Sharing Resources and Costs in Grids

Costas Courcoubetis†

InterPerf Workshop, VALUETOOLS 2008

Athens, Greece, 24 October, 2008

† Dept of Computer Science, Athens University of Economics and Business

Joint work with Richard Weber

Motivation

How do we derive sensible resource allocation policies in shared infrastructures (Grids, communication links) when participants have private information?

Motivation

How do we derive sensible resource allocation policies in shared infrastructures (Grids, communication links) when participants have private information?

Key observation: system operating policies influence strategies of agents.

Motivation

How do we derive sensible resource allocation policies in shared infrastructures (Grids, communication links) when participants have private information?

Key observation: system operating policies influence strategies of agents.

Traditional policies may fail!

Motivation

How do we derive sensible resource allocation policies in shared infrastructures (Grids, communication links) when participants have private information?

Key observation: system operating policies influence strategies of agents.

Traditional policies may fail!

Two main problems:

- The way the system resolves resource conflicts depends on declarations by agents.
- Since facility cost must be shared, agents like to free-ride.

Motivation

How do we address these issues?

Motivation

How do we address these issues?

Are auction theory and Mechanism Design appropriate? And under what assumptions on our model are these applicable?

Motivation

How do we address these issues?

Are auction theory and Mechanism Design appropriate? And under what assumptions on our model are these applicable?

What is fundamentally new in this problem?

Motivation

How do we address these issues?

Are auction theory and Mechanism Design appropriate? And under what assumptions on our model are these applicable?

What is fundamentally new in this problem?

Can we compute the optimal policies?

Our model

1. The facility already exists, has fixed size Q and known operating cost, or

Our model

1. The facility already exists, has fixed size Q and known operating cost, or
2. The facility does not exist, it will be determined by contributions of the participants.

Our model

1. The facility already exists, has fixed size Q and known operating cost, or
2. The facility does not exist, it will be determined by contributions of the participants.

How to share resources and recover costs?

Our model

1. The facility already exists, has fixed size Q and known operating cost, or
2. The facility does not exist, it will be determined by contributions of the participants.

How to share resources and recover costs?

- Easy when we know utilities of participants.

Our model

1. The facility already exists, has fixed size Q and known operating cost, or
2. The facility does not exist, it will be determined by contributions of the participants.

How to share resources and recover costs?

- Easy when we know utilities of participants.
- But in practice participants have private information! If the policy of the system depends on this information, they will disclose it to the best of their advantage.

Example: scheduling a server

- N agents will use a FCFS server, each generating a Poisson stream of exponential jobs with different mean.

Example: scheduling a server

- N agents will use a FCFS server, each generating a Poisson stream of exponential jobs with different mean.
- Each agent suffers a linear delay cost that is private information, but its distribution is public information.

Example: scheduling a server

- N agents will use a FCFS server, each generating a Poisson stream of exponential jobs with different mean.
- Each agent suffers a linear delay cost that is private information, but its distribution is public information.
- Initially each agent i contributes an amount y_i, and the server is build to size $Q = \sum y_i$ (assume linear cost).

Example: scheduling a server

- N agents will use a FCFS server, each generating a Poisson stream of exponential jobs with different mean.
- Each agent suffers a linear delay cost that is private information, but its distribution is public information.
- Initially each agent i contributes an amount y_i, and the server is build to size $Q = \sum y_i$ (assume linear cost).
- Lots of free ridding at the NE! How can we improve the system?

Example: scheduling a server

- N agents will use a FCFS server, each generating a Poisson stream of exponential jobs with different mean.
- Each agent suffers a linear delay cost that is private information, but its distribution is public information.
- Initially each agent i contributes an amount y_i, and the server is build to size $Q = \sum y_i$ (assume linear cost).
- Lots of free ridding at the NE! How can we improve the system?
- Relate the scheduling of the server to the contribution of the agents! E.g., processor sharing with coefficients proportional to the y_is.

Example: scheduling a server

- N agents will use a FCFS server, each generating a Poisson stream of exponential jobs with different mean.
- Each agent suffers a linear delay cost that is private information, but its distribution is public information.
- Initially each agent i contributes an amount y_i, and the server is build to size $Q = \sum y_i$ (assume linear cost).
- Lots of free ridding at the NE! How can we improve the system?
- Relate the scheduling of the server to the contribution of the agents! E.g., processor sharing with coefficients proportional to the y_is.
- Which is the optimal scheduling policy?

A simple example of full information

- Discrete time t,
- Facility of size Q (or to be determined), fixed cost c per time slot,
- In slot t agent i has utility $\theta_{i,t}\sqrt{x_i}$, $\theta_{i,t}$ iid, $U[0,1]$.

A simple example of full information

- Discrete time t,
- Facility of size Q (or to be determined), fixed cost c per time slot,
- In slot t agent i has utility $\theta_{i,t}\sqrt{x_i}$, $\theta_{i,t}$ iid, $U[0,1]$.

