The Shared Assignment Game

and Applications to Pricing in Cloud Computing

Yoram Bachrach, *Microsoft Research*

Gideon Blocq, *Technion*

Peter Key, *Microsoft Research*

Microsoft

Cloud Computing

Computational Tasks

Have value for task completion

Require resources (Cores, Memory, Bandwidth)

Compete for resources

- How much is a task or resource "worth"
  - Can we use to price & allocate?

Cloud Resources

Public – prices

Private – virtual currency, allocation, "fairness

The Shared Assignment Game (SAG)

- Model
- Full Information Setting
- Private Information Setting
  - Manipulations:
    - Splits
    - Bluffs
- Application to Cloud Pricing

Shared Assignment Game

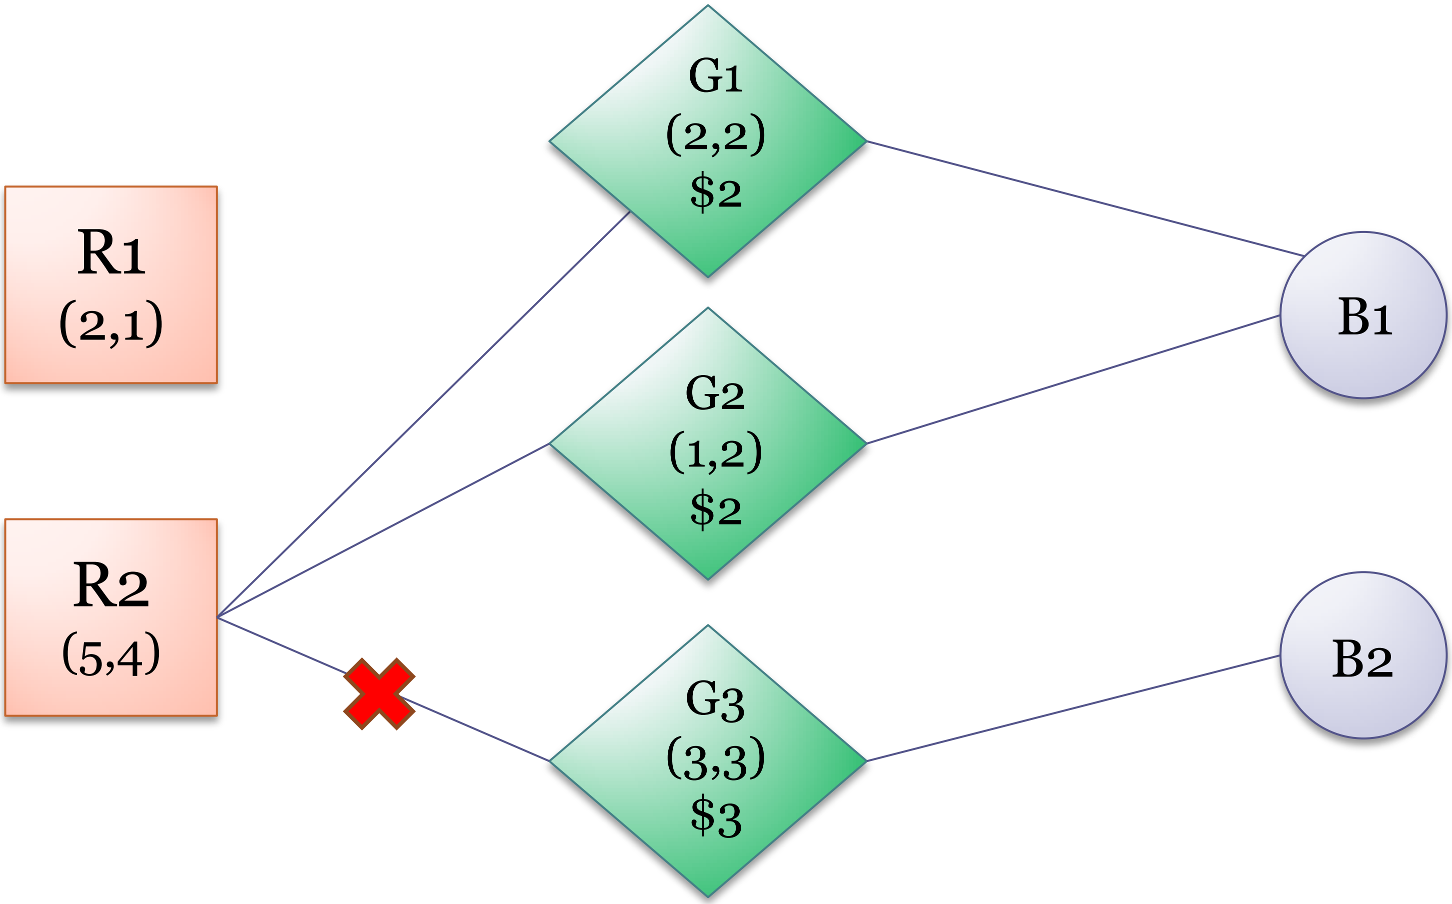

- Goods can be tasks in cloud scenarios or workers in companies.
- How should we assign sellers to buyers? (Optimal?)
- What should the buyers pay for their goods?

The SAG Model

- There are L resource and N agents (B buyers and R sellers).
- Buyer $b,$ for good g has
  - demand $d_g = \left[{d_g}^1, \ldots {d_g}^L\right]$
  - value (worth) $w_g(b)$

- Seller j has
  - capacity vector $C_j = [{C_j}^1, \ldots {C_j}^L]$

