Incentive Schemes for p2p

Costas Courcoubetis
AUEB
joint research with
Richard Weber
Panos Antoniadis

supported partly by IST proj. MMAPPs

Contents

p2p and public goods

- Public good:
  - non-rivalrous (one peer's consumption does not reduce the amount available to others)
  - positive externalities (a peer benefits from the presence of other peers because of cost sharing)
- p2p: content, coverage, connectivity have PG aspects
- Major problem: free-riding
- Our goal: design optimal incentives for contribution

A non-excludable public good

- n agents bargain to provision a public good
- Q = quantity of public good, all agents enjoy it
- $c(Q)$ = cost of public good, agent i pays p_i

$$\theta_i u(Q) - p_i$$ = agent's i net benefit

- θ_i iid, has distribution F

- Examples:

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

$$Q \in \{0,1\}, \quad u(Q) = Q, \quad c(Q) = cQ$$

Allocations

- For each $\theta = (\theta_1, \dots, \theta_n)$
  - what quantity $Q(.)$?
  - what contributions $p_1(.), \dots, p_n(.)$?
- Feasible: $c(Q(\theta)) \leq \sum p_i(\theta)$
- incentive compatible: $E_{\theta_{-i}}[\theta_i u(Q(\theta_i, \theta_{-i}) - p_i(\theta_i, \theta_{-i})] \geq E_{\theta_{-i}}[\theta_i u(Q(\hat{\theta}_i, \theta_{-i}) - p_i(\hat{\theta}_i, \theta_{-i})]$
- Individually rational: $E_{\theta_{-i}}[\theta_i u(Q(\theta_i, \theta_{-i}) - p_i(\theta_i, \theta_{-i})] \geq 0, \quad \forall \theta_i$

Allocations (2)

- First-best: maximizes Social Welfare (SW) under complete information (is trivially feasible)

$$\max_{Q(.)} \sum \theta_i u(Q(\theta)) - c(Q(\theta))$$
$$= \max_{Q} u(Q) \sum \theta_i - c(Q)$$

- Second-best: maximizes SW under incomplete information, i.e.,
  - subject to
    - feasibility
    - incentive compatibility
    - individual rationality

Example

Build a bridge

$$Q \in \{0,1\}, \quad u(Q) = Q, \quad c(Q) = cQ, \quad \theta_i \text{ iid uniform on } [0,1]$$

($c \leq 1$)

- First-best policy:

$$\max_{Q(.)} \sum \theta_i u(Q(\theta)) - c(Q(\theta))$$

- Solution (n=2):

$Q(\theta) = 1$ if $\theta_1 + \theta_2 > c$, use any $p_1 \leq \theta_1, p_2 \leq \theta_2$, s.t. $p_1 + p_2 = c$

$Q(\theta) = 0$ if $\theta_1 + \theta_2 \leq c$

$$p_i(\theta) = \frac{\theta_i}{\theta_i + \theta_j} c$$

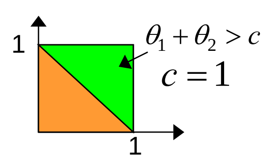

Example (2)

- Why should agents declare their actual θs ?
- If $p_i(\theta) = \frac{\theta_i}{\theta_i + \theta_j} c$ then agent with highest θ_i gains by declaring less -> SW loss
- Which is the best allocation policy? (second-best)
- Impossibility Theorem (Myerson-Satterhwaite (1983))
  - Second Best (SB) < First Best (FB)

More on second best policies

- Problem:

$$\max_{Q(.),\ p_1(.),\dots,p_n(.)} E[\sum \theta_i u(Q(\theta)) - c(Q(\theta))]$$

- subject to
  - feasibility $c(Q(\theta)) \le E[\sum p_i(\theta)]$
  - individual rationality $E_{\theta_{-i}}[\theta_i u(Q(\theta_i,\theta_{-i}) - p_i(\theta_i,\theta_{-i})] \ge 0, \quad \forall \theta_i$
  - incentive compatibility ….

