

Incentives for Large Peer-to-Peer Systems

Costas Courcoubetis* and Richard Weber†

September, 2005

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 social welfare in a way that is incentive compatible, rational and feasible. Although it is unfortunately almost never possible to calculate or implement a truly 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 a fixed contribution (payment) 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 LANs, 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.

Keywords: file sharing, incentive mechanisms, mechanism design, peer-to-peer, peering, wireless LANs.

1 Introduction

The design of a peer-to-peer (P2P) system poses many interesting questions. If the quantity of the service it provides, say Q , is a design parameter, what should it 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 LAN (WLAN) resources when peers roam in some geographic location away from their home LAN. Now Q is ‘the availability of connectivity for roaming peers’.

Let us think of the service offered by a peer-to-peer 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?

*Department of Computer Science, Athens University of Economics and Business, Patission 76, Athens 10434, GR

†Centre for Mathematical Sciences, University of Cambridge, Wilberforce Road, Cambridge CB3 0WB, UK

Each of the peer-to-peer systems that we consider in this paper has two important characteristics. Firstly, viewed as a good, it is *non-rivalrous*, 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 it 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 consider such 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, set Q , and decide which of them 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 social welfare. While the full solution of this problem is extremely complex and not easily solved in practice, we show in Section 2 that as the number of agents becomes large, there is a good solution to our problem that takes a very simple form. We merely require each agent to pay the same fixed fee towards payment of 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 to a fixed part of the overall service. Such a simple contribution policy is easy to implement and requires no centralized implementation. The only information required by the system designer to compute the fixed fee is the distribution of the agents' valuations for the service.

Although there were results pointing to the fact that in the limit optimal incentive payments reduced to fixed and equal contributions by all peers, there were no result for how well such simple policies were performing as a function of the number of peers. Our proofs allow to obtain such tight bounds. The positive result is that the optimal policies are only a 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 in the sense that the system designer declares a fixed participation fee and the expected system size to all participants. Then peers must make a simple decision, to 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 fast. We also show how to extend the system to handle multiple constraints on how the cost should be covered, and also to reflect a dependence on the number of final participants. We apply our approach to model file sharing and peering wireless LANs, and we derive the simple optimization problems from which one can compute the optimal incentive policies. We also show that the existence of a weakly feasible incentive compatible mechanism implies the existence of a strongly feasible one.

We must emphasize that 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 very simple models for our P2P systems, in which peers are identified by a single parameter, their valuation of the service. One can think of it as the 'first order term' of a more detailed model which would take into account more accurate

models of cost, utility, and peer interactions. But we expect such more complex models 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 2 we describe a model for an excludable public good and how to solve the optimization problem of finding a social welfare maximizing feasible incentive compatible mechanism. Section 3 presents our asymptotically optimal scheme and states the main theorem concerning it. In 3.2 we work through a numerical example that illustrates the ideas, in 3.3 we discuss other schemes that also require each participant to make the same payment, and 3.4 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 4 and 5 contain our applications to P2P file sharing systems and WLANS. Section 6 looks at questions of stability and convergence when some information must be learned. Our conclusions are in Section 7.

Some problem formulation and longer proofs are placed in the Appendix. Appendix A contains the set up of the optimization problem of maximizing social welfare, and B justifies the fact that it can be solve using Lagrangian methods. Appendix C contains proof of a technical lemma. The proofs of our main theorems 1 and 2 are in D and E. Appendix F contains Theorem 3, which is a significant theoretical result stating even when the mecahnism may use exclusions, the existence of a weakly feasible incentive compatible scheme implies the existence of a strongly feasible incentive compatible one. Appendix G contains detailed calculations for Section 3.3.

1.1 Related work

Our work has been motivated by that of Hellwig [13] and Norman[16], 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 simpler arguments that 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 5.

Golle et. al. [11] 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. [5] 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 [4], [15], [14]. Antoniadis et. al. [1] have attempted to compare different incentive schemes, one of them being the simple contribution scheme analyzed in this paper. Similarly, [3] and [8] contain

more elaborate applications to file sharing and WLANs using our simple contribution scheme. Regarding the asymptotic results, all this work referenced [7], 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 is tight.

From a practical perspective, a P2P system designer has to deal with the fact that he is unable to rely on trusted software¹ 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 [10], [18], [12]. 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 the free riders by reducing their reputation and restricting their downloads. The MMAPPS Consortium [6] 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.² This P2P application relies on central control exercised by a special peer sub-group enforcing specific minimum contribution rules, and excluding peers that are found to contribute less based on their IP addresses.

2 An Excludable Public Good Model

Consider an excludable public good as described in Section 1. 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 towards the cost of providing the good. Here $\theta = (\theta_1, \dots, \theta_n)$ 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 . The reader should 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 provided through a single parameter, θ_i . We could also easily refine this model to assume that agents are not of the same type (i.e., 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 this case 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 $\theta = (\theta_1, \dots, \theta_n)$, 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

¹Kazaa 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.

