Incentivizing Participation in Resource-sharing Networks

Richard Weber

Statistical Laboratory

10 October, 2007

Sharing Files

A file sharing system Peers contribute files to a shared library of files which they can access over the Internet.

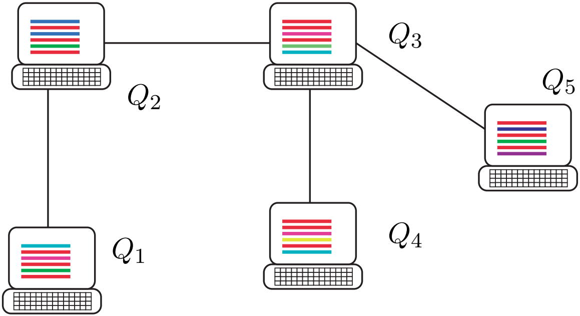

Peer i shares Q_i files.
The benefit to peer j is $\theta_j u(Q_1 + \cdots + Q_5)$.

Sharing WLANS

A sharing of wireless LANS system Peers share their wireless Local Area Networks so that they can enjoy Internet access via one another's networks whenever they wander away from their home locations.

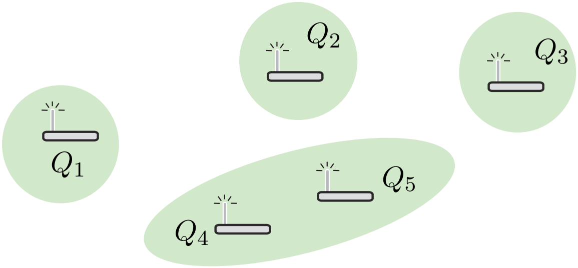

Peer i makes his WLAN available for a fraction Q_i of the time.
The benefit to peer j is $\theta_j u(Q_1 + \cdots + Q_5)$.

Grid Computing

A sharing of computing resources Participants share computing resources, statistically multiplexed through time. Peer i is present a fraction α_i of the time.

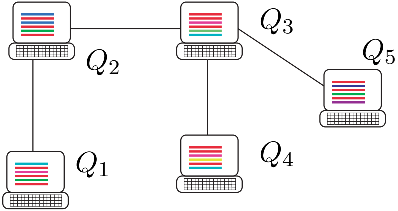

If the set $S \subset \{1, 2, 3, 4, 5\}$ of peers need resources at the same time, and peer $i \in S$, then he receives a share $y_i(\boldsymbol{\alpha}, \boldsymbol{\theta}, S)$ and obtains benefit $\theta_i u(\alpha_i, y_i)$, where

$$\sum_{j \in S} y_j(\boldsymbol{\alpha}, \boldsymbol{\theta}, S) \leq Q_1 + \cdots + Q_5 \,.$$

Free Rider Problem

All our example systems can suffer from the **free rider problem**. Each peer would like to contribute only a very few files, very little wireless LAN resources, or only a little computing resource, and yet fully to enjoy the system that others have provided. This is a typical problem in the provision of shared resources.

Free Rider Problem

All our example systems can suffer from the **free rider problem**. Each peer would like to contribute only a very few files, very little wireless LAN resources, or only a little computing resource, and yet fully to enjoy the system that others have provided. This is a typical problem in the provision of shared resources.

- Free riding is rational for peers, but leads to inefficiency.
- We should like to provide incentives for peers to act 'honestly', and enable systems of appropriate size to be built.
- Peers should benefit from the system in proportion to their contributions.

Example: A Bridge

A bridge may or may not be built. There are 2 potential users.

Mathematical Bridge, Queens' College, Cambridge

Example: A Bridge

A bridge may or may not be built. There are 2 potential users.

Mathematical Bridge, Queens' College, Cambridge

If it is built (at costs $\$1$) then user i benefits by $\$\theta_i$. Knowing θ_1 and θ_2, we should build the bridge if $\theta_1 + \theta_2 > 1$.

Example: A Bridge

A bridge may or may not be built. There are 2 potential users.

Mathematical Bridge, Queens' College, Cambridge

If it is built (at costs $\$1$) then user i benefits by $\$\theta_i$. Knowing θ_1 and θ_2, we should build the bridge if $\theta_1 + \theta_2 > 1$.

If we build the bridge we must charge for the cost. Suppose we decide to charge user i a fee of $\theta_i/(\theta_1 + \theta_2)$. The problem is that user i will have an incentive to under-report his true value of θ_i.

Example: A Bridge

A bridge may or may not be built. There are 2 potential users.

Mathematical Bridge, Queens' College, Cambridge

If it is built (at costs $\$1$) then user i benefits by $\$\theta_i$. Knowing θ_1 and θ_2, we should build the bridge if $\theta_1 + \theta_2 > 1$.