A lemma for IC

- Let $V_i(\theta_i) = E_{\theta_{-i}}[u(Q(\theta_i,\theta_{-i}))], \quad P_i(\theta_i) = E_{\theta_{-i}}[p_i(\theta_i,\theta_{-i})]$
- A necessary and sufficient condition for IC is

$$P_i(\theta_i) = P_i(0) + \theta_i V_i(\theta_i) + \int_o^{\theta_i} V_i(y)dy$$

- Given IC, the system is IR iff

$$P_i(0) \le 0$$

- Then $E_\theta[\sum p_i(\theta)] = E_\theta[\sum P_i(\theta_i)]$

$$= E_\theta[\sum_1^n (\theta_i - \frac{1-F(\theta_i)}{f(\theta_i)})u(Q(\theta))]$$

The SB problem

- Solve

$$\max_{Q(.)} E_\theta[\sum_1^n \theta_i u(Q(\theta)) - c(Q(\theta))]$$

- subject to

$$E_\theta[\sum_1^n (\theta_i - \frac{1-F(\theta_i)}{f(\theta_i)})u(Q(\theta)) - c(Q(\theta))] \ge 0$$

The Lagrangian

- A Lagrangian formulation

$$\int [\sum_1^n \theta_i u(Q(\theta)) - c(Q(\theta))]\, dP_n$$

$$+ \lambda \int [\sum_1^n (\theta_i - \frac{1-F(\theta_i)}{f(\theta_i)})u(Q(\theta)) - c(Q(\theta))] dP_n$$

$g(\theta)$

- Calculation of Q(.): point wise maximization

$$Q(\theta) =$$

$$\arg\max_Q \left[\sum_1^n \theta_i u(Q) - c(Q) + \lambda \left(\sum_1^n g(\theta_i) u(Q) - c(Q) \right) \right]$$

Back to the bridge construction

- Simple calculation of optimal Q(.)

$$Q(\theta) = 1 \ \text{ iff } \ \sum_1^n \theta_i - \frac{\lambda}{1+\lambda} \sum_1^n \frac{1 - F(\theta_i)}{f(\theta_i)} \geq c$$

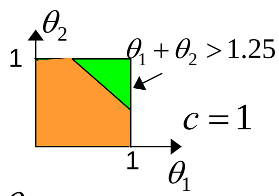

- Incentive payments: extremely complex!
  - functions of complete vector $\theta_1, \ldots, \theta_n$
  - involve money transfers between agents
  - no known simple approximation

The payments for n=2

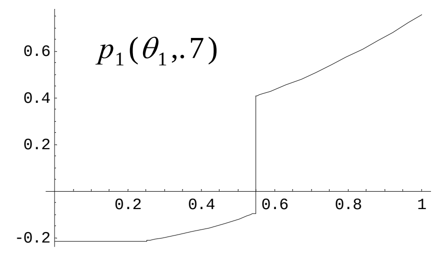

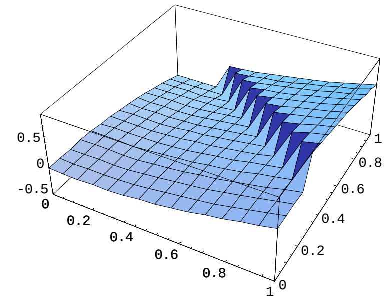

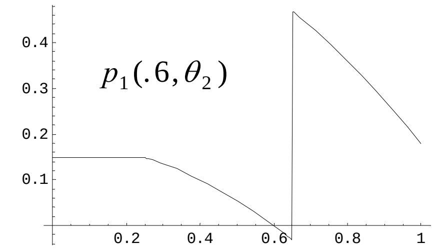

$$Q = 1 \text{ if } \ \theta_1 + \theta_2 \geq 1.25$$

Some remarks

- Optimal incentive policies are impractical to evaluate in most situations
  - Need for good approximations
- Existing results for specific models suggest that as $n \to \infty$

$$\frac{SB}{FB} \to 0$$

- If **exclusions** are possible, then

$$\frac{SB}{FB} \to \alpha > 0$$

  - Incentive payments converge to fixed contributions
  - can we obtain a general theorem?