²<http://www.neo-modus.com/>

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)–(2) that follow the expectation is taken over $\boldsymbol{\theta}$ and in (3)–(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:

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

subject to a ‘weak feasibility constraint’, which says that the expected payments must at least cover the expected cost

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

an ‘individual rationality’ constraints, which says each agent can expect positive net benefit:

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

and ‘incentive compatibility’ constraints, such that each agent i does best by declaring his true θ_i rather than ‘free-riding’ by declaring some other θ'_i :

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

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 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 be 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 weakly feasible³, individually rational and incentive compatible.

³In fact, without actually making the problem any more difficult, we can strengthen the constraint of weak feasibility to strong feasibility, in which the expectation operator is removed from (2) thus requiring that cost is covered for each realization of preference parameters, not just on the average. Crampton, et. al. [9] give an argument for a case when exclusions are not allowed. However, the method of their proof does not generalize to mechanism designs in which a possible control is to exclude agents. Although some authors have assumed that the result is still true in such circumstances, we believe it has actually been an open question. We give what we believe to be the first proof of this result, and provide a novel proof that is quite different to the constructive proof of [9]. See Appendix F.

Let us define

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

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

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

This is the same as a model of Norman [16], but he takes $c(n, Q) = Qc(n)$.

In Appendix B 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

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

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

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

where

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

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 (e.g., this would follow from the assumption that 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 $Q(\boldsymbol{\theta})$ which maximizes (8) is decreasing in λ .

3 The Asymptotically Optimal Mechanism

3.1 A Scheme of Equal Contributions

The full solution of our problem is, in general, very complex. However, in Appendix D we prove Theorem 1 below, which shows that, when n is large, a nearly optimal solution can be achieved with a simple mechanism design. First we make an assumption.

Assumption 1 *Suppose that*

$$u(Q) = A Q^\alpha, \quad (10)$$

$$c(n, Q) = B n^\delta Q^\beta, \quad (11)$$

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

We shall also prove results under the following weaker assumption.

Assumption 2 Suppose $u(Q) = \Theta(Q^\alpha)$ and $c(n, Q) = h(n)c(Q)$, where $c(Q) = \Theta(Q^\beta)$, where $0 < \alpha \leq 1$, $\beta \geq 1$, 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 ,

$$A_1 Q^\alpha \leq u(Q) \leq A_2 Q^\alpha, \quad (12)$$

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

The assumption has an important consequence in bounded the growth rate of the optimal value from above and below, as we see in the following lemma.

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

$$\xi(x) = \max_Q \{x u(Q) - c(n, Q)\}. \quad (14)$$

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 C

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

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

subject to

$$n[1 - F(\theta)]\theta u(Q) - c(n, Q) \geq 0. \quad (16)$$

Let the optimal value be Φ_n^* .

Suppose we take as a feasible solution to \mathcal{P} the decision variables $\pi_i(\theta) = 1\{\theta_i \geq \theta^*\}$ and $Q(\theta) = Q^*$. 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 and Φ_n^* are $\Theta(n^\gamma)$, where $\gamma = (\beta - \delta\alpha)/(\beta - \alpha)$.

The intuition behind 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 can expect it to be very nearly optimal to adopt the mechanism above, i.e., to take $Q(\theta) = Q^*$ and set a fixed fee of $\theta^* u(Q^*)$, thus including those peers for which $\theta_i \geq \theta^*$.

3.2 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 social welfare that could be achieved if we were to have full information about $\theta_1, \dots, \theta_n$ and were not restricted by constraints of weak feasibility, individual rationality and incentive compatibility. Given that $\sum_{i=1}^n \theta_i$ take some value T , the 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 value of the social welfare is thus

$$(1/9)(\text{var}(T) + (ET)^2) = (1/9)(n/12 + (n/2)^2) = n^2/36 + n/108 \approx 0.02778 n^2.$$

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

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

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

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

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

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 $\theta_1 > \lambda/(1 + 2\lambda)$,

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

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

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

The minimizing value of Lagrange multiplier is

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

(with $\lambda_n \rightarrow \lambda = 1/2$, 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

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

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 social welfare per capita increases with n and there is advantage in having more participants to share the cost. 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.




Figure 1: Plot of Φ_n/Φ_n^* against n .

This illustrates the statement in Theorem 1 that $\Phi_n/\Phi_n^* = 1 + O(1/n)$.

Finally, let us compare what happens if exclusions are not allowed. Let us compute the second-best social welfare value, say Φ_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 $E[(\sum_i \theta_i)^2] = n^2/4 + n/12$. We find that in our example that, when exclusions are not allowed

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

Note that the social welfare that can be obtained per capita grows as \sqrt{n} , but is vanishingly small relative to the first-best level of social welfare per capita, which grows linearly in n . find this by appropriate calculations of The optimal fee structure is

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

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.

3.3 Other Equal Contribution 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. Two obvious such mechanisms suggest themselves. 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 social welfare using this scheme is $3n^2/128 - n/384 + O(1)$. So the imposition of strict feasibility occasions a loss compared to $\Phi_n^* = 3n^2/128$. Further details of these calculations are in Appendix G.

