Efficiency Loss in a Resource Allocation Game: Random Users

Shie Mannor

McGill University

(with Jia-Yuan Yu)

February 2006

Agenda

Allocation of divisible resources
(e.g., bandwidth in a network)

Simple “mechanisms”
(based on market clearing) - a reminder

Efficiency loss, w.r.t. social optimum
(in the presence of selfish users) - a reminder

What happens when there are random users
(based on medium-size sample)

Outline

1. The basic mechanism - single link and fixed capacity
2. Network case, fixed capacities
3. Elastic capacities (just briefly)
4. Random users

PART I: Single Link, Fixed Capacity

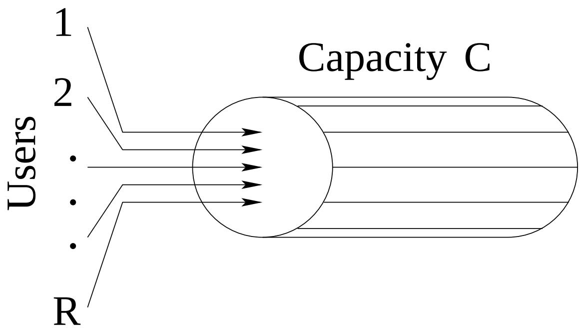

Rate $d_r \quad \longrightarrow \quad$ utility $U_r(d_r)$

U_r: concave, strictly increasing, nonnegative

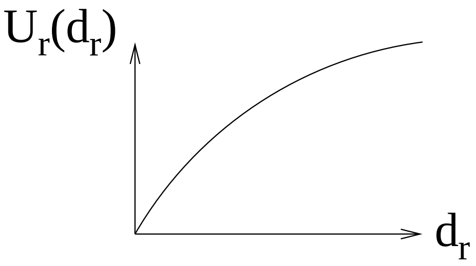

The Social Optimum

$$\begin{aligned} \text{maximize} \quad & \sum_r U_r(d_r) \\ \text{subject to} \quad & \sum_r d_r \le C \\ & \boldsymbol{d} \ge 0 \end{aligned}$$

An “easy” problem.

The Shubik Pricing Mechanism

User r submits a bid w_r.

Receives bandwidth: $\frac{w_r}{w_1 + \cdots + w_R} C$

Example:

$$w_1 = 2 \qquad d_1 = 2C/5$$
$$w_2 = 3 \qquad d_2 = 3C/5$$

All bandwidth is allocated

Unit price of bandwidth: $\mu = \frac{w_1 + \cdots + w_R}{C}$

$$d_r = \frac{w_r}{\mu}$$

Users as Price Takers

Given price μ, user r solves:

$$\max_{w_r \geq 0} \; U_r\left(\frac{w_r}{\mu}\right) - w_r$$

Theorem 1 (Existence of Competitive Equilibrium; Kelly, 1997)
There exist $\boldsymbol{w}$ and μ such that:

(a) w_r *is optimal for user* r *given the price* μ*.*

(b) $(w_1 + \cdots + w_R)/\mu = C$*.*

The resulting allocation is socially optimal.

Users as Price Anticipators

Suppose users *know* the price setting procedure.

Given $(w_s, s \neq r)$, user r solves:

$$\max_{w_r \geq 0} \; U_r\left(\frac{w_r}{w_r + \sum_{s \neq r} w_s} C\right) - w_r$$

This is now a *game*, where the strategy of user r is the bid w_r.

Nash equilibrium exists

Example

$C = 1, \;\; U_1(d_1) = 2d_1, \;\; U_2(d_2) = d_2.$

Social optimum: $d_1 = 1, d_2 = 0$

Price-taking equilibria: $\mu = 1$

Price-anticipating users:

(a) $\mu > 1 \implies w_2 = 0 \implies w_1 = \;?$

(b) $\mu = 1, w_2 > 0$; user 2 will reduce w_2, and reduce the price

(c) $\mu < 1, w_1 > 0, w_2 > 0$: inefficient

Nash Equilibrium

Theorem 2 (Hajek & Gopalakarishnan, 2002) *Assume* $R > 1$.