- An assignment $f: B \rightarrow R \cup \{\perp\}$ maps buyers to sellers.

The SAG Model

- For coalition $S \subseteq \mathcal{N}$, let A_S be the set of feasible assignments.
- Let $G_b(a_S)$ be the set of goods that buyer b gets at a_S.
- Social Welfare as $W_{a_S} = \sum_{b \in B} \sum_{g \in G_b(a_S)} w_g(b)$.
- Now construct a cooperative game G=($\mathcal{N}$,v):
  - For a coalitions S, $v(S) = \max_{a_S \in A_S} W_{a_S}$.
  - i.e. $v(S)$ is the total worth from optimal matching in coalition $S \subseteq \mathcal{N}$.
- Assumption: $\forall i, k$, there exists a seller j for which ${C_j}^k \geq {d_i}^k$.

How to divide the “spoils”?

- The Core?
  - May not exist for SAG (Lemma 1)
  - Can be unfair 

- We focus on the Shapley value as a fair solution concept.

$$\varphi_i(v) = \sum_{S \subseteq \mathcal{N} \setminus \mathrm{i}} \frac{|S|!\,(N-|S|-1)!}{N!}\,[v(S \cup i) - v(S)]$$

✓ It follows from a set of fairness axioms
✓ Examines contribution of each agent to each coalition.
✓ Sum of total payments is zero,
  ! But some buyers might get paid without receiving any goods.