Resource sharing problem: At each time t allocate the resource to maximize value to agents, obtain payments to cover the cost.

$$\underset{\{x_i\}}{\text{maximize}} \sum_{i=1}^n \theta_{i,t}\sqrt{x_i} \text{ such that } \sum_{i=1}^n x_i \leq Q.$$

A simple example of full information

- Discrete time t,
- Facility of size Q (or to be determined), fixed cost c per time slot,
- In slot t agent i has utility $\theta_{i,t}\sqrt{x_i}$, $\theta_{i,t}$ iid, $U[0,1]$.

Resource sharing problem: At each time t allocate the resource to maximize value to agents, obtain payments to cover the cost.

$$\underset{\{x_i\}}{\text{maximize}} \sum_{i=1}^{n} \theta_{i,t}\sqrt{x_i} \text{ such that } \sum_{i=1}^{n} x_i \leq Q.$$

Solution:

$$x_i = \frac{\theta_{i,t}^2}{\sum_{k=1}^{n} \theta_{k,t}^2} Q,\; V_{i,t} = \frac{\theta_{i,t}^2}{\sqrt{\sum_{k=1}^{n} \theta_{k,t}^2}} \sqrt{Q},\; V_t = \sqrt{\sum_{k=1}^{n} \theta_{k,t}^2}\sqrt{Q}\,.$$

A simple example of full information

- Discrete time t,
- Facility of size Q (or to be determined), fixed cost c per time slot,
- In slot t agent i has utility $\theta_{i,t}\sqrt{x_i}$, $\theta_{i,t}$ iid, $U[0,1]$.

Resource sharing problem: At each time t allocate the resource to maximize value to agents, obtain payments to cover the cost.

$$\underset{\{x_i\}}{\text{maximize}} \sum_{i=1}^{n} \theta_{i,t}\sqrt{x_i} \text{ such that } \sum_{i=1}^{n} x_i \leq Q.$$

Solution:

$$x_i = \frac{\theta_{i,t}^2}{\sum_{k=1}^{n} \theta_{k,t}^2} Q,\ V_{i,t} = \frac{\theta_{i,t}^2}{\sqrt{\sum_{k=1}^{n} \theta_{k,t}^2}}\sqrt{Q},\ V_t = \sqrt{\sum_{k=1}^{n} \theta_{k,t}^2}\sqrt{Q}\,.$$

If $E[V_t] > c$, fix payments $p_{i,t} = c/n$. This may not work if $\theta_{i,t} = \theta_i$ for all t (it may not be ex-post rational).

Subtle issues

- What do we really know about the agents?

Subtle issues

- What do we really know about the agents?
  - The true value of the $\theta_{i,t}$s?

Subtle issues

- What do we really know about the agents?
  - The true value of the $\theta_{i,t}$s?
  - The distributions F_i of the $\theta_{i,t}$s (the types of the agents) ?

Subtle issues

- What do we really know about the agents?
  - The true value of the $\theta_{i,t}$s?
  - The distributions F_i of the $\theta_{i,t}$s (the types of the agents) ?
  - The distribution of the types of the agents?

Subtle issues

- What do we really know about the agents?
  - The true value of the $\theta_{i,t}$s?
  - The distributions F_i of the $\theta_{i,t}$s (the types of the agents) ?
  - The distribution of the types of the agents?
- If an agent declares its F_i, can we police it? i.e., punish him if his subsequent declarations $\theta_{i,t}$, $t = 1, 2, \ldots,$ are not consistent with F_i?

Subtle issues

- What do we really know about the agents?
  - The true value of the $\theta_{i,t}$s?
  - The distributions F_i of the $\theta_{i,t}$s (the types of the agents) ?
  - The distribution of the types of the agents?
- If an agent declares its F_i, can we police it? i.e., punish him if his subsequent declarations $\theta_{i,t}$, $t = 1, 2, \ldots,$ are not consistent with F_i?
- In a variation of the model, agents are on/off with a certain probability. How do we obtain this information?

The general model

The general model

- N agents share common resource Q **over time** $t = 1, 2, \ldots$.

The general model

- N agents share common resource Q **over time** $t = 1, 2, \ldots$.
- At time t agent i has utility $\theta_{i,t} u(x_i)$, where
  - x_i is the amount of resource allocated, $u(.)$ convex increasing,
  - $\theta_{i,t}$ is its *personalization parameter*, iid, realized from distribution F_i (e.g., $U[0, Y_i]$).

The general model

- N agents share common resource Q **over time** $t = 1, 2, \ldots$.
- At time t agent i has utility $\theta_{i,t} u(x_i)$, where
  - x_i is the amount of resource allocated, $u(.)$ convex increasing,
  - $\theta_{i,t}$ is its *personalization parameter*, iid, realized from distribution F_i (e.g., $U[0, Y_i]$).
- System operates according to game G (see next slide).