There exists a unique Nash equilibrium w.

The resulting allocations d_r *are the unique socially optimal solution for modified utilities:*

$$\hat{U}_r(d_r) = \left(1 - \frac{d_r}{C}\right) U_r(d_r) + \left(\frac{d_r}{C}\right) \left(\frac{1}{d_r} \int_0^{d_r} U_r(z)\, dz\right)$$

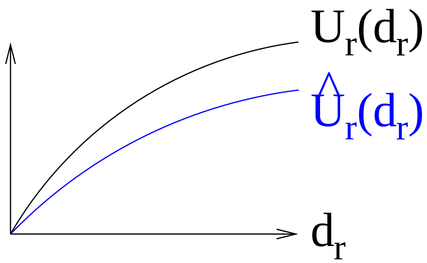

Efficiency Loss

Theorem 3 (Johari and Tsitsiklis, 2004) *The efficiency loss is no more than 25%:*

$$\begin{aligned} \sum_r U_r(d_r^G) &\geq \frac{3}{4} \sum_r U_r(d_r^S) \\ \texttt{(Nash eq. utility)} &\geq \frac{3}{4} \times \texttt{(socially optimal utility)} \end{aligned}$$

Furthermore, this bound is tight.

Worst case:

Many users, linear utility functions, one “dominant” user

The Worst Case

$$C = 1 \qquad\qquad d_1^G \to 1/2$$

$$U_1(d_1) = d_1 \qquad \Longrightarrow \qquad \sum_{r>1} d_r^G \to 1/2$$

$$U_r(d_r) \approx d_r/2 \qquad\qquad \mu \to 1/2$$

$$R \to \infty$$

$$\sum_r U_r(d_r^S) = 1 \qquad\qquad \sum_r U_r(d_r^G) \to 3/4$$

Main question: Is this a pathological case?

PART II: Networks

Link j has capacity C_j
Each user is identified with a path

Social optimum:

$$\begin{array}{rl} \text{maximize} & \sum_r U_r(d_r) \\ \text{subject to} & \text{capacity constraints} \end{array}$$

The Pricing Mechanism

User r submits a bid w_r^j, at each link j

Receives bandwidth at that link: $x_r^j(\boldsymbol{w}) = \frac{w_r^j}{w_1^j + \cdots + w_R^j} C_j$

User r sends as much as possible: $d_r = \min_{j \in \text{path}} x_r^j$

User r payoff: $U_r(d_r(\boldsymbol{x}_r(\boldsymbol{w}))) - \sum_j w_r^j$ Concave!