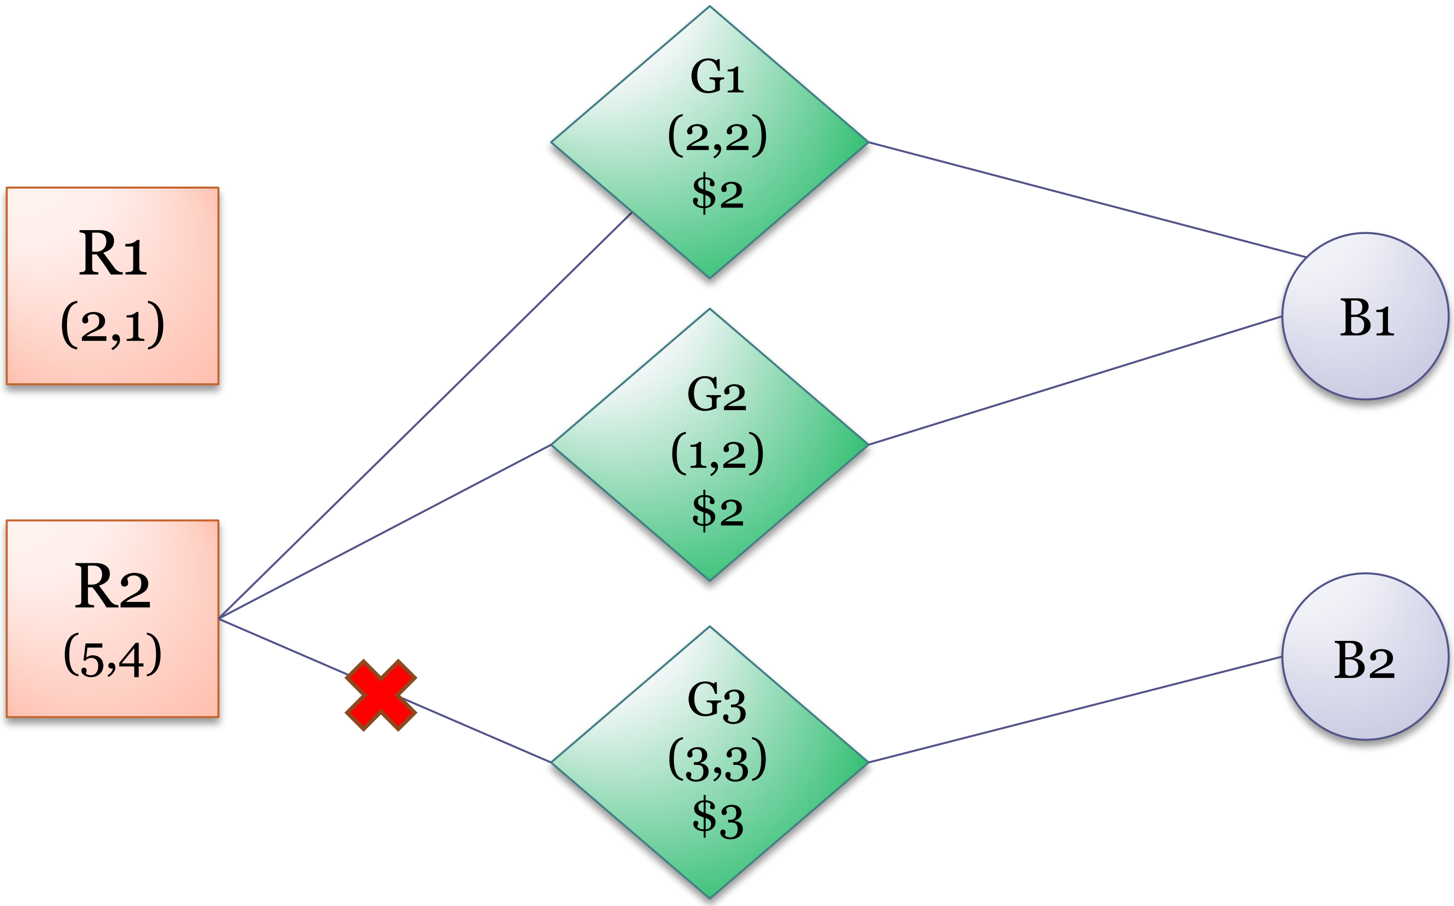

Utility Sellers	Utility Buyers	Discount	Total
$\varphi_{R1} = 0, \varphi_{R2} = \frac{30}{12}$	$\varphi_{B1} = \frac{12}{12}$	$\varphi_{B2} = \frac{6}{12}$	\$2+\$2=\$4

Private Information and Cheating

Worths and demands of the goods may be private information.

For Now, assume the following:

- Each buyer b is interested in a single good,
  - With demand ${d_b}^k$ for resource k
  - Receives its good from a single seller.

We consider two types of manipulations.