The general model

- N agents share common resource Q **over time** $t = 1, 2, \ldots$.
- At time t agent i has utility $\theta_{i,t} u(x_i)$, where
  - x_i is the amount of resource allocated, $u(.)$ convex increasing,
  - $\theta_{i,t}$ is its *personalization parameter*, iid, realized from distribution F_i (e.g., $U[0, Y_i]$).
- System operates according to game G (see next slide).
- Aim: design G so that at the Nash equilibrium the expected net benefit of the agents is maximized while recovering a cost c/period in the long run.

The general model

- N agents share common resource Q **over time** $t = 1, 2, \ldots$.
- At time t agent i has utility $\theta_{i,t} u(x_i)$, where
  - x_i is the amount of resource allocated, $u(.)$ convex increasing,
  - $\theta_{i,t}$ is its *personalization parameter*, iid, realized from distribution F_i (e.g., $U[0, Y_i]$).
- System operates according to game G (see next slide).
- Aim: design G so that at the Nash equilibrium the expected net benefit of the agents is maximized while recovering a cost c/period in the long run.

The size Q of the system may be determined as part of the game, for a given cost $c(Q)$.

The general model

- N agents share common resource Q **over time** $t = 1, 2, \ldots$.
- At time t agent i has utility $\theta_{i,t} u(x_i)$, where
  - x_i is the amount of resource allocated, $u(.)$ convex increasing,
  - $\theta_{i,t}$ is its *personalization parameter*, iid, realized from distribution F_i (e.g., $U[0, Y_i]$).
- System operates according to game G (see next slide).
- Aim: design G so that at the Nash equilibrium the expected net benefit of the agents is maximized while recovering a cost c/period in the long run.

The size Q of the system may be determined as part of the game, for a given cost $c(Q)$.

An agent should profit more by participating in this system than by building his own facility.

The game G

The game G

1. System designer posts operating rules of the facility,

The game G

1. System designer posts operating rules of the facility,
2. Participants chose contracts (*), i.e., decide on disclosing certain information. They may make some initial payment depending on contract chosen.

The game G

1. System designer posts operating rules of the facility,
2. Participants chose contracts (*), i.e., decide on disclosing certain information. They may make some initial payment depending on contract chosen.
3. The facility operates in discrete slots $t = 1, 2, \ldots$. At each time t the agents derive value by accessing the facility and may make further payments to cover the running cost. They may also be asked to disclose some private information regarding their utility of the service at time t (**).

The game G

1. System designer posts operating rules of the facility,
2. Participants chose contracts (*), i.e., decide on disclosing certain information. They may make some initial payment depending on contract chosen.
3. The facility operates in discrete slots $t = 1, 2, \ldots$. At each time t the agents derive value by accessing the facility and may make further payments to cover the running cost. They may also be asked to disclose some private information regarding their utility of the service at time t (**).
4. The resource sharing and the payment policies take into account the information provided in (*),(**).

Example: allocating a single item

We wish to share a single machine between 2 agents. Each day agent i has utility $\theta_{i,t}$, where $F_1 = U[0,1]$ and $F_2 = U[0,2]$, considered known to system operator. How de we allocate the machine and take payments to cover the cost c?

Example: allocating a single item

We wish to share a single machine between 2 agents. Each day agent i has utility $\theta_{i,t}$, where $F_1 = U[0,1]$ and $F_2 = U[0,2]$, considered known to system operator. How de we allocate the machine and take payments to cover the cost c?

Lets first consider a simple intuitive policy (A1):

Example: allocating a single item

We wish to share a single machine between 2 agents. Each day agent i has utility $\theta_{i,t}$, where $F_1 = U[0,1]$ and $F_2 = U[0,2]$, considered known to system operator. How de we allocate the machine and take payments to cover the cost c?

Lets first consider a simple intuitive policy (A1):

- Charge a price p for using the machine.

Example: allocating a single item

We wish to share a single machine between 2 agents. Each day agent i has utility $\theta_{i,t}$, where $F_1 = U[0,1]$ and $F_2 = U[0,2]$, considered known to system operator. How de we allocate the machine and take payments to cover the cost c?

Lets first consider a simple intuitive policy (A1):

- Charge a price p for using the machine.
- If one agent wants to use it, he gets it.

Example: allocating a single item

We wish to share a single machine between 2 agents. Each day agent i has utility $\theta_{i,t}$, where $F_1 = U[0,1]$ and $F_2 = U[0,2]$, considered known to system operator. How de we allocate the machine and take payments to cover the cost c?

Lets first consider a simple intuitive policy (A1):

- Charge a price p for using the machine.
- If one agent wants to use it, he gets it.
- If both agents want to use it, agent 2 gets it.

Example: allocating a single item

We wish to share a single machine between 2 agents. Each day agent i has utility $\theta_{i,t}$, where $F_1 = U[0,1]$ and $F_2 = U[0,2]$, considered known to system operator. How de we allocate the machine and take payments to cover the cost c?

