Mean Field Equilibria of Dynamic Auctions

Ramesh Johari

Stanford University

Background

- Background in networks; at Stanford since 2004
- Recent area of interest: design and analysis of online markets
- Just returned from 18 month leave at oDesk (online marketplace for remote work);
  last 9 months as director of data products
- This tutorial: Two examples of how mean field models help us understand complex dynamic markets
- Today: Dynamic auctions
- Tomorrow: Dynamic matching market

Outline

- A motivating example: dynamic auctions with learning
- A mean field model
- Mean field equilibrium
- Characterizing MFE
- Using MFE: dynamic revenue equivalence, reserve prices
- Other models: budget constraints, unit demand bidders
- Open problems

PART I: A MOTIVATING EXAMPLE

Sponsored search markets

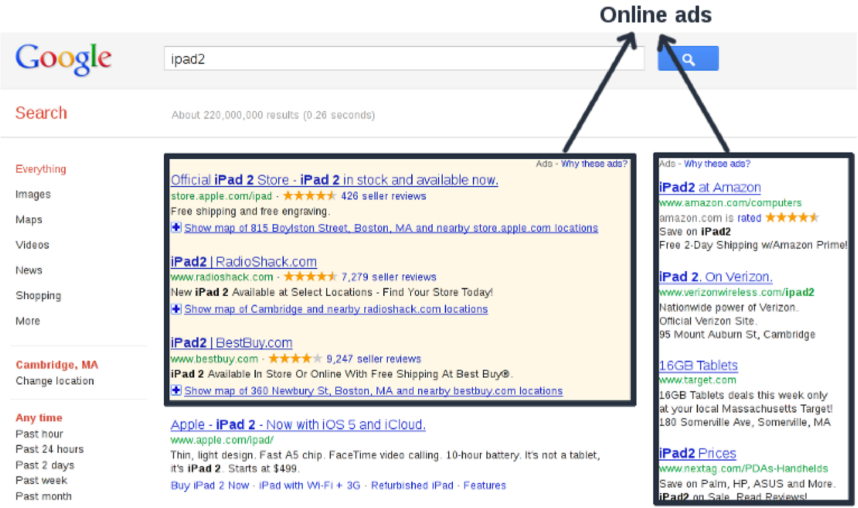

Sponsored search markets

Advertisers bid on various keywords to get their ads placed on the search page.

On each query, an auction occurs among the relevant advertisers, and winners get their ads placed.

Cost-per-click (CPC): The good being auctioned is a click, i.e., advertisers pay only if a user clicks on their ad.

Advertisers care for **conversion** – how an ad click converts into sales or profit.

Sponsored search markets

There is a **mismatch** between the good being auctioned and what the advertisers value.

This creates a **dynamic incentive**:

Bidders must simultaneously **estimate** their conversion rates while bidding on keywords.

Repeated auctions with learning

Here we consider a simple abstraction:

- N bidders
- Bidder i has a valuation $v_i \in [0, 1]$ that is *unknown* to her
  - Think of this as the conversion rate.
- v_i distributed according to prior F_i (independent across bidders)
- Bidders compete in a sequence of second price auctions

What should a bidder do?

First suppose there is a single period.

Dominant strategy: bid expected value (according to current belief).

What should a bidder do?

What about multiple periods?

There is now a **value for learning**:

Agents will tend to *overbid* above expected valuation, because learning about their value might help them in future periods

What should a bidder do?

But the amount to overbid depends critically on
what a bidder believes about her competitors.

The classical solution concept is *perfect Bayesian equilibrium* (PBE): A bidder optimizes with respect to:

- her beliefs over all that is unknown, given the history so far; and
- her prediction of how others will behave in the future, in response to her action today.

Challenge 1: PBE is implausible

There seems to be a **"law of large numbers of rationality"**:

Complex beliefs and forecasting become uncommon even with relatively small numbers of players (5-10).

Therefore PBE seems to be a highly implausible model of agent behavior, even in settings with fairly sophisticated agents.

Challenge 2: PBE is intractable

The dynamic optimization problem of an agent has a very high dimensional state space:

An agent optimizes given beliefs over all that is unknown.

Even computing best responses is prohibitive, let alone equilibria!

Moral

This is a bad place to be:

One does not want theory to be both intractable and implausible.

