Incentives for Large Peer-to-Peer Systems

Costas Courcoubetis and Richard Weber

***Abstract*—We consider problems of provisioning an excludable public good amongst n potential members of a peer-to-peer system who are able to communicate information about their private preferences for the good. The cost of provisioning the good in quantity Q depends on Q, and may also depend on n, or on the final number of participating peers m. Our aim is to maximize the expected social welfare in a way that is incentive compatible, rational and budget-balanced. Although it is unfortunately almost never possible to calculate or implement a truely optimal mechanism design, we show that as the number of participants becomes large the expected social welfare that can be obtained by the optimal design is at most a factor $1 + O(1/n)$ or $1 + O(1/\sqrt{n})$ greater than that which can be obtained with a very simple scheme that requires only payment of a fixed contribution from any agent who joins the system as a participating peer. Our first application is to a model of file sharing, in which the public good is content availability; the second concerns a problem of peering wireless local area networks, in which the public good is the availability of connectivity for roaming peers. In both problems, we can cope with the requirement that the payments be made in kind, rather than in cash.**

***Index Terms*—File sharing, incentive mechanisms, mechanism design (MD), peering, peer-to-peer (P2P), wireless local area networks (WLANs).**

I. Introduction

The design of a peer-to-peer (P2P) system poses many interesting questions. If the quantity of the service it provides is Q, and this is a design parameter, what should Q be? How should peers contribute to the cost of providing Q and how can the "free-rider problem" be avoided? In this paper, we consider how these questions might be answered for models of two possible P2P systems. The first is a file sharing system, in which Q is "the number of distinct files shared." The second is a system for sharing wireless local area networks (WLANs) resources when peers roam in some geographic location away from their home local area network (LAN). Now, Q is "the availability of connectivity for roaming peers."

Let us think of the service offered by a P2P system as an economic good, and imagine that the agents (or peers) know their own differing preferences for the quantity (or quality) of service Q that the system can provide. It is reasonable that a peer who values the system more should make a greater contribution to its cost, or share more of his own resources. But how do we force each peer to truthfully tell us how much he values the system?

Each of the P2P systems that we consider in this paper has two important characteristics. First, viewed as an economic good, it is *nonrivalrous*, meaning that one agent's consumption of the good does not decrease its utility to another agent (at least to a reasonable first approximation). For instance, content availability in a file sharing system is not reduced by peers downloading files. Similarly, the probability for obtaining wireless connectivity by a roaming agent randomly located in an area partially covered by WLANs is not affected by the fact that other such agents request similar service. A second characteristic is that the good is *excludable*, meaning that it is possible to prevent particular agents from having access to the good (e.g., by requiring a password to access the system). Economists often call such goods "excludable public goods" or "club goods." A critical aspect in our modeling approach is that we adopt a public good model for describing the value obtained by the peers from using the system.

In economics, our problem is known as the mechanism design (MD) problem. We need to elicit truthful information from the agents, or peers, regarding their valuation of the service, choose Q, and decide which peers are allowed to participate in using the system, and how much each should contribute to covering the cost of building the system at level Q. This is to be done to produce the greatest possible expected social welfare. While the full solution of this problem is extremely complex and not easily solved in practice, we show in Section II that as the number of agents becomes large, there is a good solution to this problem that takes a very simple form. We merely require each agent to pay the same fixed fee toward the total cost, and exclude agents that are unwilling to do so. In the cases we consider, this fee need not be paid in cash, but can be paid in kind, i.e., by contributing resources to provide a fixed part of the overall service. Such a simple contribution policy is easy to implement and requires no centralized implementation. The only information which the system designer needs to compute the fixed fee is the distribution of the agents' valuations for the service.

Although there are prior results pointing to the fact that, in the limit, optimal incentive payments reduce to fixed and equal contributions by all peers, there are no result to say how such simple policies perform as a function of the number of peers n. We prove results that allow one to obtain tight bounds. The positive result is that the optimal policies obtain expected social welfare that is only a multiplicative factor $(1+O(1/n)$ or $(1+O(1/\sqrt{n})$ better than a very simple class of policies which are easy to compute. These policies are simple: the system designer has only to declare a fixed participation fee and the expected system size to all participants. Each peer then makes a simple decision, to

Manuscript received September 1, 2005; revised October 1, 2005. This work was supported in part by the IST FET Coordination Action ACCA under Grant IST-6475.

C. Courcoubetis is with the Athens University of Economics and Business, Athens 10434, Greece.

R. Weber is with the Statistical Laboratory, Centre for Mathematical Science, University of Cambridge, Cambridge CB3 0WB, U.K. (e-mail: r.r.weber@statslab.cam.ac.uk).

Digital Object Identifier 10.1109/JSAC.2006.872885

participate or not. One may think of more complex equal contribution policies, where the system size and the fee may depend on the number of final participants, requiring more complex calculations both from the system designer and from the peers. Our results show that the gain of using such policies is small and the per capita gain tends to zero quickly. We also show how to extend the system to handle multiple constraints on how the cost of providing Q should be covered, and also to reflect a dependence on the number of final participants. We apply this to model file sharing and peering WLANs, and, we derive the simple optimization problems from which one can compute the optimal incentive policies. We also show that given an incentive compatible mechanism that equates the expected cost of providing Q with the expected total payments that peers make (the so-called ex-ante budget balance constraint) there exists an equally good mechanism under which the actual cost of providing Q is exactly covered by the total payment that the participating peers make (ex-post budget balance).

Our results must be taken with a degree of salt in designing a practical system. In order to obtain an asymptotic solution to the notoriously difficult MD problem, we have had to use a simple model for our P2P systems, in which each peer's valuation of the service is identified by a single parameter. One can think of it as the "first-order term" of a more detailed models that more accurately take account of cost, utility, and peer interactions. However, such more complex models are likely to lead to intractable game design problems. Even in our simple setup, the single parameter MD problem is not analytically solvable (although, we have managed to solve it completely for a particular case in Example 1). Hence, showing that such simple policies are asymptotically optimal has great value by suggesting that simple practical policies for incentives may suffice in practice.

The paper is organized as follows. In Section II, we describe a model for an excludable public good and how to solve the optimization problem of finding an expected social welfare maximizing budget-balanced incentive compatible mechanism. Section III presents our asymptotically optimal scheme and states our main theorem concerning it. In Section III-B, we work through a numerical example that illustrates the ideas. In Section III-C, we discuss other schemes that also require each participant to make the same payment, and Section III-D presents an extension to a model in which there are k "types"' of peer, and k constraints, which impose conditions that peers of the same type must cover a certain aspect of the cost. Sections IV and V contain our applications to P2P file sharing systems and WLANS, respectively. Section VI looks at questions of stability and convergence when some information must be learned. Our conclusions are in Section VII.

Some problem formulation and longer proofs are placed in the appendix. Appendix I contains a derivation of the optimization problem of maximizing expected social welfare, and Appendix II justifies the fact that it can be solve using Lagrangian methods. Appendix III contains proof of a technical lemma. The proofs of our main theorems 1 and 2 are in Appendix IV and V. Appendix VI contains Theorem 3, which is a significant new theoretical result stating even when the mechanism may use exclusions, the existence of an ex-ante cost-covering incentive compatible mechanism implies the existence of an equally good ex-post cost-covering incentive compatible mechanism. Appendix VII contains detailed calculations for Section III-C.

A. Related Work

Our work has been motivated by that of Hellwig [1] and Norman [2], who have also investigated asymptotic properties of optimal solutions to the MD problem in the public goods context. Our contribution is to focus on the form of the limiting solution and obtain results by arguments that are simpler and also permit some extensions. We show that, depending upon certain assumptions regarding the forms of the utility and cost functions, the system obtains expected social welfare such that an optimal incentive mechanism could obtain no more than a multiplicative factor of $1+O(1/n)$ or $1+O(1/\sqrt{n})$ better. Previous research did not obtain such exact bounds on the performance of the limiting policy, nor was it able to handle multiple constraints. Our proof technique is much simpler and focuses exactly on that aspect. We also treat the case where the cost may depend on the final number of participants (instead of the number of potential participants). This type of congestion cost is in fact the basis of our WLAN model in Section V.

Golle *et al.* [3] made a first effort to model the utilities and costs associated with the participation in a P2P file sharing system and using game theoretic analysis proposed the use of micropayments for achieving the desirable equilibria. Buragohain *et al.* [4] follow a game theoretic approach and study the equilibria and corresponding efficiency achieved in a P2P file sharing system, based on a similar utility and cost model to ours, assuming that the system can enforce a level of reciprocity. Other relevant modeling references are [5]–[7]. Antoniadis *et al.* [8] have attempted to compare different incentive schemes, one of them being the simple contribution scheme analyzed in this paper. Similarly, [9] and [10] contain more elaborate applications to file sharing and WLANs using our simple contribution scheme. Regarding the asymptotic results, all this work referenced [11], an unpublished version of this current paper where the proofs did not yet obtain the $1+O(1/n)$ bound on the performance of the asymptotic, which we now know is tight.

From a practical perspective, a P2P system designer has to deal with the fact that he is unable to rely on trusted software[1] or on central entities for monitoring and accounting. Thus, a significant part of the research literature in P2P economics studies the game theoretic and implementation issues related to the effective accounting of peers' transactions (by means of reputation, credits, etc.) in such a fully distributed and untrusted environment, see for example [12]–[14]. Such approaches adhere to the principle that peers should benefit from the system in proportion to the extent that they contribute to actual downloads and uploads, and be able to identify and punish free riders by reducing their reputation and restricting their downloads. The MMAPPS

[1]Kazaa is a characteristic example of a real world application that tried to implement a reciprocative incentive mechanism by giving priority to peers that contribute more by having less downloads than uploads, which failed due to a hacked version of its software. This version, Kazaa-lite, was assigning by default the maximum of credits to its users to enhance their priority. Original Kazaa users started blocking users with large amounts of credit considering them fraudulent.

Consortium [15] have discussed the difficulty in providing and storing reliable accounting information for enforcing incentive policies in P2P systems.

An interesting system that follows similar incentive policies as the ones proposed in this paper is Direct Connect. This P2P application relies on central control exercised by a special peer subgroup enforcing specific minimum contribution rules, and excluding peers that are found to contribute less based on their IP addresses.

II. An Excludable Public Good Model

Consider an excludable public good as described in Section I. Suppose that to provide the good in quantity Q costs $c(n, Q)$. Once it is provided, the net benefit to agent i, if he is permitted to use it, is

$$\theta_i u(Q) - p_i$$

where p_i is the payment he makes toward the cost of providing the good. Here, $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_n)$ is a vector of "preference parameters," which are assumed to be independent and identically distributed samples from a distribution on [0, 1] with distribution function F. This distribution F is known to all agents, but the value of θ_i is known to agent i alone. We suppose that $u(Q)$ and $c(n, Q)$ are, respectively, concave and convex functions of Q. Notice the simplicity of our model: the value agent i obtains is $\theta_i u(Q)$, where $u(Q)$ is common to all agents. Differentiation is through a single parameter, θ_i. One can easily refine this model to assume that agents are not of the same type (i.e., not characterized by the same distribution F), but belong to some finite number of types, each characterized by a different distribution of its preference parameter. In such a model the type of a peer is common information.

Knowing n and F, a social planner wishes to design a mechanism which, as a function of the declared values $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_n)$, sets Q, and determines which agents may use the good and what fees they should pay if they do. These fees must cover the cost $c(n, Q)$. Knowing the mechanism the planner will use, each agent i declares θ_i to his best advantage. The mechanism then sets $Q(\boldsymbol{\theta})$ and decides which agents may use the good and which are to be excluded from using it. Let $\pi_i(\boldsymbol{\theta})$ be the probability with which the mechanism includes agent i given the announced preferences $\boldsymbol{\theta}$. If agent i is excluded from using the good, then $\pi_i(\boldsymbol{\theta}) = 0$. If he is allowed to use it, then $\pi_i(\boldsymbol{\theta}) = 1$ and he must pay a fee $p_i(\boldsymbol{\theta})$. If exclusion is not an option for the planner, then we simply make the restriction $\pi_i(\boldsymbol{\theta}) = 1$ for all $\boldsymbol{\theta}, i$.

The mechanism that the social planner chooses to implement defines a game among the agents, and given that the agents make rational responses, it has a Nash equilibrium. The planner's problem is to design the mechanism so that this equilibrium is a point of maximum economic efficiency, where efficiency is measured by expected social welfare. Let us now put this in mathematical terms. In (1) and (2) that follow the expectation is taken over $\boldsymbol{\theta}$ and in (3) and (4) it is taken over $\boldsymbol{\theta}_{-i}$, where this denotes all the preferences parameters apart from θ_i. The problem is to maximize expected social welfare

$$\underset{\pi_1(\cdot),\ldots,\pi_n(\cdot),Q(\cdot)}{\text{maximize}}\; E\left[\sum_{i=1}^{n} \pi_i(\boldsymbol{\theta})\theta_i u\left(Q(\boldsymbol{\theta})\right) - c\left(n, Q(\boldsymbol{\theta})\right)\right] \quad (1)$$

subject to a "ex-ante cost-covering constraint," which says that the expected payments must at least cover the expected cost

$$E\left[\sum_{i=1}^{n} \pi_i(\boldsymbol{\theta}) p_i(\boldsymbol{\theta}) - c\left(n, Q(\boldsymbol{\theta})\right)\right] \geq 0 \quad (2)$$

an "individual rationality" constraints, which says each agent can expect positive net benefit

$$E_{\boldsymbol{\theta}_{-i}}\left[\pi_i(\theta_i, \boldsymbol{\theta}_{-i})\left\{\theta_i u\left(Q(\theta_i, \boldsymbol{\theta}_{-i}\right) - p_i(\theta_i, \boldsymbol{\theta}_{-i})\right\}\right] \geq 0, \quad \text{for all } \theta_i \quad (3)$$

and "incentive compatibility" constraints, such that each agent i does best by declaring his true θ_i rather than attempting some sort of "free-riding" by declaring some other θ_i'

$$\begin{aligned} &E_{\boldsymbol{\theta}_{-i}}\left[\pi_i(\theta_i, \boldsymbol{\theta}_{-i})\left\{\theta_i u\left(Q(\theta_i, \boldsymbol{\theta}_{-i}\right) - p_i(\theta_i, \boldsymbol{\theta}_{-i})\right\}\right] \\ &\quad \geq E\left[\pi_i\left(\theta_i', \boldsymbol{\theta}_{-i}\right)\left\{\theta_i u\left(Q\left(\theta_i', \boldsymbol{\theta}_{-i}\right) - p_i\left(\theta_i', \boldsymbol{\theta}_{-i}\right)\right\}\right], \\ &\text{for all } i \text{ and } \theta_i'. \end{aligned} \quad (4)$$

It is the incentive compatibility constraint (4) that ensures that the agents declare the true values of their preference parameters. This is assumed in (1)–(3). It is natural to ask whether imposing (4) means that the expected social welfare cannot be as great as if we optimized over all possible mechanisms, including those that are not incentive compatible. However, there is a well-known "revelation principle" in the theory of mechanism design which states that any Nash equilibrium that can obtained by some mechanism can also be obtained by an incentive compatible mechanism. This justifies the restriction to incentive compatible mechanisms and leads to the following simple and useful analytic characterization.