Example

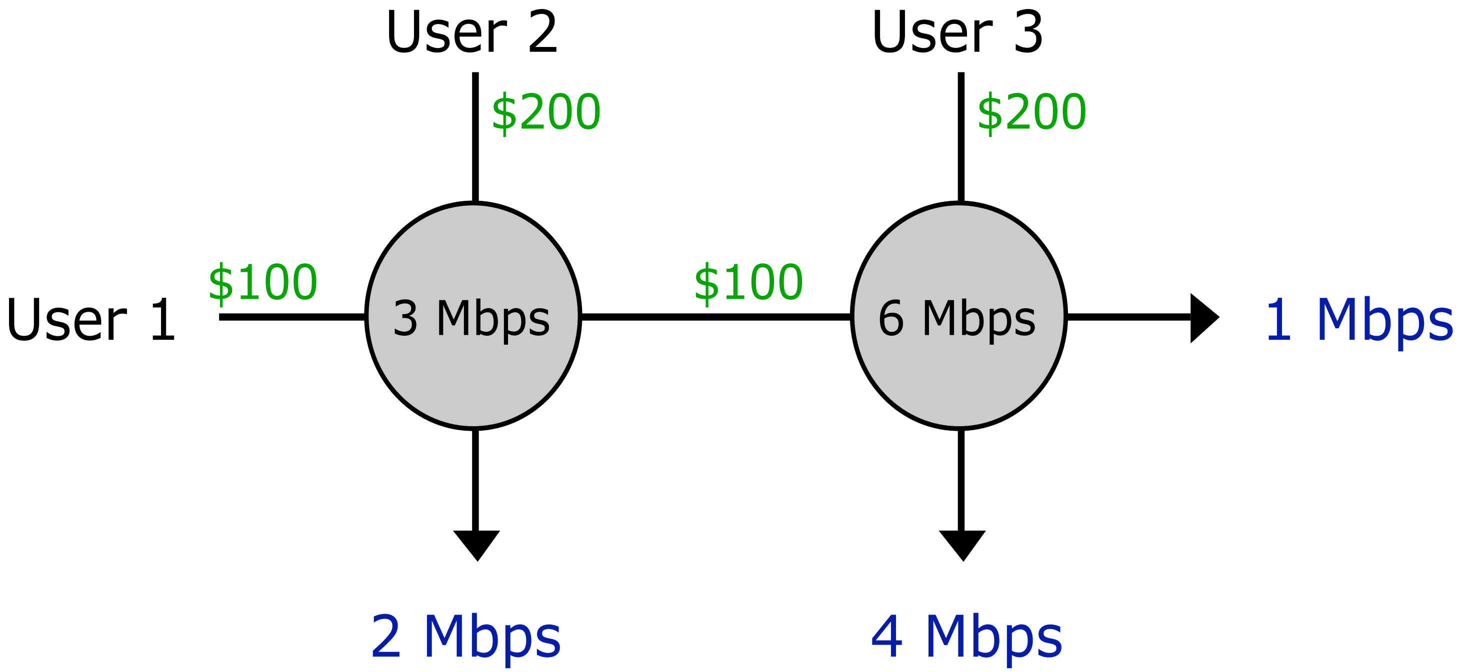
User 2
User 3
$200
$200
User 1
$100
3 Mbps
$100
6 Mbps
1 Mbps
2 Mbps
4 Mbps

Existence of a Nash Equilibrium

Theorem 4 *With “price-anticipating” users,*
there exists a Nash equilibrium $(\boldsymbol{w})$ *for an extended game.*

NE of original game maps to a NE of the extended game

Proof idea:

– Introduce “virtual user” bidding $\epsilon > 0$ at each j.
– With $\epsilon > 0$, no discontinuities, NE exists (Rosen’s theorem)
– Take the limit as $\epsilon \longrightarrow 0$
– Perturbed game allocations $x_r^j(\epsilon)$ in the limit

General Networks: Efficiency Loss

Theorem 5 (Johari and Tsitsiklis, 2004) *The efficiency loss is no more than 25%:*

$$\texttt{(Nash eq. utility)} \geq \frac{3}{4} \times \texttt{(socially optimal utility)}$$

Proof idea:

Reduce to analyzing multiple single link games, one at each $j \in J$.

PART III: Elastic Capacities

Supplying f units incurs cost \$$C(f)$.

$C(f)$ strictly increasing, strictly convex, nonnegative

For user r, consuming d_r units yields utility \$$U_r(d_r)$.

$U_r(d_r)$ strictly increasing, concave, nonnegative

An efficient allocation maximizes net monetary benefit:

$$\sum_r U_r(d_r) - C\left(\sum_r d_r\right)$$

A Pricing Mechanism

Define the *price function* $p(f) = C'(f)$.

$p(f)$ strictly increasing, convex, nonnegative

We use the following mechanism (Kelly, Maulloo, and Tan, 1998):

(1) User r submits a bid w_r.

(2) The mechanism chooses $f(\boldsymbol{w})$ to “clear the market”:

$$\begin{array}{ccccc} \sum_r w_r & = & f(\boldsymbol{w}) & \times & p(f(\boldsymbol{w})) \\ \text{Revenue} & & \text{Quantity} & \times & \text{Price} \end{array}$$

User r’s rate: $d_r = w_r / p(f(\boldsymbol{w}))$

User r’s payoff: $U_r(d_r) - w_r$

Price Anticipating Users

Suppose users anticipate the effect of their bid on $p(f(\boldsymbol{w}))$.