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 $\Phi_n = 3n^2/128 + 7n/9 + O(1)$. The strongly feasible policy does a bit better than a scheme in which Q is fixed a priori, because although it provides the same Q on average, it automatically saves bigger values of Q for the times that more users participate. Note that $mu(Q) = m\sqrt{m\phi}$ is convex in m .

An even better strongly feasible equal contribution scheme would be one in which, having learnt that m 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.

3.4 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 . Think now that Q is a vector (Q_1, \dots, Q_k) . Since \mathcal{P}^* can be solved by Lagrangian methods (cf. Appendix B), we know

there exists multipliers $\lambda_1, \dots, \lambda_k$ such that

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

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

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

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 each 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 $u_j(Q) = Q_1^{1/2}$, ($Q_2 = \dots = Q_k = 0$), in which peers of type i are required to cover exactly $(1/k)$ th of the cost, perhaps by making payments in kind.⁴ We now have from (17), and redefining $\xi(x) = \max_Q \{xQ^{\alpha} - Q^{\beta}\}$,

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

The second line follows from the first by bringing the sum on ℓ to the outside, dividing and multiplying with the term $\left[\sum_{j=1}^k (1 + \lambda_j) \nu_{j\ell} \right]$, and defining

$$T^{\ell} = \frac{\sum_{j=1}^k \sum_{i=1}^{n_j} (\theta_i + \lambda_j g_j(\theta_i))^+}{\sum_{j=1}^k (1 + \lambda_j) \nu_{j\ell}} \eta_{j\ell}, \text{ hence } ET^{\ell} = n \frac{\sum_{j=1}^k \rho_j E[(\theta_i + \lambda_j g_j(\theta_i))^+]}{\sum_{j=1}^k (1 + \lambda_j) \nu_{j\ell}} \eta_{j\ell}.$$

⁴In fact if we were to take more generally, something like $u_j(Q) = Q^{\alpha_j}$ and $c_j(Q) = Q^{\beta_j}$, the social welfare will be $\Theta(n^{\gamma})$, where $\gamma = \max_j \{\beta_j / (\beta_j - \alpha_j)\}$ and so it is only one type of good that really matters as $n \rightarrow \infty$.

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

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

Notice that the optimizing $\lambda_1, \dots, \lambda_n$ depend on ρ_1, \dots, ρ_k , but not n . The remaining details are as before. We find, under these assumptions, that $\Phi_n^* = \Theta(n^\gamma)$ and $\Phi_n/\Phi_n^* = 1 + O(1/n)$.

4 Application to File Sharing

We apply the above ideas to a problem of peer to peer 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. Imagine that 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

$$Q = N(1 - (1 - \phi/N)^n), \quad (19)$$

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

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

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 peers supply.⁵

In Figure 1, 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

$$c(n, Q) = n\phi(n, Q) = Q[1 + 0.495(Q/N) + 0.32835(Q/N)^2 + \dots]$$

⁵An alternative would be that a peer's cost is proportional to the rate at which he serves upload requests. Assuming files are equally popular this means that 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 neighbourhood of their location, this might be better modelled 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 B) 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 $c(Q) = [n(1 - F(\theta^*))]^\beta Q$.




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

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

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, $a > 0$. Let i^* be an upper bound on the number of files that any one peer can share, and $\sum_i \rho_i = 1$. The expected number of distinct files supplied will be

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

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 $n \rightarrow \infty$ is $c(n, Q) = -N \log(1 - Q/N)$. If Q/N is small, we again have $c(n, Q) = Q(1 + \frac{1}{2}(Q/N) + \frac{1}{3}(Q/N)^2 + \dots) \approx Q$.

Both of the above lead to models that are covered by Section 2. The social planner wishes to design a mechanism which maximizes social welfare, subject to its being feasible, 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

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

subject to

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

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 non-limiting version of the problem, addressed by the optimal mechanism design of solving \mathcal{P} in Section 2, the individually rationality constraint (3) is in terms of expected value, so for some θ_{-i} it can be that $\theta_i u(Q(\theta_i, \theta_{-i})) - p_i(\theta_i, \theta_{-i}) < 0$. When this happens, peer i might be tempted to defect and to not pay $p_i(\theta_i, \theta_{-i})$. 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 θ_{-i} is such that his net benefit would be positive. We are imagining that θ 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 such 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 social welfare.

5 Application to WLANS

Now we apply our ideas to wireless LANS. 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, deploying infrastructure with wide enough coverage to support this is a non-trivial task, especially from the business perspective.