As a result, we leave engineers with few tools to guide design:

How does market structure, auction format, reserve prices, etc. affect bidder behavior?

PART II: A MEAN FIELD MODEL

Bounding rationality

“Bounded rationality” models offer a way out of the impasse; but which bounded rationality approach to use?

We’ll discuss an approximation founded on the premise that there are a large number of bidders present.

This is a **mean field model**.

A formal model

We now formally describe a mean field model for dynamic auctions with learning.

Key components:

- Bidder model: learning and payoffs
- The "mean field": competitors' bid distribution

A formal model

A bidder participates in a sequence of second price auctions.

α bidders in each auction.

The bidder lives for a geometric(β) lifetime (mean $1/(1-\beta)$).

The bidder has an **unknown** private valuation $v \in [0, 1]$:

$$\mathsf{P}(\text{reward}_t = 1) = 1 - \mathsf{P}(\text{reward}_t = 0) = v$$

Learning model

Initial prior: Beta(m, n)

- (m, n) and v chosen on arrival.

Mean: $\mu(m, n) = m/(m+n)$

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

Belief update is through Bayes' rule;
let $s_k = (m_k, n_k)$ denote belief parameters after k'th auction.

Belief update

On losing the auction:

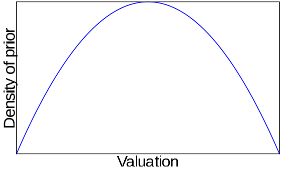

$\longrightarrow$

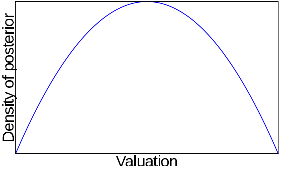

Beta(m, n) Beta(m, n)

Belief update

On winning the auction, and getting a **positive** reward:

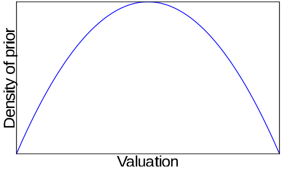

Beta(m, n)

$\longrightarrow$

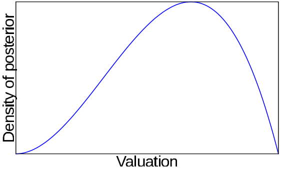

Beta$(m+1, n)$

Belief update

On winning the auction, and getting **zero** reward:

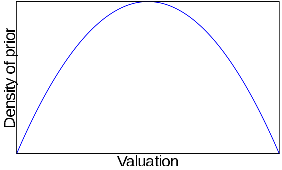

Beta(m, n)

$\longrightarrow$

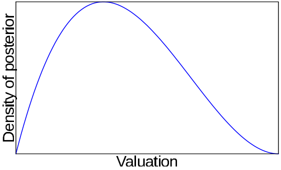

Beta$(m, n+1)$

Objective

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

(Per period payoff = reward - payment)

The “mean field” market

Suppose the *distribution* of bids in the market is g

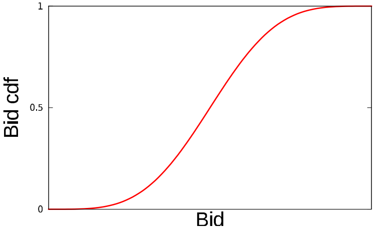

The **mean field** assumption:

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

Sponsored search: Bid landscape

Why is the mean field model reasonable?

In sponsored search, advertisers use **bid landscape** information to model the rest of the market.

Bid landscapes use the last week's data to give aggregated estimates of cost-per-click, number of clicks, and number of impressions that can be expected for a given bid.

The mean field model captures this information structure.

Sponsored search: Bid landscape

Bid simulator: operations research

Simulation based on performance from Feb 2, 2012 to Feb 8, 2012

These estimates do not guarantee similar results in the future. Learn more

Max. CPC	Estimated Impr.	Estimated Top Impr.
$5.67	17,000	53
$3.40	14,500	19
$2.86	12,100	9
$2.47	9,550	5
$2.00	7,040	3
$1.45	4,530	--
$1.00 (current)	2,040	--
Use a different bid: $		

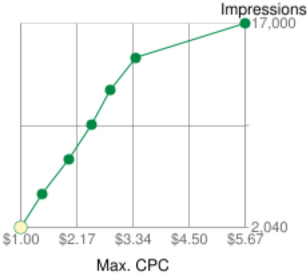

Save Cancel

Questions

- What is a reasonable notion of equilibrium for this system?
- Does it exist?
- What is the structure of bidders' optimal strategy?
- Do mean field models approximate games with finitely many players?
- How do we compute an equilibrium?