1. Splits
2. Bluffs

Splits

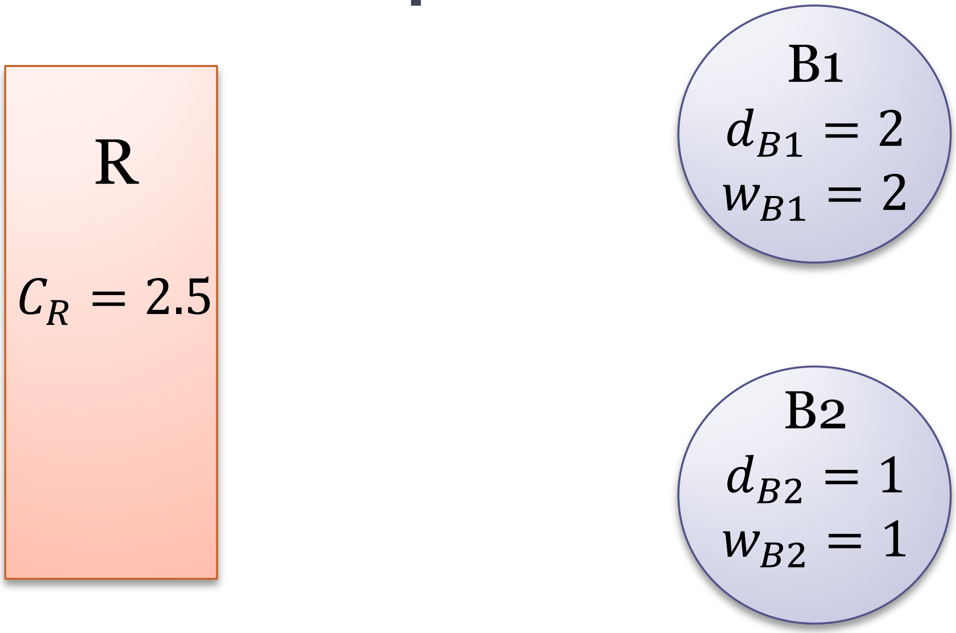

Utility Sellers	Utility Buyers	Discount	Total
$\varphi_R = \frac{7}{6}$	$\varphi_{B1} = \frac{2}{3}$	$\varphi_{B2} = \frac{1}{6}$	\$2

Shapley value is lowered

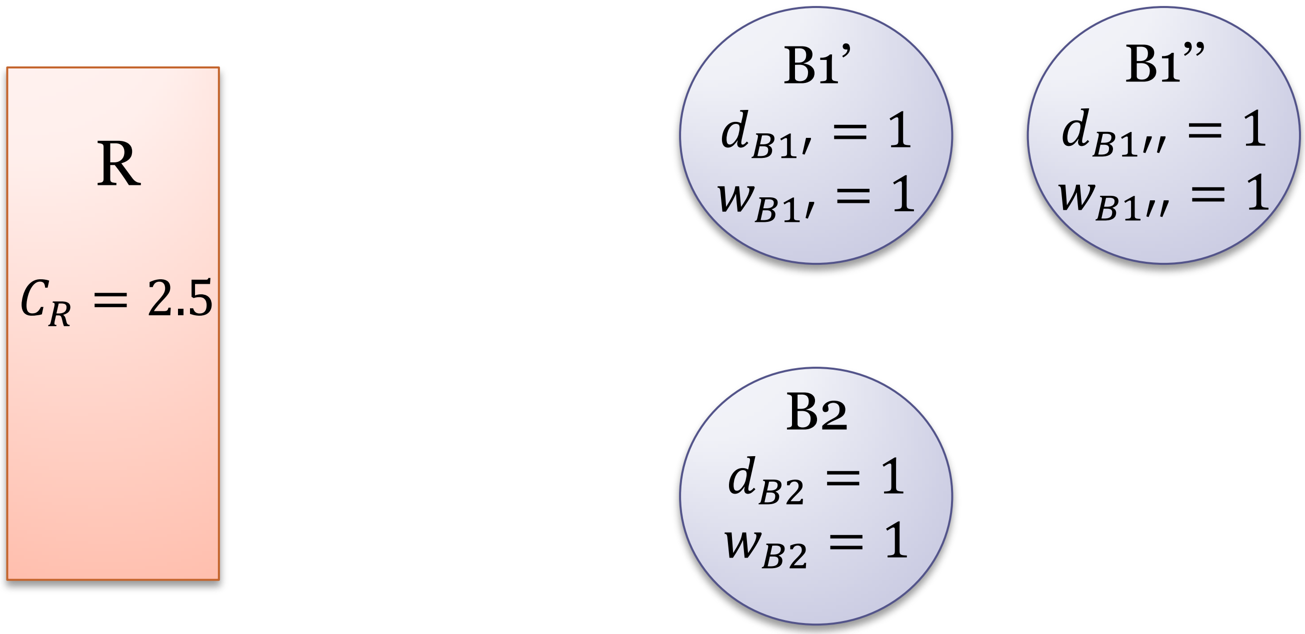

Utility Sellers	Utility Buyers	Discount/Utility	Total
$\varphi_R = \frac{5}{4}$	$\varphi_{B1'} + \varphi_{B1''} = \frac{1}{2}$	$\varphi_{B2} = \frac{1}{4}$	$2

