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

Ramon Casadesus-Masanell† Andres Hervas-Drane‡

August 4, 2009

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 the Coeditor and three anonymous Referees for suggestions that helped improve the paper substantially. We also thank Alexei Alexandrov, Jernej Copic, Albert Creus-Mir, Ig Horstmann, Rahmi İlkiliç, Jeff MacKie-Mason, Felix Oberholzer-Gee, Markus Reitzig, Alan Sorensen, and Pai-Ling Yin, and seminar participants at the Bosphorus Workshop on Economics Design, the CoCombine conference, The Wharton School of the University of Pennsylvania, London Business School, Harvard Business School, MIT, Ross School of Business (Michigan), Robert H. Smith School of Business (Maryland), 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, the Harvard EconCS Group, the 2007 NET Institute Conference, the Third International Conference on Ultrabroadband Networks, the 2007 Annual Meeting of the Academy of Management, the Summer Institute in Competitive Strategy at the Haas School of Business (Berkeley), the 2008 INFORMS Annual Meeting, and the TILEC Workshop on Media Markets at Tilburg University. The second author conducted this research at the Universitat Autònoma de Barcelona and Harvard University. We are grateful for financial support from the HBS Division of Research, Fundación BBVA, Fundación Caja Madrid, 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

‡Columbia University and Universitat Pompeu Fabra

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 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 recent 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.

Similarly to content, bandwidth is a nonexcludable good. When a peer provides upload bandwidth to the network she cannot decide who will or 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.

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 legal and fast downloads, based on traditional client-server architectures, at positive prices.[1] 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 lowering 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.* [5] show that congestion worsens as the size of p2p networks grow. Our model generates this result endogenously. In fact, the effect of network size on congestion helps explain the coexistence of multiple 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 show that the firm has an additional incentive to charge high prices, compared to what a standard model of vertical differentiation would predict. When prices are high, more consumers prefer to download from the p2p network. As a consequence, congestion increases and the value of the p2p network decreases, increasing the attractiveness of the firm's product.

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.[4] A study performed by Internet research firm CacheLogic in 2005 revealed that over 60% of total Internet traffic belonged to p2p file sharing applications.[5] This suggests that file sharing is indeed a driver for broadband demand. We believe that our results should be of interest to all participants in

[1] Note that an online store could be set up to distribute content through p2p, rather than through a client-server model. Motivated by iTunes' success, however, we assume that the firm not only has centralized governance but that it also distributes through a client-server model that provides content instantly. We show that the use of a client-server architecture for distribution is a critical source of advantage for the firm.

[2] See, for example, 'Downloading myths challenged,' BBC News.com, July 27, 2005.

[3] See Crandall and Alleman [14].

[4] See Parker [32].

[5] See Johnson *et al.* [25].

markets for digital information goods.

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 [10], 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 [13] 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 the 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 peer-to-peer file sharing. This literature asks why individuals share files in p2p networks. The pioneering work of Adar and Huberman [1] revealed that 70% of Gnutella users share no files, a phenomenon dubbed 'the tragedy of the digital commons:' because contributing files is costlier than freeriding, selfish utility maximizers 'should' freeride and this should lead to the collapse of the network. The problem has been related to the provision of public goods, as discussed by Krishnan *et al.* [26] and [28]. Our work contributes to this debate by separating the two different goods that are provided by peers who share: content and bandwidth. We focus on the provision of scarce bandwidth in open p2p networks, a rival and nonexcludable good. Separating both goods is helpful because they are different in nature as content is a pure public good (nonrival and nonexcludable).

In order to analyze the extent of freeriding and design schemes to reduce it, 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.* [20] and Antoniadis *et al.* [3] consider agents that derive utility from contributing content to the network. Feldman *et al.* [19] explicitly consider agent types which differ in their willingness to contribute. Ranganathan *et al.* [33] and Feldman *et al.* [18] model the problem as a repeated prisoners' dilemma game, where the network collapses in the absence of generosity. Cunningham *et al.* [16] assume reciprocity by a positive fraction of users, where increased sharing by some ensures a further increase in the overall provision. And Jian and MacKie-Mason [24] present a model of generalized reciprocity where peers share expecting others in the network to indirectly reciprocate. Our model is one of selfish, utility-maximizing agents. While social preferences underlie many aspects of human behavior, their role in large p2p networks where thousands of individuals interact anonymously is debatable. In the absence of peers caring about each others' utility, the question of why peer-to-peer networks such as Gnutella exhibit sharing is still open to argument.

Our approach builds on the notion of endogenous congestion to characterize the aggregate properties of p2p networks and their scalability. Asvanund *et al.* [5] analyze p2p traffic data and build a stylized model based on congestion, Krishnan *et al.* [27] propose some theoretical implications of congestion to explain sharing by selfish peers, and Butler [9] presents a conceptual framework for the scalability of online communities. To the best of our knowledge, no earlier model of p2p file sharing has bridged the aggregate properties of these networks and the local interactions between peers concerned solely about the impact of their actions on their own utility. In addition, the literature on the economics of p2p has not considered interactions between p2p and for-profit client-server distribution models.

Based on Gnutella traffic data, Asvanund *et al.* [4] examine the benefits of club formation, which can reduce freeriding by excluding peers with low levels of contribution from the network. While excludability opens several interesting issues in the context of p2p, we focus instead on open and nonexcludable networks (which have become the most widespread and successful instances of the technology), the size of which is endogenously determined as peers compare the expected utility of joining them versus available outside options.

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[6]

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 to directly communicate with others in 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.[7] 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

[6] This section draws from Casadesus-Masanell, Hervas-Drane, and Mitchell [12]. See also Casadesus-Masanell, Hervas-Drane, and Sean Silverthorne [11].

[7] 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 [34].

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 to hours. Sometimes downloads do not complete at all.