PART III: MEAN FIELD EQUILIBRIUM

Mean field equilibrium

Inspired by large markets.

In an MFE:

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

Mean field equilibrium

Optimality:

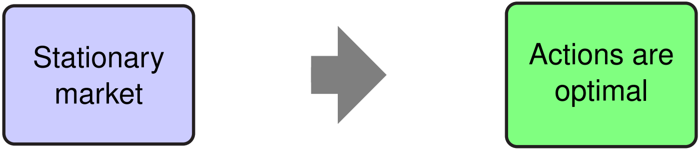

Consistency:

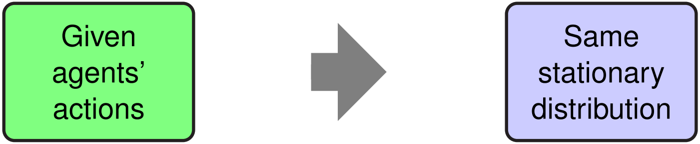

Mean field equilibrium: Dynamic auctions

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

Optimality:

Fixed bid distribution g

Strategy C is optimal

Consistency:

Given each agent follows C

Market bid distribution is g

Mean field equilibrium: Formal definition

Fix a bid distribution g.

- Let $C(\cdot|g)$ be an optimal strategy for the agent's expected lifetime profit maximization problem, given g.
- Let Φ be the steady state distribution (on valuations and states) induced by the resulting agent dynamics under the strategy $C(\cdot|g)$, and assuming other agents' bids are drawn from g. (Note that these dynamics include regeneration.)
- Let g' be the new steady state bid distribution derived by integrating the strategy $C(\cdot|g)$ against the steady state distribution Φ.

The bid distribution g is a MFE bid distribution if it is a fixed point of this map.

Mean field equilibrium: Related work

Mean field models arise in a wide variety of fields:
physics, applied math, engineering, economics, ...

Extensive work on mean field models for static games (e.g., competitive equilibrium, nonatomic games, etc.)

Mean field equilibrium: Related work

Mean field models in dynamic games:

- *Economics*: Jovanovic and Rosenthal (1988); Stokey, Lucas, Prescott (1989); Hopenhayn (1992); Sleet (2002); Weintraub, Benkard, Van Roy (2008, 2010); Acemoglu and Jepsen (2010); Bodoh-Creed (2011)
- *Control*: Glynn, Holliday, Goldsmith (2004); Lasry and Lions (2007); Huang, Caines, Malhamé (2007-2012); Gueant (2009); Tembine, Altman, El Azouzi, le Boudec (2009); Yin, Mehta, Meyn, Shanbhag (2009); Adlakha, Johari, Weintraub (2009, 2011)
- *Finance*: Duffie, Malamud, Manso (2009, 2010)
- *Dynamic auctions*: Wolinsky (1988); McAfee (1993); Backus and Lewis (2010); Iyer, Johari, Sundararajan (2012); Gummadi, Proutière, Key (2013); Bodoh-Creed (2012); Balseiro, Besbes, Weintraub (2013)

(Other names for MFE: Stationary equilibrium, oblivious equilibrium)

Mean field equilibrium: Related work

Another relevant line of literature is on *dynamic mechanism design*.

Examples: Athey and Segal (2007); Bergemann and Valimaki (2010); etc.

- In dynamic mechanism design, a *hard* optimization problem is solved (optimal dynamic allocation), and payments are structured so equilibrium behavior bidder is *simple* (truthtelling).
- But, in many real markets: repetitions of *simple* mechanisms are implemented, leading to *complex* equilibrium bidder behavior.

PART IV: CHARACTERIZING MFE

Characterizing MFE

- Optimal strategies
- Existence of MFE
- Approximation and finite games
- Computation

PART IV-A: Optimal strategies

MFE: Stationary market

Suppose the distribution of bids in the market is g

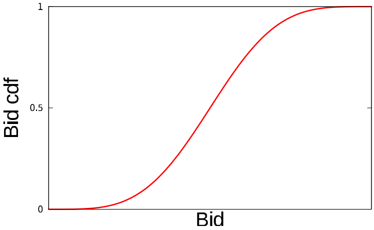

Probability of winning: $q(b|g) = g(b)^{\alpha-1}$

Expected payment: $p(b|g)$

MFE: Agent's decision problem

Let $V(s|g)$ denote the agent's maximum possible expected lifetime payoff, when her current belief is s, and the population bid distribution is g.