If we build the bridge we must charge for the cost. Suppose we decide to charge user i a fee of $\theta_i/(\theta_1 + \theta_2)$. The problem is that user i will have an incentive to under-report his true value of θ_i.

What is the best fee mechanism, $p_1(\theta_1, \theta_2)$ and $p_2(\theta_1, \theta_2)$?
Fees should incentivize users to truthfully reveal θ_1, θ_2, with

$$p_1(\theta_1, \theta_2) + p_2(\theta_1, \theta_2) = 1 \text{ or } 0\,, \text{ as bridge is built or not built}\,.$$

Economic Model

- n agents participate in providing and sharing resources.
- The total resource pool they provide is of quantity Q.
- The total cost of providing it is $c(Q)$.
- Agent i has a usage frequency α_i (exogenous and public knowledge), and a preference parameter θ_i (private knowledge).
- Agent i declares θ_i, and is required to pay $p_i(\boldsymbol{\alpha}, \boldsymbol{\theta})$.

  $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_n), \quad \boldsymbol{\theta} = (\theta_1, \dots, \theta_n).$
- This permits a system of size $Q(\boldsymbol{\alpha}, \boldsymbol{\theta})$ to be built.
- When the set of agents who are seeking to share is $S \subseteq \{1, \dots, n\}$, agent i receives a share $y_i(\boldsymbol{\alpha}, \boldsymbol{\theta}, S)$.
- Agent i obtains benefit $\theta_i u(\alpha_i, y_i)$.

The Social Planner's Problem

A social planner wishes to maximize social welfare:

$$SW = E_S \left[\sum_{i=1}^{n} \theta_i u(\alpha_i, y_i(S)) - c(Q) \right]$$

- If he knew the $\theta_1, \ldots, \theta_n$ he would choose Q optimally and share it optimally (the so-called **first-best** optimum).
- However, in practice he might have to elicit θ_i indirectly by asking agent i to volunteer paying some p_i.
- Since cost $c(Q)$ must be met by the agents' payments (possibly 'in kind') we must have $c(Q) \leq \sum_i p_i(\boldsymbol{\alpha}, \boldsymbol{\theta})$.

Peer-to-peer: the case of Nonrivalrous, Nonexcludable Resources

$$u(\alpha_i, y_i(\boldsymbol{\alpha}, \boldsymbol{\theta}, S)) = u(Q(\boldsymbol{\theta}))$$

The Mechanism Design Problem

The social planner has a **mechanism design problem**.

It is a sort of game: the planner asks agent i to declare θ_i. The agents know that as a function of what they say,

The Mechanism Design Problem

The social planner has a **mechanism design problem**.

It is a sort of game: the planner asks agent i to declare θ_i. The agents know that as a function of what they say,

- the system size (quality/quantity) will be $Q(\theta_1, \ldots, \theta_n)$;

The Mechanism Design Problem

The social planner has a **mechanism design problem**.

It is a sort of game: the planner asks agent i to declare θ_i. The agents know that as a function of what they say,

- the system size (quality/quantity) will be $Q(\theta_1, \ldots, \theta_n)$;
- agent i will required to make payment $p_i(\theta_1, \ldots, \theta_n)$;

The Mechanism Design Problem

The social planner has a **mechanism design problem**.

It is a sort of game: the planner asks agent i to declare θ_i. The agents know that as a function of what they say,

- the system size (quality/quantity) will be $Q(\theta_1, \ldots, \theta_n)$;
- agent i will required to make payment $p_i(\theta_1, \ldots, \theta_n)$;
- agent i will be excluded with probability $\pi_i(\theta_1, \ldots, \theta_n)$.

The Mechanism Design Problem

The social planner has a **mechanism design problem**.

It is a sort of game: the planner asks agent i to declare θ_i. The agents know that as a function of what they say,

- the system size (quality/quantity) will be $Q(\theta_1, \ldots, \theta_n)$;
- agent i will required to make payment $p_i(\theta_1, \ldots, \theta_n)$;
- agent i will be excluded with probability $\pi_i(\theta_1, \ldots, \theta_n)$.

Knowing that the planner will adopt this mechanism, the agents behave self-interestedly.

The planner wishes to design the mechanism so that the expected social welfare is maximized (achieving the so-called **second best** optimum).

Example

Suppose

$$u(Q) = \tfrac{2}{3}Q^{1/2}\,, \quad c(Q) = Q\,,$$

$\theta_1, \ldots, \theta_n$ are known a priori to be i.i.d. samples from $U[0, 1]$.

Example

Suppose

$$u(Q) = \tfrac{2}{3} Q^{1/2}\,, \quad c(Q) = Q\,,$$

$\theta_1, \dots, \theta_n$ are known a priori to be i.i.d. samples from $U[0,1]$.