Wireless Local Area Networks (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 neighbourhood 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 i th 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, 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. Assuming that peering agent i accepts service requests from roaming peers with probability p_i , we can express Q as a function of p_1, \dots, p_n . 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 does not peer, then $p_i = 0$. If agent i does peer 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 non-rivalrous 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. 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_i s offered by the agents. We can cast the mechanism design 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

$$c(Q) = \sigma \lambda m \sum_{i=1}^n p_i = (\sigma \lambda B/A) m Q,$$

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, 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 us take the cost to be of the more general form $m c(Q)$. Our problem \mathcal{P} now becomes

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

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 $u(Q) = A Q^\alpha$, $c(Q) = B Q^\beta$ and $\gamma = \beta/(\beta - \alpha)$. Choose units so that

$$\xi(x) = \max_Q \{x u(Q) - c(Q)\} = (A \alpha x)^{\frac{\beta}{\beta - \alpha}} (B \beta)^{-\frac{\alpha}{\beta - \alpha}} (1/\alpha - 1/\beta) = x^\gamma.$$

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

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

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 $\theta_{(1)} \geq \dots \geq \theta_{(n)}$ 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

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

subject to

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

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

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

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

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

This has a unique solution if θ has a distribution that is ‘new better than used in expectation’ (NBUE), i.e., the right hand side of (30) is decreasing in $\bar{\theta}$. 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,*

$$u(Q) = AQ^\alpha. \quad (31)$$

$$c(m, Q) = Bm^\delta Q^\beta \quad (32)$$

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*

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

The proof is in Appendix E.

6 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 $f = \theta u(Q)$. Since the fee is paid ‘in kind’ and equates to providing f files, say (in the case that the problem is one of file sharing), 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 realise that it is to their advantage to participate. The fees paid by these will provide Q_2 files where

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

Write this as $Q_2 = \phi(Q_1)$ and imagine iterating $Q_{k+1} = \phi(Q_k)$, $k = 1, 2, \dots$. 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

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

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 $Q_k \rightarrow 0$. For $Q = 100$ the 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 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 .

7 Conclusions

In this paper we have formalized an interesting connection between P2P systems, namely file sharing and peering wireless LANs, 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 that our economic model is rather crude and abstracts many practical aspects of the implementations, it captures the ‘first-order’ properties of the large externalities that such system exhibit, and which are the main cause for their large adoption. So in that sense, we have a very positive result.

There are many ways to extend this simple model but most of them may not lead to tractable solutions. For instance, personalizing agents by more than a single parameter (for instance, add a parameter that captures cost sensitivity) probably makes any analysis hopeless. One could perhaps do a more careful model of the utility function or the cost function by making it depend on more detailed congestion effects and hence capture performance aspects encountered in specific network technologies. But it is not obvious that making the problem more complex will essentially provide us with a better inside. The results one obtains by such

models are more qualitative than quantitative showing rather the form of the optimal control than computing its exact value.

In our view, interesting extensions include the file sharing case by introducing the concept of popularity and uploading cost. In this more refined model, one should model the effect of different file types (in terms of popularity), possibly using the results in Section 3.4. A conjecture is that since popular files may result in greater uploading costs the fixed contributions should be in terms of total upload rates. This will affect which files peers make available for uploading, and so influence the equilibrium of the type of content that will be available in the system. Also the wireless LANs to capture quality of service (amount of bandwidth allocated to roaming peers) and differentiating areas of low demand from areas of higher demand in terms of the amounts of peers should contribute in these different areas. Finally, we could assume that the number of participants, n , is not fixed, but that there are arrivals and departures.

An open problem of practical interest is 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 [2].

We are grateful to Robin 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. We are also grateful to Panos Antoniadis for his collaboration in the development of the models for content sharing and peering WLANs.

References

- [1] P. Antoniadis, C. Courcoubetis, and R. Mason. Comparing economic incentives in peer-to-peer networks. *Special Issue on Network Economics, Computer Networks, Elsevier*, 45(1):133–146, 2004.
- [2] P. Antoniadis, C. Courcoubetis, and B. Strulo. Incentives for content availability in memory-less peer-to-peer file sharing systems. *Special Issue on P2P Economics, ACM SIG on Ecommerce*, 5(4):11–20, July 2005.
- [3] P. Antoniadis, C. Courcoubetis, and R. Weber. An asymptotically optimal scheme for p2p file sharing. 2nd Workshop on Economics of Peer-to-Peer Systems, Harvard University, 2004.
- [4] A. Asvanund, R. Krishnan, M. D. Smith, and R. Telang. Intelligent club management in peer-to-peer networks. 1st Workshop on Economics of Peer-to-Peer Systems, University of Berkeley, 2003.
- [5] C. Buragohain, D. Agrawal, and S. Suri. A game theoretic framework for incentives in p2p systems. In *Proceedings of the Third IEEE International Conference on Peer-to-Peer Computing (P2P 2003), Linkoping, Sweden*, 2003.