By the principle of optimality for discounted dynamic programming, V must satisfy **Bellman's equation**.

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:

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

(1) Expected payoff in current auction

MFE: Agent's decision problem

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

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

(2) Future expected payoff on winning and **positive** reward:

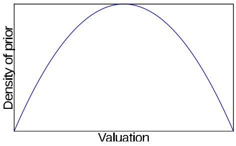

s

$\longrightarrow$

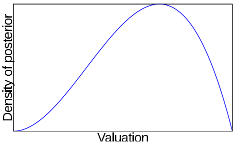

$s + e_1$

MFE: Agent's decision problem

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

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

(3) Future expected payoff on winning and **zero** reward:

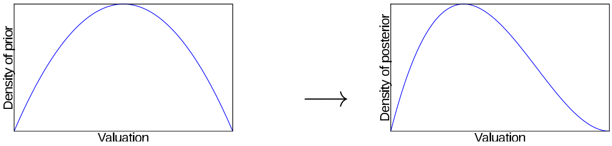

$s \longrightarrow s + e_2$

MFE: Agent's decision problem

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

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

(4) Future expected payoff on losing:

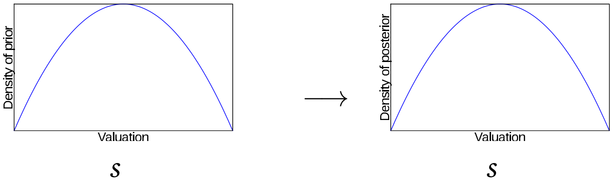

MFE: Agent's decision problem

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:

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

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

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

Simple description of agent behavior!

PART IV-B: Existence of MFE

Existence of MFE

We make one assumption for existence:

We assume that the distribution from which the value and belief of a single agent are initially drawn has compact support with no atoms.

Existence of MFE

Theorem

A mean field equilibrium exists where each agent bids her **conjoint valuation** *given her posterior.*

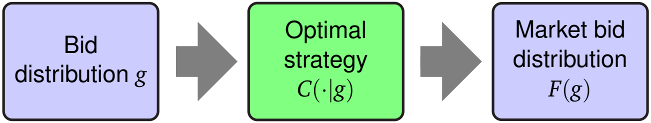

- Show: With the right topologies, F is continuous
- Show: Image of F is compact (using previous assumption)

PART IV-C: Approximation and MFE

Approximation

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

Issues:

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

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

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 $C(\cdot\|g)$, given others play $C(\cdot\|g)$	$\rightarrow$	$\mathbf{0}$

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 (from Graham and Méléard, 1994).

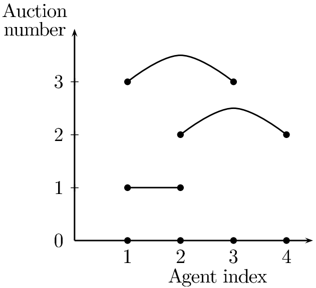

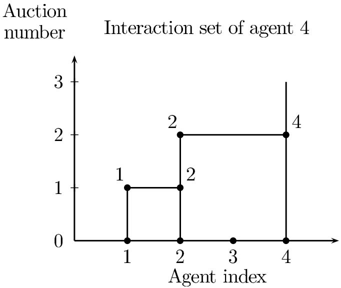

Propagation of chaos $\Longrightarrow$ As market size increases, any two agents' interaction sets become disjoint with high probability.

Approximation

Theorem

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

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

PART IV-D: Computing MFE

MFE computation

A natural algorithm inspired by **model predictive control** (or **certainty equivalent control**)

Closely models **market evolution** when agents optimize given current average estimates

MFE computation

Initialize the market at bid distribution g_0.

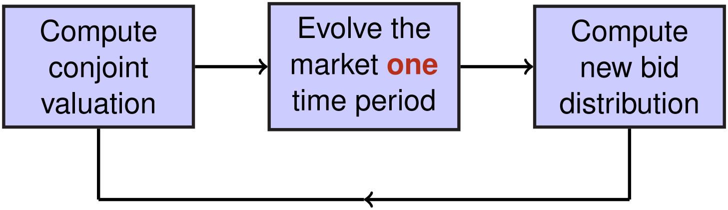

Continue until successive iterates of bid distribution are sufficiently close.

