

Peer-to-Peer File Sharing Communities

Brendan M. Cunningham, Peter J. Alexander, and Nodir Adilov

October 15, 2001

Abstract

Peer-to-peer file sharing communities present a paradox for standard public goods theory, which predicts that free-riding should preclude the success of the community. We present a model in which users choose their level of sharing, downloading, and listening in the presence of sharing costs and endogenous downloading costs. In our model, sharing emerges endogenously, largely as a byproduct of users' attempts to reduce own-costs.

*Cunningham: Department of Economics, U. S. Naval Academy, email: bcunning@usna.edu. Alexander: Federal Communications Commission, email: palexand@fcc.gov. Adilov: Department of Economics, Cornell University. We thank David Sappington for his many insightful and useful comments. The usual caveat applies.

1 Introduction

In this paper we explore peer-to-peer digital file sharing communities and present a positive model in which generalized contributions to the public good by members of the community arise. Those with only basic knowledge of file sharing may recall that the (now defunct) ‘Napster Music Community’ was one such very prominent community, in which members shared music files with each other. These types of ‘peer-to-peer’ file sharing communities are growing in size and scope, with users sharing a variety of digital file types such as music, movies, pictures and software programs.¹

We find the success of peer-to-peer file sharing communities somewhat puzzling, especially in light of Mancur Olsen’s declaration that “rational, self-interested individuals will not act to achieve their common or group interests” (Olson (1965) p.2). The supposition that in large anonymous groups individuals will rationally free-ride and thus vitiate opportunities for mutually beneficial voluntary group cooperation has long been a core assumption of standard economic theory. The data in Adar and Huberman (2000) provides limited evidence that this dynamic may be at work in peer-to-peer communities; many members of these communities fail to contribute to the public good by sharing a significant number of files. Despite this “stylized fact”, these communities continue to operate and grow in their variety.²

The explanation for the success of peer-to-peer file sharing communities from a behavioral perspective can be rationalized in a variety of ways that are consistent with observed behavior. For example, one could assert that pure or reciprocal altruism plays a deterministic role. In fact, a rather large and impressive body of literature exists suggesting these characteristics, among others, may play an important role in human sociality and willingness to cooperate (see Sugden (1984), Ledyard (1995), and Ostrom (2000)). However, none of these frameworks can theoretically accommodate cooperation

¹The model in this paper will focus upon the use of peer-to-peer communities to share music, although the descriptive power of the model is fairly broad.

²For another empirical analysis of the inner workings of peer-to-peer communities, see Asvanund, Clay, Krishnan and Smith (2001).

among extremely large, anonymous groups of individuals with no central authority guiding them (see Fehr and Gächter (2000) for a discussion of the theory and evidence regarding cooperation in human behavior).

We suggest that individual self-interest is the central explanatory element of the success of file sharing communities. Our explanation for the success of these communities, as detailed in our model, is that users share to reduce own costs. The act of sharing helps reduce search and congestion costs. In this sense, the virtual ‘commons’ are continuously and endogenously restored. Our approach is similar in spirit to the model in Becker (1976), whereby ‘simulated’ altruistic behavior emerges from individual rationality. As we suggest in our model, a dynamic of group cooperation can emerge from the interplay of a modest amount of altruism, in conjunction with a large aggregate percentage of self-interested types.³ Our results are consistent with the case in which the technology of sharing imposes (non-binding) sharing restrictions as the default behavior of users. In this scenario, some fraction of the users may be unaware that they are sharing, or simply be unable to make the adjustments required not to share.

In Section 2 of this paper, we detail the emergence of peer-to-peer file sharing communities and discuss some of the issues involved in digital file sharing. In Section 3, we develop a model of file sharing and demonstrate analytically that individual sharing can be individually rational - a behavior that is consistent with what we observe occurring within the peer-to-peer file sharing communities. In Section 4, we explore additional properties of the model and suggest possible extensions.

³In contrast to the approach explored in Golle, Leyton-Brown and Mironov (2001), our model explicitly accounts for the multi-period nature of users’ environment. The model in Golle et al. (2001) is essentially a one-shot prisoner’s dilemma, where parties defect unless an organizing authority explicitly guides them. The Folk Theorem dictates that this is not necessarily a multi-period equilibrium.

2 Digital Distribution and the Emergence of Peer-to-Peer File Sharing Communities

Prior to the introduction of digital technology, music was distributed using vinyl disks and magnetic tape. Digital distribution was facilitated by the introduction of compact disks in the 1980s. A compact disk uses a series of 0's and 1's to reproduce sound waves. Given that computer architecture employs the same digital structure, compact disks can be played using computers. Moreover, the sound waves from a compact disk can be readily transferred to the computers storage device (e.g., the hard drive), and given the appropriate connections, shared with other computer users (see Alexander (1994), (2001) for discussions of digital distribution and file sharing).

Until quite recently, this type of music file sharing was modest in scope, for several reasons. First, connections between personal computers were slow. However, the growing consumer market penetration of broadband and high-speed connections has increased aggregate bandwidth, and hence the speed at which files can be transferred. Second, and perhaps more importantly, new audio compression routines, including MP3, have greatly reduced the size of the file to be transferred. MP3, created by engineers at the German company Fraunhofer Gesellschaft, is short-hand for Motion Picture Experts Group-Layer 3. MP3 is an audio compression format that generates near compact disk quality sound at approximately 1/20 the size of the original. For example, while each minute of music on a compact disk requires the equivalent of 10 megabytes of computer storage space, an MP3 format of the same piece could be stored on 1 megabyte or less. MP3 achieves this reduction in overall file size in two ways: (1) discrete sampling of continuous sounds waves, and (2) passing the resulting samples through high and low band filters. To give a practical example of the compression savings achieved by MP3, consider that Elvis Presley's "Hound Dog" on compact disk requires 24 megabytes of hard disk space, but when converted to MP3 the storage requirement falls to 2 megabytes. On a 28.8 kilobit per second modem, the compact

disk version of “Hound Dog” would take at least one and one-half hours to download from another computer. On the other hand, if the file were first converted to MP3, it would take approximately eight and one-half minutes.