- [6] MMAPPS Consortium. Market management of peer-to-peer services - white paper. Available at <http://www.mmapps.org/>, 2004.
- [7] C. Courcoubetis and R. R. Weber. Asymptotics for provisioning problems with a large number of participants. AUEB Technical Memorandum, Available at <http://nes.aueb.gr/research/>, 2003.
- [8] C. Courcoubetis and R. R. Weber. Asymptotics for provisioning problems of peering wireless lans with a large number of participants. In *Proceedings of WiOpt'04 workshop, University of Cambridge, UK*, 2004.
- [9] P. Crampton, R. Gibbons, and P. Klemperer. Dissolving a partnership efficiently. *Econometrica*, 55:615–632, 1987.
- [10] M. Feldman, K. Lai, J. Chuang, and I. Stoica. Quantifying disincentives for collaboration in peer-to-peer networks. 1st Workshop on Economics of Peer-to-Peer Systems, University of Berkeley, 2003.
- [11] P. Golle, K. Leyton-Brown, I. Mironov, and M. Lillibridge. Incentives for sharing in peer-to-peer networks. In *Proceedings of WELCOM'01*, November 2001.
- [12] 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 *Proceedings 3rd IEEE International Conference on Peer-to-Peer Computing, Linkoping, Sweden*, 2003.
- [13] M. Hellwig. Public good provision with many participants. *Review of Economic Studies*, 70(3):589–614, July 2003.
- [14] R. Krishnan, M. Smith, Z. Tang, and R. Telang. The virtual commons: why free-riding can be tolerated in peer-to-peer networks. Workshop on Information Systems and Economics, 2003.
- [15] R. Krishnan, M. D. Smith, and R. Telang. The economics of peer-to-peer networks. *Journal of Information Technology Theory and Application*, 2003.
- [16] P. Norman. Efficient mechanisms for public goods with use exclusions. *The Review of Economic Studies*, 7:1163–1188, 2004.
- [17] P. Norman and T. Borgers. A note on budget balance under interim participation constraints: The case of independent types. <http://faculty.arts.ubc.ca/pnorman/Research/Research.htm>, June 2005.
- [18] V. Vishnumurthy, S. Chandrakumar, and E. G. Sirer. Karma: A secure economic framework for peer-to-peer resource sharing. 1st Workshop on Economics of Peer-to-Peer Systems, University of Berkeley, 2003.

Appendix

A 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 $p_i(\boldsymbol{\theta})$.⁶ Let us define

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

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

Thus $\theta_i V_i(\theta_i)$ and $P_i(\theta_i)$ 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

$$\left[\theta'_i V_i(\theta_i) - P_i(\theta_i) \right] + \left[\theta_i V_i(\theta'_i) - P_i(\theta'_i) \right] \leq \left[\theta_i V_i(\theta_i) - P_i(\theta_i) \right] + \left[\theta'_i V_i(\theta'_i) - P_i(\theta'_i) \right]$$

and so

$$(\theta'_i - \theta_i)(V_i(\theta'_i) - V_i(\theta_i)) \geq 0.$$

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.

⁶Since the agents are identical, apart from their labels, it is reasonable to suppose that the social planner can maximize welfare with a mechanism design 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.

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

$$\begin{aligned}
& \sum_{i=1}^n \int \pi_i(\theta_i, \boldsymbol{\theta}_{-i}) p_i(\boldsymbol{\theta}) dF^n(\boldsymbol{\theta}) \tag{39} \\
&= \sum_{i=1}^n \int P_i(\theta_i) dF(\theta_i) \\
&= \sum_{i=1}^n P_i(0) + \sum_{i=1}^n \int \left[\theta_i V_i(\theta_i) - \int_0^{\theta_i} V_i(\eta) d\eta \right] dF(\theta_i) \\
&= \sum_{i=1}^n P_i(0) + \sum_{i=1}^n \left\{ \int \theta_i V_i(\theta_i) dF(\theta_i) + (1 - F(\theta_i)) \int_0^{\theta_i} V_i(\eta) d\eta \Big|_0^\infty - \frac{1 - F(\theta_i)}{f(\theta_i)} V_i(\theta_i) dF(\theta_i) \right\} \\
&= \sum_{i=1}^n P_i(0) + \sum_{i=1}^n \int \pi_i(\theta_i, \boldsymbol{\theta}_{-i}) g(\theta_i) u(Q(\boldsymbol{\theta})) dF^n(\boldsymbol{\theta}). \tag{40}
\end{aligned}$$

Since the scheme is to be weakly feasible, we use (40) to deduce that our problem is one of maximizing (1) subject to

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

The maximization is with respect to a choice of the function $Q(\cdot)$ and the constants $P_1(0), \dots, P_n(0)$. 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 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 social welfare subject to use of a feasible, individually rational and incentive compatible scheme.

B 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 2, 3 and 4. The case $k > 1$ is what we need for Section 5. For simplicity of notation we drop the n from the cost $c(n, Q)$, and simply write $c(Q)$. We suppose that agents are of k types. There are n_j agents of type j , and their preference parameters are $\theta_{j1}, \dots, \theta_{jn_j}$.

