Efficiency Loss in Competitive Market Mechanisms: Asymptotics and Dynamics

Shie Mannor
McGill University

January 2007

Agenda

The big picture: control, games, and learning

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

Simple “mechanisms” based on market clearing

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

What happens usually?

Dynamics

Outline

1. Context - distributed control, games, and learning
2. The basic mechanism - single link and fixed capacity
3. Network case, fixed capacities
4. Random users
5. Dynamics

PART I: Context

Classical control:

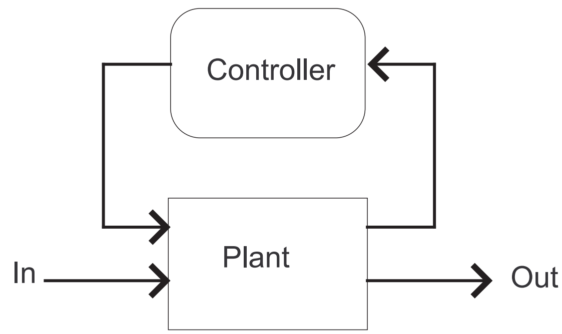

Basic tool: optimization

Sample models: Kalman filters, Markov decision processes, LQG, etc.

Focus: Cost minimization or stabilization

Multi-agent control:

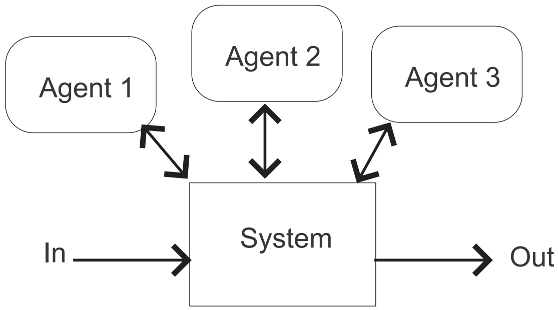

Examples: communication networks, UAV teams, deregulated power networks

- Distributed VS selfish
- Communication/information limitations
- Complexity issues

Game theory:

- Provides conceptual framework
- Self-optimizing / Selfish
- Equilibrium concepts: an agent cannot benefit by changing its action
- Try to enforce “social” behavior (taxes, incentives etc.)

Learning:

- Goal: learn a policy
- Reinforcement learning
- Minimize coordination/coupling
- Reduce “structural” complexity - self organization

PART II: Single Link, Fixed Capacity

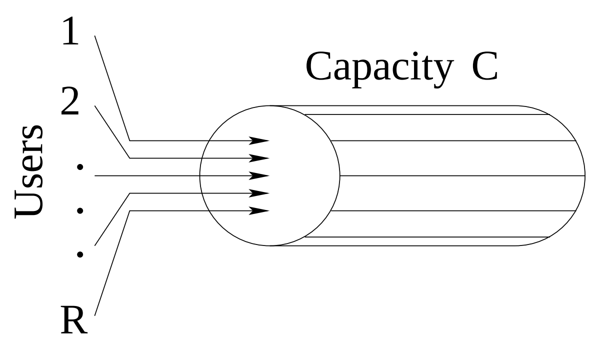

Rate d_r $\longrightarrow$ utility $U_r(d_r)$

U_r: concave, strictly increasing, nonnegative

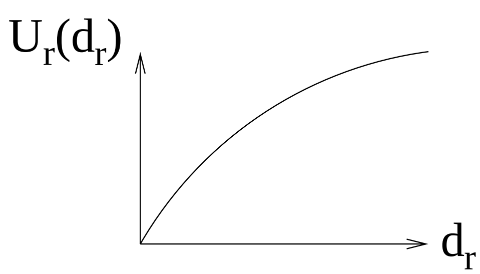

The Social Optimum

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

An “easy” problem.

But a lot of knowledge is needed.

The Kelly-Shubik Pricing Mechanism

User r submits a bid w_r (payment).

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

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: $\mu = \dfrac{w_1 + \cdots + w_R}{C}$

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

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:

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

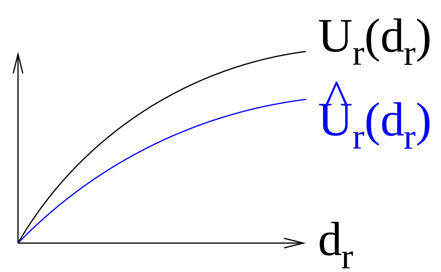

Efficiency Loss

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

$$\sum_r U_r(d_r^G) \geq \frac{3}{4}\sum_r U_r(d_r^S)$$

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

Furthermore, this bound is tight.