Lets first consider a simple intuitive policy (A1):

- Charge a price p for using the machine.
- If one agent wants to use it, he gets it.
- If both agents want to use it, agent 2 gets it.
- Choose p so that expected payment per time slot is c.

Using an auction

Can we do any better? Auction theory suggests the following scheme.

Using an auction

Can we do any better? Auction theory suggests the following scheme.

- Pick a number b and construct the prices

$$p_1 = \max\{b, \theta_{2,t} - b\},\ p_2 = \max\{2b, \theta_{1,t} + b\}\,.$$

Using an auction

Can we do any better? Auction theory suggests the following scheme.

- Pick a number b and construct the prices

$$p_1 = \max\{b, \theta_{2,t} - b\}, \; p_2 = \max\{2b, \theta_{1,t} + b\}.$$

- Ask the agents to reveal their $\theta_{i,t}$. If $\theta_{i,t} > p_i$, the machine is allocated to agent i (this cannot hold for both 1 and 2), and he pays p_i.

Using an auction

Can we do any better? Auction theory suggests the following scheme.

- Pick a number b and construct the prices

$$p_1 = \max\{b, \theta_{2,t} - b\},\ p_2 = \max\{2b, \theta_{1,t} + b\}\,.$$

- Ask the agents to reveal their $\theta_{i,t}$. If $\theta_{i,t} > p_i$, the machine is allocated to agent i (this cannot hold for both 1 and 2), and he pays p_i.
- Choose b so that expected payment per slot becomes c.

Using an auction

Can we do any better? Auction theory suggests the following scheme.

- Pick a number b and construct the prices

$$p_1 = \max\{b, \theta_{2,t} - b\},\ p_2 = \max\{2b, \theta_{1,t} + b\}\,.$$

- Ask the agents to reveal their $\theta_{i,t}$. If $\theta_{i,t} > p_i$, the machine is allocated to agent i (this cannot hold for both 1 and 2), and he pays p_i.
- Choose b so that expected payment per slot becomes c.

Solution maximizes sum of expected agent utilities conditional on recovering c on the average, is incentive compatible. Note that agent 2 wins if $\theta_{2,t} > \theta_{1,t} + b$.

Comparing the policies

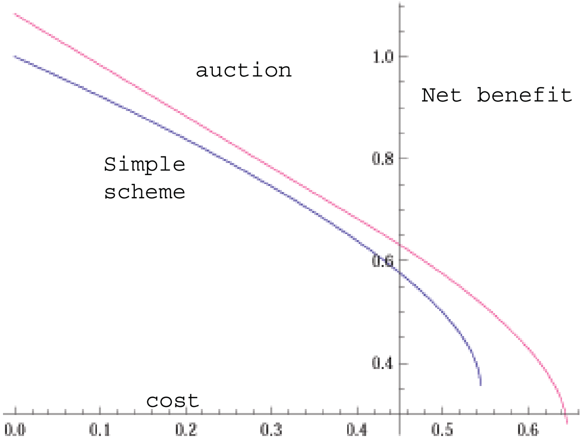

What did we learn?

What did we learn?

- The cost coverage constraint induces inefficient allocations!

What did we learn?

- The cost coverage constraint induces inefficient allocations!
- But infinite repetition simplifies the problem significantly. If we know the types of the agents and can police them we can easily achieve efficient allocation:

What did we learn?

- The cost coverage constraint induces inefficient allocations!
- But infinite repetition simplifies the problem significantly. If we know the types of the agents and can police them we can easily achieve efficient allocation:
  1. Assume that the system will run with full information at each t; charge each agent a fixed payment/round less than his average utility/round, to cover cost.
  2. Ask the value of the $\theta_{i,t}$ at each t and share the server efficiently; police the declarations according to the F_is.

What did we learn?

- The cost coverage constraint induces inefficient allocations!
- But infinite repetition simplifies the problem significantly. If we know the types of the agents and can police them we can easily achieve efficient allocation:
  1. Assume that the system will run with full information at each t; charge each agent a fixed payment/round less than his average utility/round, to cover cost.
  2. Ask the value of the $\theta_{i,t}$ at each t and share the server efficiently; police the declarations according to the F_is.
- This scheme is incentive compatible (not immediate, requires a proof).

What did we learn?

- The cost coverage constraint induces inefficient allocations!
- But infinite repetition simplifies the problem significantly. If we know the types of the agents and can police them we can easily achieve efficient allocation:
  1. Assume that the system will run with full information at each t; charge each agent a fixed payment/round less than his average utility/round, to cover cost.
  2. Ask the value of the $\theta_{i,t}$ at each t and share the server efficiently; police the declarations according to the F_is.
- This scheme is incentive compatible (not immediate, requires a proof).
- Open problem: optimal scheme if we do not know the F_is.

What did we learn?

