Peer-to-Peer File Sharing and the Market for Digital Information Goods*

Ramon Casadesus-Masanell† Andres Hervas-Drane‡

May 15, 2007

Abstract

We study competitive interaction between two alternative models of digital content distribution over the Internet: peer-to-peer (p2p) file sharing and centralized client-server distribution. We present microfoundations for a stylized model of p2p file sharing where all peers are endowed with standard preferences and show that the endogenous structure of the network is conducive to sharing by a significant number of peers, even if sharing is costlier than freeriding. We build on this model of p2p to analyze the optimal strategy of a profit-maximizing firm, such as Apple, that offers content available at positive prices. We characterize the size of the p2p network as a function of the firm's pricing strategy, and show that the firm may be better off setting high prices, allowing the network to survive, and that the p2p network may work more efficiently in the presence of the firm than in its absence.

*We thank Alexei Alexandrov, Jernej Copic, Albert Creus, Ig Horstmann, Jeff MacKie-Mason, Pai-Ling Yin, and seminar participants at the Bosphorus Workshop on Economics Design, the CoCombine conference, London Business School, Harvard Business School, MIT, the CRES-Second Annual Foundations of Business Strategy Conference, the 2006 NetEcon workshop, the 2006 NYU Summer Workshop on the Economics of Information Technology, the 5th Annual International Industrial Organization Conference and the Harvard EconCS Group seminar. We are grateful for financial support from the HBS Division of Research, Fundación BBVA, the Public-Private Sector Research Center at IESE Business School, the Real Colegio Complutense, and the FPU program of the Spanish Ministry of Education and Science.

†Harvard Business School (casadesus@gmail.com)

‡Harvard Business School and Universitat Autònoma de Barcelona (andresonline@gmail.com)

1 Introduction

Since the inception of copyright law to grant intellectual property owners a temporal monopoly on their works, the ability to capture value by copyright holders has persistently been threatened by unauthorized reproduction of content. Technological innovations have not only presented new market opportunities but also new threats. The radio, the cassette player, the video recorder, and the compact disc have allowed the industry to deliver additional value and meet new demand. But backed by fair use provisions, these same technologies have also been employed to replicate and distribute content without the consent of copyright holders.

In recent years, advances in the digitalization of content paired with the widespread adoption of broadband Internet have shaped a new and formidable threat with the emergence of peer-to-peer (p2p) file sharing networks. Peer-to-peer file sharing has grown spectacularly in the last few years. The content industry has reacted, with limited success, by legally confronting the p2p phenomenon and slowly embracing online distribution. Apple's iTunes Store, built on a traditional client-server architecture, has emerged as the dominant player in the market for legal digital downloads.

Peer-to-peer and licensed online stores constitute two fundamentally different distribution models that 'compete' against each other. Demanders of digital content are faced with the choice of whether to download content from p2p file sharing networks or from legal sites. The ability (or even the desire) of Apple to sustain high prices for downloads is affected by the presence of p2p file sharing networks. Likewise, the success of p2p file sharing is partly determined by actions taken by Apple and the majors such as pricing per download, the proneness to embark into legal action against users of p2p, their relationships with and demands from Internet Service Providers (ISPs), and the like.

In this paper we present a simple formal model to investigate how these two systems of digital distribution interact. Our model begins with the observation that peers in p2p networks face a fundamental choice between sharing content and freeriding. Sharing entails additional costs for the peer: committing computing resources such as storage space and upload bandwidth to the network and increasing the likelihood of legal action against her. When a peer in a p2p network decides to share content, she is effectively supplying two different goods. On the one hand, she provides *content*. Obviously, the peer who shares does not benefit from the content that she is sharing as she already owns it. In the absence of social preferences (i.e., altruism, reciprocity...), providing content has no direct benefit to the peer who shares. On the other hand, by sharing, a peer also supplies *upload bandwidth* to the network and this may result in lower network congestion. Sharing results in lower congestion if upload bandwidth is a scarce resource. Based on the available empirical evidence, we assume this to be the case.[1]

Similarly to content, bandwidth is a nonexcludable good. When a peer provides upload bandwidth to the network she cannot decide who will and who will not have access to that bandwidth. But contrary to content, bandwidth is a rival good because its use by one peer prevents use by another peer. We show that the nature of p2p networks, however, warrants that the provision of bandwidth benefits all peers equally in expected terms. Indeed, when a peer decides to share content, the average number of peers connected to sharers decreases (because there are more peers to connect to) and this reduces average network congestion. In sum, peers face a trade off: by sharing they bear costs that could be avoided by freeriding, but sharing also reduces average network congestion and this benefits every peer, including the peer who shares.

Building on this insight, we construct a model where peers provide bandwidth in addition to content when they share. Specifically, we consider a finite population of agents that derive positive and homogenous utility from digital content. Peers suffer disutility from the costs associated with downloading content. These costs are proportional to the time required to complete downloads, the level of *congestion*, which in turn depends on the bandwidth provision available in the network. Peers may reduce their expected congestion by providing upload bandwidth to other peers. We model this decision as a binary choice: share content or freeride. By allowing agents to differ in their disutility of congestion (or impatience) we show that an endogenous level of sharing emerges in the network. This depends on the size of the network, the costs faced by agents, and the disutility of congestion of the population. Selfish utility-maximizing peers are better off sharing because by doing so they face less congestion. We fully characterize the congestion faced and the utility enjoyed by all participants.

[1] In fact, the content industry has begun experimenting with payment schemes that reward peers for the supply of upload bandwidth in licensed p2p distribution networks. See, for example, Peer Impact: http://www.peerimpact.com. The willingeness to reward peers for the supply of upload bandwidth is evidence that this is indeed a scarce resource.

We build on this framework to analyze the optimal strategy of a profit-maximizing firm that offers the same content available on the network. Contrary to p2p networks, online stores offer fast downloads on a traditional client-server architecture and sell content at positive prices. We derive the demand function faced by the firm and characterize its optimal pricing strategy. We show that the firm may be better off setting high prices instead of attempting to shut the p2p network down by setting low prices. Moreover, we show that the p2p network may function more efficiently in the presence of the firm than in its absence.

The model captures important stylized facts identified by the literature. First, Asvanund et al. [3] show that congestion worsens as the size of peer-to-peer networks grows. Our model endogenously generates this result. In fact, the effect of network size on congestion helps explain the coexistence of multiple different p2p networks. The model can also accommodate positive network effects when users value content variety. Second, many studies have shown that heavy users of p2p file sharing networks are more prone to purchase content online.[2] Our framework not only suggests that there is no contradiction in this observed behavior, but also sheds light on the factors that explain the demand for online content in the presence of a p2p network. Third, we provide insights on content pricing and the effectiveness of industry initiatives such as suing heavy sharers. Finally, researchers and industry analysts have long questioned the existence of applications that drive broadband demand ("killer apps").[3] Our model shows that file sharing networks strictly benefit from improvements in broadband capacity, creating value for all participants. A study performed by Internet research firm CacheLogic in 2005 revealed that over 60% of total Internet traffic belonged to p2p file sharing applications.[4] This suggests that file sharing is indeed a driver for broadband demand and helps explain why Internet service providers have not taken action to limit the spread of p2p applications and file sharing traffic load. We believe that our results should be of interest to all participants in markets for digital information goods.

[2] See, for example, 'Downloading myths challenged,' BBC News.com, July 27, 2005.
[3] See Crandall and Alleman [10].
[4] See Parker [20]. Modem transmission speeds are considered too low for effective use of file sharing applications. The disutility of congestion given the long download delays over modem connections may be too high.

1.1 Literature

This paper contributes to an emerging literature in Strategy that explores competitive interaction between organizations with different business models. While there are several formal models of asymmetric competition that exist in Strategy (differences in costs, resources endowments or information, mainly), the asymmetries that this literature wrestles with are of a different nature: firms with fundamentally different objective functions, opposed approaches to competing, or different governance structures. Casadesus-Masanell and Ghemawat [6], for example, introduce a dynamic mixed duopoly model in which a profit-maximizing competitor interacts with an open source competitor that prices at zero with the installed base affecting their relative values over time and Casadesus-Masanell and Yoffie [9] study competitive interactions between two complementors, Microsoft and Intel, with asymmetries in their objectives functions stemming from technology – software vs. hardware.

Interest in the study of competitive interactions between organizations with different business models has increased in the last few years as new technologies, regulatory changes, and new customer demands have allowed firms to implement new approaches to competing in a wide range of industries spanning from airlines (Ryanair) to furniture (Ikea) and from circus (Cirque du Soleil) to software (open source projects). In fact, many of the fastest-growing firms in the recent past appear to have taken advantage of opportunities sparked by globalization, deregulation, or technological change to 'compete differently' and to innovate in their business models.

To assess the sustainability of competitive advantage of new business models it is critical to understand how they interact with those of other players. So far, the literature has studied interactions between new and traditional business models. This paper studies the competitive interactions between two new business models: p2p file sharing and profit-maximizing client-server digital distribution.

Our model of p2p (Sections 4 and 5) also contributes to a growing literature on the economics of p2p. This literature asks why individuals share files in p2p networks. Because contributing files is costlier than freeriding, selfish utility maximizers 'should' freeride and freeriding should lead to the collapse of p2p; however, p2p is thriving. To solve this puzzle, the literature has, for the most part, assumed that individuals are concerned with each others' wellbeing. Altruistic agents, for example, realize a direct benefit from contributing content. Golle et al. [14] and Antoniadis et al. [2] consider agents that derive utility from contributing content to the network. Feldman et al. [13] explicitly consider agent types which differ in their willingness to contribute. Cunningham et al. [12] assume reciprocity by a positive fraction of users, where increased sharing by some ensures a further increase in the overall provision. More recently, Jian and MacKie-Mason [17] present a model of generalized reciprocity where peers share expecting others in the network to indirectly reciprocate. While social preferences underlie many aspects of human behavior and may indeed play a role in p2p file sharing, we question the adequacy of models built on the assumption that peers care about each others' utility in a setting where millions of individuals interact anonymously.

Closer to our approach, Asvanund et al. [3] analyze p2p traffic data and build a stylized model around the notion of congestion and Krishnan et al. [18] propose some theoretical implications of congestion to explain sharing. To the best of our knowledge, there is no earlier model of p2p file sharing that models congestion endogenously and considers peers concerned solely about the impact of their actions on their own utility. Moreover, the literature on the economics of p2p has not considered interactions between p2p and client-server distribution models.

