Mean Field Equilibria of Dynamic Auctions with Learning

Ramesh Johari, Stanford University

Joint work with
Krishnamurthy Iyer, Stanford University, and
Mukund Sundararajan, Google Inc.

Auction markets

Many marketplaces employ auctions to conduct trade:

- Sponsored search
- Online markets
- Licensing, etc.

Dynamic auctions with learning

Inspired by auction settings where agents *do not know* their valuation for an item a priori

Through repeated participation, agents **learn** their preferences for goods in the market

Example: In sponsored search advertising, an advertiser only learns the value of an ad based on conversion to a sale *after* a user clicks on the ad

Sponsored search advertising

Online ads

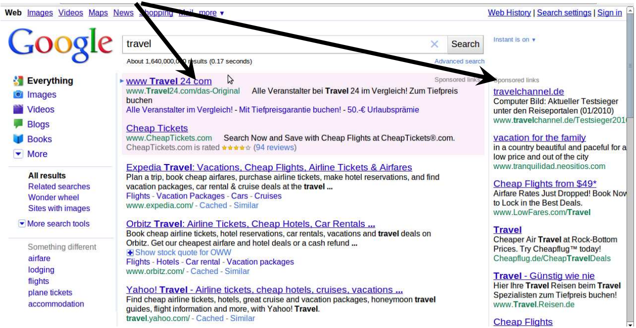

Why study learning?

Major challenges:

- What strategies are “optimal” for bidders?
- Can we characterize market behavior, and in particular the distribution of bids in the market?
- What auction format should the market operator use?
- Should the market operator subsidize learning?

Auction markets

In this talk

1. We use a mean field model to characterize agents' behavior in presence of learning.
2. We establish existence and approximation theorems for MFE.
3. We use MFE to study market design: the impact of auction format and reserve prices on the auctioneer's revenue.

Outline

Outline

Finite market

Finite number of agents

Sequence of second price auctions:

- α agents (sampled uniformly) per auction

Geometric lifetimes with parameter β:

- After an auction, participating agents leave independently with probability $1 - \beta$.
- Each departing agent is replaced by a new agent.

[Aside: Second price auctions]

In a (static) second price auction:

The highest bidder wins, and pays the second highest bid.

Exercise: It is a *dominant strategy* to bid your true valuation.

Private valuation

Agent i's private valuation $v_i \in [0, 1]$: **unknown**, independent.

Valuation determines the reward $x_{i,t}$:

$$x_{i,t} = \begin{cases} 1 & \text{with probability } v_i; \\ 0 & \text{otherwise.} \end{cases}$$

- e.g., in sponsored search:

  reward = sale after a click-through

Observing reward $x_{i,t}$ informs an agent about her valuation v_i.

Belief over valuation

Initial prior: Beta(m, n)

Density: $f_{(m,n)}(x) \;\propto\; x^{m-1}(1-x)^{n-1}$

Mean: $\mu(m, n) = \frac{m}{m+n}$

Variance: $\sigma^2(m, n)$ decreasing in m and n

Belief update

On losing the auction:

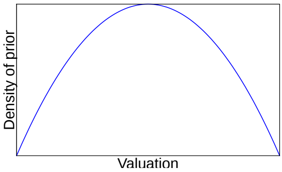

Beta(m, n)

$\longrightarrow$

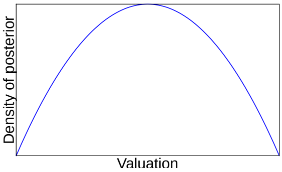

Beta(m, n)

Belief update

On winning the auction, and getting positive rewards:

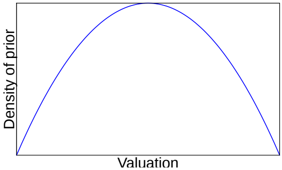

$\longrightarrow$

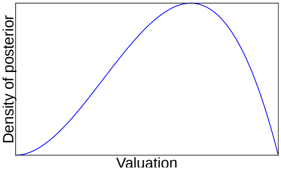

Beta(m, n) Beta$(m + 1, n)$

Belief update

On winning the auction, and getting zero rewards:

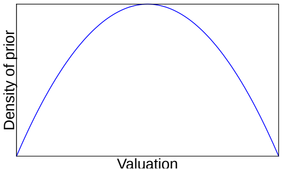

Beta(m, n)

$\longrightarrow$

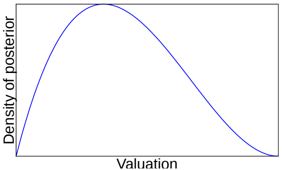

Beta$(m, n+1)$

Belief update

s_k = Belief parameters after k^{th} auction

Belief update: after k^{th} auction,

$$s_k = \begin{cases} s_{k-1} & \text{if the agent does not win;} \\ s_{k-1} + e_1 & \text{if the agent wins and } x_k = 1; \\ s_{k-1} + e_2 & \text{if the agent wins and } x_k = 0, \end{cases}$$

where $e_1 = (1, 0)$ and $e_2 = (0, 1)$.