- The cost coverage constraint induces inefficient allocations!
- But infinite repetition simplifies the problem significantly. If we know the types of the agents and can police them we can easily achieve efficient allocation:
  1. Assume that the system will run with full information at each t; charge each agent a fixed payment/round less than his average utility/round, to cover cost.
  2. Ask the value of the $\theta_{i,t}$ at each t and share the server efficiently; police the declarations according to the F_is.
- This scheme is incentive compatible (not immediate, requires a proof).
- Open problem: optimal scheme if we do not know the F_is.
- Lets look at some other fun problems!

Building a facility from scratch

A different model: a facility of size Q and cost $c(Q) = Q$ (per slot) is formed by initial contributions of agents. These are incentivized to contribute because their contribution will affect the amount of resources they will get at run time. Probably a good model for virtual Grid infrastructures.

Building a facility from scratch

A different model: a facility of size Q and cost $c(Q) = Q$ (per slot) is formed by initial contributions of agents. These are incentivized to contribute because their contribution will affect the amount of resources they will get at run time. Probably a good model for virtual Grid infrastructures.

- Agents are 'on-off' with probability α_i (public information).

Building a facility from scratch

A different model: a facility of size Q and cost $c(Q) = Q$ (per slot) is formed by initial contributions of agents. These are incentivized to contribute because their contribution will affect the amount of resources they will get at run time. Probably a good model for virtual Grid infrastructures.

- Agents are 'on-off' with probability α_i (public information).
- $\theta_{i,t} = \theta_i \forall t$ (private information).
- Control policy is $x_i(\boldsymbol{\theta}, S)$, $S=$ conflicting set of agents.

Building a facility from scratch

A different model: a facility of size Q and cost $c(Q) = Q$ (per slot) is formed by initial contributions of agents. These are incentivized to contribute because their contribution will affect the amount of resources they will get at run time. Probably a good model for virtual Grid infrastructures.

- Agents are 'on-off' with probability α_i (public information).
- $\theta_{i,t} = \theta_i \forall t$ (private information).
- Control policy is $x_i(\boldsymbol{\theta}, S)$, $S=$ conflicting set of agents.
- System planner posts {agent's contribution, operating policy} as a function of the θ_is that the agents will declare.

Building a facility from scratch

A different model: a facility of size Q and cost $c(Q) = Q$ (per slot) is formed by initial contributions of agents. These are incentivized to contribute because their contribution will affect the amount of resources they will get at run time. Probably a good model for virtual Grid infrastructures.

- Agents are 'on-off' with probability α_i (public information).
- $\theta_{i,t} = \theta_i \forall t$ (private information).
- Control policy is $x_i(\boldsymbol{\theta}, S)$, $S=$ conflicting set of agents.
- System planner posts {agent's contribution, operating policy} as a function of the θ_is that the agents will declare.
- Agents declare their θ_is and system runs according to posted policy.

A large n result

A large n result

- Assume $c(Q) = Q$ and n very large.

A large n result

- Assume $c(Q) = Q$ and n very large.
- Then S will always be near its typical value and $x_i(\boldsymbol{\theta}, S)$ become $x_i(\boldsymbol{\theta})$.

A large n result

- Assume $c(Q) = Q$ and n very large.
- Then S will always be near its typical value and $x_i(\boldsymbol{\theta}, S)$ become $x_i(\boldsymbol{\theta})$.
- The allocations should satisfy $\sum_i \alpha_i x_i(\theta) \leq Q$.

A large n result

- Assume $c(Q) = Q$ and n very large.
- Then S will always be near its typical value and $x_i(\boldsymbol{\theta}, S)$ become $x_i(\boldsymbol{\theta})$.
- The allocations should satisfy $\sum_i \alpha_i x_i(\theta) \leq Q$.

It turns out that the solution of the Mechanism Design problem implies a simple 'effective bandwith' tariff for type i agents:

- System guarantees resource y for a contribution of $\alpha_i y$.
- Agent i indirectly declares its θ_i by selecting y to maximize $\max_y\{\theta_i u(y) - \alpha_i y\}$.
- No information on F_i required!

Analysis of proportional sharing

How do sharing policies affect the incentives for agents to contribute? Consider the simple case of 2 identical participants $\theta_1 = \theta_2 = 1$, $c(Q) = Q$.

Analysis of proportional sharing

How do sharing policies affect the incentives for agents to contribute? Consider the simple case of 2 identical participants $\theta_1 = \theta_2 = 1$, $c(Q) = Q$.

Agent i contributes q_i. Average value of using the system for agent 1 is

$$\alpha_1(1-\alpha_2)u(x_1^{\{1\}}) + \alpha_1\alpha_2 u(x_1^{\{1,2\}}) - q_1\,.$$

Analysis of proportional sharing

How do sharing policies affect the incentives for agents to contribute? Consider the simple case of 2 identical participants $\theta_1 = \theta_2 = 1$, $c(Q) = Q$.