Shapley value is raised

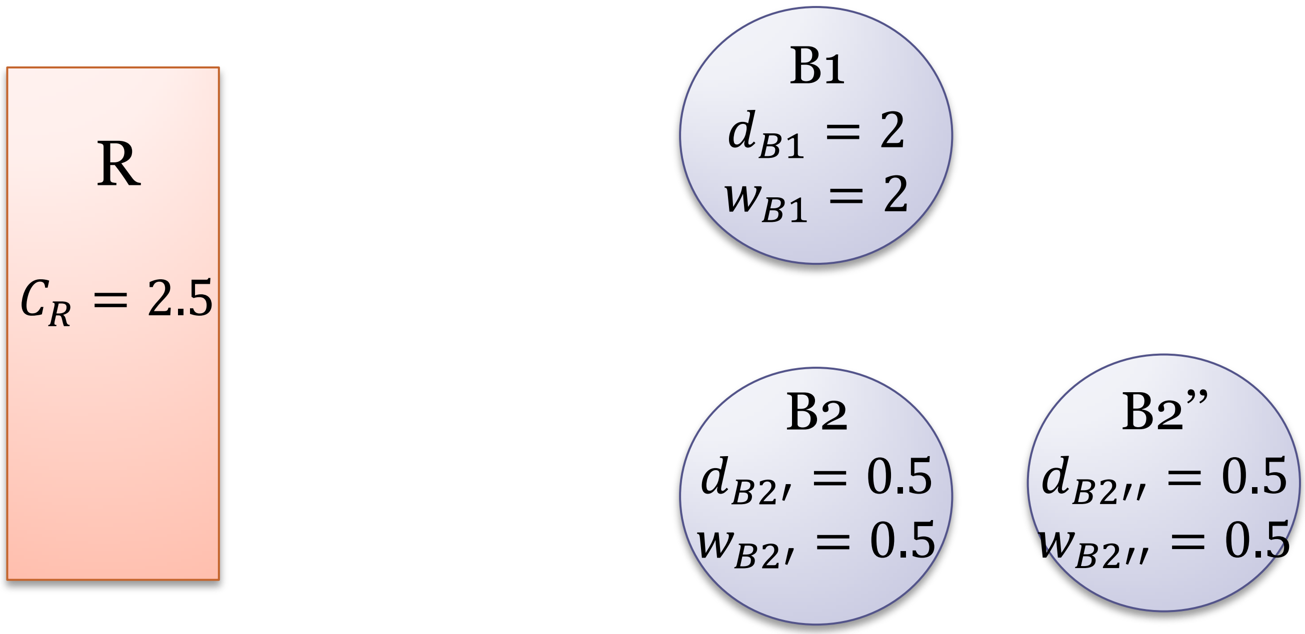

$$\boldsymbol{\varphi_{B2'} + \varphi_{B2''} = \frac{1}{4} > \frac{1}{6}}$$

What are the bounds?

Splits

Have a New game $H = (\mathcal{N}', w)$, where | $\mathcal{N}$'|=N+1.

Define $R_b \triangleq \{resources\ \ with\ enough\ capacity\ for\ b\}$

$$R_b \stackrel{\text{def}}{=} \{r \in R\ |\forall k, C_j{}^k \leq\ d_b{}^k\}$$

Theorem 1

$$\varphi_{b'}(w) + \varphi_{b''}(w) \geq \left[\frac{1}{N} + \frac{N-1}{N(N-|R_b|)}\right]\varphi_b(w)$$

Thus when $|R_b|$=1, the lower bound is equal to $\frac{2}{N}$.

But. cannot execute half a task! Hence both b' and b'' need to be allocated.

⇒the actual lower bound is much lower.

Can find examples where the bounds are very low and tight.

Splits

However an agent also stands to gain from a split:

Theorem 2

$$\varphi_{b'}(w) + \varphi_{b''}(w) \leq \left\lceil \frac{N(N - |R_b|)}{|R_b| + 1} \right\rceil \varphi_b(w)$$

This bound is tight.

Summary:

- Can gain a lot from splitting (Thm 2)
- But takes a risk (Theorem 1) it also takes a risk by doing so.
- A small change in the system or a wrong estimation may severely lower the agent's utility.