The paper is organized as follows. Section 2 describes the phenomenon that we study. Section 3 introduces the building blocks of our model of peer-to-peer file sharing. In Section 4 we present a simple approximation to the average congestion in an arbitrary peer-to-peer network. Section 5 derives the equilibrium network configuration and studies its properties. In Section 6 we introduce a profit maximizing firm that competes for users against the p2p network and analyze interdependencies that arise between both models of digital distribution. All the proofs are in the appendix.

2 Peer-to-Peer vs. iTunes[5]

The technology enabling peer-to-peer networks became mainstream in 1999 with the release of a music file sharing application called Napster. Contrary to a client-server model, in which all communication takes place through a central server, peer-to-peer architectures allow every computer

[5]This section draws from Casadesus-Masanell, Hervas-Drane, and Mitchell [8]. See also Casadesus-Masanell, Hervas-Drane, and Sean Silverthorne [7].

to directly communicate with others pertaining to the same network (running the same software) without having to go through intermediaries. This network topology increases scalability and robustness for a wide range of applications and is an active area of research.[6] File sharing has so far been the most disruptive and widespread application of p2p architectures.

A file sharing network allows participants to offer digital content for download to other connected peers, enabling content exchange to take place on a large scale. Because the value of file sharing networks depends on individual contributions and a high proportion of users consume network resources without contributing their own, congestion is one main problem in p2p. Depending on the resources available, downloading a full music album can take anywhere from minutes or hours to several days, or not complete at all in some cases.

Peer-to-peer file sharing and iTunes constitute two different paradigms for digital content distribution over the Internet. Contrary to p2p networks, Apple and its partners appear to be motivated by profit maximization. iTunes offers downloads on a traditional client-server architecture at positive prices, $0.99/song in the U.S. The client-server architecture allows Apple to manage congestion. A full music album can be downloaded from iTunes over a broadband connection in less than three minutes, and a song in just a few seconds.

Content distributed through p2p networks has several advantages over licensed content distributed through iTunes. Digital Rights Management (DRM) restrictions render licensed content an inferior good compared to unlicensed content; only the latter can be played, transferred, or replicated on any multimedia device without limitations. Furthermore, DRM restrictions may compromise future playback compatibility. For example, Apple restricts to five the number of PCs (or Macs) where a user can play content purchased on iTunes, and only on iTunes software. Digital encoding quality makes for an ambiguous case between both systems, and it is not infrequent for the same content to be available at higher quality in p2p networks.

But peer-to-peer distribution also has disadvantages. iTunes offers not only faster content download but it is also well-integrated with the iPod. In addition, iTunes metadata (file naming and tagging) is superior to that of files distributed through p2p. Moreover, iTunes is legal while

[6] Applications of p2p networks include: file sharing (Napster, BitTorrent and eMule), distributed computing (SETI@home and Folding@home) and voice over IP (Skype). These applications all create an overlay network over the host network (generally the Internet). For a detailed review of the technology see Schollmeier [21].

users of p2p networks are open to legal action against them. The following table summarizes the main dimensions on which p2p and iTunes differ:

p2p	**iTunes**
Free content	Pay for content
No DRM restrictions	DRM restrictions
Bad metadata	Good metadata
Congestion (slow downloads)	No congestion (fast downloads)
More difficult to use	Easy to use
Mostly illegal	Legal
Heavy use of upload bandwidth (sharers)	Little use of upload bandwidth

Both models have grown rapidly in the past few years. As of 2006, it is estimated that over 10 million users participate in p2p file sharing networks worldwide at any given instant. BigChampagne.com found that over 90% of the content exchanged was copyrighted. And according to CacheLogic.com, over 60% of Internet traffic in Europe and the U.S. was accounted for by p2p file sharing. In a 2005 survey, half of the experts consulted believed that file sharing on p2p networks will still be easy even a decade from now.[7] Apple, on the other hand, in 2006 claimed more than 80% share of U.S. legal music downloads.[8] And according to the Recording Industry Association of America (RIAA), in 2005 iTunes outsold several large traditional music retailers such as Tower Records, Borders, and Sam Goody.[9] Presently, iTunes offers more than 3 million songs for download and sells over 2,000 songs per minute, on average.

Copious amounts of airtime has been given to p2p's effect on music industry players. Some industry participants feel that p2p file sharing is destroying the industry. The RIAA, for example, claims that piracy cost the industry $4.2 billion each year on a worldwide basis.[10] The RIAA has been at the centre of the battle against p2p networks, first starting with lawsuits against MP3.com in 1997, then Diamond Media in 1998 for their MP3 player and then Napster in 1999. Lawsuits against other p2p networks have continued (Kazaa, Morpheus and Grokster in 2001) but to little avail. In 2003, the RIAA launched lawsuits against 261 individuals and over 2,000 other individuals in the following two years. Meanwhile, numerous consumer groups and online activists, such as the Electronic Frontier Foundation, feel that p2p represents unmatched opportunities the industry has failed to understand.

[7] See Pew Internet & American life project 2005 survey on 'The future of the Internet.'
[8] See 'France Poised To Soften Controversial iTunes Bill,' CNNMoney.com, June 21, 2006.
[9] See 'iTunes outsells traditional music stores,' CNET News.com, November 21, 2005.
[10] RIAA website, www.riaa.com.

3 The model

Consider a population of M agents that obtain utility from the consumption of digital information goods. They all value content equally and differ only in their disutility of congestion. We model the formation of a peer-to-peer file sharing network as a two stage process. In the first stage, agents choose (simultaneously) whether or not to join the network. Agents who choose to belong to the network can either share their content or freeride. Sharers offer content for download by other peers while freeriders do not. While sharing content is costly, some sharing is required for the network to survive as downloads can only be realized from other peers. We will refer to agents in the network as *peers* and those outside as *outsiders*. We let $N \leq M$ denote the number of peers. $M - N$ is the number of outsiders. In the second stage peers interconnect and downloads are realized. The utility of a peer that freerides is given by

$$u_i^f = u_d - (c_n + \rho_i)t_d, \tag{1}$$

and that of a peer who shares content by

$$u_i^s = u_d - (c_n + c_s + \rho_i)t_d, \tag{2}$$

where $i \in \mathbf{N} = \{1, 2, ..., N\}$. Outside utility is normalized to zero.

The utility derived from content once a download has been completed is u_d. We assume that u_d is common across all content and peers. This simplifies the analysis and lets us focus on the role of congestion as a determinant of peers' willingness to contribute content (and bandwidth). The assumption amounts to stating that content is not a scarce resource. As will become clear shortly, the scarce resource in our model is bandwidth.[11]

The time required to complete a download, t_d, is endogenous and depends on the level of congestion which, in turn, depends on how many peers share content. Lower bandwidth provision implies higher levels of congestion resulting in higher download time.

[11] If every sharer contributes sufficiently many files, a peer will always find some content they value at u_d when downloading from any given sharer. The specific content may differ depending on the sharer, though. See Appendix 2 for further details on content availability in p2p networks.

Every peer suffers a positive cost c_n of using the p2p network. This captures the opportunity costs of computing resources employed, the bandwidth for signalling traffic required to remain connected to the network until a download completes, and possible legal action against the peer. In addition to c_n, sharers (but not freeriders) bear cost c_s. This is the cost of offering content for download on a public p2p network, including the use of additional computing resources such as storage space and upload bandwidth and the increased likelihood of legal action against sharers (over and above that faced by freeriders).

Parameter $\rho_i \geq 0$ reflects the disutility of congestion experienced by peer i. The larger ρ_i is, the higher the disutility from an increase in the time required to complete a download. Hence ρ_i can be interpreted as impatience or how much peer i values quick access to content. Without loss of generality we choose indexes i so that $\rho_i \leq \rho_{i+1}$ for all i. All other costs being equal, peers prefer to obtain content immediately avoiding congestion delays.

In the second stage, after peers have decided whether to share or to freeride, interconnections take place. Let $\mathbf{S} \subset \mathbf{N}$ be the set of sharers (given the agents' first-stage strategies) and denote by S the number of sharers (the cardinality of $\mathbf{S}$). We assume all peers have an upload bandwidth capacity of $1/\theta$, where $\theta > 0$. Because delay (congestion) is measured by the inverse of bandwidth, this implies that a downloader *exclusively served by a sharer* downloads a unit of content in θ units of time. That is, we normalize to $t_d = \theta$ the time required for a download when a peer is served by a sharer that receives no other incoming connections. Parameter θ captures the residential capacity offered by the broadband infrastructure; the relation between file size and bandwidth available to peers. An improvement in broadband infrastructure that increases available bandwidth, decreases download times and amounts to a reduction in θ.[12] Download bandwidth capacity of peers is assumed not to be a limiting factor. If more than one downloader are connected to a given sharer, upload bandwidth is shared evenly amongst them. This can be interpreted as downloading taking place simultaneously or, alternatively, the sharer serving download queues for fractions of content by turns.

A set of links connecting peers to sharers where every peer connects to one sharer only and

[12] An improvement in encoding efficiency reducing file sizes has the same effect. In practice these improvements tend to be modest in comparison to changes in broadband infrastructure. We focus our discussion below on the latter.

no sharer connects to herself is called a *network allocation*. A *stable* network allocation is one where no peer can be made strictly better off by connecting to a different sharer. We assume that following the first stage, a stable network allocation ensues. If a social planner were to assign peers to sharers to distribute bandwidth as equitably as possible, only stable allocations would be considered.[13] Similarly, if peers were given the choice to update their link in a random sequence, the resulting network allocations would also be stable. Clearly, if the network allocation was not stable, at least one peer would have incentives to update her link and connect to a different sharer. We assume for simplicity that all stable allocations are equiprobable.

To summarize, our model of p2p assumes:

- All peers have $1/\theta$ units of upload bandwidth capacity;
- All peers have at least $1/\theta$ units of download bandwidth capacity;
- Every peer connects to one sharer only;
- A sharer may not connect to herself;
- Upload bandwidth is allocated equably amongst all peers connected to a sharer;
- Second-stage network allocations are stable and equiprobable.