Worst case:
Many users, linear utility functions, one “dominant” user

Main question: How likely is a “bad” case?

PART III: 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

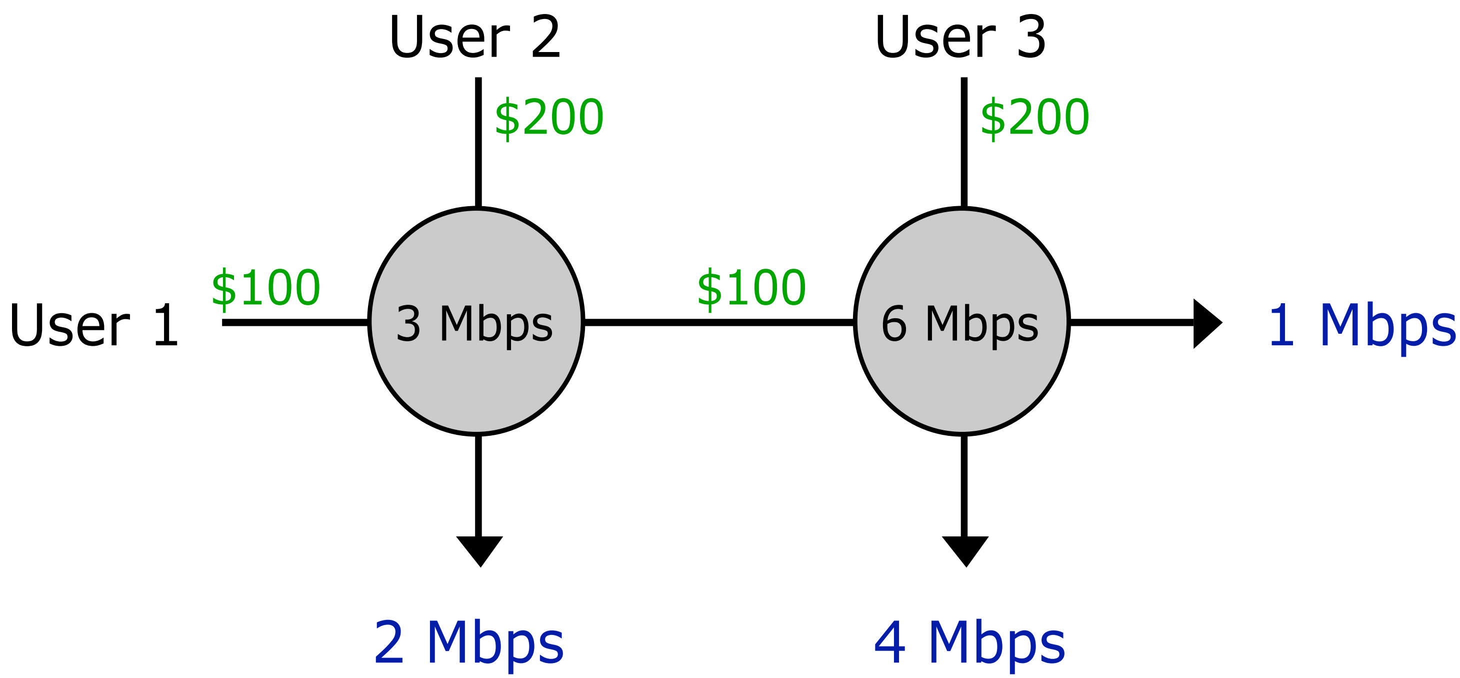

User 2
User 3
$200
$200
User 1
$100
$100
3 Mbps
6 Mbps
1 Mbps
2 Mbps
4 Mbps

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 "free bids for empty links".
– With no empty links, no discontinuities, NE exists (Rosen's theorem)
– With empty links, a NE still exists

General Networks: Efficiency Loss

Theorem 5 *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 $j \in J$.

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:

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

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":

$$\sum_r w_r \quad = \quad f(\boldsymbol{w}) \quad \times \quad p(f(\boldsymbol{w}))$$

Revenue = Quantity × Price

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: Efficiency Loss

Theorem 6 (Johari Mannor and Tsitsiklis)

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:

Users observe individual link prices

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.

Assumption 1 *The probability measure satisfies*

1. *Bounded slope at 0 for all utility functions,*

2. *Infinite slope at 0 with positive probability,*

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

   or

3. *bounded utility*

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

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).

Main Results

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

- $\text{LOE}_n \to 0$ (a.s.) for:
  - single link and inelastic supply (constant C),
  - single link and elastic supply, and
  - for networks with inelastic supply.

- Convergence rate (nearly exponential).

