

The Impact of Free-Riding on Peer-to-Peer Networks

Ramayya Krishnan[†], Michael D. Smith[†], Zhulei Tang[‡], Rahul Telang[†]

[†] *H. John Heinz III School of Public Policy and Management, Carnegie Mellon University, Pittsburgh, PA, 15213.*

[‡] *Graduate School of Industrial Administration, Carnegie Mellon University, Pittsburgh, PA, 15213.*

{rk2x, mds, rtelang, zhuleit}@andrew.cmu.edu

Abstract

Peer-to-Peer networking is gaining popularity as a architecture for sharing information goods and other computing resources. However, these networks suffer from a high level of free-riding, whereby some users consume network resources without providing any network resources. The high levels of free-riding observed by several recent studies have led some to suggest the imminent collapse of these communities as a viable information sharing mechanism.

Our research develops analytic models to analyze the behavior of P2P networks in the presence of free-riding. In contrast to previous predictions, we find that P2P networks can operate effectively in the presence of significant free-riding. However, we also show that without external incentives, the level of free-riding in P2P networks will be higher than socially optimal. Our research also explores the implications of these findings for entrepreneurs, network designers, and copyright holders.

1. Introduction

Peer-to-peer (P2P) networks have attracted a significant amount of attention in the press as a popular network architecture for the sharing of information goods. Popular P2P sites for sharing consumer

information goods, including Napster, Kazaa, and Morpheus Music City, have attracted millions of global users to share their songs, films, software, and computer games. At the same time, a variety of entrepreneurs are developing commercial applications of P2P technology. Notable examples include Allcast, Blue Falcon Networks, Kontiki, and Uprizer for the distribution of streaming media content; Groove Networks for enterprise collaboration; and Bad Blue and Nextpage for enterprise information sharing.

While P2P networks vary in their architectural design and application domain, in all P2P networks files are transferred directly between the computers of users (a.k.a. peers) connected to the network. Further, once these files have been delivered, the user accessing the file becomes a provider of that content by default. Thus, in an ideal case, the provision of content on the network will scale to match the level of demand for the content. This characteristic also means that P2P networks can be modeled in the context of the economic concept of public goods. In contrast to private goods, public goods have the characteristics of non-excludability in supply (individuals can't be excluded from consuming the product) and non-rivalry in demand (one individual's consumption does not diminish another user's value of the product) (Hardin 1968). Clean air is a typical example of a public good

In the ideal case, P2P networks will exhibit both non-excludability (information is made available to all members of the

network) and non-rivalry (consumption by one user doesn't decrease network download possibilities in the absence of free-riding). However, in the presence of free-riding, P2P networks will exhibit levels of rivalry (Asvanund et al. 2002), which distinguishes them from either public goods or club goods.

Applying the public goods literatures to P2P networks, we are motivated by two interesting observations regarding P2P network performance. First, researchers have observed high levels of free-riding in P2P networks. While users share downloaded files by default, they can, and frequently do, turn off this feature to economize on their own private allocation of bandwidth. For example, Adar and Huberman (2000) find that 70% of the users in Gnutella version 0.4 network do not contribute content to the system and the top 1% of users contribute over 50% of the total amount of content. Asvanund et al. (2002) find evidence of free-riding in the context of OpenNap networks and Asvanund et al. (2003) document free-riding in the Gnutella version 0.6 network.

Traditional economic theory predicts that the free-rider problem causes inefficient provision of public goods, and calls for a central intervention to remedy this problem. In the context of P2P networks, Adar and Huberman (2000) observe "Free-riding leads to degradation of the system performance...if this trend continues copyright issues might become moot compared to the possible collapse of such systems."

In contrast, however, the second observed characteristic of P2P networks is that they appear to persist in spite of these high levels of free-riding. For example, according to a CNN report the number of P2P sites totals nearly 38,000, up 535 percent in the past year.

These empirical facts raise several research questions. First, given the predictions of Adar and Huberman (2000) above, can P2P networks persist in spite of

high levels of free-riding? Second, how much does free-riding reduce the utility of the network to its members versus a socially optimal outcome? And third, can network designers use economic incentives to improve network performance?

Answers to these questions are important for a variety of managerial audiences including consumer P2P networks implementing incentives to reduce free-riding, copyright holders implementing incentives (such as the threat of lawsuits against individual sharers) to increase free-riding, entrepreneurs designing enterprise-level P2P content management systems, and businesses deploying these systems.