‘Napster’ refers to a particular type of file sharing software used to transfer MP3 files. The architecture of the Napster Music Community, summarized in Figure 1, consists of a series of centralized servers that direct electronic traffic and rout requests for files. In this figure, three users have installed and executed the Napster software. These users, in conjunction with the Napster server, form the constituent elements of the Napster network.

This figure illustrates one of the core features of the genus of “peer-to-peer” networks, of which Napster is a species. All sizeable transfers of files (and hence information) take place directly between the computers of individual users: users act as file clients and servers. The Napster network is unique from other species of peer-to-peer networks in its addition of a centralized server which maintains and updates a searchable index of all shared files located on individuals’ computers.

Figure 1 depicts a typical use of the Napster network, in which user 1 has queried the Napster server for a file with certain characteristics. The Napster server responds to this query with a list of other users who have served files with characteristics in the neighborhood of user 1’s query. Once user 1 has identified a desirable file and server source, in this case user 2, the user begins to transfer, or download, the file.

In this case, Napster user 2 is assumed to have chosen to share files which have appeal to user 1. Similarly, Napster user 1 is assumed to have shared files which Napster user 3 is downloading. The act of sharing is costly since any downloading from a sharer implies that the sharer is sacrificing bandwidth while the file is acquired by the downloading party. At best, we would expect users with infinite bandwidth to be indifferent between sharing and not sharing. Still, many users share, an outcome which defines the success of the Napster Music Community. Without sharing, the network collapses.

Our explanation for this behavior is that sharing is a rational, self-interested response by which users attempt to lower own costs. To motivate this result, we assume that a small but positive fraction of the peer-to-peer file sharing community increase their sharing of files in response to increased sharing by other users. We then show that other purely self-interested users will optimally share as well.

3 The Model

In this section, we construct a non-linear dynamic optimization model in which a user chooses the level of sharing, downloading, and listening in the presence of sharing costs and endogenous downloading costs. We begin by defining the following variables, dated t

- s_t – the volume of files shared by a given user
- $S(s_t)$ – the aggregate amount of sharing in the community
- h_t – the amount of time spent away from the community
- l_t – the amount of time the user spends listening to files
- M_t – the total stock of files held by the user
- d_t – the volume of files downloaded by a user.

where M_t and d_t are symmetric across all users. All variables are measured in units of time. Note $S(s_t)$ implies that aggregate sharing is a function of own sharing and our assumption of reciprocity on the part of some positive fraction of users implies that $S_s \equiv \partial S / \partial s > 0$, provided increased sharing by some does not result in a more than offsetting decline by others.

This assumption can be justified by a number of aspects of peer-to-peer file sharing networks. In particular, the default configuration of most peer-to-peer software installations will automatically share

any file which is downloaded. Provided a fraction of users fail to discover the process by which they may “un-share” files, those users will tend to share more files as they download from a larger pool of shared files.

We employ the following exogenous parameters defined as:

- φ – the fraction of MP3 files which are still of interest after 1 period
- ω – the rate at which listening and downloading are substitutes
- γ_s – the time cost per file to sharing
- $\Gamma(d_t, S)$ – the total time cost of downloading.

The rate of persistent interest in any individual MP3 file is given by $0 < \varphi < 1$. This parameter captures the fact that popular interest, individually and collectively, in any given MP3 has a finite life span. We assume that $(1 - \varphi)$ percent of files lose their interest to the user and the community each period. For simplicity we assume the rate of persistent interest is symmetric across all users.

We fix ω as the rate at which listening and sharing are substitutes. If $\omega = 0$, listening and downloading are perfect complements which occur simultaneously, in this case downloading will not consume time beyond that spent listening. If $\omega = 1$, listening and downloading are perfect substitutes and all time spent downloading will consume time beyond that spent listening. For values of ω between 0 and 1, listening and downloading exhibit varying degrees of substitutability.

The positive parameters γ_s and $\gamma_d \equiv \partial\Gamma/\partial d_t$ are the marginal time costs of sharing and downloading, respectively. These activities reduce a user’s available time since bandwidth is consumed by each activity, thereby increasing the time it takes to complete other activities. We assume that $\gamma_S \equiv \partial\Gamma/\partial S > 0$ since an increase in aggregate sharing implies a reduction in the search and congestion costs associated with downloading. By search costs, we mean that a user typically must explore multiple peer sources for a

file before a satisfactory source is found. By congestion costs, we mean that as multiple users attempt to download a given file, each will experience a lower download rate, with bandwidth consumed over a longer period. If aggregate sharing is high, a given user is less likely to encounter this problem.

The user's constraints include:

$$l_t \leq M_t + d_t. \quad (3.1)$$

which requires that the user cannot spend more time listening, in any given period, than the amount of potential music supplied by files already downloaded and those downloaded within the period. Repeat listening, while not currently modelled, could be added by scaling the right-hand side of (3.1) by a constant. Next, we have:

$$s_t \leq M_t + d_t. \quad (3.2)$$

This constraint reflects the fact that users can not share more files than those which are currently on the drive and those which are currently downloaded. The law of motion for the stock of files is given by :

$$M_{t+1} = \varphi M_t + d_t. \quad (3.3)$$

Note that if $d_t = 0$ the stock of files converges to zero at rate φ . Finally, we have:

$$l_t + \omega d_t + h_t = H - \gamma_s s_t - \Gamma(d_t, S(s_t)) \quad (3.4)$$

which requires that listening, downloading, and time away from the community are constrained by the time endowment (H) net of the time costs associated with downloading and sharing.