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, averaging at $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 back on any multimedia device and is future-proof compatible.[8] 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 users of p2p networks are generally open to legal action against them. Figure 1 summarizes the main dimensions on which p2p and iTunes differ.

Non-commercial 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

Figure 1: Non-commercial p2p vs. iTunes.

[8] See "Big Content: ludicrous to expect DRMed music to work forever," ArsTechnica.com, July 29 2009. While the music industry has gradually eliminated DRM requirements in recent years, video content largely continues subject to DRM restrictions. See "iTunes Store goes DRM-free," MacWorld.com, January 6 2009.

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.[9] Apple, on the other hand, in 2006 claimed more than 80% shares of U.S. legal music downloads.[10] 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.[11] Presently, iTunes offers more than 10 million songs for download and has sold six billion songs since going online in 2003.

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.[12] The RIAA has been at the center 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 continued to do so over the following years. More recently, the RIAA is attempting to block traffic of copyrighted content over p2p by reaching agreements with ISPs. 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.

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 p2p file sharing network as a two stage process. In the first stage, agents choose

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

(simultaneously) whether or not to join the network. In the second stage, agents interconnect and downloads are realized. 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, so the number of outsiders is given by $M - N$. 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. Sharing content is costly, but some sharing is required for the network to survive as downloads can only be realized from other peers. We assume that peers are anonymous, that they cannot coordinate their actions, and that no peer can be excluded from accessing the network.

Let u_d denote the utility derived from content once a download has been completed. We assume that u_d is common across all content and peers. This simplifies the analysis and allows us to 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.[13]

We assume that 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 signaling 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). Since both c_n and c_s are incurred over time, the total costs incurred by any given peer will depend on how long she remains actively connected to the p2p network.[14]

To introduce the disutility of congestion experienced by peer i, we use $\rho_i \geq 0$ which is a measure of how impatient peer i is. The larger ρ_i is, the higher the disutility from the time required to complete a download. 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.

[13] 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.

[14] Copyright holders willing to prosecute sharers need to *crawl* p2p networks in order to identify them, in a manner similar to how Google crawls webpages to index them. This implies that changes in peer activity (just as changes in webpage content) are not captured instantaneously.

Finally, we let t_d be the time required to complete a download. This download time is endogenous and depends on the level of congestion which, in turn, will depend on the proportion of peers sharing content in the network. Lower bandwidth provision implies higher levels of congestion resulting in higher download time.

With all the notation in place, we express the utility of a peer that freerides as

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

and that of a peer who shares content as

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

where $i \in \mathbf{N} = \{1, 2, ..., N\}$. Outside utility is normalized to zero.[15]

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 relationship between the file size and the bandwidth capacity available to peers. An improvement in broadband infrastructure that increases available bandwidth, decreases download times and amounts to a reduction in θ.[16] 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

[15] Note that we could redefine ρ_i as $\rho_i + c_n$ and omit c_n from our specification. For clarity of exposition, however, we prefer to use both c_n and ρ_i, as this notation allows us to separate the costs of using the network from agent heterogeneity. The benefits of this notation become clearest in Section 5.2 where $\rho_i \sim U\,[0, \bar{\rho}]$.

[16] 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.

fractions of content by turns.

The following two definitions are helpful in what follows:

Network allocation: A set of links connecting peers to sharers where every peer connects to one sharer only and no sharer connects to herself.

Stable network allocation: A network allocation such that 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.[17] 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 sharing in a p2p network with minimum congestion is always preferred to the outside option of not participating in 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

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

the network by allowing agents to not consume content or download from the firm instead of using the p2p network.

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-Mir on bandwidth allocation in p2p file sharing networks (see Creus-Mir *et al.* [15]).

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: $S_1 \rightarrow S_2$ (this means that S_1 connects to and downloads from S_2), $S_2 \rightarrow S_1$, $F_1 \rightarrow S_2$, $F_2 \rightarrow S_1$, and $F_3 \rightarrow S_1$. 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 [15] 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.[18]

[18] 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 $b_f = S\left(1-b_s\right)/F$. Recall that $F = N - S$. Therefore we have $b_f = S\left(1-b_s\right)/\left(N-S\right)$. 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 $b_f < S/N$.

It is interesting that sharers obtain a slightly higher bandwidth than freeriders even if they face more constraints 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 expected bandwidth of sharers and freeriders differs from S/N by, at most, 0.0012.[19] 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 is equal to S/N. Notice also that our approximation punishes (slightly) sharers and rewards (slightly) freeriders. Therefore, the approximation makes it 'harder' for peers to share. If instead of using the approximation we used the exact formula, our results would only be strengthened.

Our model is a simplified abstraction of a p2p network. There are two important aspects in which our model differs from real p2p networks. First, we assume that all nodes are homogeneous in capacity. It is well known that in addition to residential users with largely homogeneous connections, there are nodes with high-speed connections (e.g., computers in university dorm rooms) that contribute a large fraction of upload bandwidth and content. Sharers with heterogeneous upload capacities will typically end up with different expected bandwidths. Numerical simulations suggest that sharers with larger upload capacities end up obtaining more bandwidth unless upload capacities are very asymmetric. The reason is that by offering larger bandwidth, a supersharer 'frees up' other sharers (as an increased number of peers desire to connect to the supersharer) and more upload capacity is made available to her.

Consider for example a p2p network with three peers only: two sharers and one freerider. In this case, there are two stable allocations. Sharer 1 *must* download from sharer 2 (and vice versa). Given this, the freerider is indifferent between downloading from sharer 1 or from sharer 2. In both cases, the freerider obtains an effective bandwidth of 1/2. Depending on the sharer to which the freerider connects to, sharers may end up with an effective bandwidth of 1 or 1/2, so their expected bandwidth is $(1+1/2)/2 = 3/4$. Suppose now that sharer 1 is a supersharer that has $\lambda \in (1, 2]$ units of upload bandwidth available. In this case, both the second sharer and the freerider will connect to sharer 1. This is the only stable network configuration. Sharer 1 obtains an effective bandwidth of 1, and sharer 2 and the freerider obtain $\lambda/2 > 1/2$. Clearly, the supersharer benefits *the most* from

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

offering larger upload bandwidth. The freerider is also better off, but sharer 2 is better off only if $\lambda > 3/2$. Note also that if $\lambda > 2$, the contribution of the supersharer is so significant that sharer 2 and the freerider end up with effective bandwidths larger than that obtained by the supersharer.

