

NETWORK EFFECTS, ECONOMIC EFFICIENCY,
AND USAGE-BASED PRICING FOR
INTERNET ACCESS

by

MATTHEW L. COOPER

B.S., Kansas State University, 2008

B.S., Kansas State University, 2008

A REPORT

submitted in partial fulfillment of the requirements for the degree

MASTER OF ARTS

Department of Economics
College of Arts and Sciences

KANSAS STATE UNIVERSITY
Manhattan, Kansas

2011

Approved by:

Major Professor
Dr. Tracy M. Turner

Abstract

This paper attempts to shed some light on the issue of net neutrality by examining the extent to which Internet usage is efficiently allocated under current conditions. I discuss the unique features of Internet usage which make it a good that markets will tend to provide at an inefficient level. I then discuss alternative pricing regimes that will move the market for Internet usage to efficiency. I conclude with a discussion of the current economic research on the topic.

Table of Contents

List of Figures	iv
1. Introduction.....	1
2. Market Structure and Pricing	4
<i>Pricing</i>	6
<i>Two-Sided Markets</i>	7
3. Economic Theory: The Efficiency Aspects of Broadband Internet.....	8
4. Economic Literature: Impacts of Price Structure.....	13
References.....	19

List of Figures

Figure 1. Basic Architecture of the Internet.....	5
---	---

1. Introduction

The Internet is a highly efficient conduit of information transfer. So efficient, in fact, that businesses have emerged which take advantage of that efficiency and are able to make money by offering goods and services online. There are some firms offering them for a direct fee which they charge to the user. This type of transaction is very similar to traditional buyer/seller relationships where the price is determined through negotiations between the two sides. The Internet also generates revenue for businesses because many people use it: owners of the websites people use sell space on their pages to advertisers. However, in order for users to view these pages in the first place, data must be transmitted over a network which someone owns, adding a third party to the transaction.

When an Internet user streams video from YouTube.com, that usage can exert a negative effect on other users on his or her local network. This is a result of the traditional way in which internet commerce has been organized. Broadly speaking, all users of a local network must compete to use an amount of that network's finite capacity. When the amount of data attempting to reach end users exceeds the local network capacity, everyone using the network experiences congestion in the form of slow page loads, downloads, and server time-outs. Google, YouTube's owner, has decided that the solution to internet congestion is more speed from better infrastructure. In February of 2010, it announced that it would choose a few American cities as pilot programs, for which they would establish local internet connections a thousand times faster than what current infrastructure allows. Out of 1100 cities who applied, on March 30th of this year Kansas City, Kansas was announced as Google's first test location. Google hopes to alleviate congestion by delivering more speed to subscribers without charging higher prices than current providers do (Kincaid, 2011). This is one approach to solving the problem of internet

congestion: simply building a larger pipe to move data through. The other solution, however, is to prioritize internet traffic by directly charging producers and consumers of content based on the costs and benefits created from usage.

Prioritizing this data can be done automatically and presumably with no transaction cost. However, the rules by which this prioritization takes place and the effects of these decisions have vast ramifications for all of the parties involved in internet commerce. Since the owners of the network are able to set these prioritization rules any way they wish, they may realize an incentive to behave anti-competitively and have been accused of doing so in the past. This, of course, has several stakeholders concerned. However, neither regulators in the government nor the various advocacy groups which have sprung up are able to agree on what, exactly, constitutes anti-competitiveness in the broadband connection market.

Rather than wait for an instance of consumer harm as a result of networks behaving in such a way, some stakeholders want the government to write prioritization rules which identify any deviation from a typical pricing structure as illegal, arguing that the status quo is what will allow the Internet to become more valuable to society over time. As economists continue to debate over who the real winners and losers are under the many possible pricing structures, the Federal Communications Commission (FCC) has adopted a set of rules by which it wants owners to manage its networks (Gustin, 2010). The first rule requires full disclosure regarding data prioritization. Anyone providing internet access must fully disclose network management practices because they determine whose data arrives at its destination first, affecting the delay experienced by the user.