We now consider the problem of maximizing expected social welfare subject to the constraint that the mechanism is ex-ante cost-covering, individually rational, and incentive compatible.[2] Let us define

$$g(\theta_i) = \theta_i - \frac{1 - F_i(\theta_i)}{f(\theta_i)}. \quad (5)$$

Using standard MD theory, it is shown in Appendix I that our problem reduces to single constraint problem, namely of maximizing (1) subject to the constraint

$$E\left[\sum_i \pi_i(\boldsymbol{\theta}) g(\theta_i) u\left(Q(\boldsymbol{\theta})\right) - c\left(n, Q(\boldsymbol{\theta})\right)\right] \geq 0. \quad (6)$$

[2]In fact, without loss in the value of the maximal expected social welfare, we can strengthen constraint (2) by removing the expectation operator and requiring that the cost must be covered for each $\boldsymbol{\theta}$, not just on the average. This gives an "ex-post cost-covering constraint." Crampton *et al.* [16] prove this fact under the assumption that exclusions are not allowed. Their proof generalizes to circumstances in which exclusions are allowed, provided we may require payments from excluded agents. However, that may be unreasonable or impossible. That is why we express the total payment as $\sum_i \pi_i(\boldsymbol{\theta}) p_i(\boldsymbol{\theta})$ rather than $\sum_i p_i(\boldsymbol{\theta})$. It costs us nothing to do this, since, as we show in Appendix VI, the maximal expected social welfare is just as great as it is when we are allowed to require payments from excluded agents. We can also show that it is possible to obtain just as much expected social welfare if one retains the expectation operator in constraint (2), but strengthens both (3) and (4) to ex-post constraints by removing the expection operators. This result will be published elsewhere.

This is the same as a model of Norman [2], but he takes $c(n,Q) = c(n)Q$.

In Appendix II, we show that this problem can be solved using Lagrangian methods. That is, for some $\lambda > 0$ it can be solved by maximizing a Lagrangian of

$$E\left[\sum_{i=1}^{n} \pi_i(\boldsymbol{\theta})\left(\theta_i + \lambda g(\theta_i)\right) u\left(Q(\boldsymbol{\theta})\right) - (1+\lambda)c\left(n, Q(\boldsymbol{\theta})\right)\right]. \tag{7}$$

The maximization is carried out pointwise. That is, given $\boldsymbol{\theta}$, the values of $\pi_1(\boldsymbol{\theta}), \ldots, \pi_n(\boldsymbol{\theta})$ and $Q(\boldsymbol{\theta})$ are chosen to maximize

$$A(\boldsymbol{\theta}, \lambda) u\left(Q(\boldsymbol{\theta})\right) - c\left(n, Q(\boldsymbol{\theta})\right) \tag{8}$$

where

$$A(\boldsymbol{\theta}, \lambda) = \frac{\sum_{i=1}^{n} \pi_i(\boldsymbol{\theta})\left(\theta_i + \lambda g(\theta_i)\right)}{1+\lambda}. \tag{9}$$

The fact that the coefficient $A(\boldsymbol{\theta}, \lambda)$ should be maximized means that we should take $\pi_i(\boldsymbol{\theta}) = 1$ if and only if $(\theta_i + \lambda g(\theta_i)) > 0$. Now, it is intuitively reasonable that agents with greatest preference parameters should be the ones to be included. This is ensured if we impose a restriction on the shape of the distribution function F by assuming that $g(\theta_i)$ is nondecreasing (which follows, for example, if the hazard rate, $f(x)/(1-F(x))$ is nondecreasing). Assuming this, agent i should be included if and only if θ_i exceeds some $\bar{\theta}(\lambda)$, where $\bar{\theta}(\lambda) + \lambda g(\bar{\theta}(\lambda)) = 0$. Note that $\bar{\theta}(\lambda)$ is increasing in λ, $A(\boldsymbol{\theta}(\lambda), \lambda)$ is decreasing in λ, and the $Q(\theta)$ which maximizes (8) is decreasing in λ.

III. Asymptotically Optimal Mechanism

A. A Scheme of Equal Contributions

In general, the full solution of our problem is very complex. However, in Appendix IV, we prove Theorem 1 below, which states that, when n is large, a nearly optimal solution can be achieved with a simple MD. First, we make an assumption.

Assumption 1: Suppose that

$$u(Q) = AQ^{\alpha} \tag{10}$$

$$c(n,Q) = Bn^{\delta}Q^{\beta} \tag{11}$$

where $A, B > 0$, $\delta > 0$, $0 < \alpha \leq 1$, $\beta \geq 1$, and $\alpha < \beta$.

We also consider the following weaker assumption.

Assumption 2: Suppose $u(Q) = \Theta(Q^{\alpha})$ and $c(n,Q) = h(n)c(Q)$, where $c(Q) = \Theta(Q^{\beta})$, $0 < \alpha \leq 1$, $\beta \geq 1$, and $\alpha < \beta$. That is, there are positive constants A_1, A_2, B_1, B_2, and a function h such that for all Q and n

$$A_1 Q^{\alpha} \leq u(Q) \leq A_2 Q^{\alpha} \tag{12}$$

$$B_1 h(n) Q^{\beta} \leq c(n,Q) \leq B_2 h(n) Q^{\beta}. \tag{13}$$

This assumption ensures that the growth rate of the optimal value is bounded both from above and below, as we see in the following lemma.

Lemma 1: Suppose Assumption 2 holds. Let $\gamma = \beta/(\beta-\alpha)$ and define

$$\xi(x) = \max_{Q}\left\{xu(Q) - c(n,Q)\right\}. \tag{14}$$

Then, $\xi(x) = \Theta(x^{\gamma})$ and the optimizing Q satisfies $Q(x) = \Theta(x^{1/(\beta-\alpha)})$. Moreover, $\xi'(x) = \Theta(x^{\gamma-1})$.

The proof is in Appendix III. Our main result is the following.

Theorem 1: Suppose Assumptions 1 or 2 holds. Let $\mathcal{P}$ be the problem of maximizing (1) subject to (6), with optimal value Φ_n. Let Q^* and θ^* be the optimizing decision variables in the problem $\mathcal{P}^*$, defined as

$$\underset{\theta\in[0,1], Q\geq 0}{\text{maximize}} \left\{ n\left(\int_{\theta}^{1} \eta f(\eta) d\eta\right) u(Q) - c(n,Q)\right\} \tag{15}$$

subject to

$$n\left(1-F(\theta)\right)\theta u(Q) - c(n,Q) \geq 0. \tag{16}$$

Let the optimal value be Φ_n^*.

Suppose we take as a feasible solution to $\mathcal{P}$ the decision variables $\pi_i(\boldsymbol{\theta}) = 1\{\theta_i \geq \theta^*\}$ and $Q(\boldsymbol{\theta}) = Q^*$. Then, the expected social welfare under this (suboptimal) mechanism is Φ_n^*, and this is asymptotically optimal, in the sense that $\Phi_n \Phi_n^* \leq 1 + O(n^{-1})$ (under Assumption 1), or $\Phi_n \Phi_n^* \leq 1 + O(1/\sqrt{n})$ (under Assumption 2).

Moreover, Φ_n, $\Phi_n^* = \Theta(n^{\gamma})$, where $\gamma = (\beta - \delta\alpha)/(\beta - \alpha)$.

The intuition for this result is as follows. For each $\theta \in [0,1]$, let $S(\theta)$ be the set of all agents who have preference parameters in the interval $[\theta, 1]$. Denote the size of this set by $|S(\theta)|$. Now, $E|S(\theta)| = n(1-F(\theta))$, for all θ. Suppose we had the stronger fact that $|S(\theta)| = n(1-F(\theta))$, for all θ. Since, by the remarks above, the optimal mechanism includes the set of agents with preference parameters greater than some θ, we find (using integration by parts) that $\mathcal{P}$ simplifies to $\mathcal{P}^*$. The planner includes all agents in $S(\theta)$, for some θ, and then charges each of these agents the same fixed fee ϕ. The mechanism is individually rational for all agents in $S(\theta)$ provided $\phi \leq \theta u(Q)$. So using the greatest charge consistent with this, namely, $\phi = \theta u(Q)$, the total payment is $n(1 - F(\theta) \times \theta u(Q)$ and by (16) this covers the cost of $c(n,Q)$.

Now, return to the original problem $\mathcal{P}$. The weak law of large numbers guarantees that $|S(\theta)|$ is close to $n(1 - F(\theta))$ with high probability when n is large. So we may expect it to be very nearly optimal to adopt the mechanism above, i.e., to take $Q(\boldsymbol{\theta}) = Q^*$ and set a fixed fee of $\theta^* u(Q^*)$, thus including those peers for which $\theta_i \geq \theta^*$.

B. A Numerical Example

Suppose that the preference parameters are uniformly distributed on [0, 1] and that $u(Q) = (2/3)Q^{1/2}$ and $c(n,Q) = Q$. Consider the so-called "first-best" value of maximized expected social welfare that could be achieved if we were to have full information about $\theta_1, \ldots, \theta_n$ and were not restricted by constraints of ex-ante cost-covering, individual rationality and incentive compatibility. Given that $\sum_{i=1}^{n} \theta_i$ take some value T, the expected social welfare is $Tu(Q) - c(Q)$ and this is maximized by $Q = T^2/9$, to a value of $T^2/9$. The expected social welfare is thus

$$\left(\frac{1}{9}\right)\left(\text{var}(T) + (ET)^2\right) = \left(\frac{1}{9}\right)\left(\frac{n}{12} + \left(\frac{n}{2}\right)^2\right)$$

$$= \frac{n^2}{36} + \frac{n}{108} \approx 0.027\,78 n^2.$$

Now, consider the solution of $\mathcal{P}^*$. For the uniform distribution, $g(\theta) = 2\theta - 1$, so our problem is

$$\underset{\theta^*,Q}{\text{maximize}} \left\{ n \left(\int_{\theta^*}^{1} \theta d\theta \right) \frac{2}{3}\sqrt{Q} - Q \right\}$$
$$\text{subject to } n \left(\int_{\theta^*}^{1} (2\theta - 1) d\theta \right) \frac{2}{3}\sqrt{Q} - Q \geq 0.$$

The solution of this has $\theta^* = 1/4$, $Q^* = n^2/64$ and an optimal fee of $\phi = n/48$. The expected social welfare achieved is $\Phi_n^* = 3n^2/128 \approx 0.02344 n^2$. Thus, we satisfy the constraints and obtain a value of expected social welfare which grows with n similar to the first-best, but which is asymptotically smaller by a factor $(3/128)/(1/36) = 27/32 \approx 0.84$. The expected social welfare is less, but we have satisfied the constraints (2) and (3).

Next, we compare the expected social welfare value that we have found for $\mathcal{P}^*$ (i.e., $\Phi_n^* = 3n^2/128$)) with the second-best expected social welfare value that can be obtained for $\mathcal{P}$. Considering $\mathcal{P}$, we have that its solution by Lagrangian methods is

$$\Phi_n = \inf_\lambda E \left[\max_Q \left\{ \left(\sum_{i=1}^{n} (\theta_i + \lambda g(\theta_i))^+ \right) \times \frac{2}{3}\sqrt{Q} - (1+\lambda)Q \right\} \right]$$
$$= \inf_\lambda E \left[\frac{1}{9(1+\lambda)} \left(\sum_{i=1}^{n} (\theta_i + \lambda g(\theta_i))^+ \right)^2 \right].$$

To compute this, we define $I_0(s) = s^2$ and

$$I_n(s) = E \left[\left(s + \sum_{i=1}^{n} (\theta_i + \lambda g(\theta_i))^+ \right)^2 \right].$$

Then, recalling $\theta_1 + \lambda g(\theta_1) > 0$ if and only if $\theta_1 > \lambda/(1 + 2\lambda)$

$$I_n(s) = \frac{\lambda}{1+2\lambda} I_{n-1}(s) + \int_{\theta_1 \geq \frac{\lambda}{(1+2\lambda)}} I_{n-1}(s + \theta_1 + \lambda g(\theta_1))\, d\theta_1.$$

It turns out that $I_n(s)$ is a quadratic in s, and we can solve recurrence relations for the coefficients, ultimately to give

$$I_n(s) = s^2 + \frac{n(1+\lambda)^2}{1+2\lambda} s + \frac{(1+\lambda)^3}{(1+2\lambda)^2} \left(\frac{1}{4}n^2 + \frac{1}{12}n + \left(\frac{1}{4}n^2 + \frac{5}{12}n \right) \lambda \right).$$

Thus

$$\Phi_n = \min_\lambda \left\{ \frac{1}{9(1+\lambda)} I_n(0) \right\}$$
$$= \min_\lambda \left\{ \frac{1}{9} \frac{(1+\lambda)^2}{(1+2\lambda)^2} \times \left(\frac{1}{4}n^2 + \frac{1}{12}n + \left(\frac{1}{4}n^2 + \frac{5}{12}n \right) \lambda \right) \right\}.$$

The minimizing value of Lagrange multiplier is

$$\lambda_n = \frac{-5 - 3n + \sqrt{-95 + 78n + 81n^2}}{4(5+3n)}$$

(with $\lambda_n \to \lambda = 1/2$, so in the limit excluding peers with $\theta_i \leq \lambda/(1+2\lambda) = 1/4$, which is consistent with what we found for $\mathcal{P}^*$). We find the equation shown at the bottom of the page. So

$$\frac{\Phi_n}{\Phi_n^*} = 1 + \frac{7}{9n} + O(n^{-2}).$$

Note that in this case, in which $c(n, Q)$ does not depend on n, the expected social welfare per capita increases with n and there is advantage in having more participants to share the cost. The value of the welfare per capita is growing as $3n/128$ for the approximate to the second-best, this is only 7/348 less than the welfare per capita under the optimal second-best. We find numerically as shown below. This illustrates Theorem 1's claim that $\Phi_n/\Phi_n^* = 1 + O(1/n)$.