Bluff

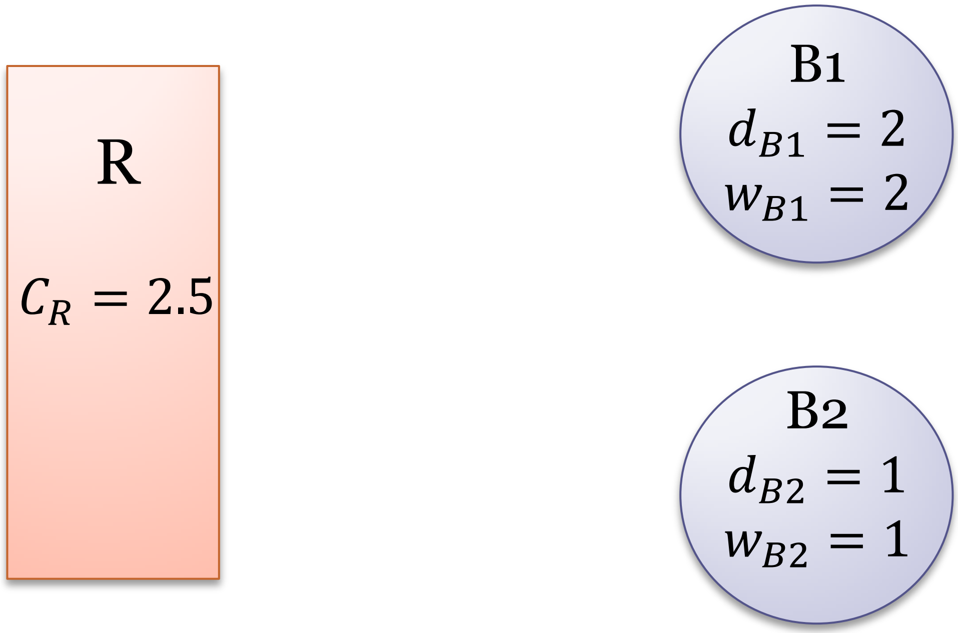

If B2 knows that B1 has a higher valuation, it can steal a 'discount' from the system.

Utility Sellers	Utility Buyers	Discount	Total
$\varphi_R = \frac{7}{6}$	$\varphi_{B1} = \frac{2}{3}$	$\varphi_{B2} = \frac{1}{6}$	$2

Bluff

Define the *lying ratio (LR),* the ratio between the Shapley value of a *unmatched* buyer b, and the maximum social welfare.

Thus, LR $= \frac{\varphi_b(v)}{v(N)}$, where b is not matched with any seller.

Theorem 3:

$$\text{LR} = \frac{\varphi_b(v)}{v(N)} \leq \frac{1}{|R_b|+1} - \frac{1}{N|R_b|}$$

- Assumes buyer having full information on the other buyers.
- Bluffing not robust to mistakes
- Moreover, when $|R_b|$ is large, LR is small

Applications to Cloud Pricing

- In current cloud pricing, you cannot specify your exact bundle per task
  - Pay for unused resources

The SAG applied to cloud pricing:

- Clients specify the size and worth of their tasks.
- Each Tasks represents a ‘good’
- Clients represent the buyers.
- Sellers represent the resources in the cloud
- The SAG optimally executes the available tasks and fairly distributes the revenue according to the Shapley value.

Simulation of the SAG vs the Fixed Pricing Scheme (FPS)

- Simulate 50 tasks with demands for two resource types
- All demands and capacities are i.i.d. Gaussian.
- Worth of task b, $w_b = \max\{ {d_b}^1, {d_b}^2\} + \alpha$, where α is offset.

Simulation of the SAG vs the Fixed Pricing Scheme (FPS)

- SAG:
- Use a greedy approximation to compute the Multi-Dimensional-Multiple-Knapsack (MDMK) problem.
- Approximate the Shapley value by sampling
- Social Welfare is sum of worths of assigned tasks
- Revenue is sum of Shapley value of servers

- FPS
- Uses a fixed, revenue optimal price
- Social value of the FPS : sum of worths of the tasks above the threshold.
- The revenue of the FPS : sum of worths of executed tasks $\times$ the fixed price.