The second and third rules deal with how, specifically, data is managed. A network cannot block any data from transmission through its wires, ensuring that a network subscriber

can access any online content he or she wishes. Also, data cannot be relegated to a slower connection speed in order to deliver other data quicker. Since network owners do create their own content and offer it online, these rules are intended to prevent anticompetitive prioritization of data. This bans a number of possible pricing strategies that block data if the sender does not pay a fee. Service providers also may not direct data to a "fast" or "slow" network speed based on who does or doesn't pay a premium. In effect, the FCC is mandating that whatever network owners do to one package of data, they must also do to everyone else's data. This would, however, still allow the network to increase price for everyone when usage is at its peak, or to offer different network speeds to consumers for a premium price.


Reactions to the rules from stakeholders are mixed, as it is unclear which pricing strategy network providers will pursue and what will happen in the market as a result. Additionally, it is unclear what the FCC has actually prevented with these mandates, as economists continue to debate the potential effects of those pricing strategies deemed anticompetitive by these regulations. This paper attempts to shed some light on the issue by examining the extent to which Internet usage is efficiently allocated under current conditions. I discuss the unique features of Internet usage which makes it a good that markets will tend to provide at an inefficient level. I then discuss alternative pricing regimes that will move the market for Internet usage to efficiency. I conclude with a discussion of the current economic research on the topic.

2. Market Structure and Pricing¹

To understand how network structure can create congestion, it is necessary to review some principles of internet technology. In order to access internet content consumers must first subscribe and log on to a local broadband network using any type of wired or wireless device capable of internet access and data exchange, such as a personal or laptop computer. Typically, there is a cost associated with access to any network. For example: subscription fees for cable internet service provision to residential and commercial users, student technology fees charged along with tuition, content providers (CP) negotiating with local service providers to hook up their file servers; are all examples of access fees charged by an internet service provider (ISP). The ISP industry is a set of telecommunications firms, cable companies, and wireless providers who have made investments in the physical infrastructure of the internet and derive revenue from charging for its use. Both consumers and providers of content gain online access through an ISP who has built commercial or residential networks in any given municipality or geographic location, often as a natural monopoly utility. This is the so called last mile connection (LMC), and to its subscribers it delivers access to the World Wide Web through various intermediary network connections. Figure 1 is a simple visualization of how these interconnections between networks are organized into the Internet.

¹ This section draws on Yoo (2006), who discusses basic internet principles and Goldratt (1985), who discusses bottlenecks.

Figure 1. Basic Architecture of the Internet



Source: Yoo (2006)

In the short run, ISPs cannot make investments in the physical network, and there is a finite amount of channel capacity (bandwidth, net bit rate, maximum throughput) available on any network. This measures the absolute maximum amount data which can move through the network during a given time span. Returning to figure 1, we see that data must transfer from one network to another several times in order to get from its origin (for example, the "File Server") to its destination (for example an "End User on a different local network). Users and originators of content pay for access to the local network, which has contractual agreements for access to intermediary networks, who also have agreements with major continental backbone networks for access to them. Intermediary and backbone networks have much higher channel capacity than the LMC. However, as content moves from backbone infrastructure to its destination, channel capacity from one network to the next is progressively smaller. Data in excess of capacity queues at these points of transfer, revealing them to be physical bottlenecks in the Internet itself. Bottlenecks are specific processes within a system which innately create delays, and they are typically characterized as being at the front of a line of uncompleted work. Currently in

broadband, data waits on a first come first serve basis at the transfer point for a place on the next network is uncompleted work, creating congestion as more data is demanded over a given amount of time. Aggregated to the LMC connection, these time delays are experienced by the consumers of content as poor quality of service: slow page load times, download interruptions and failures, and server time-outs. This congestion where data moves between networks exerts an external cost which degrades the quality of everyone's access.

Pricing

There is an important distinction between fees charged for access and fees charged for use. Pricing a service, whether the fee is usage or access based, excludes those who are not willing to pay the fee, thus decreasing congestion to some extent. A flat fee charged evenly to all individuals who wish to use a service is considered an access price, whereas a fee charged based on the intensity of use is referred to as a *usage price*. Currently, ISPs charge an access fee to their subscribers, but not to CPs from different networks when they use local network capacity (as they do when a local end user accesses the non-local CP). This is known as a *zero-price regime*, where no fee is charged to non-local CPs. This practice of maintaining a zero price regime is what is known as *net neutrality*.