Finally, the following mild assumption on the parameters is required for the results: $u_d > (c_n + c_s + \rho_i)\theta$ for all i. This ensures that a p2p network with minimum congestion is always preferred to the outside option of not pertaining to the network.

With the notation in place, we proceed to solving the model. We start by considering a fixed number of sharers (S) and network size (N) and present a simple approximation to average congestion. In Section 5 we endogenize the sharing decision, and in Section 6 we endogenize the size of the network by allowing agents to not consume content or download from the firm instead of using the p2p network.

[13] The decision rule followed by such a social planner can be implemented by a centralized algorithm that assigns links of all peers.

4 Network foundation

In this section we present an approximation to the average congestion in an arbitrary peer-to-peer network with an exogenous number of sharers and freeriders. This provides a foundation for download time t_d in the second stage of our model, a central variable of our analysis. Congestion plays a crucial role in our development as peers choose to share taking into consideration the effect that their sharing has on congestion. In Section 5 we analyze the first stage and endogenize the number of sharers and freeriders. This Section follows our work with Albert Creus on bandwidth allocation in p2p file sharing networks (see Creus et al. [11]).

To simplify the exposition, suppose initially that $\theta = 1$. Thus all peers have one unit of upload bandwidth capacity. Given a network allocation, the bandwidth obtained by peer $i \in \mathbf{N}$ can be computed as follows: if the peer is connected to a sharer to which k other peers are connected to, then peer i obtains effective bandwidth $1/(k+1)$. Freeriders are allowed to connect to every sharer. Therefore, they have S possible links available to choose from. Sharers, on the other hand, cannot connect to themselves. As a consequence, sharers have $S-1$ possible links available. This implies that, in general, the bandwidth obtained by both groups of peers will differ. To compute the *expected bandwidth* for freeriders and sharers in a network with N peers and S sharers, we begin by computing each peer's effective bandwidth in every stable network allocation. We then average these effective bandwidths assuming that every stable network allocation is equally likely. The following example illustrates our approach.

Example 1 *$N = 5$ and $\mathbf{S} = \{S_1, S_2\}$. There are three freeriding peers: F_1, F_2, and F_3. In this example there are exactly six stable network allocations.*

- *Stable network allocation 1: $S_1 \to S_2$ (this means that S_1 connects to and downloads from S_2), $S_2 \to S_1$, $F_1 \to S_2$, $F_2 \to S_1$, and $F_3 \to S_1$. No peer can be made better off by changing her connection only.*
    - *Effective bandwidths (resp.): $\frac{1}{2}$, $\frac{1}{3}$, $\frac{1}{2}$, $\frac{1}{3}$, and $\frac{1}{3}$.*
- *Stable network allocation 2: $S_1 \to S_2$, $S_2 \to S_1$, $F_1 \to S_1$, $F_2 \to S_2$, and $F_3 \to S_1$.*

– *Effective bandwidths (resp.):* $\frac{1}{2}$, $\frac{1}{3}$, $\frac{1}{3}$, $\frac{1}{2}$, *and* $\frac{1}{3}$.

- *Stable network allocation 3:* $S_1 \to S_2$, $S_2 \to S_1$, $F_1 \to S_1$, $F_2 \to S_1$, *and* $F_3 \to S_2$.

  – *Effective bandwidths (resp.):* $\frac{1}{2}$, $\frac{1}{3}$, $\frac{1}{3}$, $\frac{1}{3}$, *and* $\frac{1}{2}$.

- *Stable network allocation 4:* $S_1 \to S_2$, $S_2 \to S_1$, $F_1 \to S_1$, $F_2 \to S_2$, *and* $F_3 \to S_2$.

  – *Effective bandwidths (resp.):* $\frac{1}{3}$, $\frac{1}{2}$, $\frac{1}{2}$, $\frac{1}{3}$, *and* $\frac{1}{3}$.

- *Stable network allocation 5:* $S_1 \to S_2$, $S_2 \to S_1$, $F_1 \to S_2$, $F_2 \to S_1$, *and* $F_3 \to S_2$.

  – *Effective bandwidths (resp.):* $\frac{1}{3}$, $\frac{1}{2}$, $\frac{1}{3}$, $\frac{1}{2}$, *and* $\frac{1}{3}$.

- *Stable network allocation 6:* $S_1 \to S_2$, $S_2 \to S_1$, $F_1 \to S_2$, $F_2 \to S_2$, *and* $F_3 \to S_1$.

  – *Effective bandwidths (resp.):* $\frac{1}{3}$, $\frac{1}{2}$, $\frac{1}{3}$, $\frac{1}{3}$, *and* $\frac{1}{2}$.

The expected bandwidth of sharers is $\frac{1}{6}\left(\frac{1}{2}+\frac{1}{2}+\frac{1}{2}+\frac{1}{3}+\frac{1}{3}+\frac{1}{3}\right) = \frac{5}{12} \approx 0.417$. *The expected bandwidth of freeriders is* $\frac{7}{18} \approx 0.389$. *On average, sharers face less congestion than freeriders.*

The computational complexity of the problem increases with the number of stable allocations, which grows rapidly with N. In [11] we derive an exact expression for the expected bandwidths of both types of peers. Unfortunately, the exact formula is far too complex to be used in applied models. However, in that paper we show that S/N is a good approximation to the expected bandwidth of both sharers and freeriders. Moreover, we show that the expected bandwidth of sharers is always greater than or equal to S/N and that that of freeriders is always less than or equal to S/N.[14]

It is interesting that sharers obtain a slightly higher bandwidth than freeriders even if they face more constrains than freeriders (as they cannot connect to themselves). As N grows the difference between expected bandwidths and S/N decreases. In fact, already in a network of size $N = 10$, the

[14] It is easy to see that S/N is the 'dividing line.' Let B_u be the total upload bandwidth of all sharers. Let b_s and b_f be the expected bandwidth of sharers and freeriders respectively. Then, $B_u = b_s S + b_f F$. Notice that $S = B_u$. Therefore, $S = b_s S + b_f F$. Rearranging, we have that $b_f = S\,(1 - b_s)\,/F$. Recall that $F = N - S$. Therefore we have $b_f = S\,(1 - b_s)\,/\,(N - S)$. Dividing numerator and denominator by N we obtain $b_f = \frac{S}{N}\left(\frac{1-b_s}{1-\frac{S}{N}}\right)$. Clearly, $b_s > S/N$ implies that $b_f < S/N$.

expected bandwidth of sharers and freeriders differs from S/N by, at most, 0.0012.[15] And when $N = 100$ the difference is always less than 0.0000064. Moreover, when N is a multiple of S, the expected bandwidth of sharers and freeriders coincides and it is equal to S/N.

To sum up, all peers obtain expected bandwidth close to S/N and we take this approximation as the measure of expected bandwidth for all players. For an arbitrary value of θ, expected bandwidth can be expressed as $S/\theta N$. This allows us to easily compute *expected congestion* in a network with N peers and S sharers, which is the time required to complete downloads t_d, as the inverse of expected bandwidth. Hence, expected congestion for all peers is given by $t_d = 1/\frac{S}{\theta N} = \theta \frac{N}{S}$. It should be noted that although bandwidth depends linearly in the number of sharers, expected congestion does not. This property is crucial to our results. Technically, it ensures that our objective function is concave in S, allowing for interior equilibria in which sharing and freeriding may coexist for certain ranges of N.

Notice finally that our approximation punishes (slightly) sharers and rewards (slightly) freeriders. Therefore, we are making it 'harder' for peers to share. In Section 5 we endogenize the decision to share or to freeride and show that in equilibrium there is sharing. If instead of using the approximation we used the exact formula in [11], our results would only be strengthened.

5 Equilibrium network configurations

In this Section we analyze the first stage. This is the stage where every peer chooses whether to freeride or to share (at additional cost c_s). In other words, we now endogenize t_d. In making their decision, peers consider the effect of their choice on expected download time $\theta \frac{N}{S}$. Equations (1) and (2) imply that if expected download time was *not* affected by the sharing decision, no peer would ever share and the peer-to-peer network would not be viable.

In this Section we take N as given. This amounts to assuming that all peers in the network obtain positive utility. In general, this will depend on S and the distribution of ρs. We relax this assumption in Section 6 and let peers decide whether or not to join the network.

Let $P = \{\mathbf{F}, \mathbf{S}\}$ be a partition of $\mathbf{N}$. We refer to P as a *network configuration*.[16] $\mathbf{F}$ is the set

[15] Given N the expected congestion of sharers and freeriders will change as the cardinality of $\mathbf{S}$ varies.

[16] Notice that a network configuration can be mapped to many different network allocations.

of freeriders and $\mathbf{S}$ the set of sharers. Obviously, P constitutes an equilibrium if no $i \in \mathbf{S}$ prefers to (unilaterally) become a freerider and no $j \in \mathbf{F}$ prefers to become a sharer. Sharer i will not free ride if

$$u_d - (c_n + c_s + \rho_i)\,\theta\frac{N}{S} \geq u_d - (c_n + \rho_i)\,\theta\frac{N}{S-1}. \tag{3}$$

On the other hand, freerider j will not want to become a sharer if

$$u_d - \left(c_n + \rho_j\right)\theta\frac{N}{S} \geq u_d - \left(c_s + \rho_j\right)\theta\frac{N}{S+1}. \tag{4}$$

The following proposition characterizes the equilibrium.

Proposition 2 *Every equilibrium network configuration* $P = \{\mathbf{F}, \mathbf{S}\}$ *has the following form:* $\mathbf{F} = \{1, 2, ..., n-1\}$ *and* $\mathbf{S} = \{n, n+1, ..., N\}$ *for some* $n \in \mathbf{N}$.

The proposition says that if peer i is a sharer in equilibrium network configuration P, then peer $i+1$ must also be a sharer. Moreover, if peer j is a freerider, then peer $j-1$ must also be a freerider. Thus, the most impatient peers prefer to share while the more patient peers are better off freeriding. The reason is simple: by sharing content, peers reduce congestion and the (positive) marginal effect on peer utility implied by lower congestion is proportional to the value of ρ_i. Peers for whom the opportunity cost of time is high, are more inclined to share. This is true even though given any *fixed* level of congestion, all peers (regardless of the value of ρ) are better off freeriding than sharing.