Finally, let us make a comparison with what happens if exclusions are not allowed. Denote the second-best expected social welfare value in these circumstances by $\Phi_n^\dagger$. This is easier to compute than when finding Φ_n, because the computations can be made simply in terms of the facts that $E[\sum_i \theta_i] = n/2$ and $E[(\sum_i \theta_i)^2] = n^2/4 + n/12$ (Fig. 1). We find that in our example that, when exclusions are not allowed

$$\Phi_n^\dagger = \min_\lambda E \left(\frac{(\sum_i \theta_i + \lambda g(\theta_i))^2}{1+\lambda} \right)$$
$$= \min_\lambda E \left(\frac{((1+2\lambda)\sum_i \theta_i - n\lambda)^2}{1+\lambda} \right)$$
$$= \frac{n(\sqrt{1+3n} - 1)}{27} = O\left(n^{\frac{3}{2}}\right).$$

Note that the expected social welfare that can be obtained per capita grows as $\sqrt{n}$, but is vanishingly small relative to the first-best level of expected social welfare per capita, which grows

$$\Phi_n = \frac{n(-1 + 9n + \sqrt{-95 + 78n + 81n^2})(15 + 9n + \sqrt{-95 + 78n + 81n^2})^2}{1728(5 + 3n + \sqrt{-95 + 78n + 81n^2})^2}$$
$$= \frac{3}{128}n^2 + \frac{7}{348}n - \frac{1}{162} + \frac{47}{13122}n^{-1} + O(n^{-2})$$

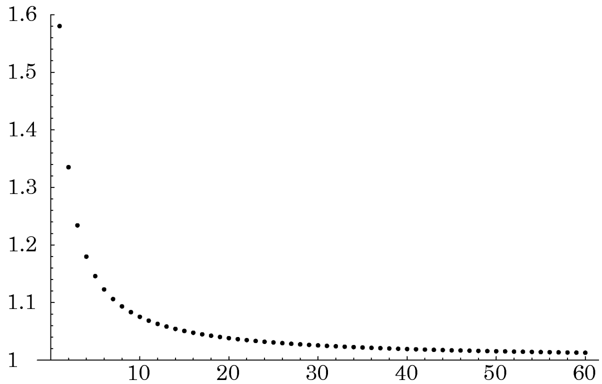

Fig. 1. Plot of Φ_n/Φ_n^* against n.

linearly in n. The fee structure and quality provision under the optimal mechanism is

$$p(\boldsymbol{\theta}) = P(\theta_i) = \frac{2}{9}\left(1 - \frac{1}{\sqrt{1+3n}}\right)\theta_i^2$$

and

$$Q(\boldsymbol{\theta}) = \frac{\left(n - \frac{2n}{\sqrt{1+3n}} - 2\left(1 - \frac{1}{\sqrt{1+3n}}\right)\sum_{i=1}^n \theta_i\right)^2}{9}.$$

Note that $EQ(\boldsymbol{\theta}) = (2n/27)(1 - 1/\sqrt{1+3n})$, which is also the value of $E[\sum_{i=1}^n P(\theta_i)]$.

There is no viable approximating second-best solution that can be obtained from $\mathcal{P}^*$ when one cannot make exclusions.

C. Other Equal-Contributions Schemes

A feature of our limiting mechanism is that all participating peers pay the same fee. One can devise other mechanisms for which this is true, of which two obvious ones are as follows. In Mechanism 1, the planner announces that he shall provide the good in quantity Q, and then share the cost $c(Q)$ amongst all those who volunteer to participate. If m out of n choose to participate, then each pays $c(Q)/m$. Those who wish to participate must make a commitment to do so, before knowing how many others will participate. In our example, it turns out that the maximized expected social welfare using this scheme is $3n^2/128 - n/384 + O(1)$. So the imposition of the ex-post cost-covering constraint occasions a loss compared to $\Phi_n^* = 3n^2/128$. Further details of these calculations are in Appendix VII.

In Mechanism 2, we charge a fee of ϕ, and then build the largest facility whose cost can be met by the number m who choose to participate, namely, Q such that $c(Q) = m\phi$. As before, peers must make a commitment to pay ϕ without knowing how many others will also participate. The optimal value in our example is $3n^2/128 + 7n/1536 + O(1)$. Compare this to $\Phi_n = 3n^2/128 + 7n/9 + O(1)$. This ex-post cost-covering scheme does a bit better than a scheme in which Q is fixed a priori, because although it provides the same Q on average it provides more Q when more peers participate.

An even better ex-post cost-covering equal-contributions scheme would be one in which, having learnt that m peers wish to participate, the planner builds a facility of size $Q(m)$ and charges each participant $c(Q(m))/m$. Potential participants know the function $Q(m)$. However, even such a mechanism cannot be more than a factor $1 + O(1/n)$ better than the simple one we propose.

D. An Extension to Multiple Constraints

Let us now consider a problem in which there are k "types" of peer and k constraints, which impose conditions that peers of the same type must cover a certain aspect of the cost. Let us suppose that there are $n_j = n\rho_j$ participants of type j, where $\sum_{j=1}^k \rho_j = 1$. Think now that Q is a vector $(Q_1, \dots, Q_k)$. Since $\mathcal{P}^*$ can be solved by Lagrangian methods (cf. Appendix II), we know there exists multipliers $\lambda_1, \dots, \lambda_k$ such that

$$\begin{aligned}\Phi_n^* = &\max_{\pi_1(\cdot),\dots,\pi_k(\cdot),Q} \\ &\times \left\{ \sum_{j=1}^k \left(n_j E\left[\pi_j(\theta_1)\left(\theta_1 + \lambda_j g_j(\theta_1)\right)\right] u_j(Q) \right.\right. \\ &\left.\left. -(1+\lambda_j)c_j(Q)\right) \right\} \\ = &\max_Q \left\{ \sum_{j=1}^k \left(E\left[\sum_{i=1}^{n_j} (\theta_i + \lambda_j g_j(\theta_i))^+\right] u_j(Q) \right.\right. \\ &\left.\left. -(1+\lambda_j)c_j(Q) \right) \right\}.\end{aligned}$$

As in the proof in Appendix IV for the case of a single constraint, we have as a bound that for any λ

$$\Phi_n \le E\left[\max_Q \left\{ \sum_{j=1}^k \left(\left[\sum_{i=1}^{n_j} (\theta_i + \lambda_j g_j(\theta_i))^+\right] u_j(Q) - (1+\lambda_j)c_j(Q) \right) \right\}\right]. \quad (17)$$

The expectation is taken with respect to the θ_i, which are i.i.d. for participants of the same type. (For simplicity we omit a second subscript on θ_i which might have been used to denote the type of peer.)

Let us suppose that type j utility and cost functions depend on different weighted sums of powers of $Q_1, \dots, Q_k$. That is, there are sets of weights $\{\eta_{j\ell}\}$ and $\{\nu_{j\ell}\}$ so that for type j

$$u_j(Q) = \sum_\ell \eta_{j\ell} Q_\ell^\alpha \text{ and } c_j(Q) = \sum_\ell \nu_{j\ell} Q_\ell^\beta.$$

This is obviously a restriction to our model, but it still includes many interesting possibilities, such as $c_j(Q) = (1/k)Q_1$, and $u_j(Q) = Q_1^{1/2}$, $(Q_2 = \cdots = Q_k = 0)$, in which peers of type i are required to cover exactly $(1/k)$th of the cost, perhaps by making payments in kind.[3] We now have from (17), and redefining $\xi(x) = \max_Q\{xQ^\alpha - Q^\beta\}$

$$\begin{aligned}\Phi_n \le & E\left[\max_Q \sum_{j=1}^k \left\{\left(\sum_{i=1}^{n_j}(\theta_i+\lambda_j g_j(\theta_i))^+\right)\sum_\ell \eta_{j\ell}Q_\ell^\alpha \right.\right.\\ &\left.\left. -(1+\lambda_j)\sum_\ell \nu_{j\ell}Q_\ell^\beta\right\}\right]\\ =& E\left[\sum_{\ell=1}^k\left[\sum_{j=1}^k(1+\lambda_j)\nu_{j\ell}\right]\max_{Q_\ell}\right.\\ &\left.\times\left\{\frac{\sum_{j=1}^k\left(\sum_{i=1}^{n_j}(\theta_i+\lambda_j g_j(\theta_i))^+\right)}{\left[\sum_{j=1}^k(1+\lambda_j)\nu_{j\ell}\right]}\eta_{j\ell}Q_\ell^\alpha - Q_\ell^\beta\right\}\right]\\ =&\sum_{\ell=1}^k\left[\sum_{j=1}^k(1+\lambda_j)\nu_{j\ell}\right]E\left[\xi(T^\ell)\right] \qquad (18)\end{aligned}$$

where

$$T^\ell = \frac{\sum_{j=1}^k\sum_{i=1}^{n_j}(\theta_i+\lambda_j g_j(\theta_i))^+}{\sum_{j=1}^k(1+\lambda_j)\nu_{j\ell}}\eta_{j\ell}.$$

The line above (18) follows from bringing the sum on ℓ to the outside, and then dividing and multiplying with by a factor $[\sum_{j=1}^k(1+\lambda_j)\nu_{j\ell}]$. We have

$$ET^\ell = n\frac{\sum_{j=1}^k \rho_j E\left[(\theta_i+\lambda_j g_j(\theta_i))^+\right]}{\sum_{j=1}^k(1+\lambda_j)\nu_{j\ell}}\eta_{j\ell}.$$

As in Appendix IV, we can bound (18) by making a Taylor expansion of $E\xi(T^\ell)$ around $\xi(ET^\ell)$, and then use

$$\Phi_n^* = \min_{\lambda_1,\ldots,\lambda_n}\sum_{\ell=1}^k\left[\sum_{j=1}^k(1+\lambda_j)\nu_{j\ell}\right]\xi(ET^\ell).$$

Notice that the optimizing $\lambda_1,\ldots,\lambda_n$ depend on $\rho_1,\ldots,\rho_k$, but not n. The remaining details are as in Appendix IV. We find, under these assumptions, that $\Phi_n^* = \Theta(n^\gamma)$ and $\Phi_n/\Phi_n^* = 1 + O(1/n)$.

IV. Application to File Sharing

We apply the above ideas to a problem of P2P file sharing. Agents are now called peers. Suppose that n peers make available various files to share with one another. What matters is the number of distinctly different files that are shared, so we must account for the possibility that more than one peer will make the same file available. Suppose that the utility obtained by peer i when the expected number of distinctly shared files is Q is $\theta_i u(Q)$, where u is concave in Q. We start by analyzing a simple model in which each peer provides the same number of files, say ϕ, choosing these randomly from amongst a set of N distinct file names. Then, the expected number of distinct files that will be available in the system is

$$Q = N\left(1-\left(1-\frac{\phi}{N}\right)^n\right) \qquad (19)$$

and so to obtain Q each peer must supply a number of files

$$\phi(n,Q) = N\left(1-\left(1-\frac{Q}{N}\right)^{\frac{1}{n}}\right). \qquad (20)$$

Suppose that each peer incurs a cost that is proportional to the number of files he contributes. For simplicity, we let the constant of proportionality be 1 (noting that we could always re-scale the utility function). Thus, the total cost is $c(n,Q) = n\phi(n,Q)$, where $n\phi(n,Q)$ is the total number of files shared by peers and this is a convex increasing function of Q, due to the duplications. Also, for any fixed Q, the cost $n\phi(n,Q)$ rapidly increases with n to the asymptote of $-N\log(1-Q/N)$. This is greater than Q, the total cost if there were no duplication in the files that the peers supply.[4]

In Fig. 2, we take $N = 1000$ and plot $n\phi(n,Q)$ against Q for $n = 1, 2, 10, \infty$. Note that for small to moderate values of Q the cost is almost linear in Q, but then increases rapidly as Q approaches N. For example, for $n = 100$, we find

$$\begin{aligned}c(n,Q) &= n\phi(n,Q)\\ &= Q\left[1+0.495\left(\frac{Q}{N}\right)+0.328\,35\left(\frac{Q}{N}\right)^2+\cdots\right].\end{aligned}$$

This justifies an approximation $c(n,Q) = Q$ when Q/N is reasonably small.