- $\text{LOE}_n \not\to 0$ for single link and inelastic supply if $U_i(x) = S_i\, x$ for S_i heavy tailed.

[Results appear in Yu and Mannor 06 and 07.]

Single Link and Inelastic Supply

Theorem 7 (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.

Uniformly Modulated Users

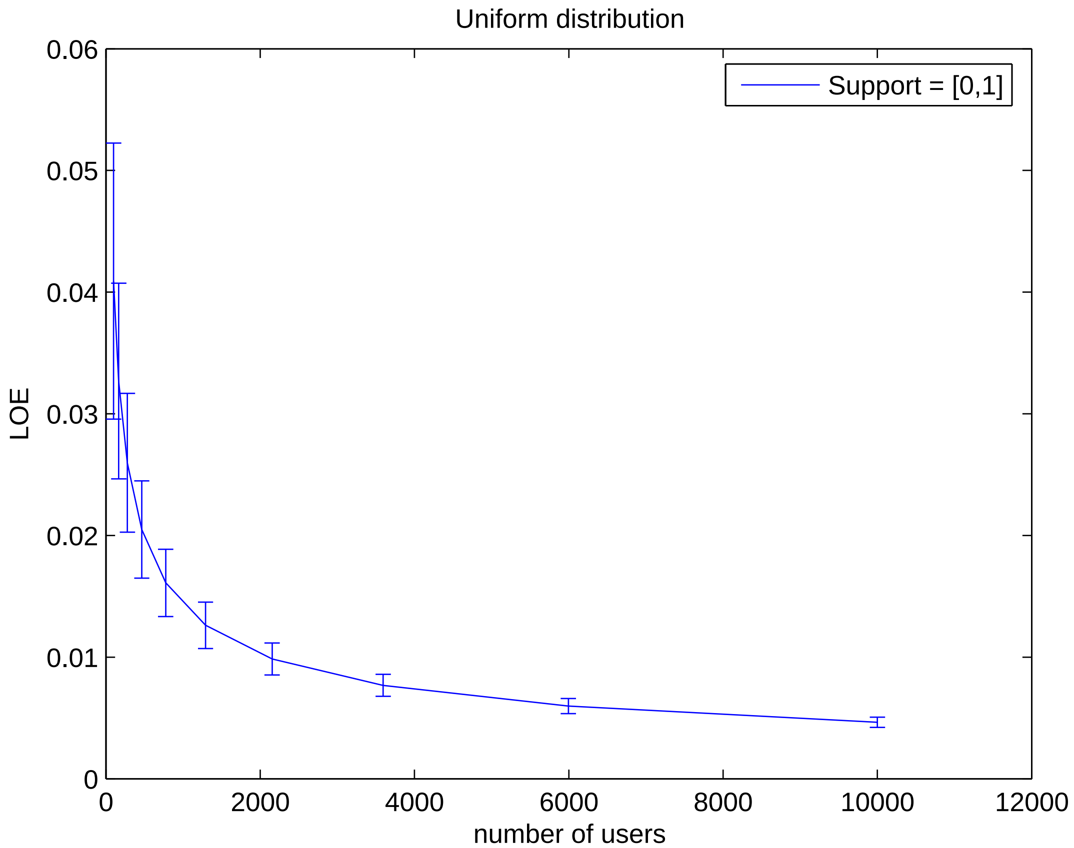

Rates of Convergence

A tighter argument leads to:

Theorem 8 *Consider a resource allocation game with random linear utility functions* $U_i(x) = S_i\, x$ *with finite slope. Then for every* $\epsilon > 0$ *there exists* C *and* t *such that*

$$\Pr\left(\text{LOE}_n > \epsilon\right) \leq Ce^{-nt}$$

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.)

- LOE $\longrightarrow 0$ for sub-linearly scaled capacity $C(n) \in o(n)$.
  Active users still tend to receive a small fraction of the scarce resource $c(n)$.

- LOE$_n \longrightarrow 0$. 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 $\text{LOE}_n \to 0$.
A "machine learning" proof:

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).