Theorem 6

There exists a Nash equilibrium $\boldsymbol{w}$*, and it is unique if* p *is differentiable.*

Proof idea:

Show that d_r is concave in w_r, for each user r.
Apply Rosen's existence theorem.

Price Anticipating Users: Efficiency Loss

Theorem 7

The efficiency loss is no more than ≈ 34%:

$$\begin{array}{c}\textit{Net monetary benefit}\\ \textit{at Nash equilibrium}\end{array} \geq (4\sqrt{2}-5) \times \begin{array}{c}\textit{Net monetary benefit}\\ \textit{at efficient allocation}\end{array}$$

Furthermore, this bound is tight.

Worst case:

Many users, linear utility functions, piecewise linear price function.

Back to Congestion Control

Faster time scale

Fixed bids w: network protocol allocates rates

Slower time scale

User observes link prices and adjusts bids

Implicit in our model:

(a) Observe individual link prices

(b) Can anticipate the effect of a bid change on prices

(c) Do not anticipate changes (reactions) in bids of other users

Main question: How much loss should we expect?

Part IV: Random Users

25% and 34% loss - good news or bad news?

How much efficiency loss should we expect?

Modeling random users - Heterogeneity.

What is the typical loss for medium (non-asymptotic) problems?

Modeling a Heterogeneous Population of Users

- Number of users is n.
- As n increases, more users enter while other users stay.
- Their utility functions are drawn i.i.d. from a set of utility functions Ω.
- We use capital U_i to denote random utility functions.
- Let v be the supremum among all $u_i'(0)$, i.e., $v = \sup_{u \in \Omega}\{u'(0)\}$.

Assumption 1 *The probability measure satisfies*

1. $v < \infty$
2. *if* $v = \infty$, *then*

$$\mathsf{Pr}(U_i'(0) = \infty) > \delta$$

or

$$\sup_{u \in \Omega} u(C) < \infty.$$

Examples Modulation satisfying Assumption 1:

- $U_i(x) = S_i\, x$, where S_i has bounded support.
- $U_i(x) = S_i\, \log(1+x)$.
- $U_i(x) = S_i\, \sqrt{x}$, (S_i need not have “bounded” support since all slopes at 0 are infinite).

Remark: We allow S_i to take a finite number of values.

Main Results

Consider the random quantity: $\text{LOE} = 1 - \frac{\sum_r U_r(d_r^G)}{\sum_r U_r(d_r^S)}$

- $\text{LOE}_n \longrightarrow 0$ (a.s.) for single link and inelastic supply (constant C).
- Convergence rate (nearly exponential).
- $\text{LOE}_n \not\rightarrow 0$ for single link and inelastic supply if $U_i(x) = S_i\, x$ for S_i heavy tailed.
- $\text{LOE}_n \longrightarrow 0$ for single link and elastic supply.
- $\text{LOE}_n \longrightarrow 0$ for networks with inelastic supply.

Single Link and Inelastic Supply

Theorem 8 (Convergence of loss of efficiency) *With Assumption 1, the loss of efficiency tends to* 0 *with probability* 1 *as the number of users tends to infinity.*

Intuition

- Assumption 1 guarantees that there is always a significant number of users that value the resource highly.
- Since the amount of resource C is fixed, with enough users with high valuation, the influence of each individual on the market outcome becomes limited.
- We end up with perfect competition and high efficiency.

Proof Outline

Hajek and Gopalakrishnan: the Nash equilibrium allocation is the solution of

$$\begin{aligned} &\textbf{GAME} \\ \max_{y_1,\ldots,y_n} \quad & \sum_{i=1}^{n} \left[\left(1 - \frac{y_i}{C}\right) U_i(y_i) + \frac{1}{C} \int_0^{y_i} U_i(z)\, \mathrm{d}z \right] \\ \text{subject to} \quad & \sum_{i=1}^{n} y_i \le C, \\ & y_i \ge 0, \quad i = 1, \ldots, n. \end{aligned} \tag{1}$$

Notice similarity with SYSTEM.