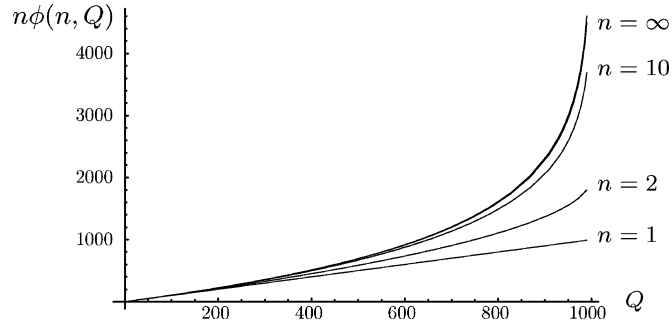

Fig. 2. $n\phi(n,Q)$ against Q, when $N = 1000$ and $n = 1, 2, 10, \infty$.

[3]In fact, if we were to take more generally, something like $u_j(Q) = Q^{\alpha_j}$ and $c_j(Q) = Q^{\beta_j}$, the expected social welfare will be $\Theta(n^\gamma)$, where $\gamma = \max_j\{\beta_j/(\beta_j-\alpha_j\}$ and so it is only one component of Q that really matters as $n \to \infty$.

[4]An alternative would be that a peer's cost is proportional to the rate at which he serves upload requests. If files are equally popular, then the total cost incurred by all the peers will be proportional to the product of the number participating peers and the number of unique files, i.e., $c(Q) = (\sum_i \pi_i)Q$. If peers can only access files held within a certain neighborhood of their location, this might be better modeled as $c(Q) = (\sum_i \pi_i)^\beta Q$, where $0 < \beta < 1$. There is a problem reproving Theorem 1 because the proof that Lagrangian methods work (proved here in the Appendix II) no longer holds. This is for future research. We would expect to be able to address a limiting problem in which $u(Q)$ is concave in Q and $c(Q) = [n(1-F(\theta^*)]^\beta Q$.

In an alternative and slightly more sophisticated model, we might imagine that the peers share different numbers of files. Suppose $n\rho_i$ of peers each share i files, each of them choosing his i files randomly from amongst a set of $N = na$ files, where $a > 0$ and $\sum_i \rho_i = 1$. Let i^* be an upper bound on the number of files that any one peer can share. The expected number of distinct files supplied will be

$$Q = na\left[1 - \prod_{i=1}^{i^*}\left(1 - \frac{i}{na}\right)^{n\rho_i}\right] = na\left[1 - e^{-\sum_i \frac{i\rho_i}{a}}\right] + O(1). \qquad (21)$$

Now, $n\sum_i i\rho_i$ is the total number of files provided by the peers, and we again assume that this is also the cost. As before, the asymptote as $n \to \infty$ is $c(n,Q) = -N\log(1 - Q/N)$. If Q/N is small, we again have $c(n,Q) = Q(1 + (1/2)(Q/N) + (1/3)(Q/N)^2 + \cdots) \approx Q$.

Both of the above lead to models that are covered by Section II. The social planner wishes to design a mechanism which maximizes expected social welfare, subject to its being cost-covering, individually rational and incentive compatible. Assuming $u(Q)$ satisfies Assumption 1 and $c(n,Q) = Q$, we can apply Theorem 1 and have an asymptotically optimal mechanism by solving the problem

$$\underset{Q,\theta}{\text{maximize}}\ nu(Q)\int_{\theta}^{1}(1 - F(\eta))\,d\eta - Q \qquad (22)$$

subject to

$$n\left[1 - F(\theta)\right]\theta u(Q) - Q \geq 0. \qquad (23)$$

Let Q^* and θ^* be the maximizing values of the decision variables. Each peer who has a preference parameter of at least θ^* is included and pays the same fixed fee of $\theta^* u(Q^*)$. Since the cost is linear in Q this fee can be paid "in kind," i.e., without monetary payments: each included peer pays his fee by contributing the same number of files: namely, $Q^*/[n(1 - F(\theta^*))]$. (Note that although our theorems assume no bound on Q, in this problem Q is bounded by N. However, this is immaterial as we expect the optimal system operates at a Q that is well away from this upper bound.)

A remark on repeated rounds: In the limiting problem, there is no reason that a peer should be other than truthful in representing himself to the system. If he knows that the expected number of unique files shared is Q and that the fee is ϕ, then peer i should join if $\theta_i u(Q) \geq \phi$. In the nonlimiting version of the problem, addressed by the optimal MD of solving $\mathcal{P}$ in Section II, the individually rationality constraint (3) is in terms of expected value, so for some $\boldsymbol{\theta}_{-i}$ it can be that $\theta_i u(Q(\theta_i, \boldsymbol{\theta}_{-i}) - p_i(\theta_i, \boldsymbol{\theta}_{-i}) < 0$. When this happens, peer i might be tempted to defect and to not pay $p_i(\theta_i, \boldsymbol{\theta}_{-i})$. However, as file sharing system is intended to last for more than one time step, we could operate a "tit-for-tat"-like protocol, that would penalize such defection, for example, by threatening to exclude peer i at a later time when $\boldsymbol{\theta}_{-i}$ is such that his net benefit would be positive. We are imagining that $\boldsymbol{\theta}$ is not fixed, but varies over time, when from time to time the peers' preference parameters are freshly sampled from F. The effect of the threat to penalize defection will be to make peer i willing to participate on occasions that he must accept a negative net benefit, knowing that on average he will benefit, as is guaranteed by (1) and (3). If every peer's preferences parameter varies over time with the distribution F, each will obtain on average $1/n$th of the maximized expected social welfare.

V. Application to WLANs

Now, we apply our ideas to WLANs. Access to the Internet is still not as ubiquitous as access to the telephone network. This greatly reduces the economic value of many new portable devices, such as PDAs, tablet computers, and smart-phones running the IP protocol. The users of these devices would benefit greatly from cost-effective Internet access that is wireless, always-on, ubiquitous and high-speed. However it a nontrivial task to deploy infrastructure with wide enough coverage, especially from the business perspective.

WLANs are an important developing infrastructure. Specifically, the IEEE 802.11 WLAN standard has grown steadily in popularity since its inception and, at least in metropolitan areas, is now well positioned to complement much more complex and costly technologies such as 3G. This is already happening. WLAN signals of networks set up by individuals for their own use already pervade many cities and such WLAN "cells" frequently cover greater areas than were originally intended at their installation. Given how easy it is to gain access to a WLAN once a potential user is within its coverage area, and leaving out the obvious security issues involved, one wonders if individuals could share such infrastructure amongst themselves to achieve ubiquitous Internet access. Sharing comes as a natural idea since WLANs provide large amounts of bandwidth that is mostly underutilized by its local users. Also the pipe that connects the local WLAN users to the Internet is usually of a broadband nature (DSL) and may also be under-used over large time periods. Existing technology allows WLAN administrators to control access to their networks and to limit the consumption of network resources by remote (roaming) users. The WLAN peering model we present next is motivated by these observations.

Suppose that n distinct WLANs are available in a given large geographical location, such as a neighborhood or a part of a city centre. The owners of the WLANs may arrange to peer with one another, and thus agent i, who is the owner of the ith WLAN can benefit when he roams in areas covered by other WLANs. When agreeing to become a peer, a WLAN owner benefits, but he also incurs some cost in providing resources to the community. We seek a mechanism, defined in terms of certain rules, to specify what quantities of resources peers must contribute and what subsidies or payments they might have to make. Our aim is that the incentives given by these rules should be such that when peers act to maximize their own benefits, expected social welfare is also maximized. To begin, we assume that there is some central authority, a "global planner," who serves as an intermediary for implementing these rules. Then, we will show that as the system gets large, the optimal rules can be approximated by simple contribution policies, alleviating the need for a central mechanism.

Let Q be the "coverage" available in the location, defined as the probability that an agent can obtain roaming service when away from his own WLAN. Assume that if agent i peers, then his WLAN accepts service requests from other roaming peers (and himself) with probability p_i. If he does not peer (is excluded), then $p_i = 0$. To express Q as a function of $p_1, \dots, p_n$, let us suppose that the total area of the location is B, the area of coverage of a typical WLAN is A, the WLANs of different peers do not overlap and that roaming peers are positioned uniformly on B. Then, we have $Q = \sum_{i=1}^{n} p_i A/B$.

If agent i peers, then his cost that is proportional to the rate of service requests that he accepts. This can be written as $\sigma m \lambda p_i$, where σ is dollars per rate of service requests accepted (which we may take as $\sigma = 1$), λ is the rate of requests generated by a typical agent, and m is the number of agents who peer, i.e., the number of i for which $p_i \neq 0$. This is a reasonable model for cost since roaming customers consume bandwidth from the WLAN.

We view coverage as a nonrivalrous public good. That is, each roaming peer benefits by the amount of coverage available, and does not reduce the probability with which other roaming peers can obtain access. That is, he benefits from, but does not consume, Q. The important issue is to provide incentives for Q to grow, while balancing the resulting costs. It can grow by having more agents participate in the peering arrangement and by increasing the p_is offered by the agents. We can cast the MD problem faced by the global planner in the formulation we used earlier. We again have that the utility of agent i is $\theta_i u(Q)$ so that the total utility is

$$\sum_{i=1}^{n} \pi_i \theta_i u(Q)$$

and total cost is

$$c(Q) = \sigma \lambda m \sum_{i=1}^{n} p_i = \left(\frac{\sigma \lambda B}{A}\right) mQ$$

where $m = \sum_{i=1}^{n} \pi_i$. The difference of this cost function with the cost functions used earlier is that now there is a multiplicative congestion factor which is proportional to the number of peers who actually participate, m, instead of the initial number of potential peers n. Thus, we need to extend our public good model to make the cost function depend on m instead of n.

Let the total cost take the form $mc(Q)$, e.g., above we would redefine $c(Q) = (\sigma \lambda B/A)Q$. Our problem $\mathcal{P}$ now becomes

$$\underset{\pi_1(\cdot),\dots,\pi_n(\cdot),Q(\cdot)}{\text{maximize}}\; E\left[\sum_i \pi_i(\boldsymbol{\theta})(\theta_i u(Q(\boldsymbol{\theta})) - c(Q(\boldsymbol{\theta})))\right] \quad (24)$$

subject to

$$E\left[\sum_i \pi_i(\boldsymbol{\theta})\left(g(\theta_i) u(Q(\boldsymbol{\theta})) - c(Q(\boldsymbol{\theta}))\right)\right] \geq 0. \quad (25)$$

Let us take $u(Q) = AQ^\alpha$, $c(Q) = BQ^\beta$ and $\gamma = \beta/(\beta - \alpha)$. Choose units so that

$$\begin{aligned}\xi(x) &= \max_Q \{xu(Q) - c(Q)\}\\ &= (A\alpha x)^{\frac{\beta}{\beta-\alpha}} (B\beta)^{-\frac{\alpha}{\beta-\alpha}} \left(\frac{1}{\alpha} - \frac{1}{\beta}\right)\\ &= x^\gamma.\end{aligned}$$

Let us solve (24), while disregarding (25). This gives

$$\Phi_n \leq E\left[\max_{m\in\{1,\dots,n\}} \left\{\frac{\left(\sum_{i=1}^{m} \theta_{(i)}\right)^\gamma}{m^{\gamma-1}}\right\}\right]$$

or, in general

$$\Phi_n \leq E\left[\max_{m\in\{1,\dots,n\}} \left\{m\xi\left(\frac{\sum_{i=1}^{m} \theta_{(i)}}{m}\right)\right\}\right]$$

where $\theta_{(1)} \geq \cdots \geq \theta_{(n)}$ are the ordered values of $\theta_1, \dots, \theta_n$.

In the limiting problem, we will admit all peers with preference parameters of at least $\bar{\theta}$. The expected number of these is $m = n(1 - F(\bar{\theta}))$ and so the problem $\mathcal{P}^*$ is

$$\underset{\bar{\theta},Q}{\text{maximize}} \left[n \int_{\bar{\theta}} \theta dF(\theta) u(Q) - n\left(1 - F(\bar{\theta})\right) c(Q)\right] \quad (26)$$

subject to

$$n\left(1 - F(\bar{\theta})\right) \bar{\theta} u(Q) - n\left(1 - F(\bar{\theta})\right) c(Q) \geq 0. \quad (27)$$

The constraint (27) says that the total payment $m\bar{\theta}u(Q)$ must cover the cost $mc(Q)$. The condition that the objective function be stationary with respect to $\bar{\theta}$ is

$$-n\bar{\theta}u(Q)f(\bar{\theta}) + nf(\bar{\theta})c(Q) = 0$$

and this implies that (27) holds with equality. Thus, in this particular problem the constraint is redundant, and we may concentrate on solving the unconstrained problem. We have

$$\Phi_n^* = n \max_{\bar{\theta}} \left\{\frac{\left(\int_{\bar{\theta}}^{1} \theta dF(\theta)\right)^\gamma}{\left(1 - F(\bar{\theta})\right)^{\gamma-1}}\right\} \quad (28)$$

or, in general

$$\Phi_n^* = n \max_{\bar{\theta}} \left\{(1 - F(\bar{\theta})\xi\left(\frac{\int_{\bar{\theta}}^{1} \theta dF(\theta)}{1 - F(\bar{\theta})}\right)\right\} \quad (29)$$

and the solution is at a $\bar{\theta}$ such that

$$\frac{\bar{\theta}}{\gamma - 1} = \frac{\int_{\bar{\theta}}^{1} (1 - F(\theta))\, dF(\theta)}{1 - F(\bar{\theta})} = E[\theta - \bar{\theta} | \theta \geq \bar{\theta}]. \quad (30)$$

The left-hand side of (30) is increasing in $\bar{\theta}$, so (30) has a unique solution if the right-hand side is decreasing in $\bar{\theta}$, i.e., if θ has a distribution that is "new better than used in expectation"

(NBUE). For example, when F is the uniform distribution $\bar{\theta} = (\gamma - 1)/(\gamma + 1)$. We now have something similar to Theorem 1.