We now further characterize the equilibrium network configurations by pinning down to the fullest possible extent the cardinality of S. Let $P = \{\mathbf{F}, \mathbf{S}\}$ be an equilibrium network configuration. Let ρ_i be the most patient sharer in $\mathbf{S}$. Equations (3) and (4) imply that

$$S \leq \frac{c_n + c_s + \rho_i}{c_s} \qquad \text{and} \qquad S \geq \frac{c_n + \rho_{i-1}}{c_s}.$$

Thus,

$$\frac{c_n + \rho_{i-1}}{c_s} \leq S \leq \frac{c_n + c_s + \rho_i}{c_s}. \tag{5}$$

Let $\mathbf{I}$ be the set of integers. The following two objects are useful in what follows:

$$G(\rho_i) = \left\{ k \in \mathbf{I} \left| \frac{c_n + \rho_{i-1}}{c_s} \leq k \leq \frac{c_n + c_s + \rho_i}{c_s} \right. \right\} \tag{6}$$

and

$$H(\rho_i) = N + 1 - i. \tag{7}$$

Correspondence G indicates the cardinality of $\mathbf{S}$ if the sharer with lowest impatience has time preference ρ_i. Function H tells us the number of peers with parameter ρ_j larger than or equal to that of peer i.

The solution to the system of equations given by G and H pins down the most patient sharer:

$$\Gamma_s = \{ i \in \mathbf{I} \, | H(\rho_i) \subset G(\rho_i) \} .$$

Because $G(\rho_i)$ is a correspondence, Γ_s may not be a singleton set. The following example illustrates this approach.

Example 3 *Assume there are* 22 *peers with time preference parameters* ρ_i *(*$i = 1..22$*) = 1, 3, 4, 5, 8, 9, 10, 12, 14, 16, 17, 18, 21, 22, 26, 27, 29, 30, 33, 34, 35, and 38. If* $c_n = 1$ *and* $c_s = 2$ *we have that* $\Gamma_s = \{13\}$*. Thus, there is one equilibrium network configuration:* $\mathbf{S} = \{13, ..., 22\}$ *(10 peers share and 12 freeride).*

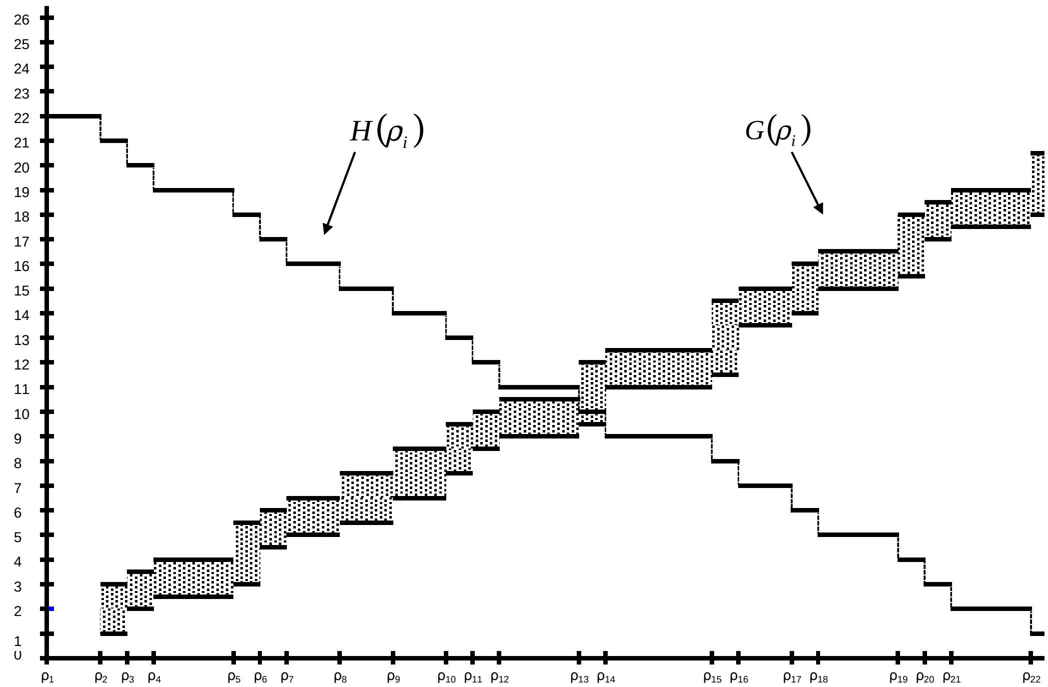

Notice that $H(\rho_i)$ does not change with the values of the parameters. Therefore, to perform comparative statics we need only look at how changes in the parameters affect the position and 'slope' of $G(\rho_i)$. When c_s falls or c_n grows, $G(\rho_i)$ shifts upwards, bringing down the most patient sharer and thus increasing sharing in the network.

The difference between the effect of c_s and c_n on congestion is as follows. When c_s increases, sharers are worse off because they bear additional cost. Freeriding now becomes more attractive. Sharers with the lowest time preference parameters ρ_i will prefer to freeride and congestion increases. When c_n increases, however, both sharers and freeriders bear additional cost. In this case, sharing becomes more attractive. Sharers do not gain from becoming freeriders as freeriders also bear c_n. Freeriders can reduce (somewhat) the negative effect of c_n on their utility by becoming sharers and thus reducing t_d. Not all freeriders will find it advantageous to become sharers but those with the largest time preference parameters ρ_i will. This is why sharing increases and congestion falls when c_n grows.

5.1 Full vs. partial sharing

Let $\mathbf{S} = \{n, n+1, ..., N\}$ be the set of sharers in an equilibrium network configuration. We refer to the case $n = 1$ as a *full-sharing network configuration* (or *full-sharing equilibrium*) and to the case $n > 1$ as a *partial-sharing network configuration* (or *partial-sharing equilibrium*). In a full-sharing network configuration all peers are sharers. In this case, congestion is minimized as the expected download time for all peers (t_d) is equal to θ. For a full-sharing network configuration to obtain, every peer must realize higher utility by sharing than by freeriding. In particular, the most patient peer (ρ_1) must be better off sharing than freeriding (given that everybody shares):

$$u_d - (c_n + c_s + \rho_1)\,\theta \geq u_d - (c_n + \rho_1)\,\theta \frac{N}{N-1}.$$

Solving for N we obtain:

$$N \leq \frac{c_n + c_s + \rho_1}{c_s}. \tag{8}$$

Therefore, if N is sufficiently small, the unique equilibrium network configuration has all peers sharing content. Notice that as the incremental cost of sharing c_s approaches zero, the maximal network size that supports full sharing grows without bound. Inequality (8) also reveals that when c_n is large, everybody shares. In this case, peers suffer more from congestion and are all better off sharing. Finally, if the most patient peer (ρ_1) is very impatient (ρ_1 large), then all peers prefer to share. When N is large, the equilibrium will typically entail partial sharing. In this case, expected download time will be larger than θ for all peers.

Although bandwidth is not a pure public good because it is rival, our model of endogenous congestion is related to work in public economics that studies the private provision of public goods. This literature proposes two approaches to modeling the decision to contribute: the self-interested approach of Bergstrom et al. [5] where individuals are concerned about the total supply of public goods, and the 'warm-glow' approach of Andreoni [1] where individuals 'feel good' when they privately provide public goods. Our approach is closer to the former. In fact, ours is a model of *negative* warm-glow (or *cold*-glow) because sharers bear additional costs (c_s) that can be avoided by freeriding. Compared to the self-interested approach our model is similar in that individuals care about the total amount of public good available and different in that the decision of how much to share is discrete (all or nothing). In this sense, ours is a model exclusively about the 'extensive margin' – the decision of whether or not to become a contributor – whereas Bergstrom et al. [5] considers both the extensive and the 'intensive margin' – the decision of how much to contribute.

5.2 Equilibrium with $\rho_i \sim U[0, \bar{\rho}]$

We have characterized the equilibrium for the general case, without specific assumptions on the distribution of ρ_is or the cardinality of $\mathbf{N}$. In order to ensure tractability when we introduce a profit maximizing firm (Section 6), we make the additional assumption that ρ_is are i.i.d. $U[0, \bar{\rho}]$. This allows us to further characterize the set of equilibrium network configurations.

Proposition 4 identifies the most patient sharer as a function of the parameters. Identifying precisely the most patient sharer will allow us to easily analyze how the different parameters affect network congestion. In particular, we are interested on the effect that N has on congestion. If congestion decreases when N grows, then the p2p network becomes gradually more valuable as the number of peers expands. If, on the contrary, network congestion grows with N, the value of the p2p network decreases with size.

Proposition 4 *Let $\rho_{s(N)}$ be the most patient sharer in equilibrium. Then, for N large,*

$$\rho_{s(N)} \simeq \frac{\bar{\rho}\left(\left(N-1\right)c_s - c_n\right)}{\bar{\rho} + Nc_s}.$$

Notice that $\rho_{s(N)}$ is increasing in N. This implies that the larger is the cardinality of $\mathbf{N}$, the lower is the proportion of sharers in equilibrium. In other words, the p2p network exhibits negative network effects (past the threshold network size of full sharing): the larger the number of peers, the lower the average utility that peers obtain. In fact, as $N \to \infty$, $\rho_{s(N)} \to \bar{\rho}$. Therefore, when the network is very large, there is essentially no sharing.

The following result is a direct implication of Proposition 4.

Corollary 5 *When N is large, the equilibrium cardinality of $\mathbf{S}$ is:*

$$S_N \simeq \frac{\bar{\rho} + c_s + c_n}{\frac{1}{N}\bar{\rho} + c_s}. \tag{9}$$

Notice first that $dS_N/dN > 0$: as N grows, the number of sharers does not decrease. However, $S_N \to \frac{\bar{\rho}+c_s+c_n}{c_s} < \infty$ as $N \to \infty$. Therefore, the proportion of sharers converges to zero as N grows. In other words, the expected download time ($\theta \frac{N}{S}$) grows without bound as N increases. This means that large p2p networks operate worse than smaller ones. The intuition lies on the effects of network size on the marginal utility of sharing. As the network grows, 'sharing' makes less of a difference on the reduction of congestion and the incentives to share fade. This problem is similar to the classic 'moral hazard in teams' (Holmstrom [16]) where team members contribute effort only to the extent that marginal benefit is larger than marginal cost. When the team grows, marginal benefit tends to decrease and freeriding grows. This effect is felt by all peers, but those with higher disutility of congestion suffer a larger decrease in utility when N grows. This has important implications for the equilibrium pricing strategy of a firm competing for customers against a p2p network (Section 6).