We assume that there exists a function mapping listening, and non-community time to utility: $u(l_t, h_t)$. This utility function has the standard property of positive but diminishing marginal utility in each argument: $u_t^l, u_t^h > 0$ where $u_t^x \equiv \partial u(\cdot) / \partial x$. We further assume that the user seeks to maximize

this function over an unbounded horizon subject to the constraints (3.1)-(3.4). In order to derive the optimal values of the state and control variables, we construct the Bellman equation:

$$\begin{aligned}
V[M_t] = & \max_{l_t, s_t, d_t} \{u(l_t, h_t) + \\
& + \lambda_t^l (M_t + d_t - l_t) + \lambda_t^s (M_t + d_t - s_t) + \\
& + \rho V[M_{t+1}]\}
\end{aligned} \tag{3.5}$$

where ρ is the discount factor and the λ^i are the relevant Lagrangian multipliers, which obey the complementary slackness conditions:

$$\lambda_t^l (M_t + d_t - l_t) = 0 \tag{3.6}$$

and

$$\lambda_t^s (M_t + d_t - s_t) = 0. \tag{3.7}$$

Three first-order conditions characterize the solution to this problem:

$$l_t : \quad u_t^l = u_t^h + \lambda_t^l \tag{3.8}$$

$$s_t : \quad -u_t^h \gamma_S S_s = u_t^h \gamma_s + \lambda_t^s \tag{3.9}$$

$$d_t : \quad \rho V'[M_{t+1}] + \lambda_t^l + \lambda_t^s = u_t^h (\omega + \gamma_d). \tag{3.10}$$

The condition (3.8) implies that the marginal utility of listening (the marginal benefit) is equal to the marginal utility of time away from the community plus the shadow price of listening (the marginal cost). The marginal benefit in (3.9) derives from the fact that downloading costs are decreasing in aggregate sharing, while aggregate sharing is increasing in own sharing. At the optimum, this benefit is equal to the marginal sharing costs and the shadow price of sharing. Finally, the marginal benefits

of downloading in (3.10) include both the discounted future value of files along with a loosening of the constraints on both listening and sharing. This benefit is balanced against the marginal loss of utility due to a reduction in non-community time.

The associated Envelope condition for the problem is:

$$V'[M_t] = u_t^h \varphi (\omega + \gamma_d) + (1 - \varphi) \lambda_t^l + (1 - \varphi) \lambda_t^s. \quad (3.11)$$

An increase in the stock of files today implies that downloading is lower, holding constant the future stock of files. It is clear from (3.11) that the value of such a marginal increase in the stock is equal to the marginal utility of the time made available by lower downloading, along with the shadow value of loosening the listening and sharing constraints.

In order to obtain a closed-form solution, we assume a Cobb-Douglas utility function which is homogeneous of degree one in listening and non-network time: $u(l_t, h_t) = l_t^\beta h_t^{1-\beta}$. The associated marginal utilities are:

$$u_t^l = \beta \left(\frac{h_t}{l_t} \right)^{1-\beta}, \quad u_t^h = (1 - \beta) \left(\frac{l_t}{h_t} \right)^\beta \quad (3.12)$$

Moreover, we assume the following functional forms for $\Gamma(\cdot)$ and $S(\cdot)$:

$$\Gamma(d_t, S(s_t)) = \frac{1}{2} (\gamma_0 d_t - \tilde{\gamma}_1 S(s_t))^2 \quad (3.13)$$

and

$$S(s_t) = \psi_1 s_t. \quad (3.14)$$

with $\tilde{\gamma}_1 > 0$ and $\psi_1 > 0$. As can be seen from (3.13) and (3.14), downloading costs are decreasing in S , which reflects our assumption that increased aggregate sharing will reduce search and congestion costs. Aggregate sharing is a strictly increasing function of an individual's sharing due to reciprocity on the

part of some users.

Combining (3.13) and (3.14), we obtain

$$\Gamma(d_t, s_t) = \frac{1}{2}(\gamma_0 d_t - \gamma_1 s_t)^2 \quad (3.15)$$

where $\gamma_1 \equiv \tilde{\gamma}_1 \psi_1$. From (3.13) and (3.14) we can conclude that $\gamma_S S_s = -\gamma_1(\gamma_0 d_t - \gamma_1 s_t)$. The coefficient γ_1 reflects the magnitude of the marginal change in the cost of downloading with respect to a marginal change in an individual's own sharing. When a user shares more files, reciprocity induces a marginal increase in aggregate sharing, thereby reducing search and congestion costs. This implies that the marginal cost of finding and downloading a file is lower. Thus, γ_1 is a measure of sharing's marginal reduction in the cost of downloading (henceforth MRCD). Similarly, γ_0 is a measure of downloading's marginal bandwidth cost.

The Appendix provides a number of propositions and proofs concerning the existence and features of optimal behavior under the functional forms described above. There are four possible steady state vectors of the state and control variables which satisfy the first-order and Envelope conditions.⁴ Each vector corresponds to one of the four combinations of the weak inequalities $\lambda^l \geq 0, \lambda^s \geq 0$.

Proposition (5.2) in the Appendix indicates that in the current framework, the listening constraint must bind in order to attain utility maximization. If the listening constraint does not bind, a user is unnecessarily downloading a volume of files which are not used to generate utility in the steady state. By reducing downloading in this situation, the user may free time which can be used for utility-generating activities. Hence, an optimal steady state will require a binding listening constraint. Moreover, the Appendix provides proofs that if the parameter values in the model are such that the conditions (5.14) and (5.18) hold, users will optimally download a significant volume of files and optimally share a positive

⁴If a variable takes on its steady value, we remove the time subscript t since that variable will remain unchanged over time.

volume of files without encountering a binding sharing constraint.

