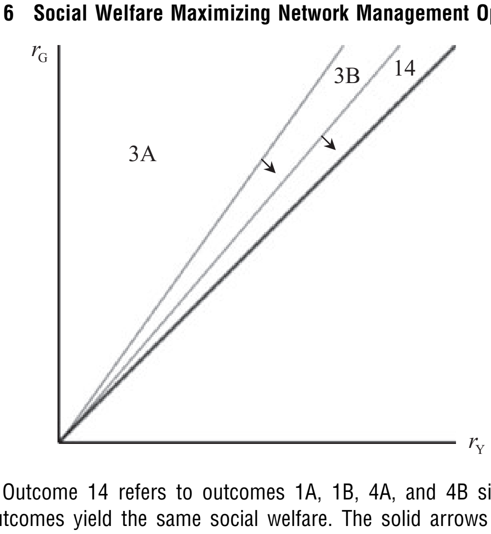

Notes: Outcome 14 refers to outcomes 1A, 1B, 4A, and 4B since these four outcomes yield the same social welfare. The solid arrows represent the changing directions of the separating lines when the BSP's capacity μ increases (utilization rate decreases).

Figure 7 Comparison between the BSP's and the Social Planner's Preferences for Network Management

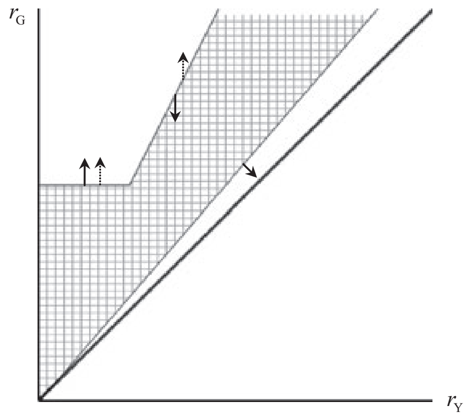

Notes: The shaded area represents the parameter space where the BSP's choice deviates from the social planner's choice. The arrows represent the changing directions of the separating lines when the BSP's capacity μ increases (utilization rate decreases). The dotted arrows represent the changing directions of the separating lines when $V_H - V_L$ increases.

Proposition 3 is a direct result of comparing the results in Proposition 1 and Proposition 2. The shaded region in Figure 7 represents the parameter space where the BSP has the incentive to deviate from the social optimum. Our research suggests that as far as the social planner is concerned, there is no universally applicable optimum policy, since the characteristics of the demand and supply sides of the BSP might differ in every local broadband market. Therefore, the social planner needs to carefully gauge different market conditions (e.g., consumer valuation for content, broadband usage patterns, local broadband infrastructure capacity, and the revenue generation rates of content providers) and accordingly adopt different network regulatory policies. This is especially true in markets where empirical evidence would suggest that the preferences of the BSP and the social planner are at odds with each other (in such cases, the parameters would fall within the shaded regions of Figure 7).

Consumers' valuation for content (V_H and V_L) does not affect social planner's choice, since these two variables affect the social welfare level in the same way. In other words, the total consumer valuations for content across all consumer groups in all outcomes are the same, all equal to $\alpha V_H + (1 - \alpha)V_L$. This is different from the BSP's choice where V_H and V_L play a critical role. This difference (along with other factors) goes to play a part in the BSP's deviation from the socially optimum choice of network management strategies. The dotted arrows in Figure 7 represent the changing directions of the separating lines when $V_H - V_L$ increases.

The FCC has stated in its recent report and order that "the novelty of Internet access and traffic management questions, the complex nature of the Internet, and a general policy of restraint in setting policy for Internet access service providers weigh in favor of a case-by-case approach," and our analysis supports their stance. The FCC has also proposed to "further develop the scope of reasonable network management on a case-by-case basis, as complaints about broadband providers' actual practices arise." However, the FCC does not provide any guidance as to what might be constituted as "reasonable network management" vs. "unreasonable discrimination." This article is a contribution toward that end. Specifically, we derive the conditions under which it makes sense for the FCC to step in and advocate specific network management practices.