Our analytic models show that P2P networks can persist in spite of free-riding. Even in the absence of external incentives to encourage sharing, some network users will find it individually optimal to share their content. The intuition is that sharing content serves to offload traffic at other servers the sharing peer is interested in accessing. At the same time, however, our models show that without external incentives, the level of free-riding will be above the socially optimal levels. Aggregate utility for network users could be improved by reducing free-riding. It is interesting to note that some free-riding can be socially optimal if the cost of sharing for the marginal user is greater than the benefit that user's content provides to the network. Finally, our models show that by differentiating the quality of service provided to sharers and non-sharers, network operators can achieve the socially optimal network outcome.

2. Literature and Theory

Our analysis combines different areas of research: the private provision of public goods, social networks, and peer-to-peer networks. We briefly review relevant papers from each literature below.

2.1. Private provision of public goods

There has been an extensive research literature on public goods (Samuelson 1954, 1955; Olson 1965; Hardin 1968; Clarke 1971; Groves, 1973; Groves and Ledyard 1977; Palfrey and Rosenthal 1984; Cornes and Sandler 1996; Sandler 1997). Following the definition of Mas-Colell et al., a public good is “a commodity for which use of a unit of the good by one agent does not preclude its use by other agents” (1995, p. 359).

As noted above, previous research on P2P networks suggests that content being shared in these networks have some of the properties of public goods (Asvanund et al. 2001; Krishnan et al. 2003). Specifically, most P2P implementations are non-excludable in the sense that if a resource is provided at all it is provided to the entire community. Likewise, resources provided over P2P networks are non-rivalrous to the extent that all users share. However, when some users free-ride, these same network resources become rivalrous.

Moreover, in most P2P networks, peers download content for free. Content in P2P networks is therefore subject to the free-rider problem: peers do not need to provide content in exchange of the one they desire. This characteristic in the distribution of public goods can lead to the well-known “tragedy of the commons” (Hardin 1968) where inadequate incentives lead to over-consumption and under-provision of community resources.

With respect to the private provision of public goods literature, our paper is different in the information setting from Palfrey and Rosenthal (1984, 1988) and Nitzan and Romano (1990). Our model assumes that individuals have complete information while their models assume that individuals have incomplete information about certain characteristics of others. We don't allow altruism in the basic model. We also extend Gradstein and Nitzan (1990) and Gradstein (1992) by explicitly modeling the

congestion function and variety function, which characterize the unique characteristics of peer-to-peer networks.

2.2. Peer-to-Peer papers

While peer-to-peer networks may significantly influence the distribution of information goods both in consumer markets and within businesses, few academic studies have analyzed the behavior of these networks, and most that have focus on technical issues, such as server architecture, data structure, and (Yang and Hector 2002).

One common theme in the P2P is a high incidence of free-riding — users consuming network resources without providing resources to the network. In a seminar paper, Adar and Huberman (2000) analyze free-riding in the Gnutella v0.4 network by sampling messages on the Gnutella network for a 24-hour period. The authors find that almost 70% of Gnutella users are free-riders. Moreover, the authors find that the top 1% of sharing hosts returns 50% of all responses. In more recent research, Asvanund et al (2003) find that 42% of Gnutella v0.6 users are free-riders.

Some authors have attempted to solve the free-riding problem by applying incentive design to P2P networks. The most common of these solutions apply telecommunications models for pricing network resources in the presence of scarcity (e.g., MacKie-Mason and Varian 1995; Wang, Peha and Sirbu 1996). For example, Golle et al (2001) attempt to provide a solution to free-riding in P2P networks by incorporating micro-payments. Similarly, Chandan and Hogenborn (2001) interpret P2P networking as a technology of wireless communications. They consider peer's decision as a node's decision to make a connection in the wireless network, just like in cellular systems. They argue since congestion overwhelms other network effects in a wireless network, peering technology with pricing will mitigate the congestion.

Other authors have proposed non-priced mechanisms based on reciprocity among users. For example, Vishnumurthy et al. (2003) implement a system they call KARMA. This system tracks both a user's contribution to and consumption of network resources. Each user receives a Karma score that governs their ability to consume network resources in the future. Kamvar et al. (2003) propose a similar scheme. In their system users can use their prior contributions of network resources to purchase the right to make query requests from other users.