The first best policy has expected social welfare

$$SW = E\left[\max_Q \left\{ \sum_{i=1}^{n} \theta_i u(Q) - c(Q) \right\}\right] = \frac{1}{36} n^2 + \frac{1}{108} n$$

Example

Suppose

$$u(Q) = \tfrac{2}{3} Q^{1/2}\,, \quad c(Q) = Q\,,$$

$\theta_1, \ldots, \theta_n$ are known a priori to be i.i.d. samples from $U[0,1]$.

The first best policy has expected social welfare

$$SW = E\left[\max_Q \left\{ \sum_{i=1}^{n} \theta_i u(Q) - c(Q) \right\}\right] = \frac{1}{36} n^2 + \frac{1}{108} n$$

The second-best policy (difficult to compute!) achieves about 84%

$$SW = \frac{3}{128} n^2 + \frac{7}{148} n + O(1)$$

What if exclusions are not allowed?

- First-best social welfare

$$\frac{1}{36}n^2 + \frac{1}{108}n$$

- Second-best social welfare

$$\frac{3}{128}n^2 + \frac{7}{348}n + O(1)\,.$$

What if exclusions are not allowed?

- First-best social welfare

$$\frac{1}{36}n^2 + \frac{1}{108}n$$

- Second-best social welfare

$$\frac{3}{128}n^2 + \frac{7}{348}n + O(1)\,.$$

- Second-best social welfare (without exclusions)

$$\frac{n(\sqrt{1+3n}-1)}{27}\,.$$

# Simple 'Equal Contribution' Schemes

Suppose we try a simple mechanism in which all participants pay the same fee, or make the equal contributions, perhaps in kind.

- Post a size Q and share the cost equally. Agent i participates if he has

$$\theta u(Q) - c(Q)E(1/X) \geq 0\,,$$

where X has the a priori distribution of the number who will participate. Choosing Q optimally,

$$SW = \frac{3}{128}n^2 - \frac{1}{384}n + O(1)$$

- Post a fee ϕ and then build the largest facility that can be covered by the fees. Agent i participates if he has

$$\theta u(X\phi) - \phi \geq 0\,,$$

where X has the a priori distribution of the number who will participate. Choosing ϕ optimally gives

$$SW = \frac{3}{128}n^2 + \frac{7}{1536}n + O(1)$$

- Post a fee ϕ and size Q. Agent i participates if he has

$$\theta_i u(Q) - \phi \geq 0 \,.$$

Choosing ϕ, Q optimally, and also so expected cost is covered,

$$SW = \frac{3}{128} n^2$$

- Post a fee ϕ and size Q. Agent i participates if he has

$$\theta_i u(Q) - \phi \geq 0\,.$$

Choosing ϕ, Q optimally, and also so expected cost is covered,

$$SW = \frac{3}{128}n^2$$

- This is almost as good as achieved by the second-best mechanism:

$$SW = \frac{3}{128}n^2 + \frac{7}{148}n + O(1)$$

- Post a fee ϕ and size Q. Agent i participates if he has

$$\theta_i u(Q) - \phi \geq 0\,.$$

  Choosing ϕ, Q optimally, and also so expected cost is covered,

$$SW = \frac{3}{128}n^2$$

- This is almost as good as achieved by the second-best mechanism:

$$SW = \frac{3}{128}n^2 + \frac{7}{148}n + O(1)$$

We aim to understand why this is, and generalize the type of simple mechanism that just posts a fee ϕ and size Q.

The Social Welfare Maximization Problem

Maximize expected social welfare:

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

The expectation here is over $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_n)$, where these are assumed to be i.i.d. samples from some distribution with known distribution function F and density f.

The maximization is subject to constraints as follow.

Budget Balance

The maximization is subject to an **ex-ante budget balance** constraint, that expected payments balance with the expected cost

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

Budget Balance

The maximization is subject to an **ex-ante budget balance** constraint, that expected payments balance with the expected cost

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

Surprisingly, there is no decrease in social welfare if we require the stronger **ex-post budget balance** constraint, that

$$\sum_{i=1}^{n} \pi_i(\boldsymbol{\theta}) p_i(\boldsymbol{\theta}) = c(n, Q(\boldsymbol{\theta}))\,.$$

Individual Rationality

Also **ex-ante individual rationality** constraints, that each agent can expect positive net benefit:

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

Incentive Compatibility

Also **ex-ante incentive compatibility** constraints, that each agent i does best by declaring his true θ_i rather than declaring some other θ_i':

$$E_{\boldsymbol{\theta}_{-i}}\Big[\pi_i(\theta_i,\boldsymbol{\theta}_{-i})\left\{\theta_i u(Q(\theta_i,\boldsymbol{\theta}_{-i})-p_i(\theta_i,\boldsymbol{\theta}_{-i})\right\}\Big]$$
$$\geq E_{\boldsymbol{\theta}_{-i}}\Big[\pi_i(\theta_i',\boldsymbol{\theta}_{-i})\left\{\theta_i u(Q(\theta_i',\boldsymbol{\theta}_{-i})-p_i(\theta_i',\boldsymbol{\theta}_{-i})\right\}\Big],$$

for all i and θ_i'.