Assumption 3: Suppose that, given that m peers are allowed to use the system

$$u(Q) = AQ^{\alpha} \tag{31}$$

$$c(m, Q) = Bm^{\delta}Q^{\beta} \tag{32}$$

where $A, B > 0$, $\delta > 0$, $0 < \alpha \leq 1$, $\beta \geq 1$, and $\alpha < \beta$.

Theorem 2: Suppose Assumption 3 holds and that the preference parameters are distributed according to a distribution with a density function bounded away from 0. Then

$$\frac{\Phi_n}{\Phi_n^*} = 1 + O\left(\frac{1}{\sqrt{n}}\right).$$

The proof is in Appendix V.

VI. Stability

Suppose that the social planner designs a mechanism on the basis that there are n peers. He expects that $(1-F(\theta))n$ of them will pay a fee of $f = \theta u(Q)$. Let us focus on the problem of file sharing covered in Section IV. Since the fee is paid "in kind" and equates to providing f files, the total number of files that are provided will be $Q = (1 - F(\theta))nf$.

Suppose that there are indeed n peers, but initially some of them are dubious that Q will be as large as the planner claims. Consequently, some do not participate and the number of files that is initially provided is $Q_1 < Q$. Once the peers have observed Q_1, those peers with $\theta_i > f/u(Q_1)$ will realize that it is to their advantage to participate. The fees paid by these will provide Q_2 files, where

$$Q_2 = \left(1 - F\left(\frac{f}{u(Q_1)}\right)\right) nf. \tag{33}$$

Write this as $Q_2 = \phi(Q_1)$ and imagine iterating $Q_{k+1} = \phi(Q_k)$, $k = 1, 2, \ldots$. In general, there can be more than one root to $Q = \phi(Q)$. For example, suppose $u(Q) = 0.6Q^{1/2}$, $f = 5$, $n = 120$, and θ_i is uniformly distributed on [0, 1]. Then

$$\phi(Q) = \left(1 - \frac{5}{0.6Q^{\frac{1}{2}}}\right)(120)(5). \tag{34}$$

In this example, there are two roots, $Q = 100.00$ and $Q = 320.87$. One can easily prove that if Q_1 exceeds the smaller root, then Q_k tends to the larger root as k tends to infinity. Otherwise $Q_k \to 0$. For $Q = 100$ the expected social welfare is 10, whereas for $Q = 320.87$ it is 184.4. Thus, the greater Q, to which the system converges, is also the root for which a greater number of peers participate and the greater expected social welfare is achieved. Similar properties hold for the more general case of $u(Q) = AQ^a$.

An interesting issue is how stability is affected by agents departing and new agents arriving. Another issue is the optimal choice of the fixed fee as a function of the system size for quickly reaching the equilibrium. Finally, might like to analyze the effect of errors in the estimation of the actual content Q.

VII. Conclusion

In this paper, we have formalized an interesting connection between P2P systems, namely, file sharing and peering WLANs, and public good theory. We have shown that simple incentive policies of the form of fixed contributions can suffice to control the overall system to a nearly optimal size. Even though our economic model is rather crude and abstracts many practical aspects of the implementations, we have captured some "first-order" properties of the large externalities that such system exhibit. These externalites are one of the main reason that P2P systems are being widely adopted.

There are many ways to extend this simple model. For instance, we might personalize agents by more than a single parameter, e.g., adding a parameter that captures cost sensitivity. However, the analysis of such extensions is probably too complicated. One could perhaps model the utility or cost function more carefully, by introducing a dependence on more detailed congestion effects, and so try to capture performance aspects encountered in specific network technologies. But it is not obvious that making the problem more complex will provide more insight. The results one obtains by such models are more qualitative than quantitative, showing the form of the optimal control rather than computing its exact value.

It would be interesting to extend the file sharing model to one in which files have different popularities and costs of uploading. This might be done using results in Section III-D. We would expect to find that the peers should be required to make equal contributions, but that these contributions should be measured in terms of total upload rates, rather than just as a number of files. Note, however, that this could affect the peers' choices as to which files they make available for uploading, and hence the equilibrium of the type of content that will be available in the system.

It would be nice to extend the WLAN model so that instead of assuming that roaming peers are uniformly distributed and that they make requests for connectivity at a constant rate λ, the rate of requests differs by geographic region. To do this, we might again think along the lines of Section III-D, and say that peers of the same type are those that are based in the same region. We would expect an optimal mechanism to be one in which peers based in the same region make equal contributions, but where this contribution differs from region to region. However, there is difficulty in continuing along the lines of Section III-D since Lemma 3 does generalize to circumstances in which the cost depends on the number of peers who actually participate, m, rather than on n, the maximum number of peers who might participate.

There are practical issues regarding the implementation of our results. Even simple exclusion schemes may be hard to enforce in a system with cheap pseudonyms. Making sure that peers make available valid files (to avoid uploading from other peers) may not always be an easy task in such a loosely designed system. We are currently investigating several of the above implementation aspects. An approach where peer contributions are kept at a minimum possible implementable level (peers are required to share a fixed number of files only during the time they download files) is in [17].

APPENDIX I
DERIVATION OF THE PROBLEM

In this appendix, we show that the constraints of individual rationality and incentive compatibility reduce to (6). We give a streamlined explanation of some fairly standard arguments.

Suppose agent a_i pays $p_i(\boldsymbol{\theta})$.[5] Let us define

$$V_i(\theta_i) = \int \pi_i(\theta_i, \theta_{-i}) u\left(Q(\theta_i, \boldsymbol{\theta}_{-i})\right) dF^{n-1}(\boldsymbol{\theta}_{-i}) \quad (35)$$

$$P_i(\theta_i) = \int \pi_i(\theta_i, \theta_{-i}) p_i(\theta_i, \boldsymbol{\theta}_{-i}) dF^{n-1}(\boldsymbol{\theta}_{-i}). \quad (36)$$

Thus, $\theta_i V(\theta_i)$ and $P(\theta_i)$ are the expected utility and expected payment of peer i when his preference parameter is θ_i. We have the following.

Lemma 2: (a) It is necessary and sufficient for incentive compatibility that (i) $V_i(\theta_i)$ is nondecreasing in θ_i, and (ii)

$$P_i(\theta_i) = P_i(0) + \theta_i V_i(\theta_i) - \int_0^{\theta_i} V_i(\eta) d\eta. \quad (37)$$

(b) Given incentive compatibility, a necessary and sufficient condition for individual rationality is $P_i(0) \leq 0$.

Proof: The incentive compatibility condition says that θ_i must maximize $\theta_i V_i(\theta_i') - P_i(\theta_i')$ with respect to θ_i'. This implies that for $\theta_i' \neq \theta_i$, we must have

$$[\theta_i' V_i(\theta_i) - P_i(\theta_i)] + [\theta_i V_i(\theta_i') - P_i(\theta_i')] \leq [\theta_i V_i(\theta_i) - P_i(\theta_i)] + [\theta_i' V_i(\theta_i') - P_i(\theta_i')]$$

and so

$$(\theta_i' - \theta_i)\left(V_i(\theta_i') - V_i(\theta_i)\right) \geq 0.$$

This implies that (i), that $V_i(\theta_i)$ must be nondecreasing in θ_i. We also have

$$\theta_i V_i'(\theta_i) - P_i'(\theta_i) = 0.$$

So, by integrating, we find a second condition, (ii)

$$P_i(\theta_i) = P_i(0) + \theta_i V_i(\theta_i) - \int_0^{\theta_i} V_i(\eta) d\eta. \quad (38)$$

Thus, (i) and (ii) are necessary for incentive compatibility. It is also easy to check that they are sufficient.

Since the scheme is to be incentive compatible, we can deduce from (38) that the expected sum of the payments is given by

$$\sum_{i=1}^n \int \pi_i(\theta_i, \boldsymbol{\theta}_{-i}) p_i(\theta) dF^n(\boldsymbol{\theta}) \quad (39)$$

$$= \sum_{i=1}^n \int P_i(\theta_i) dF(\theta_i)$$
$$= \sum_{i=1}^n P_i(0) + \sum_{i=1}^n \int \left[\theta_i V_i(\theta_i) - \int_0^{\theta_i} V_i(\eta) d\eta\right] dF(\theta_i)$$
$$= \sum_{i=1}^n P_i(0) + \sum_{i=1}^n \left\{ \int \theta_i V_i(\theta_i) dF(\theta_i) + (1 - F(\theta_i)) \times \int_0^{\theta_i} V_i(\eta) d\eta \Big|_0^\infty - \frac{1 - F(\theta_i)}{f(\theta_i)} \times V_i(\theta_i) dF(\theta_i) \right\}$$
$$= \sum_{i=1}^n P_i(0) + \sum_{i=1}^n \int \pi_i(\theta_i, \theta_{-i}) g(\theta_i) u\left(Q(\boldsymbol{\theta})\right) dF^n(\boldsymbol{\theta}). \quad (40)$$

Since the scheme is to be ex-ante cost-covering, we use (40) to deduce that our problem is one of maximizing (1) subject to

$$-\sum_{i=1}^n P_i(0) \leq \sum_{i=1}^n \int \pi_i(\boldsymbol{\theta}) g(\theta_i) u\left(Q(\boldsymbol{\theta})\right) dF^n(\boldsymbol{\theta}) - \int c\left(n, Q(\boldsymbol{\theta}) dF^n(\boldsymbol{\theta}\right). \quad (41)$$

The maximization is with respect to a choice of the function $Q(\cdot)$ and the constants $P_1(0), \ldots, P_n(0)$. The individual rationality of (3) holds if and only if $P_i(0) \leq 0$. So we must take $P_i(0) \leq 0$. These enter only through their sum, which may therefore be taken to be zero. A way to understand the role of the $P_i(0)$ is the following. In the definition of the incentive payments (38), the last two terms represent the maximum incentive compatible payment that can be extracted from agent i when his preference parameter is θ_i. Then, the first term on the right-hand side of (41) is the maximum total incentive compatible payment that can be collected from all the agents. If at the constrained expected social welfare optimum, the right-hand side of (41) is strictly positive, then we do not need to ask for the maximum possible payment and can achieve the optimum with less. In this case, the negative amount $\sum_i P_i(0)$ is the money we can give back (after collecting the maximum amount) to the agents. It is up to the system planner how to redistribute this money (or not collect it in the first place). ■

We are to maximize (1) subject to (6) by pointwise choice of $Q(\cdot)$. From this we can calculate $V_i(\theta_i)$, and then the payments from (38) and (36). Provided $V_i(\theta_i)$ turns out to be nondecreasing, we have then solved the problem of maximizing expected social welfare subject to use of a cost-covering, individually rational and incentive compatible scheme.

APPENDIX II
JUSTIFICATION FOR USE OF LAGRANGIAN METHODS

We prove that the problem $\mathcal{P}$ (of finding the second best optimum) can be solved by Lagrangian methods. The special case $k = 1$ gives the result that we need in Sections II–V. The case $k > 1$ is used for Section III-D. For simplicity of notation, we drop the n from the cost $c(n, Q)$, and simply write $c(Q)$. We

[5]Since the agents are identical, apart from their labels, it is reasonable to suppose that the social planner can maximizes welfare with a MD that does not treat agents differently because of their labels. This would mean that V_i and P_i do not depend on i. However, we will not make this simplification, and so that we have a problem that is correct even if agents are not statistically identical.

suppose that agents are of k types. There are n_j agents of type j, and their identical and independently disctributed preference parameters are $\theta_{j1},\dots,\theta_{jn_j}$.

Lemma 3: Define $\mathcal{P}$ as the problem

$$\text{maximize } E\left[\sum_{j=1}^{k}\left[\sum_{i=1}^{n_j}\pi_{ij}(\boldsymbol{\theta})\theta_{ij}u_j\left(Q(\boldsymbol{\theta})\right)-c_j\left(Q(\boldsymbol{\theta})\right)\right]\right]$$

with respect to $Q(\boldsymbol{\theta})$, $\pi_{ij}(\boldsymbol{\theta})$, with $0\le\pi_{ij}(\boldsymbol{\theta})\le 1$ and subject to

$$E\left[\sum_{i=1}^{n_j}\pi_{ij}(\boldsymbol{\theta})\theta_{ij}u_j\left(Q(\boldsymbol{\theta})\right)-c_j\left(Q(\boldsymbol{\theta})\right)\right]\ge 0 \text{ for all } j.$$

Then, there exists a Lagrange multiplier $\lambda_1,\dots.\lambda_k$ such that an optimal solution to $\mathcal{P}$ can be found by maximizing the Lagrangian

$$E\left[\sum_{j=1}^{k}\sum_{i=1}^{n_j}\pi_{ij}(\boldsymbol{\theta})\left(\theta_{ij}+\lambda_j g_j(\theta_{ij})\right)u_j\left(Q(\boldsymbol{\theta})\right)\right.$$
$$\left.-(1+\lambda)c_j\left(Q(\boldsymbol{\theta})\right)\right]. \quad (42)$$

with respect to $Q(\boldsymbol{\theta})$, $\pi_{ij}(\boldsymbol{\theta})$, with $0\le\pi_{ij}(\boldsymbol{\theta})\le 1$.