It is important to note that each of these schemes have limitations when applied to many common P2P settings. In the case of incentives based on monetary payments, most P2P networks are collections of widely dispersed, essentially anonymous individuals, making person-to-person payments extremely difficult. Likewise incentive schemes based on persistent identifies are complicated by the anonymity of users and the ease with which users can modify their online identity.

Still, these mechanisms represent useful steps to address the widespread free-riding problem in P2P networks and their viability may be enhanced based on future development of P2P network technologies. In particular, several recent papers have focused on developing persistent identity and reputation mechanisms for distributed communities. Notable papers in this literature include Avery et al. (1999) and Dellarocas (2003). Papers applying these concepts to a P2P environment include Lai et al. (2003), Moreton and Twigg (2003), Cohen (2003), Kung and Wu (2003), and Woodward and Parkes (2003).

2.3. Social Networks and Virtual Communities

Our research is also related to the literature on social networks (Sharp 1997, Kraut et al. 1998, Kollock 1999, and Faraj and Wasko 2002) and computer-supported social networks (CSSNs) in the sense that

P2P networks have elements of both social and computer-supported interaction. However, our research differs from these literatures in terms of both the underlying assumptions about individual behavior and the analytic methods employed.

Computer networks that link people as well as machines are referred to as CSSNs (Wellman et al. 1996). The economies of communication and coordination are different in CSSNs than in face-to-face communication. Examples of such networks include Usenet, BBSs, chat, MUDs, and list servers. Recent research in this area has addressed a variety of topics including the productive or supportive relationships on-line with network members, the structure and composition in such networks, and the implications of the changes those networks bring to the societies (e.g., Wellman et al. 1996, Smith and Kollock 1999, and DiMaggio et al. 2001).

Significantly, this literature includes several analyses of the provision of resources that share similarities to public goods (see Wasko 2003 for an excellent review of this literature). Faraj and Wasko (2002) investigate the motivation for and ability of knowledge exchange using three newsgroups on the Usenet. They also test the relationship between social capital and knowledge acquisition and contribution. Butler (2001) presents a resource-based model for social structures to analyze the interacting effects of membership size on benefit provision process by sampling from an e-mail-based Internet listservs. Lakhani and von Hippel (2000) explore the provision of online technical support among users of Apache open source software.

3. A Model of Contribution of Content in Peer-to-Peer Networks

3.1. Static Symmetric Contribution Game With Complete Information

In this section we provide a general overview of our model and an intuition for our basic results. A more detailed

development of the model is available at <http://www.heinz.cmu.edu/~mds/p2pa>.

In this baseline model, we consider a network in which each user has a single endowment of unique content (we relax this assumption below to allow for duplication of content). Each user independently decides whether to share their content based on their private utility. Sharing implies a cost by reducing the user's private bandwidth available for downloading. However, sharing reduces the traffic other peers place on other nodes in the network, thus increasing the user's private utility. In effect, the sharing decision serves as a means to redistribute traffic in our model. It provides a rational incentive for some users to share content even without altruism or other external incentives. Each user compares the benefit and cost to formulate her optimal strategy.

In our formal model, we consider a static game where n users join the network simultaneously with their unique endowment. In each period, each user independently decides whether or not to share content, and each user demands one unit of content randomly from the other $n-1$ users. If the user does not share, she can still download content from other users. We refer to downloading without sharing as *free-riding*. Sharing incurs a constant lump-sum cost c , which is common knowledge to the players. We note that this cost could either derive from explicit costs, such as the cost to obtain network capacity, or from implicit sources such as the cost due to legal risks from sharing. We discuss the implication of these dual sources of costs below. One could also imagine a cost related to the opportunity cost of time where cost would be a function of network congestion. We plan to explore cost related to congestion in future work.

In this setup, each user gets some utility from the content she downloads. This utility depends on the congestion in the network and the variety of content available on the network.

Each user's strategy is a decision of whether to share or not. By definition of Nash equilibrium, the equilibrium of this game is a set of user strategies such that each strategy is optimal given other users' strategies. We represent the number of peers who choose to share as k .