The first condition, (5.14), simply ensures that the user's taste for listening to files (as indicated by the parameter β) is large relative to the downloading cost (as captured by the parameter ω). If this is not the case, optimal behavior dictates that a user will not be very active in the peer-to-peer community, as measured by either downloading or sharing activity. In essence, (5.14) is a sufficient condition for actual entrance into the peer-to-peer community. The second condition, (5.18), guarantees that the MRCD parameter (γ_1) is sufficiently large such that the user need not share all available files in order to attain the maximal benefit from sharing. Thus, this condition ensures that an interior optimum is feasible for the user.⁵

We will proceed to discuss utility-maximizing behavior under the assumption that (5.14) and (5.18) hold. The Appendix demonstrates that the following actions will maximize utility in the steady state:

$$s = \frac{\gamma_0 \gamma_1 d - \gamma_s}{\gamma_1^2} \quad (3.16)$$

$$M = \frac{d}{1 - \varphi} \quad (3.17)$$

$$d = \left[H + \frac{1}{2} \left(\frac{\gamma_s}{\gamma_1} \right)^2 \right] f(\gamma_0, \gamma_1, \gamma_s, \omega, \beta, \varphi, \rho) \quad (3.18)$$

and

$$h = \frac{1 - \beta}{\beta} \left[\frac{1 - \rho\varphi}{1 + \rho(1 - \varphi)} \left(\omega + \frac{\gamma_0 \gamma_s}{\gamma_1} \right) + 1 \right] \frac{2 - \varphi}{1 - \varphi} d \quad (3.19)$$

$$l = \frac{\beta}{1 - \beta} h \quad (3.20)$$

where f is decreasing in γ_0, γ_s and ω ; f is increasing in γ_1, β and ρ .

⁵Interestingly, in this model a low level of γ_1 is associated with a binding sharing constraint, as shown in Proposition 5.4 when the condition (5.17) holds. Thus, if sharing exhibits a low marginal impact on the cost of downloading, users in this model may find it in their best interest to share a large fraction of downloaded files in order to offset this low marginal payoff. This case is not explored in the text for expositional purposes only.

A number of implications of the model are clear from these results. First, in the steady state outcome, the stock of files is entirely determined by the level of downloading. Since this stock is constant from period to period, but $(1 - \varphi)$ percent of files are effectively “lost” to the user (and the community) each period, a user must download $(1 - \varphi)$ percent of the stock each period in order to replace the fraction of the stock of files which are lost. Thus, one can calculate the stock of files by multiplying the level of downloading by the inverse of the rate of decay; downloading and the stock of files only differ by a scaling factor greater than one (see (3.17)).

In addition, the level of sharing is almost entirely determined by the level of downloading. The elasticity of sharing with respect to downloading is one (see (3.16)). This result arises since the costs associated with a large volume of downloaded files (and hence high steady state stock of files) are partially defrayed by higher levels of sharing. Thus, our model predicts high levels of sharing conditional upon high levels of downloading.

The optimal level of downloading, as described by (3.18) is fundamentally a function of most of the parameters of the model. Only four of these parameters exhibit an unambiguous impact on downloading. If listening to files carries a higher weight in the utility function (i. e., a larger β), a user will download a larger volume of files. Inspecting (3.16), we see that a user will share more as well. If the discount factor is higher (i. e., a larger ρ), a user places a larger weight on the future. Since downloading fundamentally drives the steady state stock of files, which in turn constrains the future utility of a user, a higher discount factor leads to more downloading (and, from (3.16), higher levels of sharing). Optimal behavior also requires that an increase in the marginal cost of downloading (i. e., a higher value of ω or γ_0) is matched with a decline in downloading. If a high marginal cost of downloading is only due to a high level of substitution between listening and downloading (i. e., a high ω) then sharing will simultaneously drop in response to a rise in the marginal cost of downloading (see (3.16)).

On the other hand, if the marginal cost of downloading rises due to an increase in the “direct” band-

width cost of downloading, as measured by the parameter γ_0 , sharing may actually rise as downloading declines.⁶ While both ω and γ_0 represent the marginal cost of downloading, the source of variation in the marginal cost of downloading critically effects whether sharing exhibits an inverse relationship to this variation. Sharing may not display an inverse relationship with respect to the direct marginal cost γ_0 since an increase in this cost can be offset to some extent by an increase in sharing. The total cost of downloading is $TCD = \omega d + .5(\gamma_0 d - \gamma_1 s)^2$. Since $\partial TCD / \partial \omega = d$, an increase in the substitutability of listening for downloading can not be offset by sharing, however, since $\partial TCD / \partial \gamma_0 = d(\gamma_0 d - \gamma_1 s)$, an increase in sharing will reduce the marginal impact of an increase in γ_0 on the total cost of downloading. Hence, users have an incentive to share more as the marginal bandwidth cost of downloading.

Finally, our model admits the possibility that an increase in the marginal cost of sharing (γ_s), or a decline in sharing's MRCD (γ_1) may induce either a fall or rise in sharing and / or downloading. A decline of sharing in response to such a change in the environment is intuitively plausible. The possibility that sharing might rise in response to such events is a relatively surprising prediction.

This indeterminacy comes from the fact that these parameters have ambiguous effects on the expressions in (3.16) and (3.18). Consider that (3.9) can be re-written as:

$$u^h \gamma_1 (\gamma_0 d - \gamma_1 s) = u^h \gamma_s + \lambda_s. \tag{3.21}$$

The benefit of sharing is the product of the marginal utility of non-community time and the marginal downloading time which is saved by sharing additional files (this expression lies on the LHS of (3.21)). At the optimum, this benefit is equated to the marginal cost of sharing, which is determined by the marginal utility of non-community time multiplied by the marginal bandwidth cost of sharing files plus the shadow price of sharing (see the RHS of (3.21)).

⁶In other words, (3.16) indicates that a drop in downloading caused by a rise in γ_0 may be directly offset by the rise in γ_0 itself

If the marginal bandwidth cost of sharing increases, the marginal benefit of sharing on the LHS of (3.21) may be increased by downloading a greater number of files. In essence, the user might counteract some of the increase in direct sharing costs by sharing (and downloading) a greater number of files, thereby leveraging reciprocity effects in order to offset a direct increase in the cost of sharing.

Similarly, if the MRCD increases, the user may not require a high level of sharing and downloading in order to induce, via reciprocity, lower levels of downloading costs. Thus, sharing and downloading could decline in response to an increase in the parameter γ_1 . The actual impact of γ_s and γ_1 on sharing may only be resolved through simulations of the model under various parameter values. We leave this task for future research.