Objective

Maximize the **total expected payoff** over the lifetime

Objective

Maximize the **total expected payoff** over the lifetime

(Per period payoff = reward - payment)

Outline

How should a bidder behave?

If private valuation is **known**:

No learning $\implies$ auctions decouple

$\implies$ Bid truthfully each time period

How should a bidder behave?

If private valuation is **known**:

No learning $\implies$ auctions decouple

$\implies$ Bid truthfully each time period

What if private valuation is **unknown**?

How should a bidder behave?

If private valuation is **unknown**:

Value for learning $\Longrightarrow$ agents overbid

But how much?

How should a bidder behave?

If private valuation is **unknown**:

Value for learning $\Longrightarrow$ agents overbid

But how much?

Hard to determine

Dynamic game

Standard tool to analyze dynamic games with *incomplete information* is *perfect Bayesian equilibrium* (PBE).

PBE assumes agents are completely **rational**:

- Agents track **each** competitor as long as they stay in the market, and play optimally.
- Agents maintain consistent beliefs about evolution of the **entire market**, and update them using Bayes' rule.

Two issues

PBE suffers from two issues:

1. **Intractability**: often showing equilibrium exists is difficult.
   - No structural insight
2. **Implausibility**: Not a good model of agent behavior in practice

Outline

Mean field equilibrium

Inspired by **large** markets

Weintraub et al. ('08), Lasry & Lions ('07), Huang et al. ('07), ...

In an MFE,

Agents do not track individual competitors

Each agent plays against a "stationary" market

Mean field equilibrium

Inspired by **large** markets

- Weintraub et al. ('08), Lasry & Lions ('07), Huang et al. ('07), ...

In an MFE,

- Agents do not track individual competitors
- Each agent plays against a "stationary" market

Sponsored search

Advertisers use **bid landscape** information to model the rest of the market.

Mean field equilibrium

Optimality:

Stationary market → Actions are optimal

Consistency:

Given agents' actions → Same stationary distribution

Mean field equilibrium

Mean Field Equilibrium = Optimality + Consistency

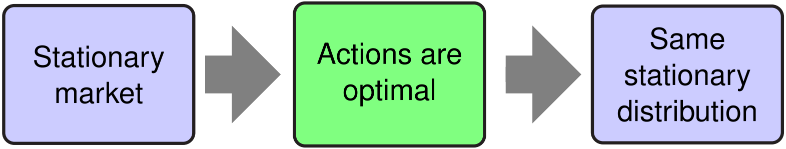

OPTIMALITY

MFE: Stationary market

Suppose the distribution of bids in the market is g

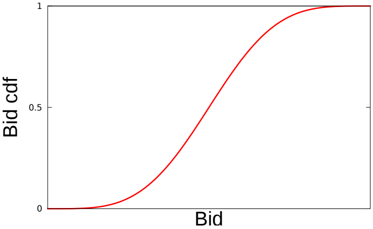

For a fixed agent, in each of her auctions, bids of other $\alpha - 1$ agents are sampled i.i.d. from g.

MFE: Stationary market

Suppose the distribution of bids in the market is g

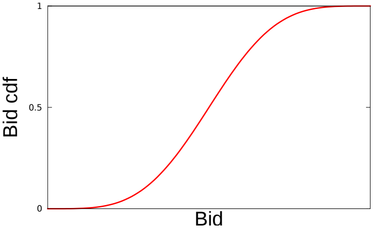

Probability of winning: $q(b|g) = g(b)^{\alpha-1}$

Expected payment: $p(b|g)$

MFE: Agent's decision problem

Expected payoff in k^{th} auction: $q(b|g)\mu(s_k) - p(b|g)$.

MFE: Agent's decision problem

Expected payoff in k^{th} auction: $q(b|g)\mu(s_k) - p(b|g)$.