Proof for finite v.

Complementary slackness conditions for social optimum:
$u_i'(x_i^{(n)}) = \lambda^{(n)}$ for all active users (i such that $x_i^{(n)} > 0$),
$u_i'(0) \leq \lambda^{(n)}$ for all inactive users ($x_i^{(n)} = 0$),
$\sum_{i=1}^{n} x_i^{(n)} = C$,

Complementary slackness conditions for Nash equilibrium:

$\left(1 - \frac{y_i^{(n)}}{C}\right) u_i'(y_i^{(n)}) = \mu^{(n)}$ for all active users (i such that $y_i^{(n)} > 0$),
$u_i'(0) \leq \mu^{(n)}$ for all inactive users,
$\sum_{i=1}^{n} y_i^{(n)} = C$.

- The sequence $\mu^{(n)}$ is monotone increasing because adding users can only increase the price. Moreover, $\mu^{(n)} \leq v < \infty$, so $\lim_{n\to\infty} \mu^{(n)}$ exists.

- Suppose $\mu^{(n)} < v - \epsilon$,

  $\Rightarrow$ infinite number of active users as $n \to \infty$,

  $\Rightarrow$ active users receive non-negligible y_i,

  – but, C is finite! $\Rightarrow$ Contradiction.

$\Rightarrow$ $\mu^{(n)} \to v$ and $y_i \to 0$,

$\Rightarrow$ the optimality conditions coincide in the limit $\Rightarrow$ LOE $\to$ 0.