4 Conclusion

In this paper, we have explored peer-to-peer file sharing communities; a curious phenomenon equivalent to a toll highway with no toll booths. Some researchers (see Adar and Huberman (2000)) have interpreted the limited empirical evidence collected from peer-to-peer communities as evidence of tragedy of the cyber-commons. In contrast, our model admits the possibility of thriving communities where generalized peer-to-peer file sharing is apparently the norm.

Our explanation for the success of these communities, as detailed in our model, is that users share to reduce own costs. The act of sharing helps reduce search and congestion costs. In this sense, the virtual ‘commons’ are continuously and endogenously restored. We have no doubt that other factors may be inducing sharing (e.g., file sharing software that defaults to sharing status and requires user intervention to change), but we believe that we provide a useful starting point for such analysis from an economic perspective. Moreover, our analysis implies that a user’s proclivity to share may be heavily affected by that same user’s propensity to download from the common pool of files. Given this result, empirical attempts to assess “free riding” and other such behavior on peer-to-peer networks must carefully collect

sufficient data to cover the multiple ways in which users are constrained by the environment as they make numerous choices regarding community interactions.

While we have provided a potentially useful starting point in this type of analysis, we remain open to the idea that the model may be usefully extended in a variety of ways. Certainly, the recent experimental work described in Fehr and Gächter (2000) suggests that ‘fairness’ (among other factors) may play a functional role in determining certain social outcomes (in the case we explore, perhaps the old adage about ‘honor among thieves’ might apply). In fact, we view empirical and experimental work in this area as a very promising next step. We think the present work provides a substantial economic foundation from which tests and experiments may be reasonably constructed.

5 Appendix

We begin by combining the expression for downloading costs, (3.15) with the time constraint (3.4). This yields the following:

$$l_t + h_t = H - \gamma_s s_t - \omega d_t - \frac{1}{2}(\gamma_0 d_t - \gamma_1 s_t)^2. \quad (5.1)$$

We conclude that the marginal effect of sharing on the total time available for the utility-yielding activities (l and h) is:

$$\frac{\partial(l_t + h_t)}{\partial s_t} = -\gamma_s + \gamma_1(\gamma_0 d_t - \gamma_1 s_t). \quad (5.2)$$

From this expression, it is clear that if downloading is very low, a marginal increase in sharing will always reduce the time available for utility-yielding activities. In this case, the optimal behavior entails zero sharing. However, if downloading is sufficiently high, positive sharing will arise in order to offset downloading costs. Also note that $\partial^2(l_t + h_t)/\partial s_t^2 < 0$ so that an interior optimal level of sharing will be globally optimal.

More formally, the optimal level of sharing is determined by the following conditions:

$$s_t^* = \begin{cases} 0 & \text{if } d_t < \frac{\gamma_s}{\gamma_0 \gamma_1} & \text{(Case I)} \\ \frac{\gamma_0 \gamma_1 d_t - \gamma_s}{\gamma_1^2} & \text{if } \frac{\gamma_s}{\gamma_0 \gamma_1} \leq d_t \text{ and } \lambda_t^s = 0 & \text{(Case II)} \\ M_t + d_t & \text{if } \frac{\gamma_s}{\gamma_0 \gamma_1} \leq d_t \text{ and } \lambda_t^s > 0 & \text{(Case III).} \end{cases} \quad (5.3)$$

Note that in Case II, the optimal level of sharing may be found by solving (3.9) for the optimal level of sharing or by setting (5.2) equal to zero and solving for sharing. From (5.3), we can see that sharing is optimally non-decreasing in downloading.

Combining (5.3) and (5.1) we have:

$$h_t = \begin{cases} H - \omega d_t - \frac{1}{2}\gamma_0^2 d_t^2 - l_t & \text{in Case I} \\ H + \frac{1}{2}\left(\frac{\gamma_s}{\gamma_1}\right)^2 - \left(\frac{\gamma_0\gamma_s}{\gamma_1} + \omega\right) d_t - l_t & \text{in Case II} \\ H - \left(\gamma_s \frac{2-\varphi}{1-\varphi} + \omega\right) d_t - \frac{1}{2}\left(\gamma_0 - \frac{2-\varphi}{1-\varphi}\right)^2 d_t^2 - l_t & \text{in Case III.} \end{cases} \quad (5.4)$$

Next, we observe that the ratio of listening to time away from the community can be determined from (3.8) since the choice of these variables has no effect on the RHS of the time constraint (5.1).

Thus,

$$l_t = \begin{cases} \frac{\beta}{1-\beta} h_t & \text{if } \lambda_t^l = 0 \\ M_t + d_t & \text{if } \lambda_t^l > 0 \end{cases} \quad (5.5)$$

from which we can conclude that for all six of the possible permutations suggested by (5.4) and (5.5), the level of time away from the community may be written as a function of current downloading and the current stock of files. By (3.3), $d_t = M_{t+1} - \varphi M_t$, so we can conclude that $h(M_t, M_{t+1})$. Combining this result with (5.5), we can write $l(M_t, M_{t+1})$ and define $\tilde{u}(M_{t+1}, M_t) \equiv u(h(M_t, M_{t+1}), l(M_t, M_{t+1}))$. The function \tilde{u} will be used to develop the FE form of our problem (for more on the FE form, see Stokey, Lucas and Prescott (1989)).

Proposition 5.1 *There exists a solution to the problem defined by (3.5).*

Proof. First, observe that the FE form of our problem is as follows:

$$V[M_{t+1}] = \max_{M_{t+1} \in \Omega(M_t)} \{\tilde{u}(M_t, M_{t+1}) + \rho V[M_{t+1}]\} \quad (5.6)$$

where

$$\Omega(M_t) = [\varphi M_t, \varphi M_t + \bar{d}] \quad (5.7)$$

with \bar{d} defined such that when $d_t = \bar{d}$, all available time is absorbed by downloading and sharing (leaving no time left for utility-yielding activities). Thus, \bar{d} represents a strict upper-bound on downloading.