The costs of using the network (c_n and c_s) play a crucial role in determining congestion and thus the viability of p2p file sharing networks.[17] In fact, given a large $N < \infty$, congestion may be low if c_s is small and/or c_n is large. Differentiating N/S_N with respect to c_s we see that congestion worsens as the cost of sharing increases. Clearly, as c_s grows, less and less peers find it attractive to share. As a consequence, congestion increases and all peers in the network are worse off. Actions targeting an increase in c_s (such as suing sharers) reduce the attractiveness of the network. An increase in the cost of pertaining to the network achieves a similar effect. In this case, the derivative of congestion N/S_N with respect to c_n is negative. However, although there is less congestion when c_n grows, all peers wind up worse off as $du_i^f/dc_n < 0$ and $du_i^s/dc_n < 0$ except for $\rho_i = \bar{\rho}$ for whom $du_i^s/dc_n = 0$. That is, in the case of $\rho_i \sim U[0, \bar{\rho}]$ the positive effect of lower congestion is more than offset by the negative effect of larger cost c_n. Therefore, increases in c_s or c_n end up hurting peers.

Corollary 5 also helps explain why p2p 'industry structure' is characterized by the presence of multiple, independent file sharing networks. The model suggests that as network size grows and

[17] Anecdotal evidence suggests that users of file sharing networks experiment in setting the p2p software parameters that determine the bandwidth allocated to uploads. That is, the decision to share is, to some extent, continuous. The aggregate equilibrium upload bandwidth predicted by our model can also be interpreted as the steady state of a process where peers decide how much upload bandwidth to offer.

congestion worsens, peers are better off forming new networks with fewer peers (initially) and faster download speeds. The number of coexisting networks must then be a function of population size and the scalability of p2p technology. Although we do not pursue this research question here, it is worth pointing out that to study the equilibrium number of p2p networks and their sizes one would want to work with a model that captures the positive network effects resulting from having higher content variety in larger networks.[18]

The evolution of p2p technology has improved its scalability. As a consequence, there is a trend towards increased concentration: fewer, larger p2p networks in operation. One key factor driving this trend are incentive schemes that promote sharing to lower congestion. In essence, these mechanisms redistribute bandwidth away from freeriders towards peers who share. The effect of these incentives is similar to that of altruism or reciprocity between peers; in both cases the marginal utility of sharing is increased. Our model, however, is one where sharers and freeriders are all assumed to obtain the exact same bandwidth S/N (there is no redistribution). We have analyzed extensions where sharers are favored and freeriders punished. As expected, congestion is reduced and scalability improves. Unfortunately, a model with such extensions becomes intractable when we introduce the firm (Section 6). The power of our model of p2p is in showing that *even without such incentive mechanisms* there is sharing in equilibrium.[19]

It is easy to see that the utility of peers increases as residential broadband infrastructure improves (lower θ). The availability of larger bandwidth reduces congestion and allows downloads to complete faster. The model suggests that file sharing plays an important role as a driver of demand for broadband. The evidence seems to confirm this.[20] File sharing may very well be the 'killer app' that broadband has reportedly been missing. Recent data on Internet traffic shows that file sharing has continued to increase and currently accounts for more than 60% of total traffic. This

[18] Asvanund et al. [3] show that the variety of content available in p2p networks is increasing and concave in network size. Positive network effects can be captured in our model by assuming that peers value content variety, a function of network size: $u_d(N) = N^\alpha$ with $\alpha \in (0, 1)$. The tradeoff between content variety and congestion determines the optimal network size. We do not follow this approach as the model with content variety becomes intractable when we introduce a profit maximizing firm that competes against the network (Section 6).

[19] Napster, the first file sharing network deployed on a large scale, operated with no incentive schemes to promote sharing. Nonetheless, the network continued to operate up until its forced closedown with usage peaks of over one million simultaneous users. The Napster protocol was reverse-engineered after the shutdown and continued to operate in parallel on several smaller networks.

[20] See, for example, Om Malik, 'P2P, the only killer broadband application,' GigaOM.com, August 29, 2005.

phenomenon seems to affect data networks of operators worldwide, even as the capacity of broadband infrastructure available to residential users varies substantially from country to country.[21] ISPs report that a small number of residential customers are responsible for a big proportion of total traffic. Our model is also consistent with this observation because sharers (a small number in large networks) generate higher volumes of traffic than freeriders.

6 The firm

We next introduce an online firm that sells digital information goods also available on the peer-to-peer network. To the firm, the network is a competitor because peers that choose to download files from the network could otherwise become paying customers. In the absence of altruism towards artists, it is an interesting question why consumers would pay to purchase content online. Empirical studies have shown that preferences for content available on p2p networks are similar to those for products available in traditional distribution channels.[22] Furthermore, the variety of content available on major file sharing networks after many years of operation is still unmatched by licensed online stores. This suggests that preferences for content do not explain why individuals choose to purchase from online stores.

Because content is free on p2p networks, for the firm to persuade users of digital content to pay positive prices, it must offer added benefits that file sharing networks cannot match. In our view, the most important advantages of licensed online distribution are: (1) it offers lower download time and (2) it is legal. Online distributors such as Apple's iTunes Store offer content on a traditional client-server architecture. Contrary to the p2p network, consumption of content acquired from the firm can be enjoyed almost immediately at the moment of purchase. A song can be downloaded from iTunes over a broadband connection in a few seconds. By adding servers and bandwidth, Apple can manage congestion. Furthermore, the industry seems to be engaged in an assertive push towards content streaming.[23] To capture this fact, we let the firm offer immediate consumption.

[21] See the data published by CacheLogic, Parker [20], the traffic measures on Japanese ISPs performed by Nissho Electronics co., and the data on traffic of Spanish ISPs disclosed by Telefónica during the Internet Global Congress 2006.

[22] The media industry has recognized the value of the Internet as a tool to learn about consumer preferences and it is increasingly using the web to better estimate demand. See Bhattacharjee et al. [4].

[23] New players in this market such as Google and CinemaNow offer content streaming. This allows for real-time

In this case, the utility of buyers becomes:

$$u_i = u_d - p. \tag{10}$$

Notice that (10) is the natural adaptation of (1) and (2) to the case of immediate consumption, as the expected download time t_d falls down to zero. Thus, the terms $c_n + \rho_i$ and c_s do not appear in (10). On the other hand, the firm charges positive prices for content.

Because we are considering pure information goods, the firm has zero marginal cost. All infrastructure and running costs related to the service are fixed or sunk and independent of activity level. We also assume that the p2p network and the firm offer the same content. Individuals face the choice of purchasing from the firm or downloading content off the network for free. Additionally, individuals now have the outside option of not consuming content at all. The outside utility is normalized to zero.

We modify the timing accordingly. In the first stage, the firm chooses the price at which content is sold. In the second stage, individuals choose to either purchase from the firm, enter the network, or stay outside. Those who enter the network may share or freeride. In the third stage, agents in the network interconnect and downloads take place.

The following proposition characterizes the demand faced by the firm.

Proposition 6 *The firm faces the following demand function:*

$$q = \begin{cases} M & \textit{if} \quad p \leq \theta(c_n + c_s) & \textit{(full coverage)} \\ (1 - \frac{p - \theta(c_n + c_s)}{\theta\bar{\rho}})M & \textit{if} \quad \theta(c_n + c_s) < p \leq \frac{\theta(c_n+c_s)(\bar{\rho}+Mc_s)}{Mc_s} & \textit{(high coverage)} \\ \max[(1 - \frac{p}{\theta(\bar{\rho}+Mc_s)})M, 0] & \textit{if} \quad \frac{\theta(c_n+c_s)(\bar{\rho}+Mc_s)}{Mc_s} < p < u_d & \textit{(low coverage)} \\ \max[(1 - \frac{p}{\theta(\bar{\rho}+Mc_s)})M, 0] & \textit{if} \quad p = u_d & \textit{(outsiders only)} \\ 0 & \textit{if} \quad p > u_d & \textit{(no demand)} \end{cases}$$

As expected, demand is downward sloping. Moreover, individuals with high disutility of congestion prefer to buy from the firm than to obtain content from the network for free. These individuals benefit the most from fast downloads. As they choose to purchase, the network becomes smaller

viewing of content purchased. Streaming of audio and video content over broadband under client-server architectures is a proven, well-established technology. Reliable content streaming over peer-to-peer architectures, however, remains a theoretical construct known to present several technical complications. See Habib and Chuang [15] and Pai and Mohr [19].

and the proportion of peers who are better off sharing increases (see Section 5.2). As a consequence, congestion falls. Hence the size and efficiency of the network is affected by the presence of the firm. Proposition 6 shows that both models of digital distribution are interdependent.

The following figure shows the shape of the demand curve and illustrates how the size of the p2p network is determined by the price charged by the firm.