Agent i contributes q_i. Average value of using the system for agent 1 is

$$\alpha_1(1-\alpha_2)u(x_1^{\{1\}}) + \alpha_1\alpha_2 u(x_1^{\{1,2\}}) - q_1 \,.$$

We analyze 4 sharing disciplines:

Analysis of proportional sharing

How do sharing policies affect the incentives for agents to contribute? Consider the simple case of 2 identical participants $\theta_1 = \theta_2 = 1$, $c(Q) = Q$.

Agent i contributes q_i. Average value of using the system for agent 1 is

$$\alpha_1(1-\alpha_2)u(x_1^{\{1\}}) + \alpha_1\alpha_2 u(x_1^{\{1,2\}}) - q_1 \,.$$

We analyze 4 sharing disciplines:

- **Acting alone:** $x_i^{\{i\}} = x_i^{\{1,2\}} = q_i$.

Analysis of proportional sharing

How do sharing policies affect the incentives for agents to contribute? Consider the simple case of 2 identical participants $\theta_1 = \theta_2 = 1$, $c(Q) = Q$.

Agent i contributes q_i. Average value of using the system for agent 1 is

$$\alpha_1(1-\alpha_2)u(x_1^{\{1\}}) + \alpha_1\alpha_2 u(x_1^{\{1,2\}}) - q_1 \,.$$

We analyze 4 sharing disciplines:

- **Acting alone:** $x_i^{\{i\}} = x_i^{\{1,2\}} = q_i$.
- **Equal sharing:** $x_i^{\{i\}} = q_1 + q_2$ and $x_i^{\{1,2\}} = \frac{1}{2}(q_1 + q_2)$.

Analysis of proportional sharing

How do sharing policies affect the incentives for agents to contribute? Consider the simple case of 2 identical participants $\theta_1 = \theta_2 = 1$, $c(Q) = Q$.

Agent i contributes q_i. Average value of using the system for agent 1 is

$$\alpha_1(1-\alpha_2)u(x_1^{\{1\}}) + \alpha_1\alpha_2 u(x_1^{\{1,2\}}) - q_1 \,.$$

We analyze 4 sharing disciplines:

- **Acting alone:** $x_i^{\{i\}} = x_i^{\{1,2\}} = q_i$.
- **Equal sharing:** $x_i^{\{i\}} = q_1 + q_2$ and $x_i^{\{1,2\}} = \frac{1}{2}(q_1 + q_2)$.
- **Proportional sharing:**

$$x_i^{\{i\}} = q_1 + q_2 \,, \quad x_i^{\{1,2\}} = \frac{q_1}{q_1 + q_2}(q_1 + q_2) \,.$$

Analysis of proportional sharing

How do sharing policies affect the incentives for agents to contribute? Consider the simple case of 2 identical participants $\theta_1 = \theta_2 = 1$, $c(Q) = Q$.

Agent i contributes q_i. Average value of using the system for agent 1 is

$$\alpha_1(1-\alpha_2)u(x_1^{\{1\}}) + \alpha_1\alpha_2 u(x_1^{\{1,2\}}) - q_1 \,.$$

We analyze 4 sharing disciplines:

- **Acting alone:** $x_i^{\{i\}} = x_i^{\{1,2\}} = q_i$.
- **Equal sharing:** $x_i^{\{i\}} = q_1 + q_2$ and $x_i^{\{1,2\}} = \frac{1}{2}(q_1 + q_2)$.
- **Proportional sharing:**

$$x_i^{\{i\}} = q_1 + q_2 \,, \quad x_i^{\{1,2\}} = \frac{q_1}{q_1 + q_2}(q_1 + q_2) \,.$$

- **s-Proportional sharing:**

$$x_i^{\{i\}} = q_1 + q_2 \,, \quad x_i^{\{1,2\}} = \frac{q_1^s}{q_1^s + q_2^s}(q_1 + q_2) \,.$$

Results for $\alpha = .8$, $u(x) = 10 - 1/x$

scheme	social welfare	values of q_1, q_2
stand alone	$r\alpha - 2\sqrt{\alpha}$	$\sqrt{\alpha}$
	6.21115	0.894427
central planner	$r\alpha - \sqrt{2\alpha(1+\alpha)}$	$\sqrt{\alpha(1+\alpha)/2}$
$s = \frac{1}{2}(1 + 1/\alpha)$	6.30294	0.848528
proportional division	$r\alpha - \frac{\sqrt{\alpha}\,(3+5\,\alpha)}{2\,\sqrt{1+3\,\alpha}}$	$\frac{1}{2}\sqrt{\alpha(1+3\,\alpha)}$
$s = 1$	6.30225	0.824621
equal division	$r\alpha - \frac{3}{2}\sqrt{\alpha(1+\alpha)}$	$\frac{1}{2}\sqrt{\alpha(1+\alpha)}$
$s = 0$	6.2	0.6