Second, our model assumes that every peer connects to one sharer only. While this was the case in earlier generations of p2p file sharing networks (such as Napster), new generations (such as Kazaa and BitTorrent) allow for multiple links. Unfortunately, relaxing the single link assumption renders the problem intractable. To provide some insight on the effect of multiple links, consider the following case. Let peers hold $L = S - 1$ simultaneous links in a p2p network with only one freerider and $S = N - 1$ sharers. This construction ensures that every sharer is connected to all other sharers in all stable allocations. Hence the upload bandwidth contributed by each sharer is always accessed by the remaining L sharers. This implies that the freerider obtains the same bandwidth across all her links in all stable allocations,

$$\hat{E}_F = \frac{L}{L+1}.$$

We can easily calculate the bandwidth obtained by sharers by taking into account that $S/N \cdot \hat{E}_S + (N-S)/N \cdot \hat{E}_F = S/N$ and $N = S + 1$, thus

$$\hat{E}_S = \frac{1 + L + L^2}{(1+L)^2}.$$

As in the case of a single link, $\hat{E}_S > \hat{E}_F$, and sharers continue to do better than freeriders. Moreover, the difference in expected bandwidths between the multiple and the single link cases (given that $S = N - 1$) is small. To see this, note that expected bandwidths in the single link case are given by our approximation S/N, which yields $S/(S+1)$. Then, substituting $L = S - 1$ in order to compare both cases we obtain $E_F \approx E_S \approx (L+1)/(L+2)$, and we have that $|\hat{E}_F - E_F| < 1/(2 + 3L + L^2)$ and $|\hat{E}_S - E_S| < 1/((1+L)^2(2+L))$. Already when $N = 10$, for instance, $|\hat{E}_F - E_F| < 0.0076$ and $|\hat{E}_S - E_S| < 0.00069$.

While asymmetric upload capacities or multiple links affect the exact values of E_S and E_F, what is important for our argument below is that upload capacity is a scarce resource. 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. In Section 5 we endogenize the decision to share or to freeride and show that in equilibrium there is sharing.

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.

We use the following definitions in our analysis:

Network configuration: A partition $P = \{\mathbf{F}, \mathbf{S}\}$ of $\mathbf{N}$, where $\mathbf{F}$ is the set of freeriders and $\mathbf{S}$ is the set of sharers.[20]

Equilibrium network configuration: A network configuration where no $i \in \mathbf{S}$ prefers to (unilaterally) become a freerider and no $j \in \mathbf{F}$ prefers to (unilaterally) become a sharer.

To characterize the set of equilibrium network configurations, notice that sharer i will not freeride if

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

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

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 1 *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* $n \in \mathbf{N}$.

Proof. All proofs are in the Appendix. ■

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:

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

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 2 *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). See Figure 2.*

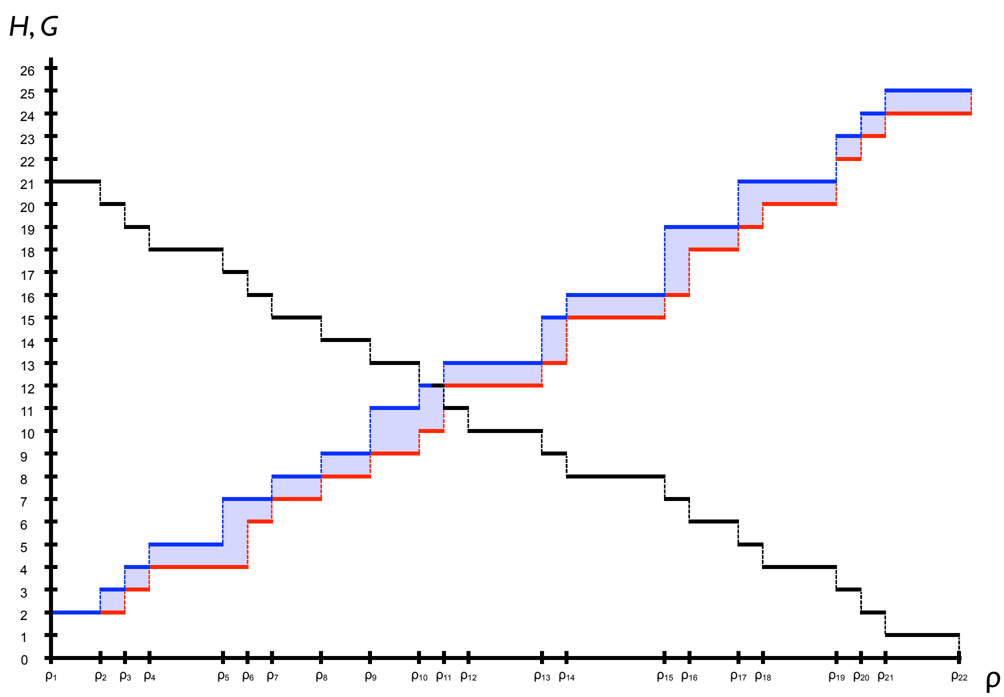

Figure 2: $H(\rho_i)$ and $G(\rho_i)$ – Equilibrium number of sharers.

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. We should point out that, depending on the values of the parameters, there may be no freeriders in the network. Likewise, when c_s is large, no peer is willing to share and the p2p network does not get off the ground.

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.