Usage pricing means that all parties in the transaction pay some kind of fee for local network capacity.² Tracking and charging for intensive use involves determining the origin or the destination of content. Only recently has it become technologically feasible for ISPs to perform this task. Identifying specific senders or receivers of content requires that computer

² Many papers in the literature call this "price discrimination" (for example, Economides & Tag, p.32, 2009), however I call it usage-based pricing or usage pricing. This is because price discrimination is based on willingness to pay, not cost differences.

software instantaneously match individual address codes (IP addresses) with their owners.³ This can allow ISPs to charge CPs for the congestion costs created as users access their file servers. It has been proposed that greater network efficiency (less congestion, less delay, etc) can be achieved if usage pricing is allowed, i.e.: packets are prioritized according to who pays or doesn't pay a fee, rather than on a first come first serve basis.

Two-Sided Markets⁴

A two sided market is one in which either a monopoly or oligopoly platform charges both the provider and the user of a service for access to one another (Rochet & Tirole, 2009). A movie theatre is an example of a *two-sided market*. It serves as a platform on which movie studios are able to access an audience. The studio bears a cost to produce, the theatre a cost to show, all driven by convincing people to go to the theatre. They put a premium on tickets to highly-anticipated film releases, recognizing that a viewer's willingness to pay for the ticket has increased. They may also recognize congestion as theatre parking lots become overcrowded when hundreds of cars sit while their owners watch four hours of the final Star Wars sequel. When moviegoers interested in other types of films show up and fail to find a parking space, they choose to drive to a less crowded theatre to watch a movie. The key insight here is that the price to enter a two sided market such as a movie theatre is dependent on the costs associated with both access and use of the platform, and also on the willingness of the two sides to pay for access to one another using the platform.

³Internet Protocol (IP) addresses are unique to each LMC connection.

⁴ See Rochet and Tirole (2006, 2003), Musacchio, Shwartz, and Warland (2009), Lee & Wu (2006), Economides and Tag (2009), Armstrong (2006) for in-depth discussions of Two-Sided Markets.

The LMC is also a two-sided market. The ISP owns the LMC and, as already noted, has the possibility to charge either access-based or usage-based fees to the end users and CPs.

Applying this framework to broadband connection, the current regime of zero pricing means that fees are not collected from both sides of the market. For usage of the LMC, economic theory suggests that switching to a price based on usage will yield efficiency gains.

3. Economic Theory: The Efficiency Aspects of Broadband Internet⁵

Markets are efficient when the price of a good reflects the true social marginal costs and benefits of the last unit's consumption. If this is the case, there will be no benefits or costs experienced by other individuals as a result of a person's consumption, all participants will have perfect information about the true costs and benefits of the good, and no agent will be able to use market power to affect prices. In essence, the price signal must work, in that it optimizes the tradeoff between benefits and costs from additional use. As a preview of the following discussion, we will see that the market for Internet use is not efficient due to usage externality issues.

In order to determine if a market will be efficient, we must understand the extent to which a good is public or private (O'Sullivan et al., 2010). Access to the network itself has so called 'public good' aspects. A pure public good must be non-rival and non-excludable, whereas a private good is both rival and excludable. The excludability of a good depends on the provider's ability to physically keep someone from accessing it or gaining its benefits from use. A good's rivalry is determined by the way a good's use affects its value; does one's enjoyment of

⁵ This section is drawn from Ch. 4(Public/Private) and Ch. 20(Club Goods), Rosen (2001)

a good affect the enjoyment or potential enjoyment derived by others? Despite a literal interpretation of the words public and private, designation as one or the other gives no indication of whether the public or private sector provides or produces the good. Instead, it depends on the intrinsic characteristics of the good itself relating to how it may be distributed and used. For instance, one classic example of a pure public good is National Defense. Non-Excludable, it is difficult to deny the benefits of protecting the borders to someone inside the United States. Non-Rival, as one's benefiting from national defense neither takes away nor adds anything to the benefit another receives from it; i.e.: the additional cost of defending US borders created by one more person is zero. Thus, the marginal cost to society of protecting one more American from invasion is equal to zero. In contrast, a private good would be exemplified with public elementary education. It is excludable, as non-payers can be physically stopped from consuming it. A teacher can only work intensely with one student at a time, so in this sense it is rival. As a result, elementary education is a private good which is publicly provided in many countries.⁶