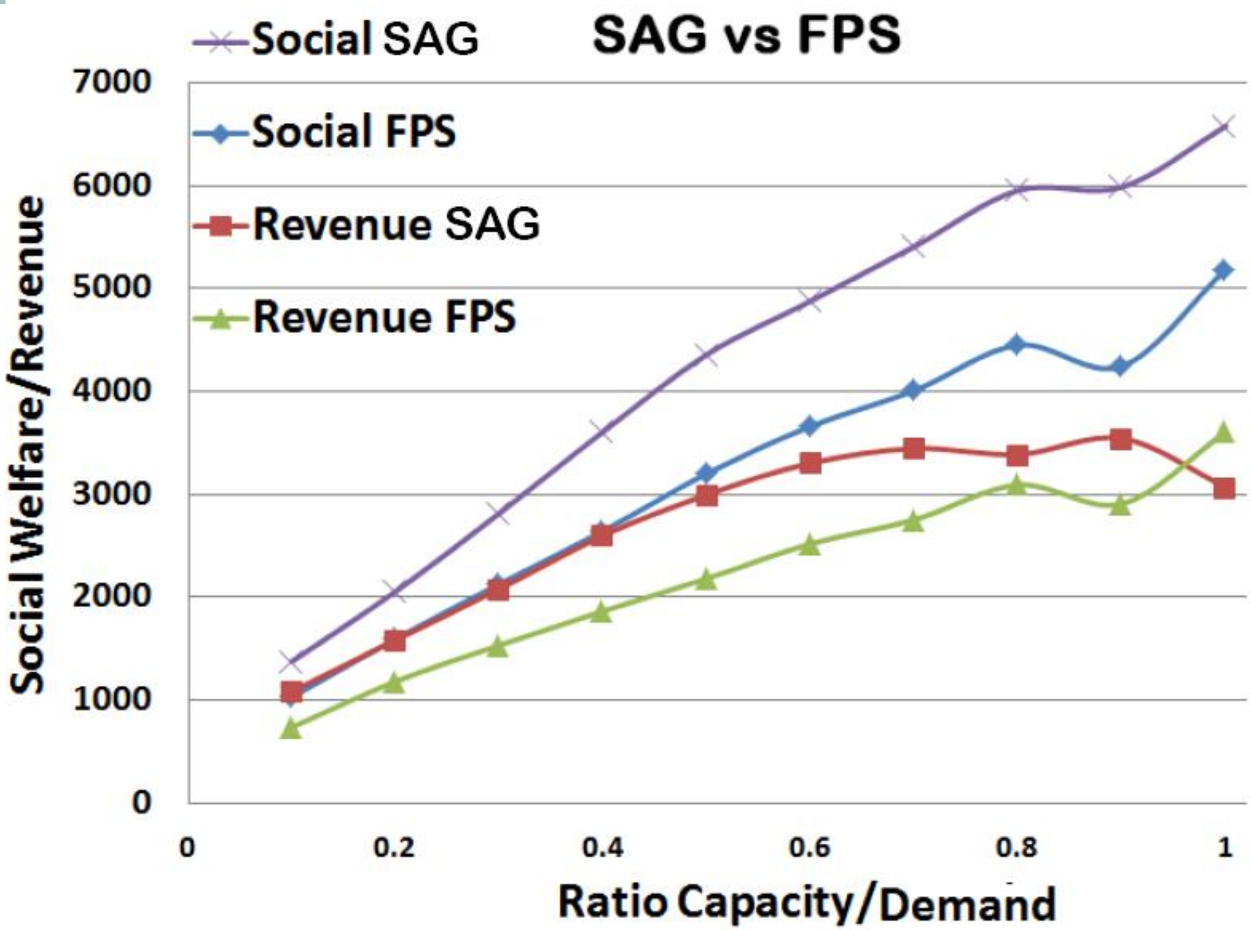

The revenue under the SAG is higher than the FPS
For a high demand ratio the two converge.

Conclusions

- Introduced an extension of the Assignment game, where sellers have resource constraints.

- Core can be empty
- Shapley value can fairly calculate the division of utility.

- Manipulation (Splits and Bluffs) when private information
  - Upper and lower bounds of their impact on the Shapley value.
  - Risk/reward trade-off for buyer

- Simulated SAG against the FPS and showed that it increases revenue.