Proof: Let us rewrite this as the problem of maximizing

$$E\left[\sum_{j=1}^{k}\left[\sum_{i=1}^{n_j}x_{ij}(\boldsymbol{\theta})-c_j\left(Q(\boldsymbol{\theta})\right)\right]\right] \quad (43)$$

with respect to $x_{ij}(\boldsymbol{\theta})$, $Q(\boldsymbol{\theta})$, subject to

$$Q(\boldsymbol{\theta})\ge 0,\quad x_{ij}(\boldsymbol{\theta})\ge 0 \quad (44)$$

$$x_{ij}(\boldsymbol{\theta})-\theta_{ij}u_j\left(Q(\boldsymbol{\theta})\right)\le 0,\quad \text{for all } i,j,\boldsymbol{\theta} \quad (45)$$

and

$$-E\left[\sum_{j=1}^{k}\left[\sum_{i=1}^{n_j}x_i(\boldsymbol{\theta})\frac{g(\theta_{ij})}{\theta_{ij}}-c_j\left(Q(\boldsymbol{\theta})\right)\right]\right]\le 0. \quad (46)$$

Assuming that $u_j(Q)$ is concave and $c_j(Q)$ is convex in Q, the objective function (43) is a concave function of the decision variables, and (44)–(46) define a region that is convex in the decision variables, $x_{ij}(\boldsymbol{\theta})$, $Q(\boldsymbol{\theta})$. These are sufficient conditions for the problem to be solvable by maximizing a Lagrangian. That is, there exist $\lambda_1,\dots,\lambda_k$ such that we can solve the problem by maximizing

$$E\left[\sum_{j=1}^{k}\left[\sum_{i=1}^{n_j}x_i(\boldsymbol{\theta})\left(1+\lambda_j\frac{g(\theta_{ij})}{\theta_{ij}}\right)-(1+\lambda_j)c_j\left(Q(\boldsymbol{\theta})\right)\right]\right]$$

with respect to $Q(\boldsymbol{\theta})$, and $x_{ij}(\boldsymbol{\theta})$, subject to (45). This is equivalent to maximizing (42) with respect to $Q(\boldsymbol{\theta})$, $\pi_{ij}(\boldsymbol{\theta})$, subject to $0\le\pi_{ij}(\boldsymbol{\theta})\le 1$. ■

Appendix III
Proof of Lemma 1

Proof: First, we can suppose $h=1$ (or we can absorb it into the constants B_1 and B_2). Note that

$$\max_Q\{xA_1Q^\alpha-B_2Q^\beta\}=C_1x^{\frac{\beta}{(\beta-\alpha)}}\le\xi(x)$$
$$\le C_2x^{\frac{\beta}{(\beta-\alpha)}}=\max_Q\{xA_2Q^\alpha-B_1Q^\beta\}$$

for constants

$$C_1=(A_1)^{\frac{\beta}{\beta-\alpha}}B_2^{-\frac{\alpha}{\beta-\alpha}}\zeta,\quad C_2=(A_2)^{\frac{\beta}{\beta-\alpha}}B_1^{-\frac{\alpha}{\beta-\alpha}}\zeta,$$
$$\zeta=\left(\frac{\alpha}{\beta}\right)^{\frac{\alpha}{(\beta-\alpha)}}-\left(\frac{\alpha}{\beta}\right)^{\frac{\beta}{(\beta-\alpha)}}$$

where $\zeta>0$. Hence, $\xi(x)=\Theta(x^\gamma)$.

Now, choose η_1 and η_2 such that for all $\eta\notin[\eta_1,\eta_2]$, we have $(A_2\eta^\alpha-B_1\eta^\beta)<C_1$. This is clearly possible, since $(A_2\eta^\alpha-B_1\eta^\beta)$ is a concave function of η which is equal to 0 at $Q=0$ and approaches $-\infty$ as $Q\to\infty$. Then, if $Q\le\eta_1x^{1/(\beta-\alpha)}$ or $Q\ge\eta_2x^{1/(\beta-\alpha)}$, we have

$$xu(Q)-c(n,Q)\le xA_2Q^\alpha-B_1Q^\beta$$
$$\le(A_2\eta^\alpha-B_1\eta^\beta)x^{\frac{\beta}{(\beta-\alpha)}}<C_1x^{\frac{\beta}{(\beta-\alpha)}}$$
$$\le\xi(x)$$

and so Q cannot be optimal. Hence, the optimizing Q, say $Q(x)$, is $\Theta(x^{1/(\beta-\alpha)})$.

Note that by differentiation through (14), we have $\xi'(x)=u(Q(x))$, and then $u(Q(x))=\Theta(Q(x)^\alpha)=\Theta(x^{\alpha/(\beta-\alpha)})=\Theta(x^{\gamma-1})$. ■

Appendix IV
Proof of Theorem 1

Proof: Suppose $c(n,Q)=h(n)c(Q)$, and for the moment take $h(n)=1$. Let us first make the strong Assumption 1: that $u(Q)=AQ^\alpha$ and $c(Q)=BQ^\beta$, where $\alpha\le 1\le\beta$, $\alpha<\beta$ (so u and c are concave and convex, respectively). Define the function ξ by

$$\xi(x)=\max_Q\left\{xu(Q)-c(Q)\right\}.$$

Note that this is just the definition of a real-valued function of a variable x (and has nothing to do with the π_i). It is a convex function of x. Let $\gamma=\beta/(\beta-\alpha)$ and choose units so that

$$\xi(x)=(A\alpha x)^{\frac{\beta}{\beta-\alpha}}(B\beta)^{-\frac{\alpha}{\beta-\alpha}}\left(\frac{1}{\alpha}-\frac{1}{\beta}\right)=x^\gamma.$$

Since $\mathcal{P}^*$ can be solved by Lagrangian methods, we know there exists a λ such that

$$\Phi_n^*=\max_{\pi_1(\cdot),Q}\left\{E\left[n\pi_1(\theta_1)\left(\theta_1+\lambda g(\theta_1)\right)\right]u(Q)\right.$$
$$\left.-(1+\lambda)c(Q)\right\}$$
$$=\max_{\pi_1(\cdot),\dots,\pi_n(\cdot),Q}\left\{E\left[\sum_{i=1}^{n}\pi_i(\theta_i)\left(\theta_i+\lambda g(\theta_i)\right)\right]u(Q)\right.$$
$$\left.-(1+\lambda)c(Q)\right\}$$
$$=\max_Q\left\{E\left[\sum_{i=1}^{n}\left(\theta_i+\lambda g(\theta_i)\right)^+\right]u(Q)-(1+\lambda)c(Q)\right\}$$

Let $T=\sum_{i=1}^n T_i$, where

$$T_i=\pi_i(\theta_i)\frac{(\theta_i+\lambda g(\theta_i))}{(1+\lambda)}=\frac{(\theta_i+\lambda g(\theta_i))^+}{(1+\lambda)}.$$

Note that the T_i are independent identically distributed (i.i.d.) random variables and that, since $\theta>g(\theta)$, it follows that $T_i\in$

$[0, \theta_i] \subset [0, 1]$. Let $\bar{T}_1 = ET_1$. Note that $\bar{T}_1$ depends on λ, but since

$$\Phi_n^* = \min_\lambda \max_Q \left\{ n \int_0^1 (\theta + \lambda g(\theta))^+ \, dF(\theta) A Q^\alpha - (1+\lambda) B Q^\beta \right\}$$
$$= \min_\lambda (1+\lambda) \left[n \frac{\int_0^1 (\theta + \lambda g(\theta))^+ \, dF(\theta)}{1+\lambda} \right]^\gamma$$

the optimizing λ does not depend on n. Hence, $\bar{T}_1$ does not depend on n.

Recall that for any λ, and so for this λ

$$\Phi_n \le \max_{\pi_1(\cdot),\ldots,\pi_n(\cdot),Q(\cdot)} E\left[\sum_{i=1}^n \pi_i(\boldsymbol{\theta}) (\theta_i + \lambda g(\theta_i)) \, u(Q(\boldsymbol{\theta})) - (1+\lambda) c(Q(\boldsymbol{\theta})) \right]$$
$$= \max_{Q(\cdot)} E\left[\sum_{i=1}^n (\theta_i + \lambda g(\theta_i))^+ \, u(Q(\boldsymbol{\theta})) - (1+\lambda) c(Q(\boldsymbol{\theta})) \right].$$

We now have, since $ET = n\bar{T}_1$

$$\Phi_n^* = (1+\lambda)\xi(ET)$$
$$\Phi_n \le (1+\lambda) E\xi(T).$$

Consider first $\gamma \ge 2$. Then, by a Taylor expansion of $\xi(T)$ around $n\bar{T}_1$, we have

$$\Phi_n \le (1+\lambda)\xi(n\bar{T}_1) + (1+\lambda)E(T - n\bar{T}_1)\xi'(n\bar{T}_1) + (1+\lambda)\frac{1}{2}E\left[(T - n\bar{T}_1)^2 \xi''(T^*)\right]$$

for some T^* depending on T. The middle term on the right-hand side is 0. Since T^* lies between $n\bar{T}_1$ and T, and so certainly $T^* \le n$, we have $\xi''(T^*) \le \gamma(\gamma-1)n^{\gamma-2}$. Hence

$$\Phi_n \le \Phi_n^* + (1+\lambda)\frac{1}{2}(n\sigma^2)\gamma(\gamma-1)n^{\gamma-2}$$

where σ^2 is the variance of T_i, which is some fixed quantity, independent of n. Using the fact that that $\Phi_n^* = (1+\lambda)(n\bar{T}_1)^\gamma$, and $\bar{T}_1$ does not depend on n, we have

$$\frac{\Phi_n}{\Phi_n^*} = 1 + O\left(\frac{1}{n}\right). \tag{47}$$

Now, consider $\gamma \in (1, 2)$. Pick k such that $k\gamma \ge 2$ and note that since $k > 1$

$$\frac{\Phi_n}{\Phi_n^*} \le E\left(\frac{T}{n\bar{T}_1}\right)^\gamma \le \left[E\left(\frac{T}{n\bar{T}_1}\right)^{k\gamma} \right]^{\frac{1}{k}}.$$

Since $k\gamma \ge 2$, we can apply the result in the first part of the proof and deduce that there is some B such that

$$\frac{\Phi_n}{\Phi_n^*} \le \left(1 + \frac{B}{n}\right)^{\frac{1}{k}} \le 1 + \frac{B}{kn}.$$

Thus, (47) holds in this case also.

Let us now turn to the result that holds under the weaker Assumption 2. The difference is that we cannot use a Taylor expansion as far as second order, but must be content with a first-order expansion

$$\Phi_n \le (1+\lambda)\xi(n\bar{T}_1) + (1+\lambda)E\left[(T - n\bar{T}_1)\xi'(T^*)\right]$$
$$\le (1+\lambda)\xi(n\bar{T}_1) + (1+\lambda)E|T - n\bar{T}_1|\xi'(n)$$

for some T^* depending on T, where we have $T^* < n$ and so $\xi'(T^*) \le \xi'(n)$. Now, Assumption 2 implies Lemma 1 which gives that $\xi'(n) = O(n^{\gamma-1})$ and $\Phi_n^* = \Omega(n^\gamma)$. Together with $E|T - n\bar{T}_1| = O(\sqrt{n})$ this gives $\Phi_n/\Phi_n^* = 1 + O(1/\sqrt{n})$. ■

APPENDIX V
PROOF OF THEOREM 2

Proof: Let $\bar{\theta}_{(1)} \ge \cdots \ge \bar{\theta}_{(n)}$ be points such that $F(\bar{\theta}_{(i)}) = 1 - i/(n+1)$. By a Taylor expansion around $\bar{\theta}_{(1)}, \ldots, \bar{\theta}_{(n)}$, we have that for some $\theta^*_{(1)}, \ldots, \theta^*_{(n)}$, with $\sum_{i=1}^m \theta^*_{(i)}$ certainly bounded above by m

$$m\xi\left(\frac{\sum_{i=1}^m \theta_{(i)}}{m}\right) = \max_Q \left\{ \sum_{i=1}^m \theta_{(i)} u(Q) - mc(Q) \right\}$$
$$= m\xi\left(\frac{\sum_{i=1}^m \bar{\theta}_{(i)}}{m}\right) + \sum_{i=1}^m \left(\theta_{(i)} - \bar{\theta}_{(i)}\right) \xi'\left(\frac{\sum_{i=1}^m \theta^*_{(i)}}{m}\right)$$
$$\le \max_m \left\{ m\xi\left(\frac{\sum_{i=1}^m \bar{\theta}_{(i)}}{m}\right) \right\} + \sum_{i=1}^n \left|\theta_{(i)} - \bar{\theta}_{(i)}\right| \xi'(1)$$

and hence, because we are ignoring the cost-covering constraint

$$\Phi_n \le E\left[\max_{m \in \{1,\ldots,n\}} \left\{ m\xi\left(\frac{\sum_{i=1}^m \theta_{(i)}}{m}\right) \right\} \right]$$
$$\le \max_m \left\{ m\xi\left(\frac{\sum_{i=1}^m \bar{\theta}_{(i)}}{m}\right) \right\} + \sum_{i=1}^n E\left|\theta_{(i)} - \bar{\theta}_{(i)}\right| \xi'(1)$$

where the final line follows from (48). Note first that

$$\int_{\bar{\theta}_{(i)}}^{\bar{\theta}_{(i-1)}} \theta dF(\theta) \ge \bar{\theta}_{(i)} \int_{\bar{\theta}_{(i)}}^{\bar{\theta}_{(i-1)}} dF(\theta)$$
$$= \bar{\theta}_{(i)} \left(F\left(\bar{\theta}_{(i-1)}\right) - F\left(\bar{\theta}_{(i)}\right) \right) = \frac{1}{n+1}\bar{\theta}_{(i)}.$$