- Stopping criterion: total variation distance is below tolerance ϵ.

Performance

Algorithm converges within 30-50 iterations in practice, for reasonable error bounds ($\epsilon \sim 0.005$)

Computation takes $\sim$ 30-45 mins on a laptop.

Overbidding

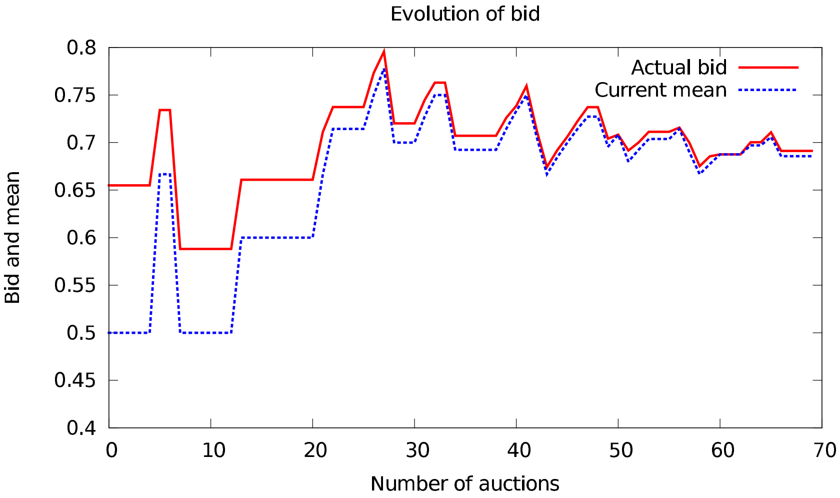

Discussion

In the dynamic auction setting, proving convergence of this algorithm remains an open problem.

However, we have proven convergence of similar algorithms in two other settings:

- Dynamic supermodular games (Adlakha and Johari, 2011)
- Multiarmed bandit games (Gummadi, Johari, and Yu, 2012)

Alternate approach: Best response dynamics (Weintraub, Benkard, Van Roy, 2008)

PART V: USING MFE IN MARKET DESIGN

Auction format

The choice of auction format is an important decision for the auctioneer.

We consider markets with repetitions of a **standard auction**:

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

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

Repeated standard auctions

Added complexity due to strategic behavior:

For example, the static first-price auction naturally induces **underbidding**.

This is in conflict with overbidding due to learning.

Repeated standard auctions

Added complexity due to strategic behavior:

For example, the static first-price auction naturally induces **underbidding**.

This is in conflict with overbidding due to learning.

We show a **dynamic revenue equivalence** theorem:

Maximum revenue over all MFE of repeated **second-price** auction.	=	Maximum revenue over all MFE of any repeated **standard** auction.

Repeated standard auctions

Added complexity due to strategic behavior:

For example, the static first-price auction naturally induces **underbidding**.

This is in conflict with overbidding due to learning.

We show a **dynamic revenue equivalence** theorem:

Maximum revenue over all MFE of repeated **second-price** auction.	=	Maximum revenue over all MFE of any repeated **standard** auction.

All standard auction formats yield the **same** revenue!

Dynamic revenue equivalence

Maximum revenue over all MFE of repeated **second-price** auction.	=	Maximum revenue over all MFE of any repeated **standard** auction.

Proof in two steps:

1. $\leq$: **Composition** of conjoint valuation and static auction behavior.
2. $\geq$: technically challenging (constructive proof).

Reserve price

Setting a reserve price can increase auctioneer's revenue.

Effects of a reserve:

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

Reserve price

Setting a reserve price can increase auctioneer's revenue.

Effects of a reserve:

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

Reserve price

Consider repeated second price auction setting.

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**: Ignores learning cost.

   Auctioneer sets a reserve r assuming bid distribution is fixed, and agents behave as in a corresponding MFE.

2. **Stackelberg**: 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 $\sim 15 - 30\%$.

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

By ignoring learning, auctioneer may incur a potentially **significant** cost.

Discussion

There is a significant point to be made here:

These types of comparative analyses are very difficult (if not impossible) using classical equilibrium concepts:

If equilibrium analysis is intractable, then we can't study how the dynamic market changes as we vary parameters.