Exclusions

- Part of the allocation policy is the exclusion capability

  Agent i is excluded if $\pi_i(\theta_1, \ldots, \theta_n) = 0$

  Agent i participates if $\pi_i(\theta_1, \ldots, \theta_n) = 1$

- Second-Best policy: solve

$$\max_{Q(.),\ \pi_1(.), \ldots,\ \pi_n(.)} E\left[\sum_1^n \theta_i \pi_i(\theta) u(Q(\theta)) - c(Q(\theta))\right]$$

such that

$$E_\theta\left[\sum_1^n \pi_i(\theta) g(\theta_i) u(Q(\theta)) - c(Q(\theta))\right] \geq 0$$

A limit theorem

Suppose

- $u(Q) = AQ^{\alpha}$, and $c(n,Q) = Bh(n)Q^{\beta}, \beta \geq 1 > \alpha$
- and Q^*, θ^* maximize

$$P = \max_{\theta \in [0,1], Q \geq 0} nu(Q) \int_{\theta}^{1} yf(y)dy - c(n,Q)$$

subject to

$$\overbrace{n[1-F(\theta)]}^{\text{\# of participants}} \overbrace{\theta u(Q)}^{\text{fixed fee}} - c(n,Q) \geq 0$$

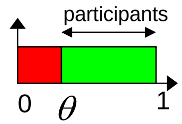

Then the simple mechanism $\pi_i(\theta) = 1\{\theta_i \geq \theta^*\}, \; Q(\theta) = Q^*,$ $p_i(\theta) = \theta^* u(Q^*)$, achieves $P \leq SB \leq (1 + O(n^{-1/5}))P$

Why large systems are simpler

- Why size helps?
  - in a large network it is hard to get people pay more than a minimum
- As the number of peers gets larger
  - a peer feels that his own declaration will have a negligible effect on the final system size
  - hence his strongest incentive is to only reduce his payment
  - therefore he declares the minimum possible theta which corresponds to the minimum fixed fee by agreeing to participate.

Example

$$u(Q) = 0.6Q^{1/2}, \quad c(Q) = Q, \quad \theta_i \text{ uniform in } [0,1]$$

$$\max_{\theta \in [0,1], Q \geq 0} \left(n \int_{\theta}^{1} ydy \right) 0.6Q^{1/2} - Q$$

$$s.t. \quad n[1-\theta]\theta \; 0.6Q^{1/2} - Q \geq 0$$

Solution: $\theta^* = 1/4, \quad Q^* = 0.0126n^2, \quad SW = 0.006328n^2$

- satisfaction of cost coverage constraint: reduction of SW by 43%

Incentives in p2p

- P2p systems exhibit strong public good aspects (externalities)
- Implication: "free market" solution is inefficient
  - each peer maximizes own net benefit
  - actions affect others
  - hence private optimum differs from social optimum
- Need regulation: use prices or rules to influence behaviour
  - incentives for each peer reflect the effect it has on others
  - example of a rule: downloads = uploads

The role of information

- Problem: optimal design requires information on user types
  - under full info: personalized price/rule for each peer
  - "first-best" policy
  - Existing approaches based on heuristics
    - reciprocity based punishments/rewards
- How can the system/planner/network manager get the required information to design optimal contribution rules?
  - if lucky, can gather personalized data about users
  - otherwise, users **must be given incentives** to reveal relevant information to planner
- Mechanism Design: set prices/rules to encourage users to act truthfully, maximize social welfare
  - for large n, use simple rules!

Applications

- File Sharing
  - public good = content availability
  - that is, number of total distinct files shared
- P2P WLANs
  - peers share wireless access to the internet
  - public good = coverage

File sharing

- Q : expected number of distinct files
- peer i :
  - utility = $\theta_i u(Q)$,
  - cost = f_i = number of files provided to the system
  - f_i randomly chosen from N files
- Rewrite equations in terms of F