Lemma 3 *Define \mathcal{P} as the problem*

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

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

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

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

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

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

Proof. Let us rewrite this as the problem of maximizing

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

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

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

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

and

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

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

$$E \left[\sum_{j=1}^k \left[\sum_{i=1}^{n_j} x_i(\boldsymbol{\theta}) \left(1 + \lambda_j \frac{g(\theta_{ij})}{\theta_{ij}} \right) - (1 + \lambda_j) c_j(Q(\boldsymbol{\theta})) \right] \right], \quad (47)$$

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 $0 \leq \pi_{ij}(\boldsymbol{\theta}) \leq 1$. ■

C 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

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

for constants

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

$$\zeta = (\alpha/\beta)^{\alpha/(\beta-\alpha)} - (\alpha/\beta)^{\beta/(\beta-\alpha)} > 0.$$

Hence $\xi(x) = \Omega(x^\gamma)$.

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 $Q \rightarrow \infty$. Then if $Q \leq \eta_1 x^{1/(\beta-\alpha)}$ or $Q \geq \eta_2 x^{1/(\beta-\alpha)}$, we have

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

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

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

D Proof of Theorem 1.

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 $\alpha \leq 1 \leq \beta$, $\alpha < \beta$ (so u and c are concave and convex respectively). Define the function ξ by

$$\xi(x) = \max_Q \{x u(Q) - c(Q)\}.$$

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

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

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

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

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 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

$$\begin{aligned} \Phi_n^* &= \min_\lambda \max_Q \left\{ n \int_0^1 (\theta + \lambda g(\theta))^+ dF(\theta) AQ^\alpha - (1 + \lambda)BQ^\beta \right\} \\ &= \min_\lambda (1 + \lambda) \left[n \frac{\int_0^1 (\theta + \lambda g(\theta))^+ dF(\theta)}{1 + \lambda} \right]^\gamma \end{aligned}$$

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 ,

$$\begin{aligned} \Phi_n &\leq \max_{\pi_1(\cdot), \dots, \pi_n(\cdot), Q(\cdot)} E \left[\sum_{i=1}^n \pi_i(\theta) (\theta_i + \lambda g(\theta_i)) u(Q(\theta)) - (1 + \lambda)c(Q(\theta)) \right] \\ &= \max_{Q(\cdot)} E \left[\sum_{i=1}^n (\theta_i + \lambda g(\theta_i))^+ u(Q(\theta)) - (1 + \lambda)c(Q(\theta)) \right] \end{aligned}$$

So we have, since $ET = n\bar{T}_1$,

$$\Phi_n^* = (1 + \lambda)\xi(ET)$$

$$\Phi_n \leq (1 + \lambda)E\xi(T)$$

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

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

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 $T^* \leq n$, we have $\xi''(T^*) \leq \gamma(\gamma - 1)n^{\gamma-2}$. Hence

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

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

$$\Phi_n/\Phi_n^* = 1 + O(1/n). \quad (48)$$

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

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

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

$$\Phi_n/\Phi_n^* \leq (1 + B/n)^{1/k} \leq 1 + B/kn.$$

Thus (48) 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

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

for some T^* depending on T , where we have $T^* < n$ and so $\xi'(T^*) \leq \xi'(n)$. 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 $\Phi_n/\Phi_n^* = 1 + O(1/\sqrt{n})$. \blacksquare

E Proof of Theorem 2.

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

$$m\xi \left(\frac{\sum_{i=1}^m \theta_{(i)}}{m} \right) = \max_Q \left\{ \sum_{i=1}^m \theta_{(i)} u(Q) - mc(Q) \right\} \quad (49)$$

$$\begin{aligned} &= m\xi \left(\sum_{i=1}^m \bar{\theta}_{(i)} / m \right) + \sum_{i=1}^m (\theta_{(i)} - \bar{\theta}_{(i)}) \xi' \left(\sum_{i=1}^m \theta_{(i)}^* / m \right) \\ &\leq \max_m \left\{ m\xi \left(\sum_{i=1}^m \bar{\theta}_{(i)} / m \right) \right\} + \sum_{i=1}^n |\theta_{(i)} - \bar{\theta}_{(i)}| \xi'(1) \end{aligned} \quad (50)$$

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

$$\Phi_n \leq E \left[\max_{m \in \{1, \dots, n\}} \left\{ m\xi \left(\frac{\sum_{i=1}^m \theta_{(i)}}{m} \right) \right\} \right] \quad (51)$$

$$\leq \max_m \left\{ m\xi \left(\frac{\sum_{i=1}^m \bar{\theta}_{(i)}}{m} \right) \right\} + \sum_{i=1}^n E |\theta_{(i)} - \bar{\theta}_{(i)}| \xi'(1) \quad (52)$$

where the final line follows from (50). Note first that

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

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 (52) is bounded by

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

Finally, the the second term on the right of (52) 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 $f(x)$, we have that for all θ, θ' ,

$$|\theta - \theta'| \leq |F(\theta) - F(\theta')|/a.$$

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 k th 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

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

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

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

■

F Weak Feasibility \implies Strong Feasibility

We prove an extension of the result of Crampton, et. al. in [9]. That result has also been recently reproved by Norman [17] in a rather simpler manner (in which he speaks of “ex ante” and “ex post” budget balance, rather than weak and strong feasibility).