The interaction between ISPs, CPs, and consumers as they access broadband networks creates cross-group externalities which can positively or negatively impact the number of producers and consumers who choose to enter, ie: pay to access the platform (Economides & Tag, 2009). Of particular note from the literature is the presence of network effects or externalities on all agents in the market, that there is a "...marginal effect that an additional user of a network has on existing users" (Asvanund et al., p.1, 2002). The positive network effects result from the nature of online interface, that in order for there to be value in it there must be someone to interface with. This means that initially the marginal willingness to pay for

⁶ Education does have non-rival aspects.

connection to online networks increases with the number of users.⁷ As the platform on which CPs and users meet, ISPs have an incentive not only to attract more subscribers, but also to encourage the development of innovative content for users to interface with. Since both ISP and CP firms can only grow as Internet usage expands, increasing network subscribers and attracting users to content has been the primary focus of their operations. This 'eyes-on-the-screen' approach to allocating access exclusively ignores the rivalry of broadband usage.

The ISP's ability to exclude potential access by denying connection to non-payers fits the definition of a private good. The network effects associated with interface gives broadband connections both rival and non-rival features. These effects are positive when each additional user contributes a net benefit to the network by providing an additional user to interface with. Since network capacity is fixed in the short run, the positive effects of additional users become smaller and smaller relative to the negative effects of congestion to everyone on the network. Eventually, any additional use has a negative effect on market efficiency when usage is too high. Therefore, access to a broadband network is a so called *club good*: a private good with non-rival characteristics when marginal consumption creates a congestion cost to the network (Bruce, p.74, 2001). This is the negative network externality present in the market.

Content available through accessing a local network is diverse. Not all use contributes the same to congestion. Confronting congestion issues by charging higher access fees would overcharge cursory use and undercharge use which requires a high information bit-rate. While charging flat access fees excludes those not willing to pay, the price will likely not reflect the congestion costs inflicted by subscribers' usage. There is also a notable effect of time on

⁷ O'Sullivan et al. (2010) discusses principles at length, Lee & Wu (2009), Rochet & Tirole (2006) both discuss application to internet.

congestion, as it appears to increase at regular intervals during the day, indicating peak periods. Since high-rate information contributes the most to congestion, it would be efficient to charge for that use, especially during times of peak usage (Crew et.al., 1995).

The problem is also experienced in our movie theatre. Even though attending a theatre is not the only way to watch a movie, and movie tickets are not the only way to capitalize on a movie, it still displays the same network effects as broadband internet. Let us suppose, for instance, that a comedy might be more enjoyable if there are many people in a theatre laughing. Romantic films may be more enjoyable with fewer people present in the room. Scary movies might have both effects, depending on the preference of the viewer. The viewer might enjoy being scared and becomes more afraid with fewer people in the theatre. Someone else might avoid being scared, preferring to watch horror films with many people around.

Like film genres, the internet's content can also be categorized according to its network effects. E-mail, for instance, increases in value as the number of people to e-mail increases. In this case, access itself would have positive network effects, whereas use itself would not. Using e-mail demands very little space on the network and does not contribute to congestion, quickly sending text from origin to receiver. Online gaming requires peer-to-peer network connections and are governed by network effects. Under the current system, there is an access fee when the game is purchased, and often an additional usage fee charged per month or per year in order to play. All games require a set number of players connected to one server sending data at a high bit-rate sustained for the entirety of play, which can last several hours. During the course of play, congestion is felt in the form of lags between user commands and the server's response. The user's command originates a data package which is sent over the broadband network to the game's server where it is processed. The games response to user action is then sent back to the

user through the network. A split second's interruption of play can and will cause a player to lose, diminishing the willingness of that person to play. As one subscriber plays online games, he or she uses bandwidth that another subscriber to that network infrastructure cannot until the game is over. When networks approach their maximum bit-rate, especially during peak hours, users compete for bandwidth and exert congestion costs on one another.