$$Q(F) \approx N(1-e^{-F/N}), \text{ where } F = \sum f_i$$

- Compute optimum fixed contributions as before

The function *F(Q)*

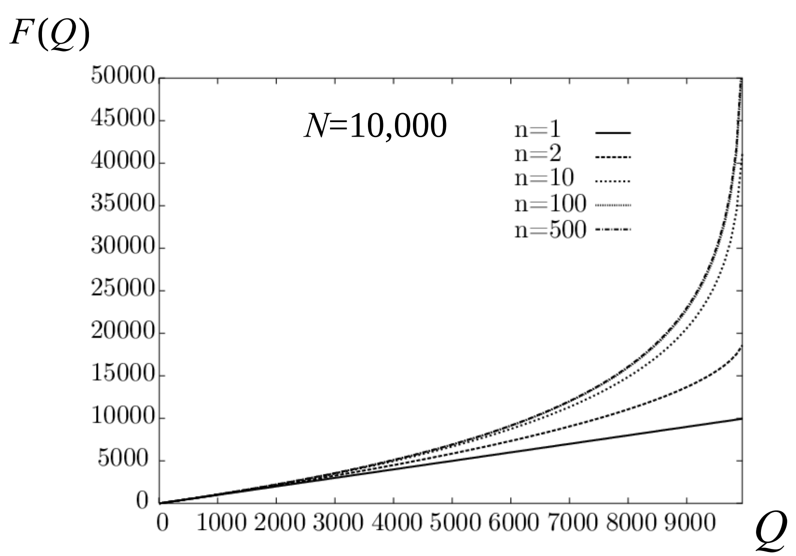

Stability

- Assume contribution f^* fixed
- Participation decision: based on prior expectation regarding total content availability *F*
- Will F^* be reached?

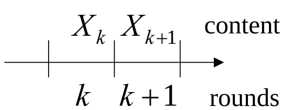

$$X_{k+1} = b(X_k, f^*)$$

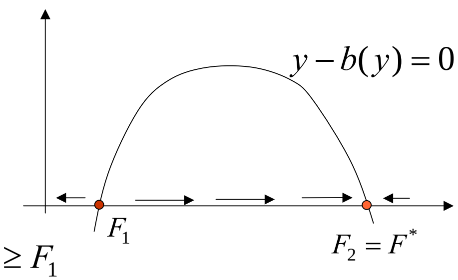

stability if $X_0 \geq F_1$

Group formation (1/3)

- Assume peers of different sub-types
- Type A: $\theta_i^A \sim [0,0.5]$ (e.g. ISDN users)
- Type B: $\theta_i^B \sim [0.5,1]$ (e.g. DSL users)
- Do they gain more by
  - forming 2 distinct groups vs forming a larger group?
  - being distinguished by the system in the larger group?

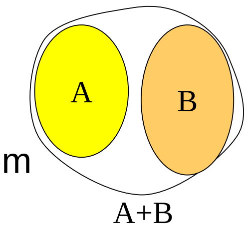

Group formation (2/3)

- Group A: $\theta_i^A \sim [0,0.5]$ (e.g. ISDN users)
- Group B: $\theta_i^B \sim [0.5,1]$ (e.g. DSL users)

Assume that the percentage of each group in the mix is 50% (n=1000)

Welfare	Group A	Group B	Total
Distinct groups	3296	35156	38452
Indistinguishable	6976 (+ 111%)	44792 (+ 27%)	51768
Distinguishable	31249 (+ 848%)	31250 (-11%)	62500

Group formation (3/3)

- How to provide better incentives for both types to combine and reveal their types?
  - reduce cost of heavy users by limiting upload rates
  - reduce fees of heavy users
- Offer sets of tariffs (versioning)
  - allow self-selection
- Model difference in cost for uploading
  - higher-cost peers benefit in a larger group when types can be distinguished

WLANs

- *L* geographic areas
  - coverage: fraction of area covered by hotspots
  - area *l* has coverage Q_l

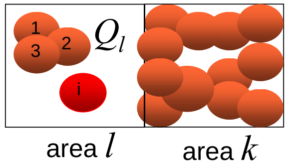

- Peer *i* of area *l*
  - obtains benefit $\theta_i^l u^l(Q_1,\dots,Q_L)$
  - contributes payment (coverage) f_i^l
- Feasibility: $Q_l \le \sum_{i=1,n_l} f_i^l$, for all areas l
- Problem: find optimum incentive scheme, maximize efficiency
- Our Theorem holds
  - use fixed contribution schemes

Conclusions

- In p2p systems with strong PG aspects
  - Fixed contribution schemes can be efficient
  - Result to tractable optimization problems
  - Provide useful insight to many interesting questions
  - Information regarding user types may be strategic
- Open issues:
  - more complex cost structures
  - adaptation
  - multiple round games
  - use existing data to tune economic model