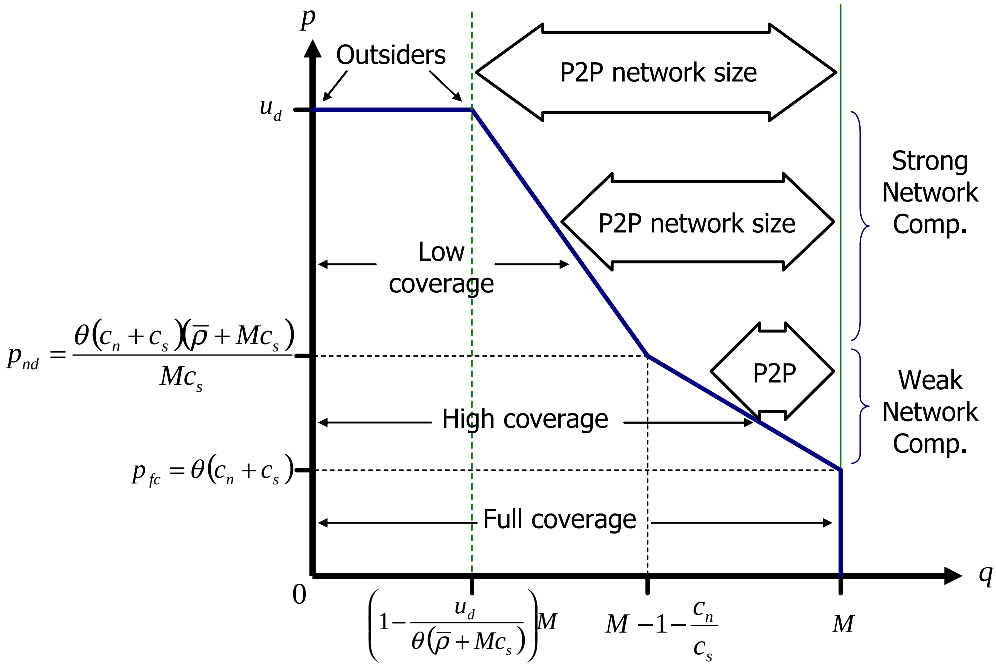

Full market coverage is obtained at any price below or equal p_{fc}. The agent who suffers congestion less ($\rho_i \simeq 0$), purchases at this price out of indifference (and all other agents prefer to purchase) rendering the network empty. Any price above p_{fc} will ensure that some agents prefer the network. Above this price, both distribution models coexist. Demand is characterized by a non-derivability at price p_{nd}. This price separates two ranges over which the behavior of congestion in the network differs. Below p_{nd} full sharing holds. In this case, congestion is not affected by peers entering or exiting to purchase. Above p_{nd} partial sharing holds. Here congestion varies with network size; the smaller the network, the lower the level of congestion. This is true although the population of peers that remain in the network are on average more patient and, thus, less prone to sharing; the effect induced by the reduction in network size is always stronger, resulting in a larger proportion of peers sharing content.[24] Finally, u_d is the maximum price that the firm will ever charge. Any

[24] Consider a reduction in price. In the full sharing price range, peers who switch to purchase are not affecting

higher price will be met with no demand.

Given the demand function derived in Proposition 6, the problem of the firm is to set p to maximize profits. The final proposition characterizes the firm's optimal pricing strategy.

Proposition 7 *Let*

$$p_{fc} := \theta(c_n + c_s) \qquad \text{(full market coverage)}$$
$$p_{hc} := \tfrac{1}{2}\theta(\bar{\rho} + c_n + c_s) \qquad \text{(high market coverage)}$$
$$p_{lc} := \tfrac{1}{2}\theta(\bar{\rho} + Mc_s) \qquad \text{(low market coverage)}$$
$$p_{oc} := u_d \qquad \text{(outsiders only coverage)}$$

The optimal pricing strategy is given by:

	$u_d \leq 2\theta(c_n + c_s)$	$u_d > 2\theta(c_n + c_s)$
$\bar{\rho} \leq c_n + c_s$	p_{fc} *if* $M < M_d$ p_{oc} *if* $M \geq M_d$	p_{fc} *if* $M < M_a$ p_{lc} *if* $M_a \leq M < M_c$ p_{oc} *if* $M \geq M_c$
$\bar{\rho} > c_n + c_s$	p_{hc} *if* $M < M_b$ p_{lc} *if* $M_b \leq M < M_c$ p_{oc} *if* $M \geq M_c$	p_{hc} *if* $M < M_b$ p_{lc} *if* $M_b \leq M < M_c$ p_{oc} *if* $M \geq M_c$

where $M_a = \frac{4(c_n+c_s)-\bar{\rho}}{c_s}$, $M_b = \frac{(c_n+c_s)(2\bar{\rho}+c_n+c_s)}{\bar{\rho}c_s}$, $M_c = \frac{2u_d-\theta\bar{\rho}}{\theta c_s}$ *and* $M_d = \frac{u_d^2-\theta\bar{\rho}(u_d-\theta(c_n+c_s))}{\theta c_s(u_d-\theta(c_n+c_s))}$.

Equilibrium profits are:

$$\pi_{fc} = M\theta(c_n + c_s) \qquad \text{(full market coverage)}$$
$$\pi_{hc} = \frac{M\theta(\bar{\rho}+c_n+c_s)^2}{4\bar{\rho}} \qquad \text{(high market coverage)}$$
$$\pi_{lc} = \tfrac{1}{4}M\theta(\bar{\rho} + Mc_s) \qquad \text{(low market coverage)}$$
$$\pi_{oc} = u_d\left(1 - \frac{u_d}{\theta(\bar{\rho}+Mc_s)}\right)M \qquad \text{(outsiders only coverage)}$$

The following figure illustrates the four different equilibrium market coverage levels. In each

the congestion experienced by those stay in the network. But in the partial sharing price range, peers leaving are (indirectly) reducing congestion by reducing the size of the network. This effect ensures that fewer peers will leave the network in response to a price reduction when price is above p_{nd}.

case, the graph on the top shows demand (black line), marginal revenue (discontinuous blue line), equilibrium price (horizontal dotted line), and equilibrium profits (yellow area). The graph on the bottom shows profit as a function of quantity sold.

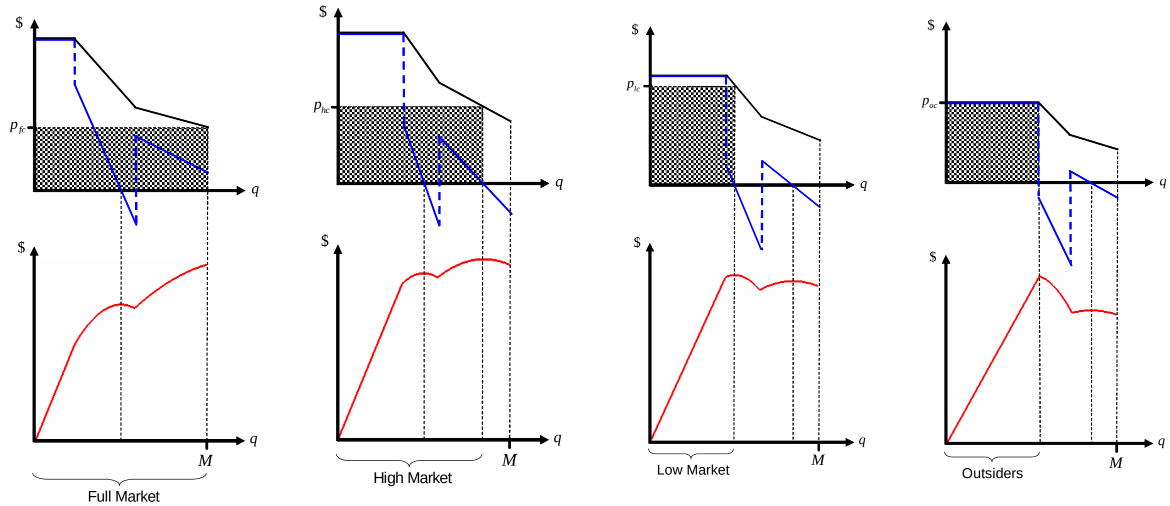

Full Market Coverage

High Market Coverage

Low Market Coverage

Outsiders Coverage

This result characterizes the optimal pricing strategy of the firm and shows that market size M is a critical parameter. The firm will only quote a low price and cover the entire market if M is sufficiently small and agents suffer little from congestion (low $\bar{\rho}$). In fact, when $\bar{\rho}$ is low, the relative benefit of the firm vs. the network is low (as peers are all patient) and the firm must set low prices to gain share.

In general, the firm will not cover the whole market. The bigger the market, the more profitable it is for the firm to target agents with high disutility of congestion by quoting high prices. If the market is sufficiently large, it is optimal for the firm to set $p = u_d$ and serve outsiders exclusively. Intuitively, as the size of the network increases, so does congestion. Therefore, when M is large the surplus that the firm can extract by targeting agents who most suffer congestion is larger than that obtained by quoting a low price and covering a larger chunk of the market.

The presence of the network always decreases firm profits. The effect of the network on the firm can be likened to a low quality firm competing against a vertically differentiated competitor. The firm has strong incentives to offer high quality service to remain competitive (by investing to minimize congestion) and to quote a high price.

The result also highlights other strategic considerations. Profits under all market configurations are increasing in market size. As M grows, so does congestion and this benefits the firm. Furthermore, profits are increasing in the cost of sharing (c_s) and nondecreasing in the cost of using the network (c_n). Because lager c_s and/or c_n result in lower utility levels for users of p2p (see Section 5.2), the firm will want to take actions to increase c_s and c_n, such as suing peers.

Proposition 7 also shows that profits decrease as θ falls. That is, as residential broadband infrastructure improves (lower θ), the network becomes a better competitor. To take actions that affect the availability of residential bandwidth (upload mainly) may be unfeasible, but selective degradation of file sharing traffic will serve the same purpose. Prioritization schemes favoring commercial traffic will strengthen the competitive position of the firm and weaken the network. Such schemes can be implemented with private contracts between telecommunications operators and content providers or through vertical integration.

7 Concluding remarks

Decentralized peer-to-peer file sharing networks and for-profit centralized online stores constitute alternative distribution models for digital information goods. In this paper we have presented a simple formal model to analyze some aspects in which p2p and online stores interact. This paper is a first step towards improving our understanding of competition between two distribution models that have been enabled by one technology, the Internet.

Although the two models have emerged only recently, it is more likely than not that they will endure. iTunes is legal and has grown spectacularly since its inception. Because peer-to-peer file sharing activity is mostly illegal, one is tempted to believe that legal attacks against p2p with the goal to shut them down will continue and eventually succeed. However, due to their decentralized nature, p2p networks have proven difficult to block. We expect our analysis to stay relevant going forward.

The effects of p2p file sharing on content providers are significant, and can be compared to those of cassette recording in earlier analog technological generations. The cassette recorder allowed individuals to generate unauthorized copies and to illegally share copyrighted content lowering potential revenues to content providers. Analog content sharing was subject to quality degradation and required physical exchange, which mainly confined the process to relatively small social networks. By eliminating these restrictions, peer-to-peer file sharing technology has increased the accessibility and attractiveness of unauthorized content replication. The threat of p2p is not different in nature, but of much larger scale as it does not require the exchange of a physical support: p2p networks allow individuals who have never met and who may be located far apart to exchange digital content as easily as if they where close next-door friends.