The user will choose their strategy to maximize their utility function. Formally, the utility function is the benefit the user gains from using the network minus the cost of sharing if they share. As above, the benefit from using the network is the product of the likelihood that a user can download the content she desires ($f(n,k)$), and the value a user derives from content available on the network that they could potentially download ($v(k)$).

Congestion in the network occurs at different stages when a user tries to download content from the network. For example, congestion can be measured as login congestion, query congestion, download attempt congestion and download speed congestion (Asvanund et al. 2002). When a user joins the network, she automatically intensifies login congestion and query congestion of the network irrespective of whether she contributes. However, download attempt congestion and download speed congestion are influenced by available content in the network. Contributors alleviate network congestion (increase the likelihood of a successful download) by providing more content. However, the marginal congestion decreases with the amount of content.

Thus the likelihood of being able to download content is a function of both the number of peers on the network and the number of sharers. The likelihood of being able to download content will increase in the number of sharers, but at a decreasing rate, thus $f_k > 0$ and $f_{kk} < 0$. Further, we note that a network with no sharers will have a zero probability of a successful download, thus, $f(n,k) = 0$ when $k = 0$.

Similarly, the value a user derives from the content on the network that they could

potentially download is an increasing function of the number of sharers only. This term is not a function of n because peers who join the network but do not share (i.e., free-riders) do not change the variety of content on the network. This value is increasing in the number of sharers, but at a decreasing rate, thus $v_k > 0$ and $v_{kk} < 0$. This implies that the marginal benefit a user gains from an additional sharer declines in the number of sharers.

Finally we assume that the users' net benefits from using the network are concave in the number of network members, and that in some limit of n , users prefer a larger network than smaller one because their probability of getting desired content from a larger network is higher. As above, this would be true if at some point the marginal benefit of adding an additional user to the network declined and is consistent with Asvanund et al. (2002).

3.2. Characterization of Equilibrium

Under this game setup we have three unique pure strategy equilibria depending on the size of the cost of sharing (c). These equilibria are determined based on whether a peer's marginal benefit to sharing justifies the cost of sharing. Note that because of the concavity of the benefit function, the marginal benefit to sharing for the next peer will decline as more peers share.

This makes sense intuitively. Peers obtain direct benefits to sharing because their sharing reduces the congestion function, and the marginal benefit is smaller the more peers are already sharing. (Note that since they are sharing their own content, they do not obtain direct benefits from increases in the content variety function.)

From this, one can show that the first class of equilibria is such that if sharing cost is sufficiently small, there is an equilibrium where everyone shares. Under this equilibrium, the gain to the n th peer from sharing is larger than their cost of sharing and likewise for the 1st through $n-1$ th users. While this equilibrium is theoretically

possible, in practice it is rarely seen. As noted above, almost all of the common peer-to-peer networks exhibit high levels of free-riding.

A second class of equilibria occurs when costs are sufficiently high such that no one shares. Such an equilibrium would occur if contingent on one peer sharing, a second peer would gain less benefit from sharing than their cost of sharing. In this situation, the second peer would never share and the network would never form since $k=1$ is not a stable equilibrium because a peer would be better off not sharing if they are the only contributor.

It is interesting to note that one could reasonably characterize the Recording Industry Association of America (RIAA) as adopting a strategy of intentionally increasing the cost of sharing for individual participating in consumer P2P networks. As noted above, the risk from legal prosecution could be incorporated into our cost function. By exacting high judgments from P2P users who share content, the RIAA may be intentionally trying to raise the implicit cost of sharing on all P2P users and thereby decreasing the level of sharing on the network, if not eliminate it altogether. However, early published results suggest that this strategy has yet to take hold for the RIAA. The number of people sharing files on P2P networks appears to be holding steady (for example, Harmon and Schwartz (2003) and Kopytoff (2003)).

The third class of equilibria occurs for intermediate values of sharing costs. For values between full sharing and no sharing there exists a set of equilibria where some, but not all users will share (i.e., some users will free-ride and their free-riding behavior is sustainable in equilibrium).

To see this how this outcome might occur, consider the decision problem of the k th user who contributes her content. It must be the case that utility from the increased probability of successfully downloading content from the other $k-1$ sharers net the cost of sharing is greater than the utility

from not sharing. Likewise, consider the decision problem of the $(k^* + 1)^{\text{th}}$ user who does not contribute her content. In her case, it must be true that the utility from the increased probability of successfully downloading content from the other k sharers net the cost of sharing is not greater than the utility from not sharing. When this occurs, the k^{th} user will share and the $(k+1)^{\text{th}}$ user will not share. Again, the concavity of the benefit function ensures that this interior solution exists.