Incentive Compatibility

Also **ex-ante incentive compatibility** constraints, that each agent i does best by declaring his true θ_i rather than declaring some other θ_i':

$$E_{\boldsymbol{\theta}_{-i}}\Big[\pi_i(\theta_i,\boldsymbol{\theta}_{-i})\,\{\theta_i u(Q(\theta_i,\boldsymbol{\theta}_{-i})-p_i(\theta_i,\boldsymbol{\theta}_{-i})\}\Big]$$
$$\geq E_{\boldsymbol{\theta}_{-i}}\Big[\pi_i(\theta_i',\boldsymbol{\theta}_{-i})\,\{\theta_i u(Q(\theta_i',\boldsymbol{\theta}_{-i})-p_i(\theta_i',\boldsymbol{\theta}_{-i})\}\Big],$$

for all i and θ_i'.

This constraint does not decrease the maximum SW that we can obtain. The **revelation principle** says that any Nash equilibrium that can obtained by some mechanism can also be obtained by an incentive compatible mechanism.

The Equivalent Problem

Let us define

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

E.g., $g(\theta_i) = 2\theta_1 - 1$ when $\theta_i \sim U[0, 1]$.

The Equivalent Problem

Let us define

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

E.g., $g(\theta_i) = 2\theta_1 - 1$ when $\theta_i \sim U[0,1]$.

It can be shown that the mechanism design problem reduces to

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

subject to a single constraint:

$$E\Big[\textstyle\sum_i \pi_i(\boldsymbol{\theta}) g(\theta_i) u(Q(\boldsymbol{\theta})) - c(n, Q(\boldsymbol{\theta}))\Big] \;\geq 0\,.$$

A key idea is that this can be solved by Lagrangian methods (but this must be proved).

The Lagrangian

That is, for some $\lambda > 0$ the problem can be solved by maximizing a Lagrangian of

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

The Lagrangian

That is, for some $\lambda > 0$ the problem can be solved by maximizing a Lagrangian of

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

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

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

where

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

Assuming $\theta_i + \lambda g(\theta_i)$ is increasing in θ_i we will exclude agents with small θ_i, such that $\theta_i + \lambda g(\theta_i) < 0$.

Optimal mechanism design for the bridge

Suppose $n = 2$, $\theta_1, \theta_2 \sim U[0,1]$. The bridge costs 1 to build. First-best optimum has $SW = 1/6 = 0.166\dot{6}$.

A possible second-best mechanism builds the bridge only if $\theta_1 + \theta_2 \geq 1.25$ and obtains $SW = 9/64 = 0.140625$.

Agent i pays fee of

$$\begin{aligned} p_i(\boldsymbol{\theta}) &= (1/2)\theta_i^2 - 1/32\,, \quad && \theta_i > 1/4\,, \\ &= 0\,, && \text{otherwise.} \end{aligned}$$

However, this mechanism is not ex-post budget balanced

$$p_1(\boldsymbol{\theta}) + p_2(\boldsymbol{\theta}) \neq 1_{\{\theta_1 + \theta_2 \geq 1.25\}}$$

An ex-post budget balance design

The optimal design is not unique. There is an equally good second-best mechanism in which agent 1 pays fee of

$$p_1(\boldsymbol{\theta}) = \tfrac{1}{2}1_{\{\theta_1+\theta_2\geq 1.25\}} + \left(\tfrac{3}{32} - \tfrac{1}{2}\theta_1(1-\theta_1)\right)1_{\{\theta_1\geq .25\}} - \left(\tfrac{3}{32} - \tfrac{1}{2}\theta_2(1-\theta_2)\right)1_{\{\theta_2\geq .25\}}$$

Now there is ex-post budget balance, as

$$p_1(\boldsymbol{\theta}) + p_2(\boldsymbol{\theta}) = 1_{\{\theta_1+\theta_2\geq 1.25\}}\,,$$

but this is still unsatisfactory.

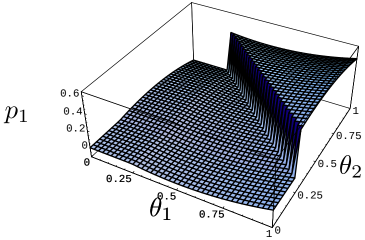

A better optimal mechanism design for the bridge

An even better second best optimal mechanism design has agent 1 paying a fee

$$p_1(\boldsymbol{\theta}) = \left(\frac{1}{3}(\theta_1 - \theta_2) + \frac{1}{2}\right)1_{\{\theta_1+\theta_2 \geq 1.25\}} .$$

This design has the advantages over the previous scheme.