Belief update: after k^{th} auction,

$$s_k = \begin{cases} s_{k-1} & \text{if the agent does not win;} \\ s_{k-1} + e_1 & \text{if the agent wins and } x_k = 1; \\ s_{k-1} + e_2 & \text{if the agent wins and } x_k = 0. \end{cases}$$

MFE: Agent's decision problem

Expected payoff in k^{th} auction: $q(b|g)\mu(s_k) - p(b|g)$.

Belief update: after k^{th} auction,

$$s_k = \begin{cases} s_{k-1} & \text{if the agent does not win;} \\ s_{k-1} + e_1 & \text{if the agent wins and } x_k = 1; \\ s_{k-1} + e_2 & \text{if the agent wins and } x_k = 0. \end{cases}$$

Geometric lifetime

MFE: Agent's decision problem

Expected payoff in k^{th} auction: $q(b|g)\mu(s_k) - p(b|g)$.

Belief update: after k^{th} auction,

$$s_k = \begin{cases} s_{k-1} & \text{if the agent does not win;} \\ s_{k-1} + e_1 & \text{if the agent wins and } x_k = 1; \\ s_{k-1} + e_2 & \text{if the agent wins and } x_k = 0. \end{cases}$$

Geometric lifetime

Expected total discounted payoff maximization

MFE: Agent's decision problem

From standard dynamic programming, optimal strategy is

- **Markovian**: bid depends only on current belief
- **Stationary**: no time dependence

CONSISTENCY

MFE: Consistency

MFE: Consistency

For consistency, we need g to be a **fixed point** of F.

MFE

Definition

A bid distribution g and a strategy C constitute an MFE if

1. Optimality: Given g, the strategy C is optimal.
2. Consistency: g is a fixed point of map F.

MFE

Definition

A bid distribution g and a strategy C constitute an MFE if

1. Optimality: Given g, the strategy C is optimal.
2. Consistency: g is a fixed point of map F.

Does an MFE exist?

Existence of MFE

Theorem

A mean field equilibrium **exists** *in the auction market.*

Existence of MFE

Theorem

A mean field equilibrium ***exists*** *in the auction market.*

Proof uses an infinite dimensional fixed point theorem.

- Show: With the **right** topologies, F is continuous.
- Show: Image of F is compact

Existence of MFE

Theorem

A mean field equilibrium **exists** *in the auction market.*

Proof uses an infinite dimensional fixed point theorem.

- Show: With the **right** topologies, F is continuous.
- Show: Image of F is compact

Existence:

- General belief models
- Multiple units per auction

OPTIMAL STRATEGY IN MFE

MFE: Agent's decision problem

Given g, agent's value function satisfies Bellman's equation:

$$V(s|g) = \max_{b \geq 0} \Big\{ q(b|g)\mu(s) - p(b|g) + \beta q(b|g)\mu(s)V(s+e_1|g)$$

$$+ \beta q(b|g)(1-\mu(s))V(s+e_2|g) + \beta(1-q(b|g))V(s|g) \Big\}$$

MFE: Agent's decision problem

Given g, agent's value function satisfies Bellman's equation:

$$V(s|g) = \max_{b \geq 0} \Big\{ \underbrace{q(b|g)\mu(s) - p(b|g)}_{(1)} + \beta q(b|g)\mu(s)V(s+e_1|g)$$
$$+ \beta q(b|g)(1-\mu(s))V(s+e_2|g) + \beta(1-q(b|g))V(s|g) \Big\}$$

(1) Expected payoff in current auction

MFE: Agent's decision problem

Given g, agent's value function satisfies Bellman's equation:

$$V(s|g) = \max_{b \geq 0} \Big\{ q(b|g)\mu(s) - p(b|g) + \underbrace{\beta q(b|g)\mu(s)V(s+e_1|g)}_{(2)}$$

$$+ \beta q(b|g)(1-\mu(s))V(s+e_2|g) + \beta(1-q(b|g))V(s|g) \Big\}$$

(2) Future expected payoff on

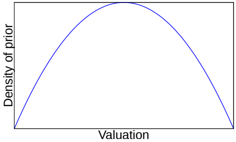

s

$\longrightarrow$

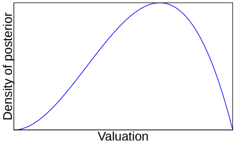

$s + e_1$

MFE: Agent's decision problem

Given g, agent's value function satisfies Bellman's equation:

$$V(s|g) = \max_{b \geq 0} \Big\{ q(b|g)\mu(s) - p(b|g) + \beta q(b|g)\mu(s)V(s+e_1|g)$$
$$+ \underbrace{\beta q(b|g)(1-\mu(s))V(s+e_2|g)}_{(3)} + \beta(1-q(b|g))V(s|g) \Big\}$$