Peer-to-peer file-sharing clients often offer incentive schemes that promote sharing to lower congestion. In essence, these mechanisms redistribute bandwidth away from freeriders towards peers who share. The work of Moulin and Shenker in [29] and [30] is helpful to analyze these allocation mechanisms: the p2p network can be understood as a production process with a unique input and output – upload bandwidth and download bandwidth. The structure of the p2p network is conducive to an equal sharing rule, while bandwidth allocation incentive schemes aim to implement a serial sharing rule. The former allocates an equal share of download bandwidth to all peers; the latter allocates download bandwidth across peers as a function of their comparative upload bandwidth contribution. As a result, serial sharing ensures freeriders obtain little download bandwidth (so there is little free lunch), but may require a central authority in the network to allocate bandwidth. Incentive schemes in p2p file sharing clients approximate this rule in a decentralized network by implementing tit-for-tat reciprocity among peers. Of course, if we allowed for such incentive schemes in our model, we would obtain even more sharing.[21] We see the result that *even*

[21]The effect of these bandwidth allocation rules is similar to that of altruism or reciprocity between peers; in

in the absence of such incentives there may be sharing in equilibrium as one main contribution of our approach.[22]

Feldman *et al.* [17] analyze the incentives for peers to share as a function of the low-level architecture of TCP-IP networks. Their simulations show that sharing can increase the download latency experienced by sharers, particularly when the level of sharing in the network is high, in turn reducing the incentives to share. This effect operates similarly to congestion in our analysis. If sharers faced increased download latency in our model, this would reduce the overall degree of sharing in the network. The impact of congestion depends on the number of peers participating in the p2p network and the marginal cost of sharing, while the impact of download latency depends on the overall degree of sharing and packet prioritization schemes. It is unclear a priori how both effects compare. We see the results in Feldman *et al.* [17] as complementary to ours. While in their paper it is the download latency experienced by sharers that averts sharing, in ours individuals freeride because of the additional costs c_s that sharers bear (uplink bandwidth, using additional hardware resources, and increased likelihood of legal action).

5.1 Full vs. partial sharing

Let $\mathbf{S} = \{n, n+1, ..., N\}$ be the set of sharers in an equilibrium network configuration. It is useful to define the following:

Full-sharing equilibrium: An equilibrium network configuration where $n = 1$, that is, all peers are sharers.

Partial-sharing equilibrium: An equilibrium network configuration where $n > 1$, that is, both sharers and freeriders are present in the network.

In a full-sharing network, 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 by

both cases the marginal utility of sharing is increased. 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).

[22] 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.

sharing (given that everybody else 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 maximum 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 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.

As mentioned above, the additional cost borne by sharers c_s reflects the cost of uplink bandwidth, the cost of the computing resources made available by sharers (such as processing capacity or hard drive space for songs that perhaps sharers are not listening to any longer but that nevertheless offer so as to liberate traffic from other sharers), and the additional prospect of legal action against them (over and above to that faced by freeriders). The opportunity cost of the uplink bandwidth for a given sharer depends on the applications she uses. One may argue that for many individuals, uplink traffic is very small (e.g., a click requesting a web page) and the congestion cost to sharing should then be small. Our model predicts that when the additional costs of sharing are small, there is ubiquitous sharing and this helps our argument that selfish individuals may prefer to share rather than freeride in anonymous, decentralized p2p networks. We note that while sharing may not affect web browsing much, increasingly widespread applications such as video conferencing and online gaming can be seriously affected by p2p. In addition, as broadband infrastructure continues to improve, the bandwidth requirements of p2p will also increase with advances in content driving the exchange of high definition video or next-generation video games. In all likelihood, p2p uplink bandwidth will remain costly for time to come.

If individuals care mostly about downlink speed, upgrading uplink speed is a pure positive network externality. A coordination problem may then arise for peers upgrading to Internet plans with large upload bandwidth capacities. For instance, Verizon's FiOS project can provide much higher symmetric bandwidth than previous generation broadband infrastructure. ISPs such as Verizon, however, must also account for the impact of p2p traffic on their network when designing Internet plans for consumers. Since customers choose among these predefined plans instead of selecting upload and download capacities separately, the severity of coordination issues is probably diminished.

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.* [6] where individuals are concerned about the total supply of public goods, and the 'warm-glow' approach of Andreoni [2] 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.* [6] considers both the extensive and the 'intensive margin' – the decision of how much to contribute.[23]

Our model is also related to that of Bliss and Nalebuff [8], where society awaits for a supplier to provide a public good in a dynamic setting. The value of the public good is discounted by the delay in its provision, and individuals have incentives to wait in order to freeride on someone else supplying the good. The individual with the lower provision cost becomes the first-mover, preferring to supply the public good rather than wait for someone else to do so. In our model, individuals with high discounting costs are the ones that provide bandwidth to the network – a similar trade-off between discounting delay and provision costs drives supply in both cases. Contrary to their

[23] 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.

approach, however, our model is static (there is no waiting game) and the total level of public good is the sum of the contributions of many individuals. Moreover, the marginal effect of an individual's contribution to the public good depends on how much others have contributed.

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 2 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 in 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 2 *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 the cardinality of $\mathbf{N}$ is, the lower the proportion of sharers in equilibrium is. 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 2.

Corollary 1 *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 [22]) 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. 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, fewer and fewer 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.

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.