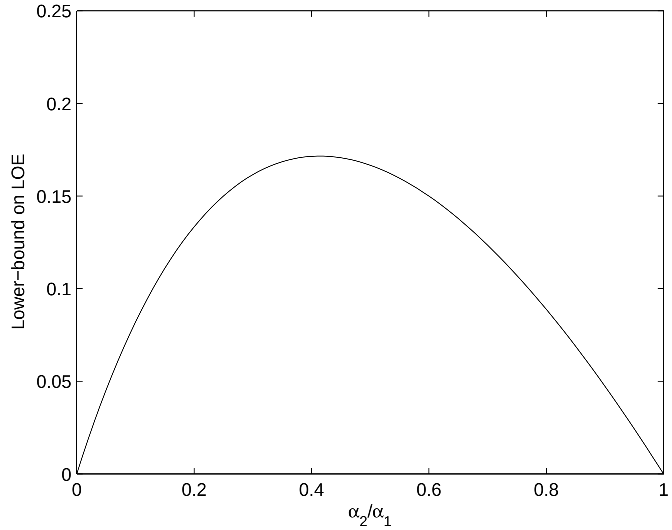

2. Show that for heavy tailed distribution $\Pr(\{\text{largest}/\text{second largest} \in \{\epsilon, 1-\epsilon\}\}) > 0$.

Pareto Modulated Users

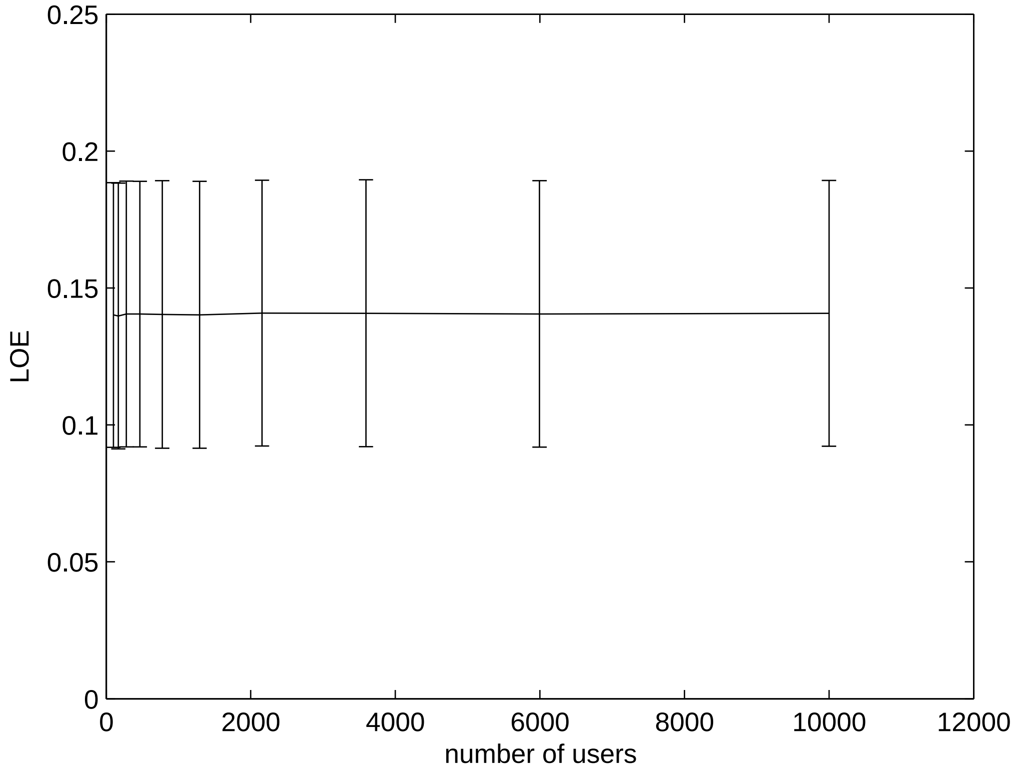

Networks

Modeling random users in a network:

1. Source-destination pairs are sampled (IID) from finitely many types.

2. For every S-D pair there is a probability measure on utility functions.

Questions:

1. Will the network converge?
2. Is there efficiency loss?

Need to look at ϵ-equilibria; there may be many

LOE $\longrightarrow$ 0 of the best ϵ-equilibrium

Can show that this NE is stable even if multiple users deviate in the same time.

PART V: Dynamics

What if we let users play best response to current price (learning)?

Each user looks at current price and optimizes bid according to:

$$w_r^{new} = \arg\max_{w_r \geq 0} U_r(w_r/\mu_{old}) - w_r$$

Minimal information requirements: only μ is common knowledge.

It turns out that:

1. “Few” users act at any time $\Rightarrow$ convergence

2. If many users act simultaneously $\Rightarrow$ may have oscillations

3. But with heterogenous users oscillations are unlikely even if all users act all the time!

4. Super fast convergence rates in practice

Classical “systems theory proofs”

Wrap-up

A control theoretic view of resource allocation

A very simple "mechanism" with low complexity/communication requirements

Mechanism is almost always efficient

A simple learning rule is usually effective

Interesting questions:
The effect of users that abandon the game
More complex dynamics/learning

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