(3) Future expected payoff on

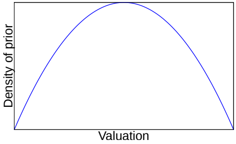

s

$\longrightarrow$

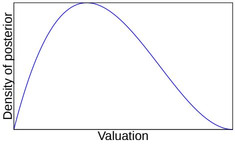

$s + e_2$

MFE: Agent's decision problem

Given g, agent's value function satisfies Bellman's equation:

$$V(s|g) = \max_{b \geq 0} \Big\{ q(b|g)\mu(s) - p(b|g) + \beta q(b|g)\mu(s)V(s+e_1|g)$$
$$+ \beta q(b|g)(1-\mu(s))V(s+e_2|g) + \underbrace{\beta(1-q(b|g))V(s|g)}_{(4)} \Big\}$$

(4) Future expected payoff on

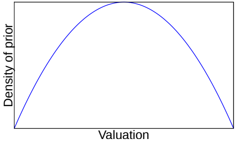

s

$\longrightarrow$

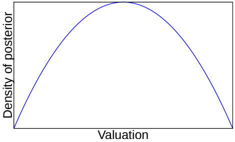

s

MFE: Agent's decision problem

Given g, agent's value function satisfies Bellman's equation:

$$V(s|g) = \max_{b \geq 0} \Big\{ q(b|g)\mu(s) - p(b|g) + \beta q(b|g)\mu(s)V(s + e_1|g) \\ + \beta q(b|g)(1 - \mu(s))V(s + e_2|g) + \beta(1 - q(b|g))V(s|g) \Big\}$$

MFE: Agent's decision problem

Given g, agent's value function satisfies Bellman's equation:

$$V(s|g) = \max_{b \geq 0} \Big\{ q(b|g)\mu(s) - p(b|g) + \beta q(b|g)\mu(s)V(s+e_1|g) + \beta q(b|g)(1-\mu(s))V(s+e_2|g) + \beta(1-q(b|g))V(s|g) \Big\}$$

MFE: Agent's decision problem

Rewriting:

$$V(s|g) = \max_{b \geq 0} \left\{ q(b|g)C(s|g) - p(b|g) \right\} + \beta V(s|g),$$

where

$$\begin{aligned} C(s|g) = \mu(s) &+ \beta\mu(s)V(s+e_1|g) \\ &+ \beta(1-\mu(s))V(s+e_2|g) - \beta V(s|g). \end{aligned}$$

MFE: Optimality

Agent's decision problem is

$$\max_{b \geq 0} \left\{ q(b|g)C(s|g) - p(b|g) \right\}$$

MFE: Optimality

Agent's decision problem is

$$\max_{b \geq 0} \left\{ q(b|g)C(s|g) - p(b|g) \right\}$$

Same decision problem as in

- **Static** second-price auction
- against $\alpha - 1$ bidders drawn i.i.d. from g
- with agent's **known** valuation $C(s|g)$.

MFE: Optimality

Agent's decision problem is

$$\max_{b \geq 0} \left\{ q(b|g)C(s|g) - p(b|g) \right\}$$

Same decision problem as in

- **Static** second-price auction
- against $\alpha - 1$ bidders drawn i.i.d. from g
- with agent's **known** valuation $C(s|g)$.

We show $C(s|g) \geq 0$ for all s

$\implies$ Bidding $C(s|g)$ at posterior s is **optimal**!

Conjoint valuation

$C(s|g)$: **Conjoint valuation** at posterior s

$$C(s|g) = \mu(s) + \beta\mu(s)V(s + e_1|g) + \beta(1 - \mu(s))V(s + e_2|g) - \beta V(s|g)$$

Conjoint valuation

$C(s|g)$: **Conjoint valuation** at posterior s

$$C(s|g) = \mu(s) + \beta\mu(s)V(s+e_1|g) + \beta(1-\mu(s))V(s+e_2|g) - \beta V(s|g)$$

Conjoint valuation = **Mean** + **Overbid**

(We show Overbid ≥ 0)

Conjoint valuation: Overbid

Overbid: $\beta\mu(s)V(s+e_1|g)+\beta(1-\mu(s))V(s+e_2|g)-\beta V(s|g)$

Conjoint valuation: Overbid

Overbid: $\beta\mu(s)V(s+e_1|g) + \beta(1-\mu(s))V(s+e_2|g) - \beta V(s|g)$

Overbid

Expected marginal future gain from **one additional observation** about private valuation

How should a bidder behave?