Results for $\alpha = .8$, $u(x) = 10 - 1/x$

scheme	social welfare	values of q_1, q_2
stand alone	$r\alpha - 2\sqrt{\alpha}$	$\sqrt{\alpha}$
	6.21115	0.894427
central planner	$r\alpha - \sqrt{2\alpha(1+\alpha)}$	$\sqrt{\alpha(1+\alpha)/2}$
$s = \frac{1}{2}(1 + 1/\alpha)$	6.30294	0.848528
proportional division	$r\alpha - \frac{\sqrt{\alpha}\,(3+5\,\alpha)}{2\,\sqrt{1+3\,\alpha}}$	$\frac{1}{2}\sqrt{\alpha(1+3\,\alpha)}$
$s = 1$	6.30225	0.824621
equal division	$r\alpha - \frac{3}{2}\sqrt{\alpha(1+\alpha)}$	$\frac{1}{2}\sqrt{\alpha(1+\alpha)}$
$s = 0$	6.2	0.6

Open problem: how do these results generalize? How much information is needed to determine the optimal s?

How bad is equal sharing?

How bad is equal sharing?

Suppose agents have types $\theta_i = 1$ and $\alpha_1 > \cdots > \alpha_n$. Then only agent 1 will have any incentive to contribute resources to the grid.

How bad is equal sharing?

Suppose agents have types $\theta_i = 1$ and $\alpha_1 > \cdots > \alpha_n$. Then only agent 1 will have any incentive to contribute resources to the grid.

The problems that are to be solved by agents 1 and 2, respectively, are to maximize with respect to q_1 and q_2

$$nb_1(q) = \alpha_1 \left[\alpha_2 Eu\left(\frac{\sum_1^n q_i}{M+2} \right) + (1-\alpha_2) Eu\left(\frac{\sum_1^n q_i}{M+1} \right) \right] - q_1 \,,$$

$$nb_2(q) = \alpha_2 \left[\alpha_1 Eu\left(\frac{\sum_1^n q_i}{M+2} \right) + (1-\alpha_1) Eu\left(\frac{\sum_1^n q_i}{M+1} \right) \right] - q_2$$

where $M=$ number of agents $3, \ldots, n$ that are present. Since $\alpha_1(1-\alpha_2) > \alpha_2(1-\alpha_1)$ it follows that $\frac{\partial}{\partial q_1} nb_1(q) = 0 \implies \frac{\partial}{\partial q_2} nb_2(q) < 0.$

Declaring activity frequencies

Declaring activity frequencies

Now the α_is are private information, iid uniform on $[0, 1]$, and $\theta_{i,t} = \theta_i = 1$.

Declaring activity frequencies

Now the α_is are private information, iid uniform on $[0, 1]$, and $\theta_{i,t} = \theta_i = 1$. Sensible if accounting of activity is costly (e.g., DSL).

Declaring activity frequencies

Now the α_is are private information, iid uniform on $[0,1]$, and $\theta_{i,t} = \theta_i = 1$. Sensible if accounting of activity is costly (e.g., DSL).

The facility is build from agent conributions. We like to compute the set of optimal tariffs $q(\omega), x(\omega)$ parametrized by ω the type of the agent, where an agent that contributes $q(\omega)$ gets $x(\omega)$ when he is active.

Declaring activity frequencies

Now the α_is are private information, iid uniform on $[0, 1]$, and $\theta_{i,t} = \theta_i = 1$. Sensible if accounting of activity is costly (e.g., DSL).

The facility is build from agent conributions. We like to compute the set of optimal tariffs $q(\omega), x(\omega)$ parametrized by ω the type of the agent, where an agent that contributes $q(\omega)$ gets $x(\omega)$ when he is active.

An agent maximizes his net benefit $f(\alpha)$, where

$$f(\alpha) = \max\left\{ \max_s \left[\alpha u(x(s)) - q(s)\right], 0 \right\}.$$

Declaring activity frequencies

Now the α_is are private information, iid uniform on $[0, 1]$, and $\theta_{i,t} = \theta_i = 1$. Sensible if accounting of activity is costly (e.g., DSL).

The facility is build from agent conributions. We like to compute the set of optimal tariffs $q(\omega), x(\omega)$ parametrized by ω the type of the agent, where an agent that contributes $q(\omega)$ gets $x(\omega)$ when he is active.

An agent maximizes his net benefit $f(\alpha)$, where

$$f(\alpha) = \max\left\{\max_s \left[\alpha u(x(s)) - q(s)\right], 0\right\}.$$

Given $f(t)$, one can recover the corresponding $x(t)$ and $y(t)$, by taking $q(s) = -f(s) + sf'(s)$ and $u(x(s)) = f'(s)$. These tariffs are incentive compatible.