1. $p_1(\boldsymbol{\theta}) \geq 0$.
2. $p_1(\boldsymbol{\theta}) = 0$ when the bridge is not built;
3. $\theta_1 - p_1(\boldsymbol{\theta}) > 0$ when it is built (ex-post individual rationality).

An fixed-fee design

Note that we might simply announce that we will charge a usage fee of $\$0.50$ and build the bridge if and only if both users are willing to pay. This produces SW of $1/8 = 0.125$. This is less than the $9/64 = 0.140625$ that is achieved with an optimal design.

The Asymptotically Optimal Mechanism

The full solution of our problem is, in general, very complex.

Our new idea (Theorem 1) is that, when n is large, a nearly optimal solution is achieved with a simple mechanism design.

Intuitively, if n is large then we would expect about $nf(t)\,d\theta$ agents to have values of θ_i in the interval $[t, t+d\theta]$. So if we build a system of size Q and charge a fee of ϕ, then agents with $\theta \geq \theta^*$ will decide to participate, where $\theta^* u(Q) = \phi$. The social welfare will be about

$$n\left(\int_{\theta^*}^{1} \theta u(Q) f(\theta)\,d\theta\right) - c(Q)\,,$$

and we will collect fees that about cover our cost if

$$n[1 - F(\theta^*)]\phi = n[1 - F(\theta^*)]\theta^* u(Q) \geq c(Q)\,.$$

The Near Optimality of the Asymptotic Mechanism

Theorem 1

Let $\mathcal{P}$ be the problem of maximizing second-best social welfare, and suppose the maximal value is Φ_n. Let Q^* and θ^* be the optimizing decision variables in the problem $\mathcal{P}^*$, defined as

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

subject to

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

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 $\Phi_n / \Phi_n^* = 1 + O(n^{-1})$ (under Assumption 1), or $\Phi_n / \Phi_n^* = 1 + O(1/\sqrt{n})$ (under Assumption 2).

Assumptions

Assumptions are that $u(\cdot)$ and $c(\cdot)$ are concave and convex respectively, and grow as powers of Q.

Assumption 1.

$$u(Q) = AQ^{\alpha}\,,$$
$$c(n, Q) = Bn^{\delta}Q^{\beta}\,,$$

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

Assumption 2. There are positive constants A_1, A_2, B_1, B_2, and a function h such that for all Q and n,

$$A_1 Q^{\alpha} \leq u(Q) \leq A_2 Q^{\alpha}\,,$$
$$B_1 h(n) Q^{\beta} \leq c(n, Q) \leq B_2 h(n) Q^{\beta}\,.$$

Numerical Example

$$u(Q) = \frac{2}{3}Q^{1/2}\,, \quad c(Q) = Q\,, \quad \theta_1, \ldots, \theta_n \text{ are i.i.d. } U[0,1].$$

Numerical Example

$$u(Q) = \tfrac{2}{3}Q^{1/2}\,, \quad c(Q) = Q\,, \quad \theta_1, \ldots, \theta_n \text{ are i.i.d. } U[0,1].$$

- Second-best social welfare is

$$\Phi_n = \frac{3}{128}n^2 + \frac{7}{348}n - \frac{1}{162} + o(1)\,.$$

Numerical Example

$$u(Q) = \tfrac{2}{3}Q^{1/2}\,, \quad c(Q) = Q\,, \quad \theta_1, \ldots, \theta_n \text{ are i.i.d. } U[0,1].$$

- Second-best social welfare is

$$\Phi_n = \frac{3}{128}n^2 + \frac{7}{348}n - \frac{1}{162} + o(1)\,.$$

- Approximate second-best social welfare is $\Phi_n^* = \dfrac{3}{128}n^2$, with $Q^* = n^2/64$ and $\phi^* = n/48$.

Numerical Example

$$u(Q) = \tfrac{2}{3}Q^{1/2}\,, \quad c(Q) = Q\,, \quad \theta_1, \ldots, \theta_n \text{ are i.i.d. } U[0,1].$$

- Second-best social welfare is

$$\Phi_n = \frac{3}{128}n^2 + \frac{7}{348}n - \frac{1}{162} + o(1)\,.$$

- Approximate second-best social welfare is $\Phi_n^* = \dfrac{3}{128}n^2$, with $Q^* = n^2/64$ and $\phi^* = n/48$.
- $\Phi_n/\Phi_n^* = 1 + 0.858n^{-1} + o(1/n)$.