If private valuation is **unknown**:

Value for learning $\implies$ agents overbid

But how much?

How should a bidder behave?

If private valuation is **unknown**:

Value for learning $\implies$ agents overbid

But how much?

Expected marginal value of one additional observation

How should a bidder behave?

If private valuation is **unknown**:

Value for learning $\implies$ agents overbid

But how much?

Expected marginal value of one additional observation

Simple description of agent behavior!

Outline

Approximation

Does an MFE capture rational agent behavior in **finite** market?

Issues:

- Repeated interactions $\implies$ agents no longer independent.
- Keeping track of history will be beneficial.

Hope for approximation only in the **asymptotic** regime

Approximation

Look at the market as an interacting particle system.

Interaction set of an agent: all agents influenced by or that had an influence on the given agent.

Approximation

Look at the market as an interacting particle system.

Interaction set of an agent: all agents influenced by or that had an influence on the given agent.

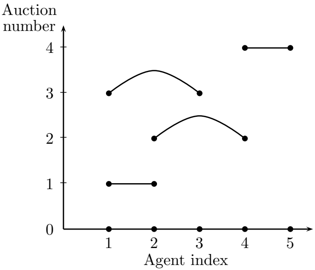

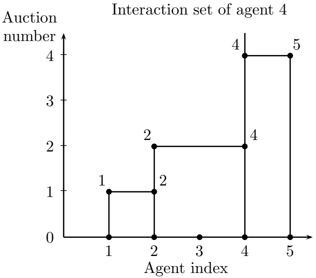

Approximation

Look at the market as an interacting particle system.

Interaction set of an agent: all agents influenced by or that had an influence on the given agent.

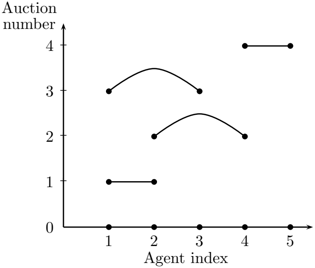

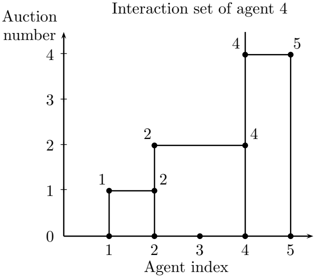

Intuition: As market size increases, any two agents' interaction sets become disjoint with high probability.

Approximation

Theorem

As the number of agents in the market **increases***, the maximum additional payoff on a* **unilateral** *deviation converges to zero.*

As the market size increases,

Expected payoff under **optimal strategy**, given others play $C(\cdot\|g)$	$-$	Expected payoff under $\mathbf{C}(\cdot\|\mathbf{g})$, given others play $C(\cdot\|g)$	$\to$	$\mathbf{0}$

Approximation

Theorem

As the number of agents in the market ***increases****, the maximum additional payoff on a* ***unilateral*** *deviation converges to zero.*

Mean field equilibrium is **good** approximation to agent behavior in finite large market.

Approximation

Theorem

As the number of agents in the market ***increases****, the maximum additional payoff on a* ***unilateral*** *deviation converges to zero.*

Mean field equilibrium is **good** approximation to agent behavior in finite large market.

Good approximation even in small markets due to **behavioral** reasons

Outline

Heuristic to compute MFE

A natural heuristic inspired by *model predictive control*.

Implicitly encodes a **learning algorithm** for the agents

- Closely models market evolution when agents optimize given current average estimates.

Algorithm

1. Initiate the market at bid distribution g_0.

Algorithm

1. Initiate the market at bid distribution g_0.
2. Given g_k, compute conjoint valuation $C(\cdot|g_k)$.

Algorithm

1. Initiate the market at bid distribution g_0.
2. Given g_k, compute conjoint valuation $C(\cdot|g_k)$.
3. Evolve the market *one time period*, assuming each agent bids her conjoint valuation.

Algorithm

1. Initiate the market at bid distribution g_0.
2. Given g_k, compute conjoint valuation $C(\cdot|g_k)$.
3. Evolve the market *one time period*, assuming each agent bids her conjoint valuation.
4. Compute the new bid distribution g_{k+1}.