In terms of describing congestion as a negative externality, we can say that the negative externality occurs because the transactions of one individual affects another person's well-being who was not a part of the original transaction. When one subscriber to a LMC decides to use any content, he or she uses network capacity which another user cannot. This has a negative effect on other local users, which increases along with usage intensity and the amount of users receiving data at a given time. The presence of externalities means that the market is inefficient in its allocation of capacity, resulting from a price which is too low. Consequently, consumption of network capacity is too high. The efficiency of markets requires that the price of use be set equal to the so called social marginal cost (SMC) of use. The SMC is equivalent to the price paid by the initial user (private marginal cost, PMC) for using an additional unit, plus the marginal damage (marginal external damage, MED) imposed on another subscriber from that use. When usage intensity is low congestion is not a factor, and the efficient price of using additional network bandwidth is close to zero. Without congestion, net neutrality is economically efficient. In contrast, usage pricing on the LMC is efficient when the network is congested, or with "bandwidth-intensive" content (Bruce, 2001).

Theory then dictates that if price for usage in the market were to be equal to the social marginal cost of one more packet of data sent over the wire, congestion would be eliminated from the network entirely at the marginal level, (Yoo, 2006). Network effects may be such that

greater use on one side of the market positively or negatively affects the other. This could be a result of the platform observing a positive network effect on one side of the market from more entrants to the other. The positive effect of additional use, in this case, increases the other side's willingness to pay, whereas a negative one would decrease it. A usage-based regime also means that users who create positive network effects ought to be subsidized by the ISP to encourage entry (Economides & Tag, 2009). Therefore, an efficient price will reflect the social marginal cost or benefit of use. This means that usage with positive network effects is charged the lowest price, and high-capacity content during peak usage hours has the highest price.

4. Economic Literature: Impacts of Price Structure

The research into this area is voluminous and draws from legal and economic disciplines dealing with government regulations and their effects. This section focuses on the key results of six recent papers. They seek to expand on most recent economic literature on net neutrality, which primarily focuses on access based rather than usage based charges. Rochet & Tirole (2006) apply two sided markets to broadband connection and clarify some of the empirical parameters with several conditions for two-sidedness. They note that the structure of pricing for access to either side of the platform affects allocation. This means that a variety of pricing structures might be efficient given the different network effects associated with usage and access. Not all Internet transactions can be called two-sided, however, as negotiation between users and CPs would efficiently allocate usage. For example, Netflix charges its own monthly subscription fee for its streaming services outside of the ISP's access price. The authors also note a seesaw principle where platforms charge one side low and the other high to capture profit margins instead of exploiting network effects. The article accounts for both usage and access fees in their

analysis and comes up with optimal pricing conditions for both sides of the market. Most of the subsequent published work on two-sided markets references the methods presented in this paper.

Economides and Tag (2009) use two-sided markets to measure how allocation changes under different pricing regimes. They focus on the question of whether neutrality or charging CPs for access will increase the content available in the market. The assumption is that since users highly value new content, more content will drive greater entry on the consumer side. This means that attracting more users will increase revenues for both CPs and ISPs, ensuring future internet growth. This takes into account the positive and negative network effects a market experiences from different types of content usage. While congestion is present as a negative externality in the market when there are strong network effects it is eliminated with the exclusionary access fees.

In the neutral case, the ISP charges zero price to CPs for access to one side of the market while charging consumers a subscription fee. Alternatively, a usage-based pricing scheme would allow ISPs to charge both sides of the market for access. The paper finds that when network effects are strong neutrality is the preferable regime. This depends on the assumption that consumers derive a great benefit from additional content, meaning that additional content has a positive effect on the amount of users who choose to enter the market. As more consumers pay the access fee, CPs and ISPs benefit from more revenue. When ISPs can be allowed to charge both sides, the amount of CPs choosing to enter is lower, and since CPs bear a share of the cost of network effects, the ISP charges a lower price for consumer access. The authors argue, however, that since more content and usage is better than less, neutrality is the preferable network regime as overall levels of both are higher.