We have here a theorem, which for $\theta_1, \ldots, \theta_n$ i.i.d., compares the solution of

$$\underset{x(\cdot)}{\text{maximize}}\, E_{\boldsymbol{\theta}}\left[\sum_{i=1}^{n} f(\theta_i, x(\boldsymbol{\theta}))\right], \quad \text{s.t. } E_{\boldsymbol{\theta}}\left[\sum_{i=1}^{n} g(\theta_i, x(\boldsymbol{\theta}))\right] = b$$

with the solution of

$$\underset{x}{\text{maximize}} \sum_{i=1}^{n} E_{\theta_i}\Big[f(\theta_i, x)\Big], \quad \text{s.t. } \sum_{i=1}^{n} E_{\theta_i}\Big[g(\theta_i, x)\Big] = b$$

The key idea that permits the solutions to be compared is that the constraint can be accommodated in a Lagrangian.

The Applications

- **File sharing**. Q is content availability. The fee $p_i = Q_i$ is paid in kind as the number of files that agent i is required to share, and $Q = \sum_i Q_i$.

The Applications

- **File sharing**. Q is content availability. The fee $p_i = Q_i$ is paid in kind as the number of files that agent i is required to share, and $Q = \sum_i Q_i$.
- **WLAN sharing**. Q is coverage — availability of WLAN service at a random geographic point. The fee p_i is paid in kind as the availability, or rate of roaming requests served to others by agent i's WLAN in i's home geographic location. For this model we can take the cost in a form $c(m, Q) = \gamma m Q = \sum_i \gamma m Q_i$ so that the cost depends on the number of peers, m, who participate.

Ex-ante and ex-post constraints

- We have solved a very complicated sort of constrained optimization problem.
- We have solved a problem with an ex-ante budget balance constraint, but remarked that that budget balance can be achieved ex-post.
- The asymptotically approximate solution is simple, but somewhat unsatisfactory, because excluded agents are required to pay. This might be hard to enforce in practice.
- Individual rationality and incentive compatibility we imposed only 'on average'. This may be enough if we see things as a 'repeated game'.
  But can we guarantee more?

Ex-ante constraints

Under **ex-ante constraints** we are to maximize expected SW subject to a 'budget balance' constraint:

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

'individual rationality' constraints, :

$$E_{\boldsymbol{\theta}_{-i}}\left[\pi_i(\theta_i, \boldsymbol{\theta}_{-i})\theta_i u(Q(\theta_i, \boldsymbol{\theta}_{-i}) - p_i(\theta_i, \boldsymbol{\theta}_{-i})\right] \geq 0\,, \tag{1b}$$

and 'incentive compatibility' constraints:

$$E_{\boldsymbol{\theta}_{-i}}\left[\pi_i(\theta_i, \boldsymbol{\theta}_{-i})\theta_i u(Q(\theta_i, \boldsymbol{\theta}_{-i}) - p_i(\theta_i, \boldsymbol{\theta}_{-i})\right]$$
$$\geq E\left[\pi_i(\theta_i', \boldsymbol{\theta}_{-i})\theta_i u(Q(\theta_i', \boldsymbol{\theta}_{-i}) - p_i(\theta_i', \boldsymbol{\theta}_{-i})\right]. \tag{1c}$$

Ex-post constraints

Under **ex-post constraints** we are to maximize expected SW subject to a 'budget balance' constraint:

$$\sum_{i=1}^{n} p_i(\boldsymbol{\theta}) - c(Q(\boldsymbol{\theta})) = 0\,, \tag{2a}$$

'individual rationality' constraints, :

$$\pi_i(\theta_i, \boldsymbol{\theta}_{-i})\theta_i u(Q(\theta_i, \boldsymbol{\theta}_{-i}) - p_i(\theta_i, \boldsymbol{\theta}_{-i}) \geq 0\,, \tag{2b}$$

and 'incentive compatibility' constraints:

$$\begin{aligned} &\pi_i(\theta_i, \boldsymbol{\theta}_{-i})\theta_i u(Q(\theta_i, \boldsymbol{\theta}_{-i}) - p_i(\theta_i, \boldsymbol{\theta}_{-i}) \\ &\quad \geq \;\; \pi_i(\theta_i', \boldsymbol{\theta}_{-i})\theta_i u(Q(\theta_i', \boldsymbol{\theta}_{-i}) - p_i(\theta_i', \boldsymbol{\theta}_{-i})\,. \end{aligned} \tag{2c}$$

Strengthening the Constraints

Theorem 2 *Given a mechanism design* $\pi_1(\cdot), \ldots, \pi_n(\cdot)$, $Q(\cdot)$, $p_1(\cdot), \ldots, p_n(\cdot)$ *satisfying the ex-ante constraints* (1a), (1b) *and* (1c)*, there is a mechanism design which achieves the same value of expected social welfare, but which additionally satisfies the ex-post budget balance constraint* (2a)*.*

It is also possible to arrange that (1b) *can be partially replaced by* (2b)*, (ex-post individual rationality), so that* $p_i(\theta_i, \boldsymbol{\theta}_{-i}) = 0$ *whenever* $\pi(\theta_i, \boldsymbol{\theta}_{-i}) = 0$*.*
This is nice since it says we need not take payments from (or pay) agents who are completely excluded from the system.