Algorithm

1. Initiate the market at bid distribution g_0.
2. Given g_k, compute conjoint valuation $C(\cdot|g_k)$.
3. Evolve the market *one time period*, assuming each agent bids her conjoint valuation.
4. Compute the new bid distribution g_{k+1}.
5. Repeat until $\|g_{k+1} - g_k\|_\infty < \epsilon$.

Performance

Heuristic converges to MFE within 30-50 iterations in practice, for reasonable error bounds ($\epsilon \sim 0.0015$)

Computation takes ~ 30 mins on a laptop.

Overbidding

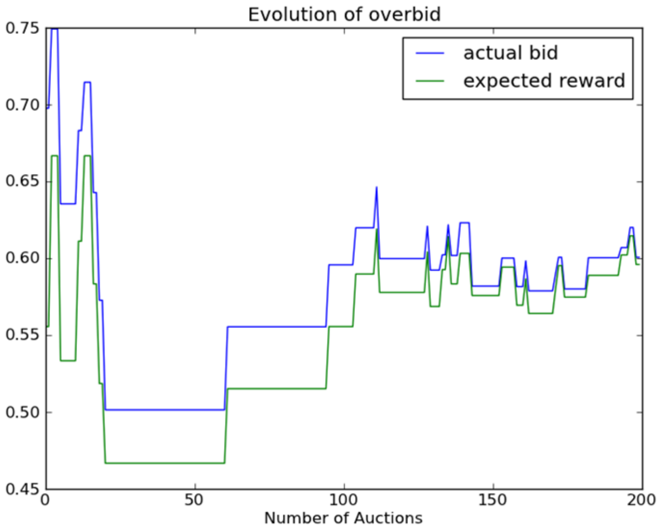

Outline

Repeated standard auctions

Standard auction:

1. Winner has the highest bid.
2. Zero bid implies zero payment.

Example: First price, all pay, third price, etc.

Theorem

MFE exists in any repeated **standard auction** *if the static auction has a symmetric Bayes-Nash equilibrium.*

Repeated standard auctions

Dynamic revenue equivalence

Expected revenue to the auctioneer is **same** irrespective of static auction format

(This is an analog of a similar result in *static* auction theory.)

Moral: Auctioneer's expected revenue not affected by choice of auction format.

Outline

Reserve price

Setting a reserve can increase auctioneer's revenue.

Effects of a reserve:

1. Relinquishes revenue from agents with low conjoint valuation.
2. Extracts more revenue from those with high conjoint valuation.

Reserve price

Setting a reserve can increase auctioneer's revenue.

Effects of a reserve:

1. Relinquishes revenue from agents with low conjoint valuation.
2. Extracts more revenue from those with high conjoint valuation.
3. Imposes a **learning cost**:
   - Precludes agents from learning, and reduces incentives to learn.

Reserve price

Due to learning cost, agents change behavior on setting a reserve.

> Auctioneer sets a reserve r and agents behave as in an MFE with reserve r.

Defines a **game** between the auctioneer and the agents.

Optimal reserve

Two approaches:

1. **Nash equilibrium**: Ignores learning cost. Auctioneer sets reserve assuming bid distribution is fixed, and agents behave as in MFE with reserve r.
2. **Stackelberg equilibrium**: Includes learning cost. Auctioneer computes revenue in MFE for each r, and sets the maximizer r_{OPT}.

We compare these two approaches using numerical computation.

Optimal reserve: Numerical findings

By definition, $\Pi(r_{OPT}) \geq \Pi(r_{NASH})$.

$\Pi(r_{OPT}) - \Pi(0)$ is greater than $\Pi(r_{NASH}) - \Pi(0)$ by $\sim 1 - 30\%$.

- Ignoring learning incurs a potentially **significant** cost.

Improvement depends on the distribution of initial beliefs of arriving agents.

CONCLUSION

Conclusion

The methodology of MFE allows for

- **Tractability**: many analytical insights possible
- **Plausibility**: conjoint valuation captures nicely the value of learning

Numerical computation feasible $\implies$ questions of practical relevance such as optimal reserve can be answered through computation.

Open problems

Proof of convergence of the heuristic

How does efficiency of such repeated one-shot mechanisms compare with optimal mechanisms?

Empirical validation

Other models:

- Unit demand bidders (eBay, Amazon, etc.)
- Budget constrained bidders

THANK YOU

Paper at http://ssrn.com/abstract=1799085