Thus, this result implies that for values of c within a certain range, a network will form where some users will be willing to share content even if others are free-riding. There are two possible reasons for this. First, in equilibrium they are better off than non-sharing. Second, when they decide to share their content, they may believe that some other users might make the same decision. It is also important to note that because each per is homogeneous in our model, there are $\binom{n}{k^*}$ such equilibria, unique up to a permutation of the players' strategies. Thus, it may be difficult to ascertain beforehand which k users might decide to share. But over time, the users will establish a mechanism where some of them may share in a period while the rest free-ride.

4. Social Optimality

In Section 3, we discussed our results that peer-to-peer networks can persist in equilibrium in spite of free-riding. In this section we discuss how the outcome in a network with free riding compares to the socially optimal outcome. Specifically, we find that social utility under an equilibrium with free-riding is below the socially optimal level of utility.

To solve for the social optimal level of sharing, first note that the players' strategies are socially optimal if the aggregate of individual utilities is maximized. That is, if the sum of their benefits minus the sum of

their costs (the number of sharers times the cost of sharing) is maximized.

There are two interesting things to note about the socially optimal outcome in our model.

First, we find that the equilibrium level of sharing in the absence of external incentives is lower than the socially optimal level of sharing. That is, content is underprovided compared with social optimum. This is an extension of typical results in the public goods literature that the private provision of public goods is socially inefficient. Without external incentives, individual network users do not internalize the externality they impose on other network users when deciding whether to share. As a result, their sharing levels are below socially optimal levels.

Since the equilibrium number of sharers and the socially optimal number of sharers differ, a natural question is how we can motivate peers to contribute more content so that social optimality can be obtained? In this vein, we note that several P2P networks have adopted external incentive mechanisms to encourage users to share. For example, the most recent *Kazaa* client tracks the sharing level of each peer and moderates a peer's ability to share based on the level of resources they provide to the network.

Second, we find that it is not necessarily the case that the social optimum does not require that everyone contribute. Indeed, the socially optimal level of contributors depends on the ratio of the cost of sharing to the number of network participants $(\frac{c}{n})$. When this ratio varies, the social optimum also varies. However, whenever full sharing is optimal to the individual, it must be also optimal to the whole society.

This finding occurs, again, because of the concavity of the benefit function. The marginal benefit of sharing to the network declines as the number of users increases. At

some point it is possible that the marginal benefit an additional sharer adds to the network does not justify the cost this peer would incur by sharing their content. In this circumstance, the socially optimal outcome would be for this user not to share.

5. Conclusions

We examine the peer-to-peer networks in a game theoretic model. Individual users face trade-offs between sharing and not sharing. In a static game with complete information, we show that in equilibrium some peers will not share their files. Thus, in contrast to prior predictions in the literature, we find that networks can sustain free-riding in equilibrium. In spite of free-riding, some peers will find it individually rational to share as a way to reduce congestion on other peers they are interested in accessing.

We then compare the outcome of our base model to the socially optimal outcome. We find that in a socially optimal outcome, some users may still be free-riders. This is because at some point it is possible that the cost of sharing for the marginal user may not justify the value this user provides to the network. However, we also find that the level of free-riding in the socially optimal case is less than the level of free-riding in the base case. In the absence of incentives on user behavior, network goods will be under-provided.

This analysis should have value for entrepreneurs and designers of P2P networks, for copyright holders seeking to protect their copyrights in the presence of P2P networks, and for businesses implementing P2P content management systems. For network designers, our results suggest that while networks can persist in spite of free-riding, the aggregate utility of the network would be improved in the presence of incentives to encourage more users to share their content. For copyright holders, our results suggest that increasing the cost of sharing can reduce the number of sharers and above a certain point lead to network collapse. One way copyright

holders appear to be increasing the implicit cost of sharing is by increasing the legal risks to individual network users from sharing copyrighted information, as noted by the recent prosecution of several college students found to be sharing copyrighted materials online. For enterprises implementing P2P network products, our results suggest that forcing all network members to share may not be the socially optimal outcome. It may be possible that the cost of sharing to an individual user (in terms of reduced bandwidth, storage, or processing power) may not justify the value this user would provide to other network members.