It is apparent that the debate on net neutrality is more nuanced than what is popularized in the literature. A strict interpretation of net neutrality would mean that all data packets should be treated equally regardless of their origin, destination, or content. Our research suggests that there are certain deviations from net neutrality that maximize social welfare under various market conditions. For example, as we know from Proposition 1, outcome 3A is the BSP's preferred choice if both the absolute magnitude of r_G and the relative magnitude of r_G compared with r_Y are high. When we evaluate this outcome from the social planner's perspective, Proposition 2 shows that outcome 3A also maximizes the social welfare, and therefore no regulatory intervention is required in this case. Such a network management policy should be considered "reasonable network management" since it is not only preferred by the BSP but also yields the socially optimal welfare level.

Some other network management options may be preferred by the BSP but detrimental to social welfare. For example, when the revenue generation capability of G (r_G) is high relative to Y (r_Y) but the absolute magnitude of r_G is low, outcome 3B is the BSP's preferred choice. However, outcome 3B yields a lower level of social welfare and therefore should be considered "unreasonable discrimination." Such a network management policy should be regulated.

5. Concluding Remarks

The debate of net neutrality and the potential regulation of net neutrality may fundamentally change the dynamics of data consumption and transmission through the Internet. Much of the extant literature on economics of net neutrality has looked at the problem from the "supply" perspective, but due to its unique position, the broadband service provider can affect the transmission of content at both the supply and demand sides. Therefore, a holistic view of the problem requires us to consider traffic management not only at the content providers' end, but also simultaneously consider the possible traffic management policies that might be designed for different types of consumers. This article analyzes both the supply and demand sides of the market in the same model. It is important to emphasize that making a network management decision for the two sides of the market independently (i.e., choosing outcome 1, 2, 3, or 4 for the content providers on the supply side and choosing outcome A or B for the consumers on the demand side) is suboptimal. The BSP and the social planner have to make their network management decision by considering both sides simultaneously (e.g., 3A or 4B). In the previous sections, we have discussed the results of our analysis in considerable detail, and therefore, we will conclude with the main "takeaway" policy implications of our research.

We find that an implementation of strict network neutrality that forbids all forms of network management (i.e., outcome 1A in our model where there is neither content provider nor consumer discrimination) is always suboptimal for the BSP and yields a lower level of social welfare in most cases. This highlights the need for a careful examination of all possible network management options and for the distinction between "reasonable network management" and "unreasonable discrimination," which is what we have attempted in this article.

Our research shows that there are deviations from the boilerplate arguments to go for or against net neutrality that only consider the supply side. It also indicates the need for a case-by-case approach that the FCC advocates using all the tools of network management available to the BSP. Characteristics of both the supply side and the demand side are important factors in evaluating the network management options. In addition to taking into account the "particular network architecture and technology" (FCC 2010a, p. 48), we suggest that policymakers should also consider other market factors such as content providers' revenue generation rates and consumers' valuation for content. In order for the regulators to implement our recommendations properly, they need to obtain information such as consumer valuation, broadband usage patterns, revenue generation rates, and so on. An ill-conceived policy resulting from lack of knowledge or information may have a negative impact on the social welfare.

Our model makes some simplifying assumptions about the actual marketplace. For example, users are categorized as H-type or L-type, but in reality, BSPs can discriminate between several classes of users. The solution to the problem will become more involved, but the core ideas should still follow. In addition, we did not consider one issue that has been raised by some commentators—that of the possibility that a select few "power users" within a network have positive externalities that help other users and the BSP in the long run (Al-Chalabi 2008). Although the veracity of such a claim can be debated, we think that future research can consider simulating the presence of such users within a network and thus examine the fallouts of their presence. In this article, we consider an advertisement-assisted revenue model for the content providers. It would be interesting to see how the results would change if the content providers can charge the consumers directly through a subscription revenue model.

Given the fact that the local broadband services market has very limited competition in the United States and the last mile infrastructure is privately owned, it is difficult for the regulator to enforce certain network management principles on the monopolistic BSP. The National Broadband Plan released by the FCC (2010b) envisages that wireless broadband technologies will usher in competition in the local broadband services, and technologies like WiMax and LTE are beginning to appear in some selected geographies. Future research can investigate how introducing competition in the local broadband services market can affect the choices of the BSP and the social planner.

Acknowledgments

The authors thank three anonymous referees, the senior editor, and the department editor for their constructive comments that helped improve the article. This research is funded by the NET Institute (www.NETinst.org) whose financial support is gratefully acknowledged.

Notes

[1]These two assumptions are consistent with empirical evidence. For example, AT&T has recently estimated that their top 5% of users (in terms of usage) account for about 40% of the total traffic (Tweney 2008).