Strengthening the Constraints

Theorem 3 *Given a mechanism design* $\pi_1(\cdot), \ldots, \pi_n(\cdot)$, $Q(\cdot)$, $p_1(\cdot), \ldots, p_n(\cdot)$ *satisfying the ex-ante constraints* (1a), (1b) *and* (1c), *there is a mechanism design which achieves the same value of expected social welfare, but which additionally satisfies the ex-post individual rationality and incentive compatibility constraints* (2b) *and* (2c).

Proof of Theorem 2. Suppose we start with some scheme which satisfies ex-ante constraints (1a), (1b) and (1c).

1. We construct a new scheme that is closer to satisfying the ex-post constraint (2a) by making an adjustment to the original scheme. This involves increasing and decreasing the payments that some pair of agents make in various circumstances.
2. The adjustment is similar to a step in the transportation algorithm. E.g., payments change as

$$\begin{array}{lcll} p_1(\boldsymbol{\theta}) & \rightarrow & p_1(\boldsymbol{\theta}) & -\epsilon \\ p_1(\boldsymbol{\theta}'') & \rightarrow & p_1(\boldsymbol{\theta}'') & +\epsilon \\ p_2(\boldsymbol{\theta}'') & \rightarrow & p_2(\boldsymbol{\theta}'') & -\epsilon \\ p_2(\boldsymbol{\theta}') & \rightarrow & p_2(\boldsymbol{\theta}') & +\epsilon \end{array}$$

3. After finitely many such steps the scheme that satisfies (2a).

Summary

- We have explored effects of externalities that arise in peer-to-peer systems by the study of a simple model or a nonrivalous nonexcludable good.
- Found that an 'equal contribution' scheme performs well as a control for large n.
  - The scheme is easy to implement and enforce.
  - Simple to compute optimal fee.
  - Can model payments that are 'contributions in kind'.
  - It is the first term of a more exact solution.
- Results can be extended to multiple constraints (subsets of peers must pay for certain portions of the total cost. E.g., perhaps peers in the same city must contribute 'in kind' to the total amount of shared WLAN that available in that city.)
- Robust (when users respond iteratively to the size of system they see produced then the system converges to an equilibrium where social welfare is maximized).

Grid Computing: the case of Rivalrous, Excludable Resources

$$u(\alpha_i, y_i(\boldsymbol{\alpha}, \boldsymbol{\theta}, S)) = \alpha_i u(y_i(\boldsymbol{\alpha}, \boldsymbol{\theta}, S))$$

The game

1. The mechanism designer posts rules.
2. Agent i declares θ_i truthfully (assuming incentive compatibility is designed into the rules). α_i is exogenous and public knowledge.
3. A system is built of size $Q(\boldsymbol{\alpha}, \boldsymbol{\theta})$, and is shared so that

$$\sum_{i \in S} y_i(\boldsymbol{\alpha}, \boldsymbol{\theta}, S) \leq Q(\boldsymbol{\alpha}, \boldsymbol{\theta}) \, .$$

4. The system is self-financing. Agent i pays $p_i(\boldsymbol{\alpha}, \boldsymbol{\theta})$ and

$$c(Q(\boldsymbol{\alpha}, \boldsymbol{\theta})) \leq \sum_{i=1}^{n} p_i(\boldsymbol{\alpha}, \boldsymbol{\theta}) \, .$$

The charge $p_i(\boldsymbol{\alpha}, \boldsymbol{\theta})$ can either occur up front, or as a usage charge. They are equivalent so far as incentivizing agents.

The mechanism design problem

The aim is to maximize the social welfare

$$E\left[\sum_{i=1}^{n}\left[\theta_i u(y_i(\boldsymbol{\alpha},\boldsymbol{\theta},S)) - p_i(\boldsymbol{\alpha},\boldsymbol{\theta})\right]\right]$$

subject to the ex-post constraints

$$c(Q(\boldsymbol{\alpha},\boldsymbol{\theta})) \leq \sum_{i=1}^{n} p_i(\boldsymbol{\alpha},\boldsymbol{\theta})$$

$$\sum_{i\in S} y_i(\boldsymbol{\alpha},\boldsymbol{\theta},S) \leq Q(\boldsymbol{\alpha},\boldsymbol{\theta})$$

and also constraints of ex ante individual rationality and incentive compatibility.