Theorem 3 *Whether or not the public good is excludable, the existence of a weakly feasible incentive compatible payment scheme implies the existence of a strongly feasible incentive compatible payment scheme.*

The constructive proofs in [9] and [17] only work for a nonexcludable good. A strictly feasible scheme is constructed by a complicated transfer of payments amongst the agents. In doing this, an agent, say i , can end up being paid by others and so his payment $x_i(\boldsymbol{\theta})$ can be negative for some $\boldsymbol{\theta}$. This cannot be avoided. If one applies the construction in a situation in which there are to be exclusions, it sometimes finishes with a scheme in which, for one or more $\boldsymbol{\theta}$ and i , we have $x_i(\boldsymbol{\theta}) \neq 0$, but $\boldsymbol{\theta}$ is such that agent i should be excluded. This would require payments to be taken from an agent who is excluded, which is usually not a practical thing to do. We have not been able to find any way to modify the arguments in [9] or [17] to cover the case of an excludable good. However, we now show how this can be done using a completely different approach.

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

Suppose we have a weakly feasible scheme, with payment functions $y_1(\boldsymbol{\theta}), \dots, y_n(\boldsymbol{\theta})$. In a scheme that makes use of exclusions θ_i will be excluded for some small values of θ_i , say for $\theta_i \leq t_\ell$, and we will have $E[y_i(\boldsymbol{\theta})|\theta_i] = 0$ for all $\theta_i \leq t_\ell$.

We will construct a strongly feasible incentive compatible scheme, with payment functions $x_1(\boldsymbol{\theta}), \dots, x_n(\boldsymbol{\theta})$. This will be such that

- (i) $x_1(\boldsymbol{\theta}) + \dots + x_n(\boldsymbol{\theta}) = c(\boldsymbol{\theta})$,
- (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 $\theta_i \leq t_\ell$.

Condition (iii) ensures that the scheme is compatible with the exclusions we wish to make. For agents who are not excluded we do permit payments to be negative.

The modification algorithm. Suppose we start with a weakly feasible incentive compatible scheme which satisfies (ii) and

$$(i)' \quad E[x_1(\boldsymbol{\theta}) + \dots + x_n(\boldsymbol{\theta})] = E[c(\boldsymbol{\theta})].$$

This scheme might have been constructed as in [9]. However, we now show how to construct it by a new method, as the explanation will be a good preparation for understanding what follows later in this proof. Pick any θ where (i) is violated because $x_1(\theta) + \dots + x_n(\theta) > c(\theta)$, and any θ' where it is violated because of $x_1(\theta') + \dots + x_n(\theta') < c(\theta')$. Note that (i)' implies that there must always be such a pair if (i) does not hold for all θ . Let e_i be the n -vector that has a 1 in the i th component and all other components 0. Let

$$\theta'' = \theta + (\theta'_j - \theta_j)e_j.$$

So θ'' is the state that is the same as θ except that the j th component is θ'_j . Let us 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 θ and θ' . Note that alterations (b) and (c) ensure that the values of $E[x_1(\theta)|\theta_1]$, $E[x_2(\theta)|\theta_2]$ and $x_1(\theta'') + \dots + x_n(\theta'')$ do not change.

- (a) $x_1(\theta) \rightarrow x_1(\theta) - \epsilon$
- (b) $x_1(\theta'') \rightarrow x_1(\theta'') + \epsilon$
- (c) $x_2(\theta'') \rightarrow x_2(\theta'') - \epsilon$
- (d) $x_2(\theta') \rightarrow x_2(\theta') + \epsilon$

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 a strongly feasible scheme.

Now suppose we have a strongly feasible incentive compatible scheme which satisfies (i) and (ii). We will modify it to produce a strongly feasible incentive compatible scheme which also satisfies (iii).

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

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

That is, states θ^i , θ^j , and θ^{ij} , are constructed from θ by, respectively, increasing the i th, j th, or both the i th and j th components of θ to t_m (for which an agent is definitely not excluded). Let $\epsilon = x_i(\theta)$. We now make the following adjustments to the payments. These are chosen carefully to preserve (i) and (ii), and remove the violation to (iii) of $x_i(\theta) \neq 0$.

- (a) $x_i(\boldsymbol{\theta}) \rightarrow x_i(\boldsymbol{\theta}) - \epsilon$
- (b) $x_j(\boldsymbol{\theta}) \rightarrow x_j(\boldsymbol{\theta}) + \epsilon$
- (c) $x_i(\boldsymbol{\theta}^j) \rightarrow x_i(\boldsymbol{\theta}^j) + \epsilon$
- (d) $x_j(\boldsymbol{\theta}^j) \rightarrow x_j(\boldsymbol{\theta}^j) - \epsilon$
- (e) $x_j(\boldsymbol{\theta}^{ij}) \rightarrow x_j(\boldsymbol{\theta}^{ij}) + \epsilon$
- (f) $x_i(\boldsymbol{\theta}^{ij}) \rightarrow x_i(\boldsymbol{\theta}^{ij}) - \epsilon$
- (g) $x_i(\boldsymbol{\theta}^i) \rightarrow x_i(\boldsymbol{\theta}^i) + \epsilon$
- (h) $x_j(\boldsymbol{\theta}^i) \rightarrow x_j(\boldsymbol{\theta}^i) - \epsilon$