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. ISPs also report that a small number of residential customers are responsible for a large proportion of total traffic. Our model is consistent with this observation because sharers (a small number in large networks) generate higher volumes of traffic than freeriders.[24]

[24] File sharing may very well be the killer app that broadband had reportedly been missing. File sharing accounts

Corollary 1 helps explain why the p2p 'industry structure' is characterized by the presence of multiple, independent file sharing networks. The model suggests that as network size grows and 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.[25]

5.3 On the empirical relevance of the analysis

To the best of our knowledge, no empirical estimates exist for the model's main parameters (c_s and c_n) or the distribution of ρ. This makes it difficult to calibrate the model and explain available data on p2p file sharing networks. In general, with non-degenerate distributions of ρ, the model predicts that a large proportion of peers share for low values of N, and that the proportion is low for large values of N. For example, suppose that $N = 100$, $c_n = 10$, $c_s = 0.5$, and that $\rho \sim U[0, 100]$. In this case, about 73.6% of peers would share. On the other hand, when $N = 1,000,000$, about 0.02% of peers share. And when $\rho \sim U[0, 100]$, $c_n = 10$, $c_s = 0.01$ and $N = 1,000,000$, about 1% of peers share.

Empirical estimates of the proportion of sharers in p2p networks are in the ballpark of 20%.[26] A reasonable estimate of the number of peers logged in peer-to-peer file sharing networks at any given time is 10,000,000.[27] However, while the number of active peers at any given time across all networks is large, N should be interpreted as the number of peers in a given p2p network engaged in the exchange of a given subset of content. For example, a peer downloading music over BitTorrent

for a large share of Internet traffic across data networks worldwide, even as the capacity of broadband infrastructure available to residential users varies substantially from country to country. (See the data published by CacheLogic, Parker [32], 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.)

[25] Asvanund *et al.* [5] 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).

[26] See for example Hughes *et al.* [23].

[27] See Johnson *et al.* [25].

does not interact with peers on eMule, Gnutella, or other file sharing networks. Similarly, a peer interested in contemporary jazz music is unlikely to be interested in modern country music. The relevant N for such a peer is not 10,000,000 but the smaller community with similar music interests she might exchange content with within a given file sharing network.

For a rough estimate of what the relevant Ns might be, consider the fact that the largest p2p network is BitTorrent with approximately 60% of worldwide p2p traffic, and that a conservative estimate of the share of peers exchanging music in any given p2p network is 33% (the real share is probably lower).[28] Thus the relevant number of simultaneous peers exchanging music in a given file sharing network is at most 2,000,000. And to the extent that music genres may be good proxies for boundaries between interest groups, the relevant Ns become much smaller. For the case of pop, with a reported market share of 7.1% of music sales, the relevant N would be $N = 2,000,000 * 0.071 = 142,000$.[29] A 20% sharing ratio would then imply that $S = 28,400$.

The model with uniformly distributed ρ predicts that a large proportion of peers share for low values of N, and that this proportion is low for large values of N, unless c_s is very small. It is easy to see that given values of N, $\bar{\rho}$, and c_n, the equilibrium S (eq. 9) can take any value (from 2 to N) for any given value of c_s. For instance, to obtain a sharing ratio of r in the network, we require

$$c_s = \frac{(1-r)\bar{\rho} + c_n}{rN - 1},$$

a very small number when N is large. Thus, to reconcile the model with the evidence, we need c_s to be small. Is this reasonable? We believe that c_s is likely to be small when N is large. The likelihood of legal prosecution is probably the most important component of c_s, and since the probability of being caught sharing decreases with the number of peers in the network, it seems reasonable that c_s could be very small when N is large.[30]

A small cost of sharing c_s is needed for reasonable levels of sharing to occur when N is large. What about the benefit from sharing? The benefit from sharing is captured by the reduction in

[28] See the Ipoque 2008/09 Internet Study.

[29] See the Taylor Research & Consulting Group data for 2006 on music sales collected by genre.

[30] A more complete model would have $c_s = c_s(N)$ with limit of $c_s(N)$ as $N \to \infty$ being zero. For example, we could have $c_s = 1/N$. In this case, if $N = 142,000$, then $c_s = 0.000007$.

congestion from becoming a sharer ($t_d(S-1) - t_d(S)$). This is a decreasing function of S:

$$t_d(S-1) - t_d(S) = \frac{N}{(S-1)S}.$$

Thus, for a sharing ratio of 20%, we have that $S = 0.2\,N$ and

$$t_d(S-1) - t_d(S) = \frac{25}{N-5},$$

which is small when N is large. For example, when $N = 142,000$, we have $t_d(28,399) - t_d(28,400) = 0.00017$. Therefore, both the cost and the benefit of sharing are small when S is reasonably large. Given this, it is not entirely clear that inequalities (3) and (4) are reasonable drivers of peer behavior when S is large: the cost/benefit trade off appears to be too insignificant for peers to pay much attention to it.

This raises the question: how relevant is our analysis if, in reality, peers are unlikely to pay much attention to inequalities (3) and (4) when S is large? The answer is that the analysis is relevant to explain why there is sharing but not to explain the exact levels of sharing that we observe in some p2p file sharing communities. While the benefit of sharing for the $28,400^{th}$ sharer of pop music is close to zero according to inequalities (3) and (4), the magnitudes involved become more significant for smaller sharing ratios (low values of S) and for the most impatient sharers (those with large ρ_i). If $S = 300$, the reduction in congestion achieved by the 300^{th} sharer is $t_d(299) - t_d(300) = 1.58$ units of time. If $S = 100$, the 100^{th} sharer is reducing congestion by 14.34 units of time for all peers. These improvements in congestion are likely to be taken into consideration by peers.

Thus, the model explains why sharing arises in the network and how it evolves with network size. The model, however, is not sufficiently detailed to explain the larger levels of sharing observed in some p2p communities. In addition to endogenous congestion, there are other factors that should be accounted for to explain real-world sharing data. All large and successful p2p networks currently in operation implement incentive schemes that drastically affect how bandwidth is redistributed across sharers and freeriders. And while our model focuses on self interest, social preferences (altruism, reciprocity, or social norms) and unawareness ('sharing' is usually the default behavior of client software) can also contribute to explain the observed levels of sharing. We see the large literature in computer science and economics on sharing with social preferences as complementary to our approach.