Define T such that

$$(Tf)(x) = \max_{\varphi x \leq y \leq \varphi x + \bar{d}} \{\tilde{u}(x, y) + \rho f(y)\}. \quad (5.8)$$

First, note that for any x and time t , $\Omega(x)$ is bounded from below by 0 and above by $\max\{\varphi x + \bar{d}, \frac{\bar{d}}{1-\varphi}\}$.

Thus, $\Omega : \mathbb{R}^1 \rightarrow \mathbb{R}^1$ is compact-valued, non-empty and continuous.

Let $f \in C(X)$ where C is the set of bounded and continuous functions. Then $\tilde{u}(\cdot) + \rho f(\cdot)$ is bounded.

Also, by the Theorem of the Maximum, (Tf) is continuous. Then $Tf \in C(X)$.

The monotonicity of T is easily verified. Assume $f_1(x) \geq f_2(x), \forall x \in X$. Therefore

$$\begin{aligned} (Tf_1)(x) &= \max_{y \in \Omega(x)} \{\tilde{u}(x, y) + \rho f_1(y)\} \\ &\geq \tilde{u}(x, y_2^*) + \rho f_1(y_2^*) \\ &\geq \tilde{u}(x, y_2^*) + \rho f_2(y_2^*) \\ &= \max_{y \in \Omega(x)} \{\tilde{u}(x, y) + \rho f_2(y)\} = (Tf_2)(x). \end{aligned} \quad (5.9)$$

In addition, T satisfies the discounting property:

$$\begin{aligned} T(f + \alpha)(x) &= \max_{y \in \Omega(x)} \{\tilde{u}(x, y) + \rho[f(x) + \alpha]\} \\ &= \max_{y \in \Omega(x)} \{\tilde{u}(x, y) + \rho f(x)\} + \rho\alpha \\ &= Tf(x) + \rho\alpha. \end{aligned} \quad (5.10)$$

Therefore, Blackwell's Sufficiency Conditions are satisfied and Tf is a contraction mapping. Since C is a complete metric space, T has a unique fixed point in the set of bounded and continuous functions. Moreover, from the Theorem of the Maximum, the optimal policy correspondence $M_{t+1}(M_t)$ is compact-valued and upper hemicontinuous. ■

Given this result, we will now proceed to solve for the optimal values of the remaining state and control variables in the steady state in which all variables take on constant optimal values. We will

drop the subscript t in order to indicate that a variable is taking on its optimal steady state value.

The following proposition is the first step in the solution.

Proposition 5.2 *In an optimal steady state, the listening constraint (3.1) must bind.*

Proof. Assume the listening constraint is not binding, so that $l < M + d = (2 - \varphi)/(1 - \varphi)d$. Note that imposing a steady state on the law of motion (3.3) yields the equality in this expression.

From (5.4), we also have:

$$l + h = \begin{cases} H - \omega d - \frac{1}{2}\gamma_0^2 d^2 & \text{in Case I} \\ H + \frac{1}{2}\left(\frac{\gamma_s}{\gamma_1}\right)^2 - \left(\frac{\gamma_0\gamma_s}{\gamma_1} + \omega\right) d & \text{in Case II} \\ H - \left(\gamma_s \frac{2-\varphi}{1-\varphi} + \omega\right) d - \frac{1}{2}\left(\gamma_0 - \frac{2-\varphi}{1-\varphi}\right)^2 d^2 & \text{in Case III.} \end{cases} \quad (5.11)$$

Choose an arbitrarily small $\epsilon > 0$ and a new level of downloading $d' \equiv d - \epsilon$, ensuring that ϵ is such that: a) if Cases I, II or III held under the optimal level of downloading d then this case still holds under d' and b) $l < (2 - \varphi)/(1 - \varphi)d'$. If b) holds, the user need not reduce the level of listening as a result of this lower level of downloading.

Inspecting (5.4), we observe that for each of the three possible cases, the total time available for listening and downloading is decreasing in downloading. Thus, relative to d , d' will necessarily permit the user to achieve higher levels of l or h (or both), thereby yielding higher utility for the user.

This demonstrates that a level of downloading such that the listening constraint does not bind will not be optimal in the steady state. ■

The only remaining task entails determining which of the Cases I, II and III hold at the optimum. The following proposition illuminates a set of parameter values such that the low level of downloading and sharing described under Case I is never a feasible optimal outcome.

Proposition 5.3 *Provided β is sufficiently high and ω is sufficiently low such that (5.14) below holds,*

there does not exist a low level of downloading (specifically, $d < \gamma_s/\gamma_0\gamma_1$) which satisfies the conditions (3.8)-(3.11). Therefore, Case I never holds and zero sharing is never optimal.

Proof. Begin by observing that if Case I holds, the sharing constraint must not bind. The law of motion (3.3) implies that $M + d = \frac{2-\varphi}{1-\varphi}d$ in the steady state. If the sharing constraint did bind in the Case I steady state (so that $\lambda^s > 0$), $0 = s = \frac{2-\varphi}{1-\varphi}d \Rightarrow d = 0$. But then, by Proposition 5.2, $l = \frac{2-\varphi}{1-\varphi}d = 0$. However, (3.9) then indicates that $\lambda^s = 0$. This is a contradiction.

Using the observation that the sharing constraint is not binding, combining (3.8), (3.10), (3.11), and (5.11) we have that in Case I:

$$(1 + \rho(1 - \varphi)) \frac{\beta}{1 - \beta} \frac{h}{l} = (1 + \rho(1 - \varphi)) + (1 - \rho\varphi)\omega + (1 - \rho\varphi)\gamma_0^2 d. \quad (5.12)$$

Using Proposition (5.2) to substitute for l in (5.12) and equating the resulting equation to Case I in (5.11), we have:

$$\begin{aligned} H = & \left\{ \frac{1-\beta}{\beta} \frac{2-\varphi}{(1+\rho(1-\varphi))(1-\varphi)} [1 + \rho(1 - \varphi) + (1 - \rho\varphi)\omega] + \omega + \frac{2-\varphi}{1-\varphi} \right\} d + \\ & + \gamma_0^2 \left\{ \frac{1-\beta}{\beta} \frac{(2-\varphi)(1-\rho\varphi)}{(1+\rho(1-\varphi))(1-\varphi)+\frac{1}{2}} \right\} d^2. \end{aligned} \quad (5.13)$$

Note that the RHS of (5.13) is increasing in the level of downloading. Therefore, an increase in β or a decline in ω or γ_0 will result in an endogenous increase in d in order to ensure that (5.13) holds.