PART VI: OTHER DYNAMIC INCENTIVES

PART VI-A: Budget constraints

Bidder model

Now suppose that a bidder faces a budget constraint B, but knows her valuation v.

The remainder of the specification remains as before.

In particular, the agent has a geometric(β) lifetime, and assumes that her competitors in each auction are i.i.d. draws from g.

Decision problem

Then a bidder's dynamic optimization problem has the following value function:

$$
\begin{aligned}
V(B, v|g) = \max_{b \leq v} \Big\{ & q(b|g)v - p(b|g) + \beta(1 - q(b|g))V(B, v|g) \\
& + \beta q(b|g)\mathsf{E}\big[V(B - b_-, v|g)|b_- \leq b\big] \Big\},
\end{aligned}
$$

where b_- is the highest bid among the competitors.

Decision problem

Some rearranging gives:

$$
\begin{aligned}
V(B,v|g) = \frac{1}{1-\beta} \max_{b \le v} \Big\{ & q(b|g)v - p(b|g) + \\
& - \beta q(b|g) \mathsf{E}\big[V(B,v|g) - V(B-b_-,v|g) | b_- \le b\big] \Big\},
\end{aligned}
$$

where b_- is the highest bid among the competitors.

Decision problem: large B

Suppose that B is very large relative to v. Then we can approximate:

$$V(B, v|g) - V(B - b_{-}, v|g)$$

by:

$$V'(B, v|g)b_{-}.$$

Decision problem: large B

Since:

$$q(b|g)\mathsf{E}\big[b_-|b_- \leq b\big] = p(b|g),$$

conclude that:

$$\beta q(b|g)\mathsf{E}\big[V(B,v|g) - V(B-b_-,v|g)|b_- \leq b\big] \approx \beta V'(B,v|g)p(b|g).$$

Decision problem: large B

Substituting we find:

$$V(B, v|g) = \frac{1 + \beta V'(B, v|g)}{1 - \beta} \max_{b \leq v} \left\{ q(b|g) \left(\frac{v}{1 + \beta V'(B, v|g)} \right) - p(b|g) \right\}.$$

As before: this is the same decision problem as an agent in a *static* second price auction, with “effective” valuation $v/(1 + \beta V'(B, v|g)$.

Optimal bidding strategy

Moral:

In a mean field model of repeated second price auctions with budget constraints (and with $B \gg v$), an agent's optimal bid is:

$$\frac{v}{1 + \beta V'(B|g)}.$$

Note that agents **shade** their bids:

This is due to the opportunity cost of spending budget now.

Large B

This model can be formally studied in a limit that captures the regime where B becomes large relative to the valuation.

See Gummadi, Proutière, Key (2012) for details;
see also Balseiro, Besbes, Weintraub (2012).

PART VI-B: Unit demand bidders

Bidder model

Now consider a setting where a bidder only wants one copy of the good, and her valuation is v.

She competes in auctions until she gets one copy of the good; discount factor for future auctions = δ.

The remainder of the specification remains as before.

In particular, the agent has a geometric(β) lifetime, and assumes that her competitors in each auction are i.i.d. draws from g.

Decision problem

Then a bidder's dynamic optimization problem has the following value function:

$$V(v|g) = \max_{b \leq v}\{q(b|g)v - p(b|g) + \beta(1 - q(b|g))\delta V(v|g)\}.$$

Decision problem

Rearranging:

$$V(v|g) = \frac{1}{1-\beta} \max_{b \leq v} \{q(b|g)(v - \beta\delta V(v|g)) - p(b|g)\}.$$

As before: this is the same decision problem as an agent in a *static* second price auction, with "effective" valuation $v - \beta\delta V(v|g)$.

Optimal bidding strategy

Moral:

In a mean field model of repeated second price auctions with unit demand bidders, an agent's optimal bid is:

$$v - \beta\delta V(v|g).$$

Note that agents **shade** their bids:

This is due to the possibility of waiting until later to get the item.

Generalization

This model has been analyzed in a much more complex setting, with many sellers and buyers, and with endogeneous entry and exit.

See Bodoh-Creed (2012) for details.

PART VI-C: Common value auctions

Bidder model

Now suppose that the valuation v is *common* to all the bidders, but *unknown* to any of them.