Also, by definition of $\bar{\theta}_{(m)}$, we have $m = (n+1)(1 - F(\bar{\theta}_{(m)}))$. So the first term on the right of (48) is bounded by

$$(n+1) \max_{\bar{\theta}} \left\{ \left(1 - F(\bar{\theta})\right) \xi\left(\frac{\int_{\bar{\theta}}^1 \theta dF(\theta)}{1 - F(\bar{\theta})}\right) \right\}$$
$$= (n+1) \max_{\bar{\theta}, Q} \left\{ \left(\int_{\bar{\theta}}^1 \theta dF(\theta) \right) u(Q) - \left(1 - F(\bar{\theta})\right) c(Q) \right\}$$
$$= \left(\frac{n+1}{n}\right) \Phi_n^*.$$

Finally, the second term on the right of (48) is $O(\sqrt{n})$. To see this, we use the assumption in the theorem statement that the density function $f(x)$ is bounded below by some $a > 0$. (If we were to have $f(x) = 0$ in the interval $[F^{-1}(0.1), F^{-1}(0.9)]$, say, then in a sample of size $2n+1$ the $k = n+1$ order statistic cannot be arbitrarily close to its mean, which lies close to $F^{-1}(0.5)$. So the claim is not true.) But assuming such a lower bound on $f(x)$, then for all θ, θ'

$$|\theta - \theta'| \le \frac{|F(\theta) - F(\theta')|}{a}.$$

Now, if θ is a random variable with distribution function F, then $F(\theta)$ has the uniform distribution on [0, 1]. This fact, combined with the above, shows that it is sufficient to prove the result for uniform random variables on [0, 1]. But the kth largest of n samples of the uniform distribution has density

$$h(x) = k\binom{n}{k} x^{n-k}(1-x)^{k-1}$$

and so it is routine to calculate that

$$\sum_{k=1}^{n} E\left|\theta_{(k)} - \bar{\theta}_{(k)}\right| \le \sqrt{n \sum_{k=1}^{n} E\left[\left(\theta_{(k)} - \bar{\theta}_{(k)}\right)^2\right]} = \sqrt{\frac{n^2}{6(n+1)}} = O(\sqrt{n})$$

where the first inequality follows from the fact that for any random variables $X_1, \ldots, X_n$

$$\frac{1}{n}\sum_{i=1}^{n} E|X_i| \le \frac{1}{n}\sum_{i=1}^{n}\sqrt{E\left[X_i^2\right]} \le \sqrt{\frac{1}{n}\sum_{i=1}^{n} E\left[X_i^2\right]}.$$

■

APPENDIX VI
STRENGTHING EX-ANTE COST-COVERING TO EX-POST COST-COVERING

We take up the issue discussed in footnote 2 and prove a generalization of a result due to Crampton *et al.* in [16], and recently reproved by Norman [18].

Theorem 3: The expected social welfare that can be obtained under constaints of ex-post cost-covering, incentive compatible and individually rationality is just as great as can be obtained when the first constraint is weakened to ex-ante cost-covering. Moreover, we need not take payments from excluded agents.

In [16] and [18], the first sentence of the theorem is proved for a nonexcludable good. An ex-post cost-covering scheme with equally great expected social welfare is constructed from an ex-ante cost-covering scheme by a complicated reorganization of the agents' payments. This can end up with an agent, say i, being paid by others and so his payment $x_i(\boldsymbol{\theta})$ can be negative for some $\boldsymbol{\theta}$; this cannot be avoided. More disturbingly, if exclusions are allowed, the result of such a construction can be a scheme in which, for one or more $\boldsymbol{\theta}$ and i, we find $x_i(\boldsymbol{\theta}) \neq 0$, but $\boldsymbol{\theta}$ is such that agent i is to be excluded. This would require taking payment from an agent who is excluded, which is usually not a practical to do. Although we have not been able to find a way to modify the proofs in [16] or [18] to avoid this, we now show that it can indeed be avoided.

Proof: Suppose that θ_i is distributed with equal probabilities over m values $\{t_1 \le t_2 \cdots \le t_m\}$. This assumption of a uniform distribution is for convenience in exposition, and other distributions can be handled by modifying the arguments below to included weighting factors against the ϵ terms that appear, or by simply thinking about repeating values, e.g., $t_1 = t_2 = t_3 < t_4 = t_5 < t_6 \cdots \le t_m$. This discrete distribution can approximate a continuous one with arbitrary accuracy. (In fact, the following proof does not even require that $\theta_1, \ldots, \theta_n$ be identically distributed, but we assume this for notational simplicity.)

Consider an ex-ante cost-covering incentive compatible scheme, with payment functions $y_1(\boldsymbol{\theta}), \ldots, y_n(\boldsymbol{\theta})$. In a scheme that uses exclusions, θ_i will be excluded for some small values of θ_i, say for $\theta_i \le t_\ell$, and we will have $E[y_i(\boldsymbol{\theta})|\theta_i] = 0$ for all $\theta_i \le t_\ell$.

We will construct an ex-post cost-covering, incentive compatible scheme, with payment functions $x_1(\boldsymbol{\theta}), \ldots, x_n(\boldsymbol{\theta})$. This will be such that

(i) $x_1(\boldsymbol{\theta}) + \cdots + x_n(\boldsymbol{\theta}) = c(\boldsymbol{\theta})$;
(ii) $E[x_i(\boldsymbol{\theta})|\theta_i] = E[y_i(\boldsymbol{\theta})|\theta_i]$, for all i;
(iii) $x_i(\boldsymbol{\theta}) = 0$ if $\theta_i \le t_\ell$.

Condition (ii) ensure that the scheme is incentive compatible. Condition (iii) ensures we do not take payments from agents who are excluded. Note that for agents who are included we permit payments to be negative.

Modifications to obtain (i): Let us initially take $x_i(\cdot) = y_i(\cdot)$, for all i, giving an ex-ante cost-covering, incentive compatible scheme which satisfies (ii) and

$$\text{(i)}' \; E[x_1(\boldsymbol{\theta}) + \cdots + x_n(\boldsymbol{\theta})] = E[c(\boldsymbol{\theta})].$$

We will show how to modify the scheme to satisfy the ex-post constraint (i). This might be done as in [16]. However, we will do it by a new method, as the explanation is a good preparation for understanding how we obtain (iii) later in the proof. For the moment, we do not worry about (iii). Pick any $\boldsymbol{\theta}$ where (i) is violated because $x_1(\boldsymbol{\theta}) + \cdots + x_n(\boldsymbol{\theta}) > c(\boldsymbol{\theta})$, and any $\boldsymbol{\theta}'$, where it is violated because of $x_1(\boldsymbol{\theta}') + \cdots + x_n(\boldsymbol{\theta}') < c(\boldsymbol{\theta}')$. Note that (i)$'$ implies that there must always be such a pair if (i) does not hold for all $\boldsymbol{\theta}$. Consider two possibilities. If we can pick this pair such that $\theta_i = \theta_i'$ for either $i = 1$ or $i = 2$, then we simply make the following alterations, gradually increasing ϵ from 0 until we have (i) holding for one or both of $\boldsymbol{\theta}$ and $\boldsymbol{\theta}'$. The value of $E[x_i(\boldsymbol{\theta})|\theta_i]$ does not change and the number of violations to (i) decreases by at least one

$$\text{(a)} \quad x_i(\boldsymbol{\theta}) \to x_i(\boldsymbol{\theta}) - \epsilon$$
$$\text{(b)} \quad x_i(\boldsymbol{\theta}') \to x_i(\boldsymbol{\theta}') + \epsilon.$$

Alternatively, if the above is not possible, we must have $\theta_1' \neq \theta_1$ and $\theta_2' \neq \theta_2$. Let e_i be the n vector that has a 1 in the ith component and all other components 0. Let

$$\boldsymbol{\theta}'' = \boldsymbol{\theta} + (\theta_2' - \theta_2)\, e_2.$$

So $\boldsymbol{\theta}''$ is the state that is the same as $\boldsymbol{\theta}$ except that the second component is θ_2'. We make the following adjustments to the payments, gradually increasing ϵ from 0 until, by (a) or (d), we have (i) holding for one or both of $\boldsymbol{\theta}$ and $\boldsymbol{\theta}'$. Note that alterations (b)–(d) ensure that the values of $E[x_1(\boldsymbol{\theta})|\theta_1]$, $E[x_2(\boldsymbol{\theta})|\theta_2=\theta_2']$ and $x_1(\boldsymbol{\theta}'')+\cdots+x_n(\boldsymbol{\theta}'')$ do not change

$$\begin{aligned}&(a)\quad x_1(\boldsymbol{\theta})\to x_1(\boldsymbol{\theta})-\epsilon\\&(b)\quad x_1(\boldsymbol{\theta}'')\to x_1(\boldsymbol{\theta}'')+\epsilon\\&(c)\quad x_2(\boldsymbol{\theta}'')\to x_2(\boldsymbol{\theta}'')-\epsilon\\&(d)\quad x_2(\boldsymbol{\theta}')\to x_2(\boldsymbol{\theta}')+\epsilon.\end{aligned}$$

The number of violations to (i) has been decreased by at least one, and this process can be continued until no such violations are left. This shows that it is possible to have an ex-post cost-covering incentive compatible scheme.[6]

So now suppose we have an ex-post cost-covering incentive compatible scheme which satisfies (i) and (ii). We will modify it to produce a ex-post cost-covering incentive compatible scheme which also satisfies (iii).

Modifications to obtain (iii): Let $\Theta=\{\boldsymbol{\theta}:\theta_j\le t_\ell \text{ for all } j\}$. This is the set of states where all agents are excluded and $c(\boldsymbol{\theta})=0$. Suppose there is a $\boldsymbol{\theta}\in\Theta$ and i such that there is a violation of (iii), i.e., $\theta_i\le t_\ell$ and $x_i(\boldsymbol{\theta})\ne 0$. Pick a j such $x_j(\boldsymbol{\theta})\ne 0$; note that this is always possible because $\boldsymbol{\theta}$ cannot be such that there is just one violation of (iii), since we have (i) and $c(\boldsymbol{\theta})=0$ for $\boldsymbol{\theta}\in\Theta$ and this implies $x_1(\boldsymbol{\theta})+\cdots+x_n(\boldsymbol{\theta})=0$. Let

$$\begin{aligned}\boldsymbol{\theta}^i&=\boldsymbol{\theta}+(t_m-\theta_i)e_i\\\boldsymbol{\theta}^{ij}&=\boldsymbol{\theta}+(t_m-\theta_i)e_i+(t_m-\theta_j)e_j\\\boldsymbol{\theta}^j&=\boldsymbol{\theta}+(t_m-\theta_j)e_j.\end{aligned}$$

That is, states $\boldsymbol{\theta}^i$, $\boldsymbol{\theta}^j$, and $\boldsymbol{\theta}^{ij}$, are constructed from $\boldsymbol{\theta}$ by, respectively, increasing the ith, jth, or both the ith and jth components of $\boldsymbol{\theta}$ to t_m (for which an agent is definately not excluded). Let $\epsilon=x_i(\boldsymbol{\theta})$. We now make further modifications to the payments, choosing these carefully to preserve (i) and (ii), and remove the violation to (iii) of $x_i(\boldsymbol{\theta})\ne 0$

$$\begin{aligned}&(a)\quad x_i(\boldsymbol{\theta})\to x_i(\boldsymbol{\theta})-\epsilon\\&(b)\quad x_j(\boldsymbol{\theta})\to x_j(\boldsymbol{\theta})+\epsilon\\&(c)\quad x_i(\boldsymbol{\theta}^j)\to x_i(\boldsymbol{\theta}^j)+\epsilon\\&(d)\quad x_j(\boldsymbol{\theta}^j)\to x_j(\boldsymbol{\theta}^j)-\epsilon\\&(e)\quad x_j(\boldsymbol{\theta}^{ij})\to x_j(\boldsymbol{\theta}^{ij})+\epsilon\\&(f)\quad x_i(\boldsymbol{\theta}^{ij})\to x_i(\boldsymbol{\theta}^{ij})-\epsilon\\&(g)\quad x_i(\boldsymbol{\theta}^i)\to x_i(\boldsymbol{\theta}^i)+\epsilon\\&(h)\quad x_j(\boldsymbol{\theta}^i)\to x_j(\boldsymbol{\theta}^i)-\epsilon.\end{aligned}$$

In detail

(a) removes the violation to (iii);
(b) makes $x_1(\boldsymbol{\theta})+\cdots+x_n(\boldsymbol{\theta})=c(\boldsymbol{\theta})$;
(c) puts $E[x_i(\boldsymbol{\theta})|\theta_i]$ back to its required value;
(d) makes $x_1(\boldsymbol{\theta}^j)+\cdots+x_n(\boldsymbol{\theta}^j)=c(\boldsymbol{\theta}^j)$;
(e) puts $E[x_j(\boldsymbol{\theta})|\theta_j]$ back to its required value;
(f) makes $x_1(\boldsymbol{\theta}^{ij})+\cdots+x_n(\boldsymbol{\theta}^{ij})=c(\boldsymbol{\theta}^{ij})$;
(g) puts $E[x_i(\boldsymbol{\theta})|\theta_i]$ back to its required value;
(h) makes $x_1(\boldsymbol{\theta}^i)+\cdots+x_n(\boldsymbol{\theta}^i)=c(\boldsymbol{\theta}^i)$ and puts $E[x_j(\boldsymbol{\theta})|\theta_j]$ back to its required value.

Unfortunately, this may create a new violations of (iii) if $x_i(\boldsymbol{\theta}^j)=0$ and $x_i(\boldsymbol{\theta}^j)+\epsilon\ne 0$, or if $x_j(\boldsymbol{\theta}^i)=0$ and $x_j(\boldsymbol{\theta}^i)-\epsilon\ne 0$. But even if this happens the number of violations within the set Θ will have been reduced by at least 1, since $\boldsymbol{\theta}^i,\boldsymbol{\theta}^j\notin\Theta$. So let us first make adjustments to violations that occur within Θ until none such are left.