Case I requires that d rise no further than $\gamma_s/\gamma_0\gamma_1$. If the parameter β is high enough while the parameters ω , and γ_0 are low enough such that:

$$\begin{aligned} H > & \left\{ \frac{1-\beta}{\beta} \frac{2-\varphi}{(1+\rho(1-\varphi))(1-\varphi)} [1 + \rho(1 - \varphi) + (1 - \rho\varphi)\omega] + \omega + \frac{2-\varphi}{1-\varphi} \right\} \frac{\gamma_s}{\gamma_0\gamma_1} + \\ & + \gamma_0^2 \left\{ \frac{1-\beta}{\beta} \frac{(2-\varphi)(1-\rho\varphi)}{(1+\rho(1-\varphi))(1-\varphi)+\frac{1}{2}} \right\} \left(\frac{\gamma_s}{\gamma_0\gamma_1} \right)^2 \end{aligned} \quad (5.14)$$

then the upper-bound on downloading required in Case I is inconsistent with utility maximization; the

user will optimally choose a level of downloading above the upper bound imposed by Case I. ■

Given this proposition, the only remaining task involves determining conditions under which of the non-binding sharing constraint case (II) or the binding sharing constraint case (III) is optimal. The following proposition provides these conditions:

Proposition 5.4 *Assume the condition (5.14) holds. If the conditions (5.17) hold, implying that γ_1 is sufficiently low, the sharing constraint will bind at the optimum and Case III emerges as utility-maximizing behavior. If the condition (5.18) holds, implying that γ_1 is relatively high, the sharing constraint will not bind and Case II emerges as utility maximizing behavior.*

Proof. The discussion above indicates that sharing maximizes the time available for listening and downloading when:

$$s = \frac{\gamma_0 \gamma_1 d - \gamma_s}{\gamma_1^2}. \quad (5.15)$$

This level of sharing will exceed the binding amount of files available whenever

$$\frac{2 - \varphi}{1 - \varphi} d < \frac{\gamma_0 \gamma_1 d - \gamma_s}{\gamma_1^2}. \quad (5.16)$$

Using this inequality along with the fact that the condition (5.14) implies that $d > \frac{\gamma_s}{\gamma_0 \gamma_1}$, it is easy to show that if both of the following conditions hold:

$$\begin{aligned} 1 &< \frac{\gamma_s}{\gamma_0 \gamma_1} \\ \gamma_s &< \gamma_1 \left(\gamma_0 - \frac{2 - \varphi}{1 - \varphi} \gamma_1 \right) \end{aligned} \quad (5.17)$$

the sharing constraint will always bind for any level of downloading and Case III always arises.

Conversely, if the inequality in (5.16) is reversed, it is easy to show that if:

$$\frac{\gamma_0}{\gamma_1} - \frac{2 - \varphi}{1 - \varphi} < 0 \quad (5.18)$$

then the sharing constraint will never bind for any level of downloading and Case II always emerges. ■

Note that the conditions (5.17) and (5.18) are not an exhaustive list of possibilities. For “middle” levels of γ_1 , Proposition 5.4 fails to discern utility maximizing behavior. Determining a behavioral outcome in this middle range is possible but not pursued in this paper.

As a final step, we derive the optimal level of downloading in the steady state in which the sharing constraint does not bind (Case II). From (5.11), Proposition (5.2), and the observation that in the steady state, $M + d = (2 - \varphi)/(1 - \varphi)$ we observe that:

$$h = H + \frac{1}{2} \left(\frac{\gamma_s}{\gamma_1} \right)^2 - \left(\frac{\gamma_0 \gamma_s}{\gamma_1} + \omega + \frac{2 - \varphi}{1 - \varphi} \right) d. \quad (5.19)$$

Moreover, the conditions (3.8), (3.10), (3.11), and the value of sharing determined in (5.3) may be combined, leading to:

$$h = \frac{1 - \beta}{\beta} \left[\frac{1 - \rho\varphi}{1 + \rho(1 - \varphi)} \left(\omega + \frac{\gamma_0 \gamma_s}{\gamma_1} \right) + 1 \right] \frac{2 - \varphi}{1 - \varphi} d. \quad (5.20)$$

These two relationships may be combined in order to arrive at the optimal level of downloading when the listening constraint binds, the sharing constraint does not bind, and the conditions (5.14) and (5.18) hold:

$$\begin{aligned} d &= \left[H + \frac{1}{2} \left(\frac{\gamma_s}{\gamma_1} \right)^2 \right] \left\{ \frac{2 - \varphi}{1 - \varphi} \frac{1 - \beta}{\beta} \left[\frac{1 - \rho\varphi}{1 + \rho(1 - \varphi)} \left(\omega + \frac{\gamma_0 \gamma_s}{\gamma_1} \right) + 1 \right] + \right. \\ &\quad \left. + \omega + \frac{\gamma_0 \gamma_s}{\gamma_1} + \frac{2 - \varphi}{1 - \varphi} \right\}^{-1} \\ &= \left[H + \frac{1}{2} \left(\frac{\gamma_s}{\gamma_1} \right)^2 \right] f(\gamma_0, \gamma_1, \gamma_s, \omega, \beta, \varphi, \rho) \end{aligned} \quad (5.21)$$

where f has been defined appropriately. Note that f is decreasing in γ_0, γ_s and ω . In addition, f is increasing in γ_1, β and ρ . Given (5.21), the optimal steady-state level of sharing may be determined from (5.3), the optimal value of the steady-state stock of files may be determined from the law of motion (3.3), the optimal level of listening is determined by the binding listening constraint (3.1), and the optimal

amount of time away from the network may be determined by (5.20). These results are presented above in the text.