The explanatory power of endogenous congestion, incentive schemes, social preferences, and unawareness is likely to be different at different stages in the life cycle of p2p networks. Our model suggests that endogenous congestion is likely to play an important role early in the life of a p2p network, before it becomes large, since the impact of sharing on congestion is always large for small N (because S must be small then). Our model also shows that as N grows, so does congestion, consistent with the available empirical evidence on congestion. With large N, however, the impact that any individual has in lowering congestion also decreases. This suggests that incentive systems in client software, social preferences, and unawareness become important factors to explain sharing. This is consistent with the observation that new p2p clients with incentive schemes, status-based user communities and non-technically-savvy users have become more prominent with the widespread adoption of p2p file sharing.

To conclude our discussion on the empirical relevance of the model, we note that fully rational peers would always take inequalities (3) and (4) into account. But when S is large, the behavior of some sharers is unlikely to be driven by these inequalities alone. Instead of full rationality, we could have assumed a threshold on the ability of peers to perceive the benefits of their sharing, such that for $S > S^*$ the inequalities no longer apply. Even with such a threshold, our main results would not change: we would still have a positive level of sharing in equilibrium, and congestion would worsen as N grows. Moreover, the results for the firm that we present in the following section continue to hold (at least qualitatively). Thus, our findings under the assumption of full rationality would be preserved.

6 The firm

We next introduce an online firm that sells digital information goods also available on the p2p 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.[31] 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. Online distributors such as Apple's iTunes Store offer content on a traditional client-server architecture and operate in agreement with intellectual property right holders. In our view, the most important advantages of licensed online distribution are: (1) it offers lower download time with no congestion, as consumption of content acquired from the firm can be enjoyed almost immediately at the moment of purchase, and (2) content is of higher quality for consumers since it is legal and offers advantages over that available on p2p, such as improved metadata and no spoof files.

We next express the utility of consumers purchasing from the firm by taking into account the above. Let u_f be the utility that consumers derive from content obtained from the firm, where $u_f \geq u_d$, and let p be the firm's price. A song can be downloaded from iTunes over a broadband connection in a few seconds, and furthermore, the industry seems to be engaged in an assertive push towards content streaming.[32] To capture this fact, we let the firm offer immediate consumption. Buyer i obtains the following net utility when purchasing content from the firm:

$$u_i = u_f - p, \tag{10}$$

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, the terms $c_n + \rho_i$ and c_s do not appear in (10). Thus the firm offers instantaneous downloads and higher quality content, but contrary to the

[31] 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.* [7].

[32] New players in this market such as Amazon and Google offer content streaming. This allows for real-time 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 [21] and Pai and Mohr [31].

network, it charges positive prices.

We assume for simplicity that the firm has zero marginal cost and 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 have the outside option of not consuming content at all. The outside utility is normalized to zero.

We modify the timing accordingly and solve for the subgame perfect Nash equilibrium. 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 not to consume. Those who enter the network may share or freeride. In the third stage, agents in the network interconnect and all downloads take place.

Let $\delta := u_f - u_d$ denote the marginal utility that consumers derive from the superior quality of the content distributed by the firm. The following proposition characterizes the demand faced by the firm.

Proposition 3 *The firm faces the following demand function:*

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

Figure 3 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.

Both models of digital distribution are interdependent. Individuals with high disutility of congestion prefer to buy from the firm rather than obtaining content from the network for free. These individuals benefit the most from fast downloads. As they choose to purchase, the network becomes smaller 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.[33]

[33] We should point out that the result that congestion falls when the network shrinks (regardless of the value of N)

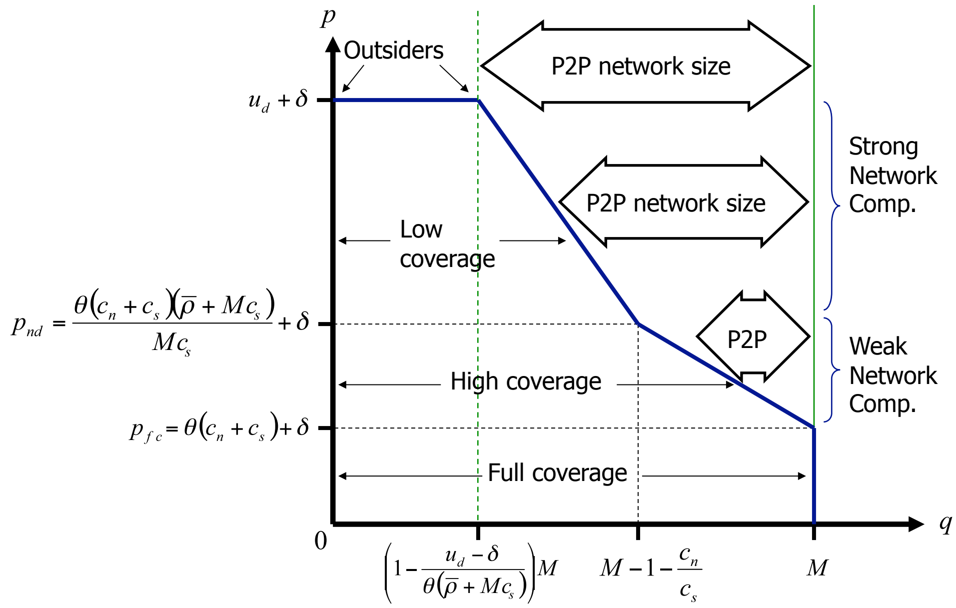

Figure 3: Demand function and size of the p2p network.

Full market coverage is obtained at any price less than or equal to p_{fc}. The individual who suffers congestion less ($\rho_i \simeq 0$) purchases at this price out of indifference (and all other individuals strictly prefer to purchase) rendering the network empty. Any price above p_{fc} ensures that some individuals 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 inclined to sharing; the effect induced by the reduction in network size is always stronger, resulting in a larger proportion of peers sharing content.[34] Finally, $p = u_d + \delta = u_f$ is the maximum price that the firm will ever charge. Any higher price will be met with no demand.