Now, look for a violation of (iii) for $\boldsymbol{\theta}\notin\Theta$. Suppose there is one, say $\boldsymbol{\theta}$, with $\theta_i\le t_\ell$ and $x_i(\boldsymbol{\theta})\ne 0$. Since $\boldsymbol{\theta}\notin\Theta$, there is a j such that $\theta_j>t_\ell$. Now, observe that there is at least one other violation of (iii) for the same agent and θ_i, since $E[x_i(\boldsymbol{\theta})|\theta_i]=0$. Suppose this is at $\boldsymbol{\theta}'$, where $x_i(\boldsymbol{\theta}')\ne 0$ and that, again since $\boldsymbol{\theta}'\notin\Theta$, there is a k such that $\theta_k'\ge t_\ell$ (where k may possibly be the same as j). Let

$$\boldsymbol{\theta}^\dagger=\boldsymbol{\theta}+(\theta_k'-\theta_k)\,e_k$$

and make the following alterations, with the same motivations as above, The effect is to retain (i) and (ii) and remove the violation to (iii) at $x_i(\boldsymbol{\theta})\ne 0$

$$\begin{aligned}&(a)\quad x_i(\boldsymbol{\theta})\to x_i(\boldsymbol{\theta})-\epsilon\\&(b)\quad x_i(\boldsymbol{\theta}')\to x_i(\boldsymbol{\theta}')+\epsilon\\&(c)\quad x_k(\boldsymbol{\theta}')\to x_k(\boldsymbol{\theta}')-\epsilon\\&(d)\quad x_k(\boldsymbol{\theta}^\dagger)\to x_k(\boldsymbol{\theta}^\dagger)+\epsilon\\&(e)\quad x_j(\boldsymbol{\theta}^\dagger)\to x_j(\boldsymbol{\theta}^\dagger)-\epsilon\\&(f)\quad x_j(\boldsymbol{\theta})\to x_j(\boldsymbol{\theta})+\epsilon.\end{aligned}$$

This does not introduce any new violations to (iii) in the states $\boldsymbol{\theta}$ and $\boldsymbol{\theta}^\dagger$ because there is already a violation $x_i(\boldsymbol{\theta}')\ne 0$, and all the other changes are to payments being made by agents j and k who are not excluded. Thus, we may repeat this until the number of violations of (iii) is zero. ■

Appendix VII
Ex-Post Cost-Covering Equal-Contributions Mechanisms

Here are details of the calculations referred to in Section III-C, where we consider fixed contribution schemes that satisfy ex-post cost-covering conditions.

A. Mechanism 1

Suppose that we build a facility of size Q and then share the cost $c(Q)$ amongst all those who volunteer to participate. They must make a commitment to do so, before knowing how many will participate.

Let X be the number of peers with preferences parameters of at least θ. The marginal θ (i.e., θ such that a peer will indifferent between participating and not participating) is given by

$$\theta u(Q)-c(Q)E\left[\frac{1}{X}\right]=0.$$

[6]This also shows that if all payments in the original scheme are nonnegative, then there is a ex-post cost-covering scheme in which at most two agents (namely, 1 and 2) might make negative payments in some states. The fact that it may be impossible to avoid negative payments completely can be seen from an example in which $n=2$ and $\theta_i\in\{t_1,t_2\}$. Suppose we have

$$(c(t_i,t_j))=\begin{pmatrix}6&7\\7&8\end{pmatrix}$$

$$((y_1(t_i,t_j),y_2(t_i,t_j)))=\begin{pmatrix}(1,1)&(1,6)\\(6,1)&(6,6)\end{pmatrix}.$$

An ex-post cost-covering scheme that respects (ii), namely, $E[x_i(\boldsymbol{\theta})|\theta_i]=E[y_i(\boldsymbol{\theta})|\theta_i]$, is

$$((x_1(t_i,t_j),x_2(t_i,t_j)))=\begin{pmatrix}(3,3)&(-1,8)\\(8,-1)&(4,4)\end{pmatrix}$$

but there is no ex-post cost-covering scheme with all payments nonnegative.

The expected social welfare will be

$$E\left[X\frac{1}{2}(1+\theta)u(Q)-c(Q)\right].$$

So, for our example, we have the problem of maximizing with respect to Q and θ

$$n(1-\theta)\frac{1}{2}(1+\theta)\frac{2}{3}\sqrt{Q}-Q \tag{48}$$

subject to

$$\theta\frac{2}{3}\sqrt{Q}-QE\left[\frac{1}{X}\right]\geq 0. \tag{49}$$

Now

$$\begin{aligned}E\left[\frac{1}{X}\right]&=\frac{1}{EX}E\left[\frac{1}{1+\frac{X-EX}{EX}}\right]\\&=\frac{1}{EX}\left(1+\frac{E(X-EX)^2}{[EX]^2}+\cdots\right)\\&=\frac{1}{n(1-\theta)}+\frac{\theta}{n^2(1-\theta)^2}+O\left(\frac{1}{n^4}\right).\end{aligned}$$

Now, use (48) to solve for Q, making the approximation $E[1/X]\approx 1/n(1-\theta)$. Substituting this in (48) and maximizing with respect to θ gives $\theta=1/4$ and a maximized value of $3n^2/128-n/384+o(1)$. Compare this to $\Phi_n^*=3n^2/128$.

B. Mechanism 2

Suppose that we charge a fee of ϕ and then build the largest facility whose cost can be covered by the fees. As before, peers must make a commitment to pay their share of the cost without knowing how many others will also participate. The marginal θ is now given by

$$\theta Eu(X\phi)-\phi=0$$

and the expected social welfare is

$$E\left[X\frac{1}{2}(1+\theta)u(X\phi)-c(X\phi)\right].$$

So we want to maximize with respect to θ and ϕ,

$$E\left[X\frac{1}{2}(1+\theta)\frac{2}{3}\sqrt{X\phi}-X\phi\right] \tag{50}$$

subject to

$$\theta\frac{2}{3}E\sqrt{X\phi}-\phi\geq 0. \tag{51}$$

Now

$$\begin{aligned}E\sqrt{X}&=\sqrt{EX}E\left[1+\frac{X-EX}{EX}\right]^{\frac{1}{2}}\\&=\sqrt{EX}E\left[1+\frac{1}{2}\frac{X-EX}{EX}-\frac{1}{8}\frac{(X-EX)^2}{(EX)^2}+\cdots\right]\\&=(n(1-\theta))^{\frac{1}{2}}\left(1-\frac{1}{8}\frac{\theta}{n(1-\theta)}+\cdots\right)\\E\sqrt{X^{\frac{3}{2}}}&=(n(1-\theta))^{\frac{3}{2}}\left(1+\frac{3}{8}\frac{\theta}{n(1-\theta)}+\cdots\right).\end{aligned}$$

So ignoring terms that become small for large n, we want to maximize

$$\frac{1}{3}\sqrt{\phi}\,(n(1-\theta))^{\frac{3}{2}}\,(1+\theta)\left(1+\frac{3}{8}\frac{\theta}{n(1-\theta)}\right)-\phi n(1-\theta) \tag{52}$$

subject to

$$\frac{2}{3}\theta\,(n(1-\theta))^{\frac{1}{2}}\left(1-\frac{1}{8}\frac{\theta}{n(1-\theta)}\right)\sqrt{\phi}-\phi\geq 0. \tag{53}$$

Now, use (53) to find ϕ as a function of θ and substitute into (52). The optimal θ tends to 1/4, as we would expect. The optimal value is $3n^2/128+7n/1536+o(1)$. Compare this to $\Phi_n^*=3n^2/128$ and $\Phi_n=3n^2/128+7n/9+o(1)$. The ex-post cost-covering scheme does a bit better because although it provides the same Q on average it provides more Q when more peers participate. The explanation is that when $u(Q)=(2/3)\sqrt{Q}$ the term in the social welfare of $X(1/2)(1+\theta)u(\phi X)$ is a convex function of X, so

$$E\left[X\frac{1}{2}(1+\theta)u(\phi X)\right]\geq EX\frac{1}{2}(1+\theta)u(\phi EX).$$

ACKNOWLEDGMENT

The authors are grateful to R. Mason for initially suggesting the connection of Norman and Hellwig's work to P2P file sharing and for many helpful discussions in the context of the IST project MMAPPS. We have had some helpful correspondence with both these authors. They are also grateful to P. Antoniadis for his collaboration in the development of the models for content sharing and peering WLANs.

REFERENCES

[1] M. Hellwig, "Public good provision with many participants," *Rev. Econ. Studies*, vol. 70, no. 3, pp. 589–614, Jul. 2003.
[2] P. Norman, "Efficient mechanisms for public goods with use exclusions," *Rev. Econ. Studies*, vol. 7, pp. 1163–1188, 2004.
[3] P. Golle, K. Leyton-Brown, I. Mironov, and M. Lillibridge, "Incentives for sharing in peer-to-peer networks," in *Proc. WELCOM*, Nov. 2001, pp. 264–267.
[4] C. Buragohain, D. Agrawal, and S. Suri, "A game theoretic framework for incentives in P2P systems," in *Proc. 3rd IEEE Int. Conf. Peer-to-Peer Comput.*, Linkoping, Sweden, Sep. 2003, pp. 48–56.

[5] A. Asvanund, R. Krishnan, M. D. Smith, and R. Telang, "Intelligent club management in peer-to-peer networks," in *Workshop on Economics of Peer-to-Peer Systems*, 2003. [Online]. Available: http://www.sims.berkeley.edu:8000/research/conferences/p2pecon/papers/s6-asvanund.pdf.

[6] R. Krishnan, M. D. Smith, and R. Telang, "The economics of peer-to-peer networks," *J. Inf. Technol. Theory and Appl.*, pp. 31–44, 2003.

[7] R. Krishnan, M. Smith, Z. Tang, and R. Telang, "The virtual commons: Why free-riding can be tolerated in peer-to-peer networks," in *Proc. Workshop Inf. Syst. Econ.*, 2003. [Online]. Available at SSRN: http://ssrn.com/abstract=450241.

[8] P. Antoniadis, C. Courcoubetis, and R. Mason, "Comparing economic incentives in peer-to-peer networks," *Special Issue on Network Economics, Computer Networks*, vol. 45, no. 1, pp. 133–146, 2004.

[9] P. Antoniadis, C. Courcoubetis, and R. Weber, "An asymptotically optimal scheme for P2P file sharing," in *Proc. 2nd Workshop on Economics of Peer-to-Peer Syst.*, 2004. [Online]. Available: http://www.eecs.harvard.edu/p2pecon/confman/papers/s6p1.pdf.

[10] C. Courcoubetis and R. R. Weber, "Asymptotics for provisioning problems of peering wireless LAN's with a large number of participants," in *Proc. WiOpt'04 Workshop*, Cambridge, U.K., 2004, pp. 217–226.

[11] ——, (2003) Asymptotics for provisioning problems with a large number of participants. AUEB Tech. Memo. [Online]. Available: http://nes.aueb.gr/research/

[12] M. Feldman, K. Lai, J. Chuang, and I. Stoica, "Quantifying disincentives for collaboration in peer-to-peer networks," in *Workshop on Economics of Peer-to-Peer Syst.*, Berkeley, CA, 2003. [Online]. Available: http://www.sims.berkeley.edu:8000/research/conferences/p2pecon/papers/s5-feldman.ps.

[13] V. Vishnumurthy, S. Chandrakumar, and E. G. Sirer, "Karma: A secure economic framework for peer-to-peer resource sharing," in *Workshop on Economics of Peer-to-Peer Systems*, Berkeley, CA, 2003. [Online]. Available: http://www.sims.berkeley.edu:8000/research/conferences/p2pecon/papers/s5-vishnumurthy.pdf.

[14] D. Hausheer, N. Liebau, A. Mauthe, R. Steinmetz, and B. Stiller, "Token-based accounting and distributed pricing to introduce market mechanisms in a peer-to-peer file sharing scenario," in *Proc. 3rd IEEE Int. Conf. Peer-to-Peer Comput.*, Linkoping, Sweden, Sep. 2003, pp. 200–201.

[15] M. Consortium. (2004) Market management of peer-to-peer services—White Paper. [Online]. Available: http://www.mmapps.org/

[16] P. Crampton, R. Gibbons, and P. Klemperer, "Dissolving a partnership efficiently," *Econometrica*, vol. 55, pp. 615–632, 1987.

[17] P. Antoniadis, C. Courcoubetis, and B. Strulo, "Incentives for content availability in memory-less peer-to-peer file sharing systems," in *Special Issue on P2P Economics, ACM SIG on Ecommerce*, vol. 5, Jul. 2005, pp. 11–20.

[18] P. Norman and T. Borgers. (2005, Jun.) A note on cost-covering under interim participation constraints: The case of independent types. [Online]. Available: http://faculty.arts.ubc.ca/pnorman/Research/Research.htm

Costas Courcoubetis received the Ph.D. degree from the University of California, Berkeley

He is Professor of Computer Science and heads the Network Economics and Services Group and the Theory, Economics and Systems Laboratory, Athens University of Economics and Business, Athens, Greece. From 1982 to 1990, he was Member of the Technical Staff in the Mathematical Sciences Research Center, Bell Laboratories. He is a coauthor of *Pricing Communication Networks: Economics, Technology and Modeling* (Wiley, 2003), with R. Weber. He has been a Consultant frequently with Bell Laboratories and the Greek regulation authorities, and has participated in many EU research projects related to network economics such as Ca$hman, M3i, and MMAPPS. His research interests include economics of communication networks, resource allocation and optimization, peer-to-peer computing, and regulation policy.

Richard Weber is Churchill Professor of Mathematics for Operational Research, Mathematics Department, University of Cambridge, Cambridge, U.K. His research interests include the economics of networks and pricing, online bin packing, queueing control, and scheduling. He has been a Consultant frequently with AT&T Research and Bell Laboratories and has participated in EU research projects related to network economics such as Ca$hman, M3i, and MMAPPS.