References

- Adar, Eytan and Bernardo A. Huberman**, “Free Riding on Gnutella,” *First Monday*, October 2000, 5 (10). http://www.firstmonday.dk/issues/issue5_10/adar/index.html.
- Alexander, Peter J.**, “New Technology and Market Structure: Evidence from the Music Recording Industry,” *Journal of Cultural Economics*, 1994, 18 (2), 113–123.
- ___, “Digital Distribution, Free Riders, and Market Structure: The Case of the Music Recording Industry,” *Review of Industrial Organization*, 2001. Forthcoming.
- Asvanund, Atip, Karen Clay, Ramayya Krishnan, and Michael Smith**, “Bigger May Not Be Better: An Empirical Analysis of Membership Rules in Peer-to-Peer Networks,” September 2001. Working Paper, H. John Heinz III School of Public Policy Management, Carnegie Mellon University.
- Becker, Gary**, “Altruism, Egoism, and Genetic Fitness: Economics and Sociobiology,” *Journal of Economic Literature*, 1976, 14 (3), 817–826.
- Fehr, Ernst and S. Gächter**, “Fairness and Retaliation: the Economics of Reciprocity,” *Journal of Economic Perspectives*, Summer 2000, 14 (3), 159–181.
- Golle, Philippe, Kevin Leyton-Brown, and Ilya Mironov**, “Incentives for Sharing in Peer-to-Peer Networks,” 2001. Working Paper, Computer Science Department, Stanford University.
- Ledyard, John**, “Public Goods: a Survey of Experimental Research,” in Alvin Roth and John Kregel, eds., *Handbook of Experimental Economics*, Princeton University Press, 1995.
- Olson, Mancur**, *The Logic of Collective Action: Public Goods and the Theory of Groups*, Cambridge, MA: Harvard University Press, 1965.
- Ostrom, Elinor**, “Collective Action and the Evolution of Social Norms,” *Journal of Economic Perspectives*, Summer 2000, 14 (3), 137–158.
- Stokey, Nancy L, Robert E Lucas, and Edward C Prescott**, *Recursive Methods in Economic Dynamics*, Cambridge, MA: Harvard University Press, 1989.

Sugden, Robert, "Reciprocity: the Supply of Public Goods Through Voluntary Contributions," *The Economic Journal*, December 1984, *94* (386), 772–787.

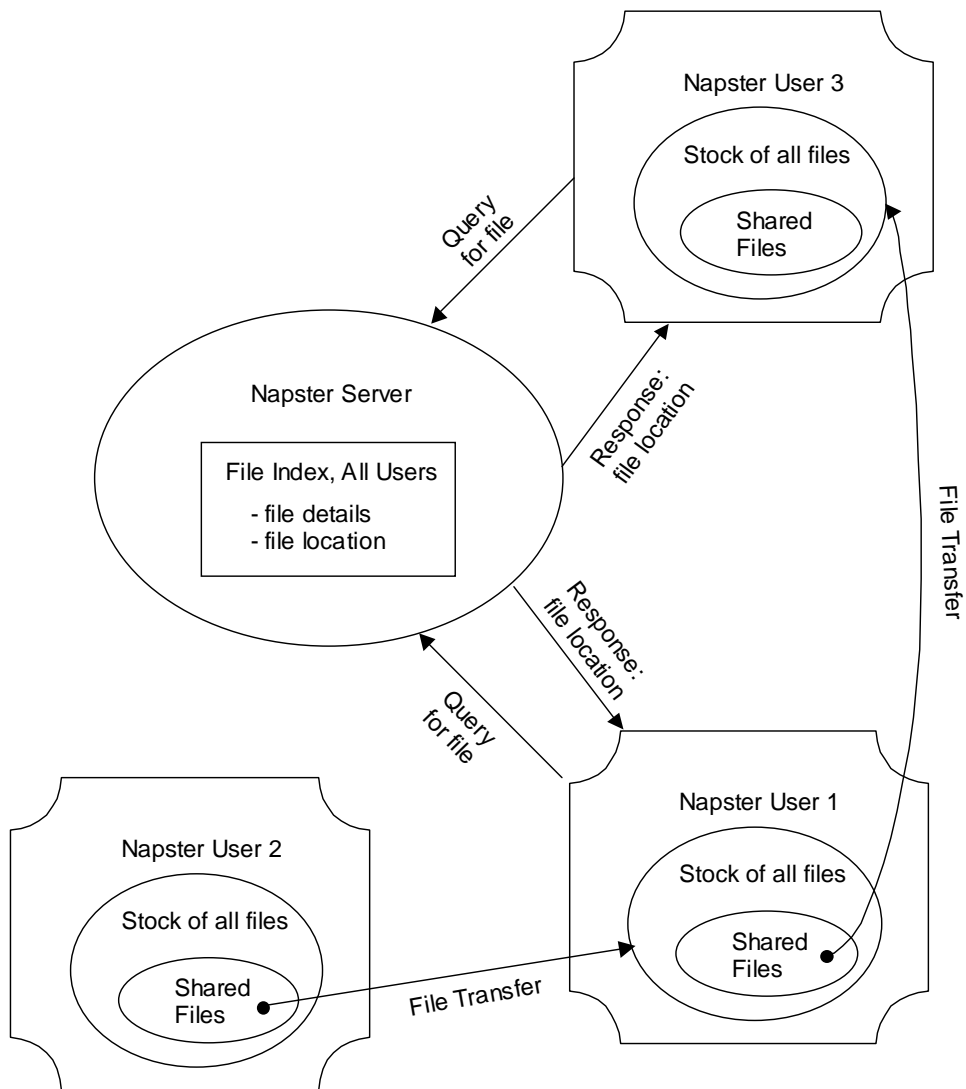


Figure 1: The Architecture of Napster