An increase in the marginal utility that consumers derive from the better quality of the content

relies on our assumption that ρ_i are uniformly distributed. For some distributions, a smaller network could result in higher congestion. However, as shown in Section 5.1, when the network is sufficiently small, congestion is always minimal, regardless of the distribution of ρ_i.

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

distributed by the firm δ has two effects: it shifts the demand curve up and it reduces the range of qs for which the p2p network has partial sharing. As the firm becomes comparatively more attractive, fewer and fewer individuals are ready to join the network. When δ is high, the number of individuals ready to join the network is so low that only full sharing occurs. For any price at which the market is uncovered, market coverage increases with δ.

Given the demand function derived in Proposition 3, the firm sets p to maximize profits. The next proposition summarizes our findings.

Proposition 4 *Profits are increasing in the size of the market M, in the marginal utility that consumers derive from the superior quality of the content distributed by the firm δ, in the spread in the distribution of impatience $\bar{\rho}$, in the costs of using the p2p network c_s and c_n, and in the inverse of the residential capacity offered by the broadband infrastructure θ.*

When market size M is large, the firm sets a high price and the marginal buyer is indifferent between joining the p2p network and not consuming at all. Intuitively, the demand lost to the p2p network N is small relative to M and the firm is better off ignoring the network by pricing at the monopoly price $p = u_f$. When M is low, the demand lost to the p2p network N is large relative to M and the firm benefits from prices below the monopoly level. In fact, when the market is very small, the firm must fight for share very aggressively; so much so, that it ends up capturing the entire market through low prices.

As δ increases, the firm's offering becomes more attractive. This allows the firm to charge higher prices and derive higher profits. An increase in δ can be implemented, for example, by improving metadata, providing complementary goods not readily available on the p2p network such as album art or artist information, reducing DRM restrictions for consumers and offering personalized content recommendations.

When $\bar{\rho}$ is low, the relative value of obtaining content from the firm vs. the network is low (as, in this case, peers become patient and do not feel much of the discomfort of congestion). As a consequence, the firm must set low prices to gain market share.

Increases in c_s and c_n render the network less attractive but have no effect on the value of purchasing from the firm. The most obvious way for the firm to increase c_s and c_n is by suing peers exchanging content in p2p file sharing networks. Targeting sharers ensures the impact is felt on c_s.

Finally, when broadband infrastructure improves (smaller θ), congestion decreases and the relative attractiveness of the p2p network improves. To take actions that affect the quality of residential bandwidth (upload mainly) may be unfeasible, but selective degradation of file sharing traffic serves the same purpose. Prioritization schemes favoring commercial traffic strengthen the competitive position of the firm and weaken the file sharing network. Such schemes can be implemented through agreements with telecommunications operators or vertical integration.

Figure 4 illustrates the four different market coverage levels that may arise in equilibrium. In each case, the graph on the top shows demand (black line), marginal revenue (discontinuous blue line), equilibrium price (horizontal dotted line), and equilibrium revenues (yellow area).[35] The graph at the bottom shows profit as a function of quantity sold.

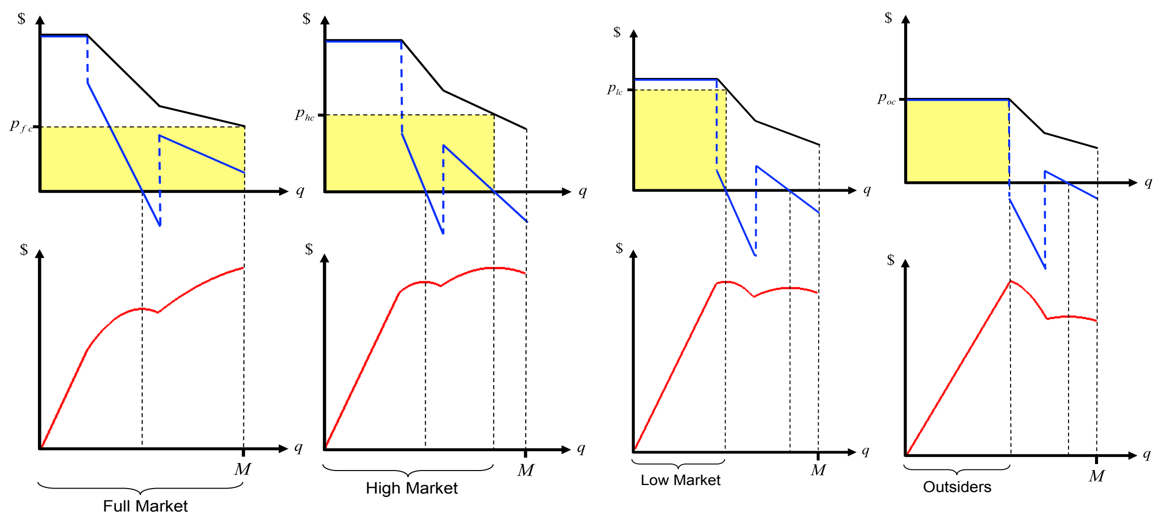

Figure 4: Firm prices and profits in the four equilibrium market coverage levels.

The effect of the network on the firm can be likened to a low quality firm competing against a vertically differentiated competitor. On the surface, the basic idea is the same: the firm sells a high-quality good to consumers with high willingness to pay for quality (those for whom congestion is costlier). Contrary to a standard vertical differentiation framework, however, the quality of the p2p network depends on its size (how many peers are in the network and the proportion of sharers vs. freeriders). When choosing prices, the firm internalizes the endogenous quality of its 'competitor.'

[35] Our assumption of zero marginal cost implies that revenue maximization is equivalent to profit maximization.