[2]Outcome B is different from the BSP's practice of charging different prices for different service packages (which correspond to different capacity levels μ in our model) from consumers. In our model, we consider a particular service package (a given capacity level μ) and the consumers who subscribe to this same service package have different usage patterns. Some consumers request more content than other consumers ($\lambda_H > \lambda_L$). The type of user discrimination modeled in this article involves distinguishing two types of consumers H and L and managing their traffic accordingly.

[3]These values reflect real-life empirical parameter values. AT&T has recently estimated that their top 5% of users (in terms of usage) account for about 40% of the total traffic, that

is, $\alpha = 0.05$ and $\alpha\lambda_H/[\alpha\lambda_H + (1-\alpha)\lambda_L] = 0.4$ (Tweney 2008). We normalize λ_L to 1, and it follows that $\lambda_H = 12.67$. We then change the value of μ such that we explore a wide range of utilization—from 10% to 80%. We further set $d = 1$ and $t = 65$ such that the fit cost is comparable to congestion cost.

References

Al-Chalabi, M. 2008. *Broadband Usage-Based Pricing and Caps Analysis*. GigaOm. September 30. Available at http://gigaom.com/2008/09/30/gigaom-white-paper-the-facts-fiction-of-bandwidth-caps/ (accessed date May 19, 2011).

Cheng, H. K., S. Bandyopadhyay, H. Guo. 2011. The debate on net neutrality: A policy perspective. *Inf. Syst. Res.* **22**(1): 60–82.

Choi, J. P., B. Kim. 2010. Net neutrality and investment incentives. *RAND J. Econ.* **41**(3): 446–471.

Economides, N., B. Hermalin. 2013. The economics of network neutrality. *RAND J. Econ.* Forthcoming.

Economides, N., J. Tåg. 2012. Net neutrality on the internet: A two-sided market analysis. *Inf. Econ. Policy* **24**(2): 91–172.

FCC. 2010a. *Report and Order: Preserving the Free and Open Internet*. Federal Communications Commission. December 23. Available at http://www.fcc.gov/document/preserving-open-internet-broadband-industry-practices-1 (accessed date October 29, 2012).

FCC. 2010b. *Connecting America: The National Broadband Plan*. Federal Communications Commission. March 16. Available at http://download.broadband.gov/plan/national-broadband-plan.pdf (accessed date May 19, 2011).

Guo, H., S. Bandyopadhyay, H. K. Cheng, Y. Yang. 2010. Net neutrality and vertical integration of content and broadband services. *J. Manag. Inf. Syst.* **27**(2): 243–275.

Guo, H., H. K. Cheng, S. Bandyopadhyay. 2012. Net neutrality, broadband market coverage and innovations at the edge. *Decis. Sci.* **43**(1): 141–172.

Helm, B. 2006. Tech giants' internet battles. *BusinessWeek*, April 27.

Hermalin, B. E., M. L. Katz. 2007. The economics of product-line restrictions with an application to the network neutrality debate. *Inf. Econ. Policy* **19**: 215–248.

Hillier, F. S., G. J. Lieberman. 1990. *Introduction to Stochastic Models in Operations Research*. McGraw-Hill, New York.

Hotelling, H. 1929. Stability in competition. *Econ. J.* **39**(153): 41–57.

McCullagh, D. 2007. Comcast really does block BitTorrent traffic after all. *CNET News*, October 19.

Newmark, C., M. McCurry. 2006. Should the net be neutral. *The Wall Street Journal*, May 24.

Turner, J. 2010. Net neutrality debate. *Citizen Economists* February 22. Available at http://www.citizeneconomists.com/blogs/2010/02/22/net-neutrality-debate/ (accessed date May 19, 2011).

Tweney, D. 2008. AT&T embraces BitTorrent, may consider usage-based pricing. *Wired*. June 6.

van Schewick, B. 2010. *Internet Architecture and Innovation*. The MIT Press, Cambridge, MA.

Waldmeir, P. 2006. The net neutrality dogfight that is shaking up cyberspace. *Financial Times*, March 23.

Supporting Information

Additional Supporting Information may be found in the online version of this article:

Appendix S1. Derivations of solutions to the BSP's profit maximization problems with different network management options

Appendix S2. Proof of Proposition 1 (BSP's network management choices)

Appendix S3. Proof of Proposition 2 (Social welfare maximizing network management options)

Appendix S4. Proof of Proposition 3 (BSP's deviation from social optimum)