The content industry has so far faced the new online paradigm as a threat more than as an opportunity. But the need to embrace digital distribution seems obvious by now; there is no way back to a world of physical distribution only. Due to this transition and the increasing value of the online channel to reach consumers, we expect ISPs to have a stronger role in shaping market structure. We also expect the content industry to reassess their revenue models. Changes towards monetizing products not subject to replication, such as the increased attention paid by major record companies to live concerts and merchandising, may be signals of a new trend.

Our formal model is necessarily partial in that it is focussed around characterizing the firm's profit-maximizing pricing strategy. More generally (but less formally), to compete effectively against p2p, online digital distribution must strive to become accessible and attractive to consumers. Online content providers are in a unique position to optimize and deliver new experiences to consumers which cannot be matched by decentralized, self-sustained peer-to-peer networks. iTunes, for example, provides a better customer experience than file sharing for similar content and this allows Apple to charge positive prices and make a profit.[25] The potential industry-wide revenue implications of p2p are still uncertain. However, our analysis suggests that there is scope for profit-maximizing online distributors and content producers to compete effectively against unauthorized file sharing.

8 Appendix – Proofs

This appendix contains all the proofs.

[25] iTunes is easy to use and it is well-integrated with the iPod, it offers a secure and simple payment process, free samples, and minimum download delay.

Proof of Proposition 2. Sharer $i \in \mathbf{S}$ will not free ride if (3) or

$$\frac{S}{S-1} \leq \frac{c_n + \rho_i}{c_n + c_s + \rho_i}$$

is satisfied. Notice that

$$\frac{d\left(\frac{c_n+\rho_i}{c_n+c_s+\rho_i}\right)}{d\rho_i} = \frac{c_s}{\left(c_n + c_s + \rho_i\right)^2} > 0. \tag{11}$$

Therefore, if (3) is satisfied for sharer $i \in \mathbf{S}$ it is also satisfied for all sharers i' with $\rho_{i'} \geq \rho_i$. Thus, the more impatient a sharer is, the less the incentive to become a freerider. A freerider $j \in \mathbf{F}$ will not want to become a sharer if (4) or

$$\frac{S}{S+1} \geq \frac{c_n + \rho_j}{c_n + c_s + \rho_j}$$

is satisfied. Notice that (11) implies that if (4) is satisfied for peer $j \in \mathbf{F}$ it is also satisfied for all peers $j' \in \mathbf{F}$ with $\rho_{j'} \leq \rho_j$. Thus, the more patient a freeriding peer is, the less the incentive to become a sharer. ■

Proof of Proposition 4. We look for $\rho_{s(N)}$ such that the set of peers with $i \geq s(N)$ all want to share. Because $\rho_{s(N)}$ is the most patient sharer, the cardinality of the set of sharers is $S = N - s(N) + 1$.

For $\mathbf{S}$ to be the set of sharers of a stable network configuration, it is necessary that the most patient sharer does not want to freeride:

$$u_d - \left(c_n + c_s + \rho_{s(N)}\right)\theta\frac{N}{S} \geq u_d - \left(c_n + \rho_{s(N)}\right)\theta\frac{N}{S-1}$$

This expression implies that

$$\rho_{s(N)} \geq \left(N - s(N)\right)c_s - c_n.$$

Therefore, for $S = N - s(N) + 1$ to be stable, $\rho_{s(N)}$ must satisfy $\rho_{s(N)} \geq \left(N - s(N)\right)c_s - c_n$.

We also need that the most impatient freerider does not want to share:

$$u_d - \left(c_n + \rho_{s(N)-1}\right)\theta\frac{N}{S} \geq u_d - \left(c_n + c_s + \rho_{s(N)-1}\right)\theta\frac{N}{S+1}$$

This expression implies that

$$\rho_{s(N)-1} \leq (N - s(N) + 1) c_s - c_n.$$

So, we need that

$$\rho_{s(N)-1} \leq (N - s(N) + 1) c_s - c_n \quad \text{and} \quad (N - s(N)) c_s - c_n \leq \rho_{s(N)}.$$

Suppose now that all ρ_is are drawn from a uniform distribution $\rho_i \sim U[0, \bar{\rho}]$. When N is large we have that $s(N) - 1 \simeq \frac{\rho_{s(N)-1}}{\bar{\rho}} N$. Furthermore, large N also implies that $\rho_{s(N)} \simeq \rho_{s(N)-1}$. Therefore $s(N) \simeq \frac{\rho_{s(N)-1}}{\bar{\rho}} N + 1 \simeq \frac{\rho_{s(N)}}{\bar{\rho}} N + 1$. Substituting in the expression above, we obtain

$$\begin{aligned}
\left(N - \frac{\rho_{s(N)}}{\bar{\rho}} N - 1\right) c_s - c_n &\leq \rho_{s(N)} \\
\left(\bar{\rho} N - \rho_{s(N)} N - \bar{\rho}\right) c_s - \bar{\rho} c_n &\leq \bar{\rho} \rho_{s(N)} \\
\frac{\bar{\rho}\left((N-1) c_s - c_n\right)}{\bar{\rho} + N c_s} &\leq \rho_{s(N)}.
\end{aligned}$$

When N is large we have that $s(N) - 2 \simeq \frac{\rho_{s(N)-2}}{\bar{\rho}} N$. Furthermore, large N also implies that $\rho_{s(N)-1} \simeq \rho_{s(N)-2}$. Therefore $s(N) - 2 \simeq \frac{\rho_{s(N)-2}}{\bar{\rho}} N \simeq \frac{\rho_{s(N)-1}}{\bar{\rho}} N$ or $-s(N) + 1 \simeq -\frac{\rho_{s(N)-1}}{\bar{\rho}} N - 1$. Now, substituting in the expression above, we obtain

$$\begin{aligned}
\rho_{s(N)-1} &\leq (N - s(N) + 1) c_s - c_n \\
\rho_{s(N)-1} &\leq \left(N - \frac{\rho_{s(N)-1}}{\bar{\rho}} N - 1\right) c_s - c_n \\
\bar{\rho} \rho_{s(N)-1} &\leq \left(\bar{\rho} N - \rho_{s(N)-1} N - \bar{\rho}\right) c_s - \bar{\rho} c_n \\
\rho_{s(N)-1} (\bar{\rho} + N c_s) &\leq \bar{\rho}\left((N-1) c_s - c_n\right) \\
\rho_{s(N)-1} &\leq \frac{\bar{\rho}\left((N-1) c_s - c_n\right)}{\bar{\rho} + N c_s}
\end{aligned}$$

So, when N is large we have that

$$\rho_{s(N)-1} \leq \frac{\bar{\rho}\left((N-1) c_s - c_n\right)}{\bar{\rho} + N c_s} \leq \rho_{s(N)}.$$

We conclude that when N is large

$$\rho_{s(N)} \simeq \frac{\bar{\rho}\left(\left(N-1\right)c_s - c_n\right)}{\bar{\rho} + Nc_s}.$$

■

Proof of Corollary 5. We have that $S = N - s(N) + 1$ and $s(N) - 1 \simeq \frac{\rho_{s(N)-1}}{\bar{\rho}}N$. Furthermore, we have just seen that

$$\frac{\rho_{s(N)-1}}{\bar{\rho}} = \frac{\left(N-1\right)c_s - c_n}{\bar{\rho} + Nc_s}.$$

Therefore,

$$\begin{aligned} S &= N - \frac{\left(N-1\right)c_s - c_n}{\bar{\rho} + Nc_s}N \\ &= N\left(1 - \frac{\left(N-1\right)c_s - c_n}{\bar{\rho} + Nc_s}\right) \\ &= N\left(\frac{\bar{\rho} + Nc_s - \left(N-1\right)c_s + c_n}{\bar{\rho} + Nc_s}\right) \\ &= N\left(\frac{\bar{\rho} + c_s + c_n}{\bar{\rho} + Nc_s}\right). \end{aligned}$$

■

Proof of Proposition 6. An agent with disutility of congestion ρ_i will purchase from the firm iff:

$$u_d - p \geq u_d - (c_n + c_s + \rho_i)t_d.$$

Because $t_d \geq \theta$ is positive, if the condition is satisfied for peer i it will also be satisfied for peer $i+1$. So, the agents who most suffer congestion are the ones for whom purchasing the content from the firm is most attractive. To solve for demand given a price p we proceed by identifying the indifferent buyer, denoted by ρ_b. The indifferent buyer obtains the same utility from purchasing content and from downloading it for free from the network. Hence, all agents with $\rho_i > \rho_b$ will strictly prefer to purchase from the firm. Note that this includes outsiders (individuals who would have experienced negative utility if in the p2p network), who choose to purchase as long as $p \leq u_d$.

If $p = u_d$, only outsiders buy from the firm, as all other agents obtain strictly positive utility in the network. If $p > u_d$ purchasing yields negative utility and the firm faces no demand. To obtain demand when $p \leq u_d$ we must solve for ρ_b, given by:

$$u_d - p = u_d - (c_n + c_s + \rho_b)t_d. \tag{12}$$

Because either full or partial sharing may hold in the network, we consider two separate cases. We begin with the latter.

Substituting $t_d = \theta\frac{N}{S}$ in (12) and taking into account that congestion under partial sharing will depend on ρ_b, as only agents such that $\rho_i \leq \rho_b$ are present in the network:

$$u_d - p \simeq u_d - \left(c_n + c_s + \rho_b^{ps}\right)\theta\frac{N\left(\rho_b^{ps}\right)}{S\left(\rho_b^{ps}\right)},$$

where

$$N\left(\rho_b^{ps}\right) = \frac{\rho_b^{ps}}{\bar{\rho}}M,$$

and

$$\begin{aligned} S\left(\rho_b^{ps}\right) &= N\left(\rho_b^{ps}\right)\left(\frac{\rho_b^{ps} + c_s + c_n}{\rho_b^{ps} + N\left(\rho_b^{ps}\right)c_s}\right) \\ &= \frac{\rho_b^{ps}}{\bar{\rho}}M\left(\frac{\rho_b^{ps} + c_s + c_n}{\rho_b^{ps} + \frac{\rho_b^{ps}}{\bar{\rho}}Mc_s}\right) \\ &= M\left(\frac{\rho_b^{ps} + c_s + c_n}{\bar{\rho} + Mc_s}\right). \end{aligned}$$

Thus,

$$u_d - p \simeq u_d - \left(c_n + c_s + \rho_b^{ps}\right)\theta\frac{\frac{\rho_b^{ps}}{\bar{\rho}}M}{M\left(\frac{\rho_b^{ps}+c_s+c_n}{\bar{\rho}+Mc_s}\right)}.$$

Solving for ρ_b^{ps} yields:

$$\rho_b^{ps} = \frac{p\bar{\rho}}{\theta(\bar{\rho} + Mc_s)}.$$

We can now use this result to identify the boundary price which separates the full and partial sharing range. For the network size to equal the full sharing boundary size,

$$\frac{\rho_b^{ps}}{\bar{\rho}} M = \frac{c_n + c_s}{c_s}.$$

Substituting ρ_b^{ps} and solving for p:

$$p_{nd} = \frac{\theta(c_n + c_s)(\bar{\rho} + Mc_s)}{Mc_s}.$$

We denote the boundary price by p_{nd} to indicate that demand exhibits a non-derivability at this point. For any price below p_{nd} only full sharing will hold in the network, hence $t_d = \theta$.