Uniformly Modulated Users

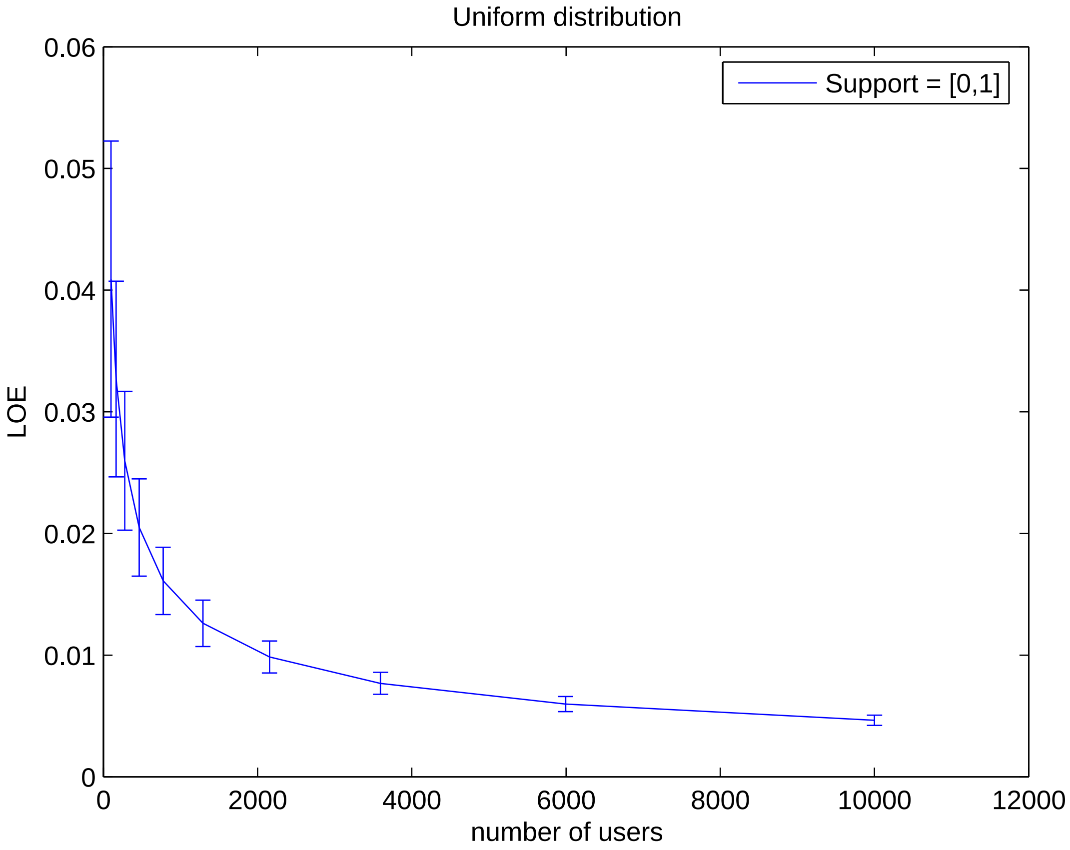

Rates of Convergence

A tighter argument leads to:

Theorem 9 *Consider a resource allocation game with random linear utility functions* $V_i(x) = S_i\, x$ *with* $v < \infty$*. For some fixed* γ *and* ϵ *such that* $0 < \gamma < \epsilon < u$*, let* $\xi \triangleq \Pr(S_i \geq u - \gamma)$ *and* $n \geq \lfloor (u-\gamma)/(\epsilon-\gamma) \rfloor / \xi$*, then*

$$\Pr\left(\text{LOE}_n > \frac{\epsilon}{u}\right) \leq \exp\left\{-2\,\frac{(n\xi - \lfloor (u-\gamma)/(\epsilon-\gamma)\rfloor)^2}{n}\right\}.$$

Main tool: Hoeffding inequality and bounding the “bad” events tightly.

Increasing Capacity

What if we allow the capacity to increase as the number of users increases? (Assume all utility functions are bounded.)

- Rate of convergence result holds for sub-linearly scaled capacity $C(n) \in o(n)$.
  Active users still tend to receive a small fraction of the scarce resource $c(n)$.

- Almost sure convergence result holds for super-linearly scaled capacity $C(n) \in \omega(n)$.
  Abundant resource; every user receives enough resource to saturate its utility function.

Increasing Capacity - the Linear Case

If $C(n) = c \cdot n$ we still get $\text{LOE}_n \to 0$.
Proof methodology:

1. Result holds if Ω is a finite set.
2. Show that there is continuity in the LOE: users should be “close” according to the metric:
$$d(U_i, U_j) = \sup_{x \geq 0} \left| U_i'(x) - U_j'(x) \right|, \quad U_i, U_j \in \Omega.$$
3. Prove that under our assumptions, Ω can be covered by finitely many functions.

Divergence

If sampling mechanism violates Assumption 1, we get divergence. In particular,

$$U_i(x) = S_i x,$$

with S Pareto ($P(x) = ab^a / x^{a+1}$). Leads to positive LOE in expectation.

Implication: if “hungry” users are likely, this leads to efficiency loss.

Result holds for all heavy-tailed distributions.

Proof Methodology

1. Bound the loss of efficiency using order statistics (largest/second largest).

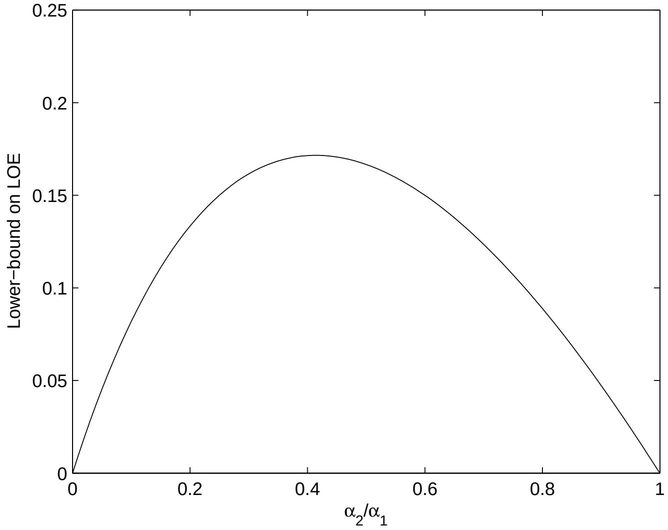

2. Show that for heavy tailed distribution $\Pr(\{\text{largest}/\text{second largest} \in \{\epsilon, 1-\epsilon\}\}) > 0$.