A special case

- Suppose $\theta_1 = \cdots = \theta_n = 1$, $\alpha_1, \ldots, \alpha_n \sim U[0,1]$, $C(Q) = Q$,

$$u(\alpha_i, y_i) = \alpha_i u(y_i)\,,$$

and the system is large so that

$$\sum_{i \in S} y_i(\boldsymbol{\alpha}, \boldsymbol{\theta}, S) \leq Q(\boldsymbol{\alpha}, \boldsymbol{\theta})$$

is nearly the same as

$$\sum_{i=1}^{n} \alpha_i y_i(\alpha_i) \leq \sum_{i=1}^{n} p(\alpha_i)\,.$$

- We offer agents a choice of tariffs $\{p(a), y(a)\}$, $0 \leq a \leq 1$, designed to be incentive compatible, in the sense that

$$\alpha_i = \operatorname*{argmax}_a \{\alpha_i u(y(a)) - p(a)\}\,.$$

A calculus of variations problem

The problem becomes:

$$\underset{y(\cdot),\,p(\cdot)}{\text{maximize}} \int_0^1 [au(y(a)) - p(a)]\,da$$

where for incentive compatibility we require

$$au'(y(a))y'(a) - p'(a) = 0$$

and as a capacity constraint we have

$$\int_0^1 ay(a)\,da \le \int_0^1 p(a)\,da\,.$$

This can be solved using Pontryagin's Maximum Principle (after some changes of variables and the use of a Lagrange multiplier with the second constraint).

The optimal control problem

Let $x_1(a)$ denote the net benefit of by an agent having $\alpha_i = a$. The formulation as an optimal control problem is

$$\underset{v(\cdot)}{\text{maximize}} \int_0^1 x_1(t)\,dt$$

subject to

$$\begin{aligned} 0 = w + \int_0^1 [tu(x_2(t)) + x_1(t) - tx_2(t)]\,dt \\ x_1'(t) = x_2(t) \\ x_2'(t) = v(t) \\ x_1(t) \geq 0 \\ x_2(t) \geq 0 \\ v(t) \geq 0 \\ w \geq 0 \end{aligned}$$

The solution

For $u(y) = \sqrt{y}$ the solution is

$$x_1(a) = \begin{cases} 0\,, & 0 \le a \le 0.1586 \\ -0.2677 + 0.5942\,a - 0.0942\log(a)\,, & 0.1586 \le a \le 1 \end{cases}$$

$$p(a) = ax_2(a) - x_1(a)$$
$$y(a) = x_2(a)^2$$

1. The social welfare is 0.1161 per agent, which is a bit less than the 0.125 that could be obtained by a central planner.
2. Note that agents with small α (less than $t^* = 0.1586$) are completely prevented from participating.

The optimal tariffs

Red lines show $p(\alpha)$ and $y(\alpha)$ (the amounts that agents will pay and receive). Most agents receive more than for what they pay. But some agents, with small values of α receive less than they contribute. However, if going-it-alone is not possible they will take up this scheme, since their net benefit is positive.

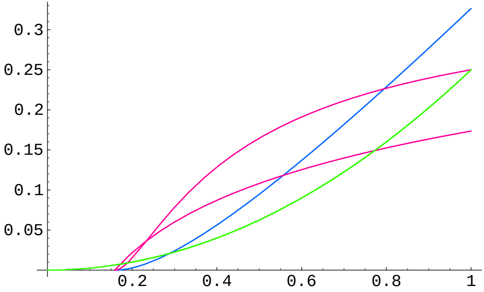

In blue we show $x_1(t)$ (an agent's net benefit), and in green $t^2/4$ (an agent's stand-alone net benefit).

The optimal tariffs

It is optimal to offer to provide $y(p)$ for a payment p.

$$y(p) = 0.3531 + 0.3531e^{-21.23\,p} - 0.7062e^{-10.61\,p}\,, \quad 0 \le p \le 0.1735\,.$$

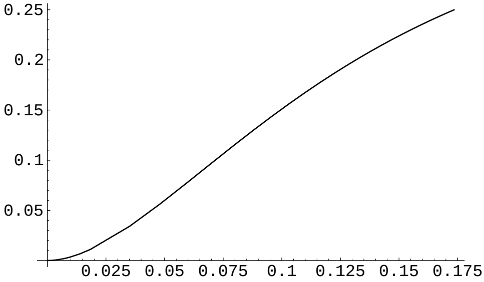

Summary

- We have modelled the sharing of rivalrous excludable good such as the computing resources shared in grid computing.
- We see that just as in the model for a nonrivalous nonexcludable good, it is optimal to set tariffs that effectively exclude from participating agents who value the system little.

  Intuitively, we need to prevent agents who value the system highly from masquerading as those who don't, so there is social benefit in excluding the latter type of agent.

References

Public good provision with many participants, M. Hellwig, *Review of Economic Studies* (2003) 70(3):589–614.

Incentives for large peer-to-peer systems, C. Courcoubetis and R. R. Weber, *IEEE Journal on Selected Areas in Communications* (2006) 24, 1034-1049.