We now solve for the indifferent buyer in the full sharing case by substituting $t_d = \theta$ in (12):

$$u_d - p \simeq u_d - \left(c_n + c_s + \rho_b^{fs}\right)\theta,$$

and solve for ρ_b^{fs}:

$$\rho_b^{fs} = \frac{p - \theta(c_n + c_s)}{\theta}.$$

The demand function for the firm is given by:

$$D = (1 - \frac{\rho_b}{\bar{\rho}})M. \tag{13}$$

Substituting ρ_b^{ps} we obtain the expression for demand in the partial sharing range:

$$D^{ps} = (1 - \frac{p}{\theta(\bar{\rho} + Mc_s)})M.$$

Note that this expression may yield negative values for higher values of p. Substituting ρ_b^{fs} in (13) we obtain demand in the full sharing range:

$$D^{fs} = (1 - \frac{p - \theta(c_n + c_s)}{\theta\bar{\rho}})M.$$

Full market coverage is obtained when $\rho_b^{fs} = 0$, which implies:

$$p_{fc} = \theta(c_n + c_s).$$

A lower price will also ensure that the market is covered. ■

Proof of Proposition 7. Given the non-derivability that the demand curve exhibits at p_{nd}, the firm faces two separate pricing ranges which describe two concave profit curves. Due to the shape of the demand curve, each profit curve lies above the other in its own price range and both intersect at the boundary price p_{nd}. Let us consider separately the optimal strategy of the firm in each range. Profits in the lower price range are given by

$$\begin{aligned} \pi_{lr} &= pD^{fs} \\ &= p(1 - \frac{p - \theta(c_n + c_s)}{\theta\bar{\rho}})M, \end{aligned} \tag{14}$$

which has a maximum at

$$p_{hc} = \frac{1}{2}\theta(\bar{\rho} + c_n + c_s).$$

We denote the maximum by p_{hc}, as high coverage of the market is obtained in this price range. Note that this price is not feasible if outside the range in which demand is well defined, that is, outside the lower price range. First, if below or equal to the price which ensures full market coverage,

$$p_{hc} \leq p_{fc},$$

the optimal price in the lower range is that which covers the market $p_{fc} = \theta(c_n + c_s)$. A lower price will not increase demand but only lower profits. Solving the inequality indicates that this is the case if and only if

$$\bar{\rho} \leq c_n + c_s. \tag{15}$$

Only in this case will a corner solution hold in the lower price range, implying full market coverage. Second, if

$$p_{lr} \geq p_{nd},$$

the optimal price for the firm is in the higher price range. As each profit curve lies above the other in its own price range, if the maximum of the lower range profit curve is attained in the higher price range it follows feasible profits in that range are effectively higher. Solving this inequality shows that this is the case if and only if

$$\bar{\rho} > c_n + c_s \text{ and } M \geq \frac{2\bar{\rho}(c_n + c_s)}{(\bar{\rho} - c_n - c_s)c_s}. \tag{16}$$

Only when neither (15) nor (16) are satisfied does the interior solution hold in the lower price range, given by p_{hc}.

We next consider the higher price range. Profits are given by

$$\begin{aligned} \pi_{hr} &= pD^{ps} \\ &= p(1 - \frac{p}{\theta(\bar{\rho} + Mc_s)})M, \end{aligned} \tag{17}$$

which has a maximum at

$$p_{lc} = \frac{1}{2}\theta(\bar{\rho} + Mc_s).$$

As market coverage is lower in this range, we denote the maximum by p_{lc}. We now consider the conditions for this price to be feasible. First, if

$$p_{lc} \geq u_d,$$

the higher feasible price for the firm is that which covers outsiders only, $p_{oc} = u_d$. A higher price will face no demand. Solving the inequality yields

$$M \geq \frac{2u_d - \theta\bar{\rho}}{\theta c_s}. \tag{18}$$

Only if this condition is satisfied does a corner solution hold in the higher price range, implying outsider coverage only. Second, if

$$p_{lc} \leq p_{nd},$$

it follows the optimal price is in the lower price range, as feasible profits must be higher in that range. Solving this inequality we have

$$M \leq \frac{2(c_n + c_s)}{c_s}. \tag{19}$$

Only when neither (18) nor (19) are satisfied does the interior solution hold in the higher price range, given by p_{lc}.

By the above construction, under condition (16) maximum profits are obtained in the higher price range, and under condition (19) in the lower price range. For the remaining cases, the price range where maximum profits are obtained needs to be determined. Profits in both price ranges as given by (14) and (17) must be evaluated under all combinations of conditions (15) and (18). Solving these systems of inequalities we obtain the firm's optimal pricing strategy. ■

References

[1] Andreoni, J. (1990), 'Impure Altruism and Donations to Public Goods: A Theory of Warm-Glow Giving,' *The Economic Journal*, Vol. 100, No. 401, pp. 464-477.

[2] Antoniadis, P., Courcoubetis, C. and Mason, R. (2004), 'Comparing economic incentives in peer-to-peer networks,' *Computer Networks*, Vol. 46, pp. 133-146.

[3] Asvanund, A., Clay, K., Krishnan, R. and Smith, M. D. (2004), 'An Empirical Analysis of Network Effects in Peer-to-Peer Music-Sharing Networks,' *Information Systems Research*, Vol. 15, pp. 155-174.

[4] Bhattacharjee, S., Gopal, R., Lertwachara, K. and Marsden, J. R. (2005), 'Using p2p sharing activity to improve business decision making: proof of concept for estimating product life-cycle,' *Electronic Commerce Research and Applications*, Vol. 4, pp. 14-20.

[5] Bergstrom, T., Blume, L. and Varian, H. (1986), 'On the Private Provision of Public Goods,' *Journal of Public Economics*, Vol. 29, pp. 25-49.

[6] Casadesus-Masanell, R. and Ghemawat, P. (2006), 'Dynamic Mixed Duopoly: A Model Motivated by Linux vs. Windows,' *Management Science* 52, no. 7: pp.1072-1084.

[7] Casadesus-Masanell, R., Hervas-Drane, A. and Sean Silverthorne (2007), 'Delivering the Digital Goods: iTunes vs. Peer-to-Peer,' *HBS Working Knowledge*. http://hbswk.hbs.edu/item/5594.html

[8] Casadesus-Masanell, R., Hervas-Drane, A. and Mitchell, J. (2006), 'Peer-to-Peer File Sharing and the Market for Digital Information Goods,' *Harvard Business School Case*, No. 706-479.

[9] Casadesus-Masanell, R. and Yoffie, D. (2007), 'Wintel: Cooperation and Conflict,' *Management Science* 53, no. 4: pp. 584-598.

[10] Crandall, R. W. and Alleman, J. (2002), 'Broadband: Should We Regulate High-Speed Internet Access?,' *AEI Brookings Joint Center for Regulatory Studies*.

[11] Creus Mir, A., Casadesus-Masanell, R. and Hervas-Drane, A. (2006), 'Bandwidth Allocation in Peer-to-Peer File Sharing Networks,' Mimeo, Harvard Business School.

[12] Cunningham, B. M., Alexander, P. J. and Adilov, N. (2004), 'Peer-to-peer file sharing communities,' *Information Economics and Policy*, Vol. 16, Issue 2, pp. 197-221.

[13] Feldman, M., Papadimitriou, C., Chuang, J. and Stoica, I. (2004), 'Free-riding and whitewashing in peer-to-peer systems,' *Proceedings of the ACM SIGCOMM 2004 workshop on Practice and Theory of Incentives and Game Theory in Networked Systems* (PINS).

[14] Golle, P., Leyton-Brown, K., Mironov, I. and Lillibridge, M. (2001), 'Incentives for Sharing in Peer-to-Peer Networks,' *Lecture Notes in Computer Science*, Vol. 2232, pp. 75+.

[15] Habib, A. and Chuang, J. (2004), 'Incentive mechanism for peer-to-peer media streaming,' *Proceedings of the Twelfth IEEE International Workshop on Quality of Service* (IWQOS), pp. 171- 180.

[16] Holmstrom, B. (1982), 'Moral hazard in Teams,' *Bell Journal of Economics*, Vol. 13, pp. 324-340.

[17] Jian, L. and MacKie-Mason, J. (2006), 'Why Share in Peer-to-Peer Networks?,' *Proceedings of the ACM EC 2006 workshop on the Economics of Networked Systems* (NetEcon).

[18] Krishnan, R., Smith, M. D., Tang, Z. and Telang, R. (2004), 'The Virtual Commons: Understanding Content Provision in Peer-to-Peer File Sharing Networks,' *Carnegie Mellon Heinz School Working Paper*, No. 2004-19.

[19] Pai, V. and Mohr, A. E. (2006), 'Improving Robustness of Peer-to-Peer Streaming with Incentives,' *Proceedings of the ACM EC 2006 workshop on the Economics of Networked Systems* (NetEcon).

[20] Parker, A. (2004), 'The True Picture of Peer-to-Peer File sharing,' http://www.cachelogic.com/research/slide1.php

[21] Schollmeier, R. (2004), 'Signaling and Networking in Unstructured Peer-to-Peer Networks,' *PhD. thesis*, Technische Universität München.