Pareto Modulated Users

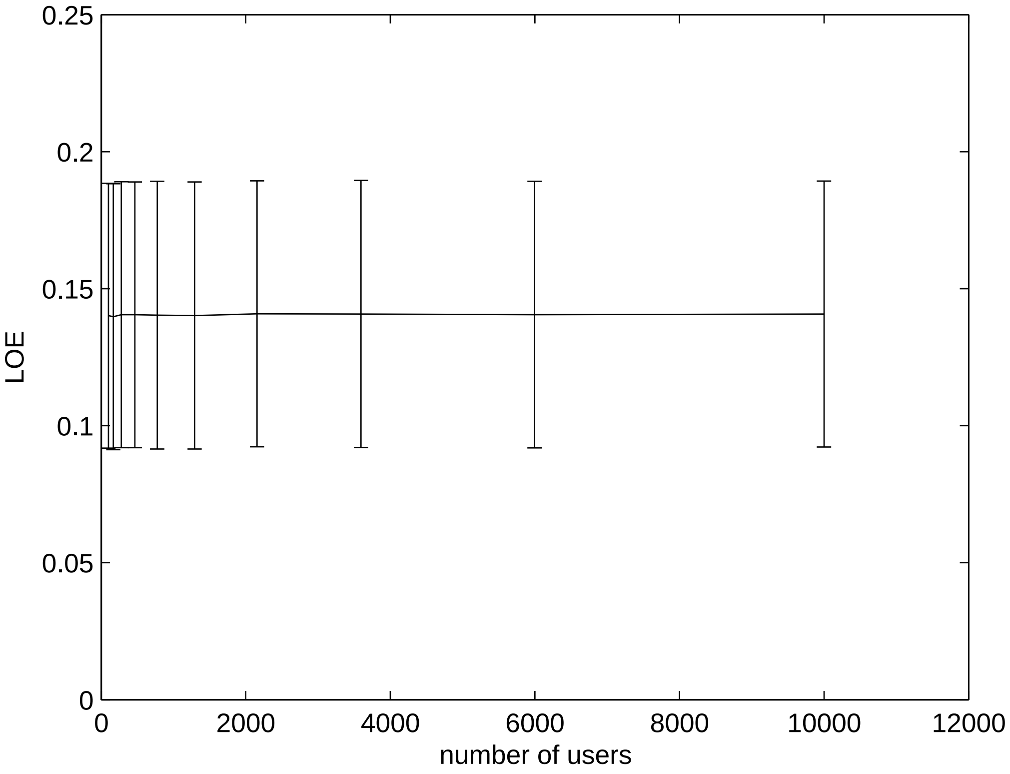

Networks

Modeling random users in a network:

1. A source-destination pair is sampled (IID).
2. For every S-D pair there is a probability measure on utility functions.

Questions:

1. Is there efficiency loss?
2. What will the network converge to?

LOE $\longrightarrow 0$ a.s.

The social case: Solve problem with maximal slope between every legit S-D pair

The game: The price becomes the maximal price guaranteeing only “optimal” S-D pairs participate.

A New Take on an Old Problem

Perfect Competition (Cournot): pure exchange economics

Dubey et al.(1980) - continuum economics model is efficient

Green (1980) - replica markets converge to efficient outcome

Our work:

1. Results for small populations
2. Heterogeneity
3. Increasing commodity
4. Probabilistic analysis

Interesting Directions

Dynamics of convergence to equilibrium (best-response and beyond)

The effect of non-myopic users (in terms of efficiency loss)

The effect of users that abandon the game