Thus the insight delivered by our model is that the firm has an additional incentive to charge high prices. The reason is that when prices are high, more consumers prefer to download from the p2p network. As a consequence of the increased number of peers in the network, congestion worsens and the value of the network decreases, in turn increasing the attractiveness of the firm.

7 Concluding remarks

We have presented a simple formal model to analyze some aspects of the interaction between p2p and online stores, two alternative models for the distribution of digital information goods. 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 were 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 relevance of the online channel to reach consumers, we expect ISPs to play 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 focused 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.[36] 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.

Proof of Proposition 1. 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

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

[36] iTunes is easy to use and it is well-integrated with the iPod, it offers a secure and simple payment process, free samples, album art, personalized recommendations, good karma, and minimum download delay.

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 2. 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

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

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

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

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

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

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

So, when N is large we have that

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

We conclude that when N is large

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

■

Proof of Corollary 1. 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

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

Therefore,

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

■

Proof of Proposition 3. An agent with disutility of congestion ρ_i will purchase from the firm iff:

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

or equivalently

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

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

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

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:

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

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,

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

Solving for ρ_b^{ps} yields:

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

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:

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

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):

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

and solve for ρ_b^{fs}:

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

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:

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

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:

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

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

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

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

Proof of Proposition 4. Direct inspection of the demand function reveals that there are four candidates to optimal price:

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

Consequently, there are four candidates to maximum profit:

$$\begin{aligned}
\pi_{fc} &= M(\theta(c_n + c_s) + \delta) && \text{(full market coverage)} \\
\pi_{hc} &= \tfrac{M}{4\bar{\rho}\theta}(\delta + \theta(\bar{\rho} + c_n + c_s))^2 && \text{(high market coverage)} \\
\pi_{lc} &= \tfrac{M(\delta + \theta(\bar{\rho} + Mc_s))^2}{4\theta(\bar{\rho} + Mc_s)} && \text{(low market coverage)} \\
\pi_{oc} &= (u_d + \delta)(1 - \tfrac{u_d - \delta}{\theta(\bar{\rho} + Mc_s)})M && \text{(outsiders only coverage)}
\end{aligned}$$

That the comparative statics within each region are as stated in the proposition, is immediate. That is, profits are strictly increasing in M, δ, c_s, and θ, and weakly increasing in $\bar{\rho}$ and c_n.

Now, because the demand function is continuous (though not differentiable at several points), the profit function is also continuous. Therefore, the comparative statics carry across the different regions of the profit function (i.e., there are no discontinuities in the comparative statics as we move from one region to the next). ■

References

[1] Adar, E. and Huberman, B. A. (2000), 'Free Riding on Gnutella,' First Monday, Vol. 5.

[2] 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.

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

[4] Asvanund, A., Bagla, S., Kapadia, M. H., Krishnan, R., Smith, M. D. and Telang, R. (2003), 'Intelligent Club Management in Peer-to-Peer Networks,' *1st Workshop on Economics of Peer-to-Peer Systems.*

[5] 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.

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

[7] 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.

[8] Bliss, C. and Nalebuff, B. (1984), 'Dragon-Slaying and Ballroom Dancing: The Private Supply of a Public Good', *Journal of Public Economics*, Vol. 25, pp. 1-12.

[9] Butler, B. S. (2001), 'Membership Size, Communication Activity, and Sustainability: A Resource-Based Model of Online Social Structures,' *Information Systems Research*, Vol. 12, No. 4, pp. 346-362.

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

[11] 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

[12] 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.

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

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

[15] Creus-Mir, A., Casadesus-Masanell, R. and Hervas-Drane, A. (2008), 'Bandwidth Allocation in Peer-to-Peer File Sharing Networks,' *Computer Communications*, Vol. 31, pp. 257–265.

[16] 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.

[17] Feldman, M., Lai, K., Chuang, J. and Stoica, I. (2003), 'Quantifying Disincentives in Peer-to-Peer Networks,' *1st Workshop on Economics of Peer-to-Peer Systems.*

[18] Feldman, M., Lai, K., Stoica, I. and Chuang, J. (2004), 'Robust Incentive Techniques for Peer-to-Peer Networks,' *Proceedings of the 5th ACM conference on Electronic commerce.*

[19] 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).

[20] 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+.

[21] 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.

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

[23] Hughes, D., Coulson, G., and Walkerdine, J. (2005) 'Free Riding on Gnutella Revisited. The Bell Tolls?,' *IEEE Distributed Systems Online*, Vol. 6, No. 6.

[24] 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).

[25] Johnson, M. Eric, McGuire, Dan, and Willey, Nicholas D. (2008), 'The Evolution of the Peer-to-Peer File Sharing Industry and the Security Risks for Users,' *Proceedings of the 41st Hawaii International Conference on System Sciences.*

[26] Krishnan, R., Smith, M. D. and Telang, R. (2003), 'The Economics of Peer-to-Peer Networks,' *Journal of Information Technology Theory and Application*, Vol. 5, No. 3, pp. 31-44.

[27] 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.

[28] Krishnan, R., Smith, M. D., Tang, Z. and Telang, R. (2007), 'Digital Business Models for Peer-to-Peer Networks: Analysis and Economic Issues,' *Review of Network Economics*, Vol. 6, No. 2, pp. 194-213.

[29] Moulin, H. and Shenker, S. (1992), 'Serial Cost Sharing,' *Econometrica*, Vol. 60, No. 5, pp. 1009-1037.

[30] Moulin, H. and Shenker, S. (1994), 'Average Cost Pricing versus Serial Cost Sharing: An Axiomatic Comparison,' *Journal of Economic Theory*, Vol. 64, pp. 178-201.

[31] 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).

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

[33] Ranganathan, K., Ripeanu, M., Sarin, A. and Foster, I. (2003), 'To Share or not to Share. An Analysis of Incentives to Contribute in Collaborative File Sharing Environments,' *1st Workshop on Economics of Peer-to-Peer Systems.*

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