Lee & Wu (2009) also acknowledge the dominion of network effects over markets for broadband access with their title "Subsidizing Creativity:...", justifying neutrality because it promotes innovation. The basis for their argument lies in the assumption, once again, that content innovation allows more entry to both sides than efficient pricing, and that current price structures already reflect the negative external costs associated with usage. The authors do not present any empirical analysis, proposing on theoretical grounds that CPs already pay for network usage and that ISPs receive enough revenue from charging its subscribers to invest in capacity expansion. For these and other reasons the article concludes that neutrality will improve the internet through increasing new investment in innovative content.

Hermailin & Katz (2007) take a look at net neutrality as a product line restriction, in other words a mandated price of zero would keep ISPs from offering product variations to subscribers. The analysis also employs two-sided markets to derive allocation results under the alternative pricing regimes and parameters indicating various network effects. Under neutrality, there is only one connection level available which must serve all subscriber/CP types. The alternative allows market entrants to choose amongst three (top, middle, and low quality) tiers at variable pricing. Overall, the results are mixed and depend on the real strengths and directions of network effects as to which regime yields greater entry. However, when only one product line is available, they find that entrants seeking low-quality are excluded. Those in the middle buy higher quality than otherwise and those demanding top-tier access end up with lower speed than they desired, resulting in less efficient online interactions for those CPs and users.

Choi & Kim (2010) also look at the potential effects of allowing ISPs to offer access tiers by assessing potential profits as incentives to invest, i.e.: enter the market. The price charged for CP access as it affects their profit potential determines that side's entry. Once again, usage-

pricing will decrease CP entry and the authors warn against stifling innovation as a side-effect. Even charging a price equal to the SMC of their use, they argue, may create barriers which innovative content might not be able to overcome. An ISP will invest if it sees a large enough profit in charging higher prices for higher access tiers. However, the authors note that there will be a negative effect for the very top tier as there may not be very many people willing to pay the price, decreasing its value to the ISP relative to the other tiers.

Kramer & Wiewiorra (2010) apply two-sided market analysis in the same way, but with a twist. Instead they build on the idea that if an ISP offers two access tiers to either side of the market, network congestion would be reduced. The intuition is that the different tiers represent a higher or lower network capacity, and that the network effects governing each side's willingness to pay will efficiently allocate market entry by both sides into one of the two tiers. The sensitivity of content to congestion plays a key role, as this parameter indicates the relative strength of network effects on the market. By analyzing all of the potential cases, they conclude that tiered access improves market efficiency over neutrality by alleviating congestion. Because congestion-sensitive content will enter the top tier, users not willing to pay the higher price for access are left with no choice but to connect to the other tier, lowering congestion in both tiers. It is clear that CP firms would be hurt by being charged a premium for the sensitivity of their content. The authors, however, argue that because congestion is eliminated from the internet and overall usage increases, the Internet as a market works better without mandated neutrality.

Another approach taken by Musacchio, Schwartz, and Walrand (2009) specifically analyzes the effects of these same conditions on ISP's and their incentives to invest in network expansion. Instead of strictly applying a two-sided market to broadband internet, they model the zero-price regime as a one-side market and a usage pricing case as two-sided. Also taken into

account is the number of ISP's available to content providers. Once again, the authors test and compare optimal values under either pricing regime. Additionally, they account for the value of content derived by consumers and the different results these cases could yield. This is accomplished by finding the marginal revenue CP's would make from advertisers per user click on a page and dividing this value by a constant representing price sensitivity of consumers. Intuitively, since advertising revenue per click is assumed to be constant in all regimes, if consumers highly value content they would be insensitive to price, and this ratio would be high. They evaluate cases representing high, middling, and low values for this parameter in both price regimes. Also as a result of this assumption, rates of return on investments for both ISP network capacity and CP content are the same in all regimes, as for both agents this is dependent on the number of users attracted to the platform.

The findings rely on the idea that market entry is maximized when investment incentives in the form of projected profits for CP's and ISP's are high. In cases where either there is an extreme value of the ratio between marginal revenue per click and the consumer's price sensitivity, unregulated monopoly pricing is shown to encourage more investment by both parties. When this ratio is mid-range, platform competition for content becomes more of a factor. As the number of platforms increases, the negative effects of any one content producer choosing not to enter is diffused amongst more platforms, giving ISPs a tendency to overcharge CPs for access in a non-regulatory regime. This has a negative net effect on content production in general, and hurts the amount of investment in innovative content available on the internet. However, as far as modeling long term investment in infrastructure or service innovation, this approach falls short. Since the test design does not allow the number of ISPs to fluctuate, these results can only really be interpreted to represent short run market outcomes after a change in

prices. It is probable that long run effects of a structural market change would not only create innovation opportunities for some, but also take them away from others. Therefore, it is difficult to make any definitive conclusions from these results.