Future research could extend our results in a variety of ways. First, it may be possible to extend our static model to take into account user incentives in a multi-period game. Second, it would be interesting to analyze what types of incentives can be best used to achieve the socially optimal outcome. Third, it may be possible to relax some of the assumptions employed in this model. For example, a model where sharing cost is a function of the number of members of the network could be explored. It also might be possible to model the impact of another tactic copyright holders are using to reduce network performance — introducing fake content into the network. It might also be possible to explore how changes in the characteristics of heterogeneous peers would impact network performance as a way to predict how changes in the profile of P2P users would impact network outcomes.

6. References

- [1] Adar, E, B. A. Huberman. 2000. Free-riding on Gnutella. *First Monday* 5(10).
- [2] Asvanund, Atip, Karen Clay, Ramayya Krishnan, Michael Smith. 2002. An Empirical Analysis of Network Externalities in Peer-To-Peer Music Sharing Networks. *Proceedings of the 23rd International Conference on Information*

Systems (ICIS), Barcelona, Spain, December 15-18.

[3] Asvanund, A., R. Krishnan, M. D. Smith, R. Telang. 2003. Intelligent Club Management in Peer-to-Peer Networks. Working Paper, H. John Heinz III School of Public Policy and Management, Carnegie Mellon University, Pittsburgh, PA.

[4] Avery, C., P. Resnick, and R. Zeckhauser. 1999. The Market for Evaluations. *The American Economic Review* **89**(3) 564-584.

[5] Butler, Brian S. 2001. Membership Size, Communication Activity, and Sustainability: A Resource-Based Model of Online Social Structures. *Information Systems Research* **12**(4) 346-362.

[6] Chandan, S., Hogendorn C. 2001. The Bucket Brigade: Pricing and Network Externalities in Peer-to-Peer Communications Networks. Telecommunications Policy Research Conference, Alexandria, VA, October 27-29.

[7] Clarke, E.H. 1971. Multipart Pricing of Public Goods. *Public Choice* **11** 19-33.

[8] Cornes, R.C. and Sndler, T. 1996. The Theory of Externalities, Public Goods, and Club Goods, Cambridge University Press (Cambridge), 2nd edition.

[9] Dellarocas, Chrysanthos. 2003. The Digitization of Word-of-Mouth: Promise and Challenges of Online Reputation Mechanisms. *Management Science* **49**(11) Forthcoming.

[10] Faraj, S., Wasko, M. 2002. The Web of Knowledge: An Investigation of Knowledge Exchange in Networks of Practice. Working Paper, University of Maryland, College Park, MD.

[11] Golle, P., K. Leyton-Brown, I. Mironov. 2001. Incentives for Sharing in Peer-to-Peer Networks. Working Paper, Stanford University, Palo Alto, CA.

[12] Gradstein, M. 1992. Time Dynamics and Incomplete Information in the Private

Provision of Public Goods. *Journal of Political Economy*, **100** 581-97.

[13] Gradstein, M., Nitzan, S. 1990. Binary Participation and Incremental Provision of Public Goods. *Social Choice and Welfare*, **7** 171-92.

[14] Groves, T. 1973. Incentives in Teams. *Econometrica* **41**: 617-31.

[15] Groves, T. and Ledyard, J. 1977. Optimal Allocation of Public Goods: A Solution to the 'Free Rider' Problem. *Econometrica*, **45** 783-809.

[16] Hardin, G. 1968. The Tragedy of the Commons. *Science* **162** 1243-48.

[17] Harmon, A., J. Schwartz. 2003. Music File Sharers Keep Sharing. *New York Times*, September 19.

[18] Kamvar, Sepandar, Beverly Yang, Hector Garcia-Molina. 2003. Addressing the Non-Cooperation Problem in Competitive P2P Systems. Working Paper, Stanford University, Palo Alto, CA.

[19] Kollock, P. 1999. The Economies of Online Cooperation: gifts and Public Goods in Cyberspace. Chapter 9, *Communities in Cyberspace*, M.A. Smith and P. Kollock eds. London, Routledge.

[20] Kopytoff, V. 2003. File sharers not changing their tune: RIAA court tactic has limited success. *San Francisco Chronicle*, September 10.