- In particular, assume $v \in \{v_H, v_L\}$.
- After each auction, each bidder in the auction observes the highest bid among the competitors (denoted B).

With common values, the bidders learn not only from observing their rewards, but also from observing B.

Bid distribution

In an MFE, the distribution of B will be correlated with the true valuation.
Denote these two distributions by g_H, g_L.
Define the log-likelihood ratio as:

$$\lambda(\cdot) = \ln\left(\frac{dg_H(\cdot)}{dg_L(\cdot)}\right).$$

We restrict attention to MFE where λ is increasing.

Belief update

In this model, there are *two sources* of belief update:

1. On seeing the price B, an agent with log-likelihood ratio s about the valuation updates it to:

$$s_{\text{update}} = s + \lambda(B).$$

2. On winning the auction, an agent updates her belief on seeing the reward as before:

$$s_{\text{update}} = \begin{cases} s+1 & \text{if } reward = 1; \\ s-1 & \text{if } reward = 0. \end{cases}$$

Decision problem

A single bidder's dynamic optimization problem becomes:

$$
\begin{aligned}
V(s) = \max_b \{ & \mathbf{E}[\mathbf{I}\{B \leq b\}(v - B)|s] \\
& + \beta \mathbf{E}[\mathbf{I}\{B \leq b\}(vV(s + 1 + \lambda(B)) + (1 - v)V(s - 1 + \lambda(B)))|s] \\
& + \beta \mathbf{E}[\mathbf{I}\{B > b\}V(s + \lambda(B))|s]\} .
\end{aligned}
$$

Optimal bidding strategy (Iyer 2012)

Following similar analysis, we can show that the optimal bid for an agent at belief s is:

$$\sup\{C(s+\lambda(b)) : C(s+\lambda(b)) > b\}$$

Informally:

The optimal bid is the **conjoint valuation** at posterior belief, assuming the price isthe **same** as the bid (i.e., that the bidder is "pivotal").

Intuition

Bidder's gain conditional on winning
$= $ Conjoint valuation at posterior belief
$= C(s+\lambda(B))$

In second price auctions, bidders bid their **maximum gain** conditional on winning:

$$\textbf{Optimal bid} = \sup\{\underbrace{C(s+\lambda(B))}_{\text{gain}} : \underbrace{C(s+\lambda(B)) > B}_{\text{win}}\}$$

[*Note:* This intuition is the same as for static second price common value auctions.]

PART VII: OPEN PROBLEMS

General theory

A similar analysis can be carried out for general *anonymous* dynamic games.

Extensions to:

- Nonstationary models (Weintraub et al.);
- Unbounded state spaces (Adlakha et al.);
- Continuous time (Tembine et al., Huang et al., Lasry and Lions, etc.).

Efficiency

There is an extensive literature in economics studying convergence of **large static double auctions** to:

- competitive equilibrium (with private values); or
- rational expectations equilibrium (with common values).

Analogously, which sequential auction mechanisms converge to dynamic competitive or rational expectations equilibria in large markets?

[*Note:* dynamic incentives such as learning or budget constraints cause an efficiency loss.]

Intractability

What does it mean to say MFE is simpler than classical equilibrium concepts?

Typical argument: curse of dimensionality.

But in the end, all concepts rely on fixed point arguments to establish existence.

Can we establish in a computational complexity-theoretic framework, that MFE is simpler?

Finding MFE

In most settings, MFE existence remains nonconstructive.

As discussed above, in some cases algorithms exist to compute MFE.

What are some other reasonable algorithms to compute MFE? In what settings can we establish uniqueness, convergence, etc.?

Interchanging limits

Our approximation theorem only holds over finite time intervals.

In general, interchanging time and number of agents is not straightforward: requires uniform convergence to mean field limit over time.

Under what conditions is this guaranteed? (See also: Glynn, 2004; Gummadi, Johari, Yu, 2012.)

Interaction models

MFE is valid with full **temporal mixing**:

Interact with a small number of agents each period, but resample i.i.d. every time period

But MFE is also valid with **full spatial mixing**:

Interact with everyone at every time period

What about more complex interaction models (e.g., random graphs that evolve over time)?

CONCLUSION

Conclusion

Modern large scale markets are highly dynamic, and present significant design challenges to engineers.

Approximation methods like MFE are both more tractable and more plausible than classical equilibrium approaches to such complex dynamic games.