The papers above are a good representation of the two sides of this price regulation debate presently in the economic literature. Most authors disagree on the extent to which positive and negative network externalities affect usage. Some base their arguments for neutrality on the assumption that new and innovative content is the main attraction to users as they pay an access fee and enter the market. Others maintain that since LMC usage is a club good, it's efficient allocation requires the price for access be set as close to the social marginal cost as possible. What is clear is that consumers pay more for internet access under neutrality, as ISPs are only allowed to recoup high capital costs from one side of the market. CPs end up with lower operational costs, and therefore higher profits than if usage pricing were allowed. Whether or not usage will increase and the market will expand under these new rules, however, is still not certain.

References

- Armstrong, M. (2006): "Competition In Two-Sided Markets," *RAND Journal of Economics*, 37: 668–691.
- Asvanund, Atip; Clay, Karen; Krishnan, Ramayya; Smith, Michael D., (2002), "An Empirical Analysis of Network Externalities in P2P Music-Sharing Networks," Working Paper/Report presented to H. John Heinz III School of Public Policy and Management, Carnegie Mellon University, Pittsburgh, PA
- Bruce, Neil (2001), Public Finance and the American Economy., 2nd ed., Boston: Addison Wesley
- Choi, Jay Pil; Kim, Byung-Cheol, (2010), "Net Neutrality and Investment Incentives," *RAND Journal of Economics*, 41: 446-471.
- Crew, Michael A.; Fernando, Chitru S.; Kleindorfer, Paul R., (1995), "The Theory of Peak-Load Pricing: A Survey," *Journal of Regulatory Economics*, 8: 215-248.
- Economides, N. and J. Tag (2009): "Net Neutrality on the Internet: A Two-Sided Market Analysis," Working Paper, Stern School of Business, New York University.
- Goldratt, Eliyahu M.; Jeff Cox. (1984) The Goal: A Process of Ongoing Improvement. Great Barrington, MA.: North River Press.
- Gustin, Sam. "FCC Passes Compromise Net Neutrality Rules." *Wired.com* Dec. 21, 2010, accessed on April 28, 2011: <<http://www.wired.com/epicenter/2010/12/fcc-order/>>

- 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.
- Jan Krämer and Lukas Wiewiorra (2010) “Network Neutrality and Congestion Sensitive Content Providers: Implications for Service Innovation, Broadband Investment and Regulation,” Karlsruhe Institute of Technology, NET Institute Working Paper #10-09.
- Kincaid, Jason. "Google Chooses Its Fiber-Networked City Of The Future: Kansas City." TechCrunch.com Mar. 30, 2011, accessed on Mar. 31, 2011: <<http://techcrunch.com/2011/03/30/google-chooses-its-fiber-networked-city-of-the-future-kansas-city/>>
- Lee, R. S. and T. Wu (2009): “Subsidizing Creativity through Network Design: Zero Pricing and Net Neutrality,” *Journal of Economic Perspectives*, 23: 61–76.
- 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,” *Review of Network Economics*, 8: 22–39
- O'Sullivan, Arthur; Sheffrin, Steven M. and Perez, Stephen J. (2010), Microeconomics., 6th ed., Upper Saddle River, New Jersey: Pearson Education, Inc.
- Rosen, Harvey S. (2001), Public Finance. : New York: McGraw-Hill/Irwin.
- Rochet, J.-C. and J. Tirole (2006): “Two-Sided Markets: A Progress Report,” *RAND Journal of Economics*, 37: 645–667.
- Yoo, Christopher S. (2006), Network Neutrality and the Economics of Congestion. *Georgetown Law Journal*, 94: 05-28; Vanderbilt Public Law Research Paper No. 05-33. Available at SSRN: <http://ssrn.com/abstract=825669>