In detail:

- (a) removes the violation to (iii);
- (b) makes $x_1(\boldsymbol{\theta}) + \dots + x_n(\boldsymbol{\theta}) = c(\boldsymbol{\theta})$;
- (c) puts $E[x_i(\boldsymbol{\theta})|\theta_i]$ back to its required value;
- (d) makes $x_1(\boldsymbol{\theta}^j) + \dots + x_n(\boldsymbol{\theta}^j) = c(\boldsymbol{\theta}^j)$;
- (e) puts $E[x_j(\boldsymbol{\theta})|\theta_j]$ back to its required value;
- (f) makes $x_1(\boldsymbol{\theta}^{ij}) + \dots + x_n(\boldsymbol{\theta}^{ij}) = c(\boldsymbol{\theta}^{ij})$;
- (g) puts $E[x_i(\boldsymbol{\theta})|\theta_i]$ back to its required value;
- (h) makes $x_1(\boldsymbol{\theta}^i) + \dots + x_n(\boldsymbol{\theta}^i) = c(\boldsymbol{\theta}^i)$ 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 $x_i(\boldsymbol{\theta}^j) + \epsilon \neq 0$, or if $x_j(\boldsymbol{\theta}^i) = 0$ and $x_j(\boldsymbol{\theta}^i) - \epsilon \neq 0$. 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 $\theta_i \leq t_\ell$ and $x_i(\boldsymbol{\theta}) \neq 0$. 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 $x_i(\boldsymbol{\theta}') \neq 0$ and that, again since $\boldsymbol{\theta}' \notin \Theta$, there is a k such that $\theta'_k \geq t_\ell$ (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 $x_i(\boldsymbol{\theta}) \neq 0$.

- (a) $x_i(\boldsymbol{\theta}) \rightarrow x_i(\boldsymbol{\theta}) - \epsilon$
- (b) $x_i(\boldsymbol{\theta}') \rightarrow x_i(\boldsymbol{\theta}') + \epsilon$
- (c) $x_k(\boldsymbol{\theta}') \rightarrow x_k(\boldsymbol{\theta}') - \epsilon$
- (d) $x_k(\boldsymbol{\theta}^\dagger) \rightarrow x_k(\boldsymbol{\theta}^\dagger) + \epsilon$
- (e) $x_j(\boldsymbol{\theta}^\dagger) \rightarrow x_j(\boldsymbol{\theta}^\dagger) - \epsilon$
- (f) $x_j(\boldsymbol{\theta}) \rightarrow x_j(\boldsymbol{\theta}) + \epsilon$

This does not introduce any new violations to (iii) in the states θ and θ^\dagger because there is already a violation $x_i(\theta') \neq 0$, 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. ■

■

G Strongly Feasible Equal Contribution Mechanisms

Here are details of the calculations referred to in Section 3.3 where we consider fixed contribution schemes that satisfy strong feasibility conditions.

G.1 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 will be such that

$$\theta u(Q) - c(Q)E[1/X] = 0.$$

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 ,

$$n(1 - \theta) \frac{1}{2} (1 + \theta) \frac{2}{3} \sqrt{Q} - Q. \quad (53)$$

subject to

$$\theta \frac{2}{3} \sqrt{Q} - QE[1/X] \geq 0. \quad (54)$$

Now

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

Now use (53) to solve for Q , making the approximation $E[1/X] \approx 1/n(1 - \theta)$. Substituting this in (53) 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 $\Phi_n^* = 3n^2/128$.

G.2 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 in our problem we want to maximize with respect to θ and ϕ ,

$$E \left[X \frac{1}{2} (1 + \theta) \frac{2}{3} \sqrt{X\phi} - X\phi \right] \quad (55)$$

subject to

$$\theta \frac{2}{3} E \sqrt{X\phi} - \phi \geq 0. \quad (56)$$

Now

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

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

$$\frac{1}{3} \sqrt{\phi} (n(1 - \theta))^{3/2} (1 + \theta) \left(1 + \frac{3}{8} \frac{\theta}{n(1 - \theta)} \right) - \phi n(1 - \theta) \quad (57)$$

subject to

$$\frac{2}{3} \theta (n(1 - \theta))^{1/2} \left(1 - \frac{1}{8} \frac{\theta}{n(1 - \theta)} \right) \sqrt{\phi} - \phi \geq 0. \quad (58)$$

Now use (58) to find ϕ as a function of θ and substitute into (57). 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 strongly feasible policy does a bit better because although it provides the same Q on average it makes provides more Q when there are more people participating. The explanation is that when $u(Q) = \frac{2}{3} \sqrt{Q}$ the term in the social welfare of $X \frac{1}{2} (1 + \theta) u(\phi X)$ is a convex function of X , so

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