[21] Kraut, R. Steinfield, C. Chan, A. Butler, B. and Hoag, A. 1998. Coordination and Virtualization: The Role of Electronic Networks and Personal Relationships, *JCMC* **3**(4).

[22] Krishnan, Ramayya, Michael D. Smith, Rahul Telang. 2003. The Economics of Peer-to-Peer Networks. Working Paper, H. John Heinz III School of Public Policy and Management, Carnegie Mellon University, Pittsburgh, PA.

[23] Kung, H. T., Chun-Hsin Wu. 2003. Differentiated Admission for Peer-to-Peer Systems: Incentivizing Peers to Contribute

Their Resources. Working Paper Harvard University, Cambridge, MA.

[24] Lai, Kevin, Michal Feldman, Ion Stoica, and John Chuang. 2003. Incentives for sharing in P2P networks”, Workshop on Economics of Peer to Peer Network. Working Paper, University of California at Berkeley, Berkeley CA.

[25] Lakhani, Karim and von Hippel, E. (2000) “How Open Source Software Works: “Free” User-to-user Assistance”. MIT Sloan School of Management Working Paper #4117.

[26] MacKie-Mason, Jeffrey K., Hal Varian. 1995. Pricing Congestible Resources. *IEEE Journal of Selected Areas in Communications* **13**(7) 1141-49.

[27] Mas-Colell, A., M.D. Whinston, J.R. Green. 1995. *Microeconomic Theory*. Oxford University Press, Oxford, New York.

[28] Moreton, Tim and Andrew Twigg. 2003. Trading in Trust, Tokens, and Stamps. Working Paper, Cambridge University, Cambridge, United Kingdom.

[29] Nitzman, S. and Romano, R.E. 1990. Private Provision of a Discrete Public Good with Uncertain Cost. *Journal of Public Economics*, **42** 357-70.

[30] Olson, M. (1965), *The Logic of Collective Action*, Harvard University Press (Cambridge, MA).

[31] Palfrey, T. and Rosenthal, H. 1984. Participation and the Provision of Discrete Public Goods: A Strategic Analysis. *Journal of Public Economics*, **24** 171-193.

[32] Palfrey, T. and Rosenthal, H. 1988. Private Incentives in Social Dilemmas: The Effects of Incomplete Information and Altruism, *Journal of Public Economics* **35** 309-332.

[33] Samuelson, P.A. 1954. The Pure Theory of Public Expenditure. *Review of Economics and Statistics* **36** 387-9.

[34] Samuelson, P.A. 1955. A diagrammatic Exposition of a Theory of Public Expenditure. *Review of Economics and Statistics* **37** 350-6.

[35] Sandler, T. 1992. *Collective Action: Theory and Applications*, University of Michigan Press, Ann Arbor, Michigan.

[36] Sharp, J. 1997. *Communities of Practice: A Review of the Literature*. March.

[37] Smith, M.A., P. Kollock. 1999. *Communities in Cyberspace*. Routledge: London.

[38] Vishnumurthy, Vivek, Sangeeth Chandrakumar, Emin Gun Sirer. 2003. KARMA: A Secure Economic Framework for Peer-to-Peer Resource Sharing. Working Paper, Cornell University, Ithaca, NY.

[39] Wang, Qiong, Jon M. Peha, Marvin Sirbu. 1997. Optimal Pricing for Integrated Services Networks, McKnight and Bailey eds. *Internet Economics*. MIT Press, Cambridge, MA, 353-376.

[40] Wasko, Molly M. 2003. Virtual Commons or Public Goods? Applying Theories of Social Dilemmas, Collective Action, and Public Goods to Electronic Networks of Practice. Working Paper, Florida State University, Tallahassee, Florida.

[41] Wellman, B. Salaff, J. Dimitrova, D. Garton, D. Gulia, M. and Haythornthwaite, C. 1996. Computer Networks as Social Networks: collaborative Work, Telework, and Virtual Community. *Annual Review of Sociology* **22** 213-38.

[42] Woodward, C. Jason, David C. Parkes. 2003. Strategyproof Mechanisms for Ad Hoc Network Formation. Working Paper, Harvard University, Cambridge, MA.

[43] Yang, B., G. Hector. 2002. Efficient Search in Peer-to-Peer Networks. Working Paper, Stanford University, Palo Alto, CA.