An optimal control formulation for $u(y) = y^{1/k}$

When n is large we wish to design $f(\cdot)$ so as to maximize

$$\int_0^1 f(s)\, ds$$

subject to a constraint that says that incoming contributions and outgoing allocations are equal, i.e.,

$$\int_0^1 [sx(s) - q(s)]\, ds = \int_0^1 [sf'(s)^k + f(s) - sf'(s)]\, ds \leq 0\,.$$

An optimal control formulation for $u(y) = y^{1/k}$

Let $z_1(t) = f(t)$, $z_2(t) = f'(t)$. Our problem becomes

$$\text{maximize} \int_0^1 z_1(t)\,dt \text{ subject to}$$

$$\begin{aligned}
0 &= w + \int_0^1 [s z_2(s)^k + z_1(s) - s z_2(s)]\,ds \\
z_1'(t) &= z_2(t) \\
z_2'(t) &= v(t) \\
z_1(t) &\geq 0 \\
z_2(t) &\geq 0 \\
w &\geq 0\,.
\end{aligned}$$

An optimal control formulation for $u(y) = y^{1/k}$

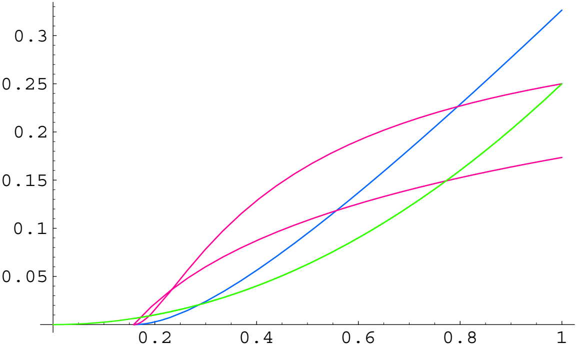

An optimal control formulation for $u(y) = y^{1/k}$

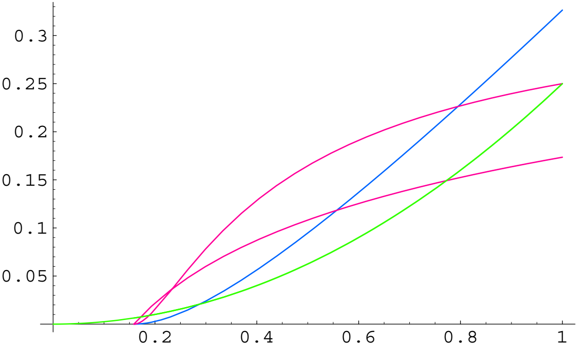

Agents with small α (less than $\alpha^* = 0.1585$) are prevented from participating.

An optimal control formulation for $u(y) = y^{1/k}$

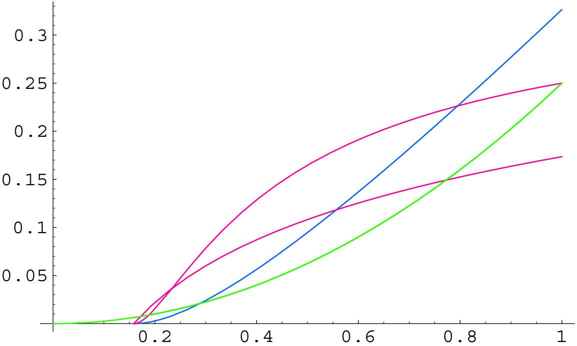

Agents with small α (less than $\alpha^* = 0.1585$) are prevented from participating. Some agents may prefer self-provisioning.

An optimal control formulation for $u(y) = y^{1/k}$

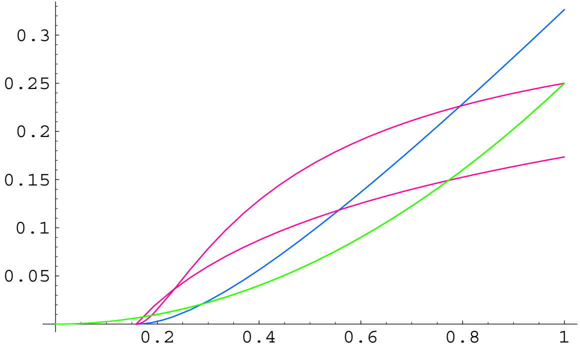

Agents with small α (less than $\alpha^* = 0.1585$) are prevented from participating. Some agents may prefer self-provisioning. Agents with large αs benefit more.

Conclusions

- In most realistic resource allocation problems there is private information to participants.

Conclusions

- In most realistic resource allocation problems there is private information to participants.
- Resource allocation policies need to take acount of incentives. But this is not a simple task!

Conclusions

- In most realistic resource allocation problems there is private information to participants.
- Resource allocation policies need to take acount of incentives. But this is not a simple task!
- Simple egalitarian sharing policies may fail and produce little incentives for participants to contribute resources.

Conclusions

- In most realistic resource allocation problems there is private information to participants.
- Resource allocation policies need to take acount of incentives. But this is not a simple task!
- Simple egalitarian sharing policies may fail and produce little incentives for participants to contribute resources.
- Many new interesting problems!!!

Conclusions

- In most realistic resource allocation problems there is private information to participants.
- Resource allocation policies need to take acount of incentives. But this is not a simple task!
- Simple egalitarian sharing policies may fail and produce little incentives for participants to contribute resources.
- Many new interesting problems!!!
- THANK YOU!!