Mean field equilibria of dynamic auctions with learning

Ramesh Johari

Stanford University

Joint work with K. Iyer and M. Sundararajan

Overview

What tools are useful for studying
dynamic systems with many interacting agents?

One possibility: *dynamic game theory.*

Traditional game theory is impractical in this regime:

- Equilibria in dynamic games make *very strong rationality assumptions*
- Equilibrium computation grows in complexity with the number of players ("*curse of dimensionality*")

Approximate approach: *mean field equilibrium*

This talk

(1) Dynamic auctions with learning

(2) Mean field

(3) Approximation

(4) Conclusions

Dynamic auctions with learning

Dynamic auctions with learning

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

  Consider N devices that compete for resources by bidding for channel use.

  Devices don't know the quality of the channel, and learn about quality each time they use it.

Dynamic auctions with learning

Model:

- Consider auction setting where bidder i has valuation $v_i \in$ [0,1] that *she does not know*
- *In other words:*

  If i wins, then P(success) = 1 – P(failure) = v_i
- Assume w.l.o.g. each successful packet transmitted is worth $1 to the bidder

Dynamic auctions with learning

Model (cont'd):

- Suppose bidders live for geometric lifetime with parameter β
- True v_i sampled from a beta distribution
- Initial state of bidder i:
  ($w_{i,0}$, $\ell_{i,0}$) = parameters of beta prior for agent i

[Assume initial state sampled from compact set with smooth density]

- (v_i, $w_{i,0}$, $\ell_{i,0}$) independent across bidders

Goal: **maximize long-run expected profit.**

Dynamic auctions with learning

- **Suppose n = k α bidders in market.**

$$n = k\,\alpha$$

Dynamic auctions with learning

- Suppose $n = k\,\alpha$ bidders in market.
- At each time t, divided into k subsets of α bidders each, *uniformly at random.*
- Each subset bids in a *second price (Vickrey) auction* for one slot.

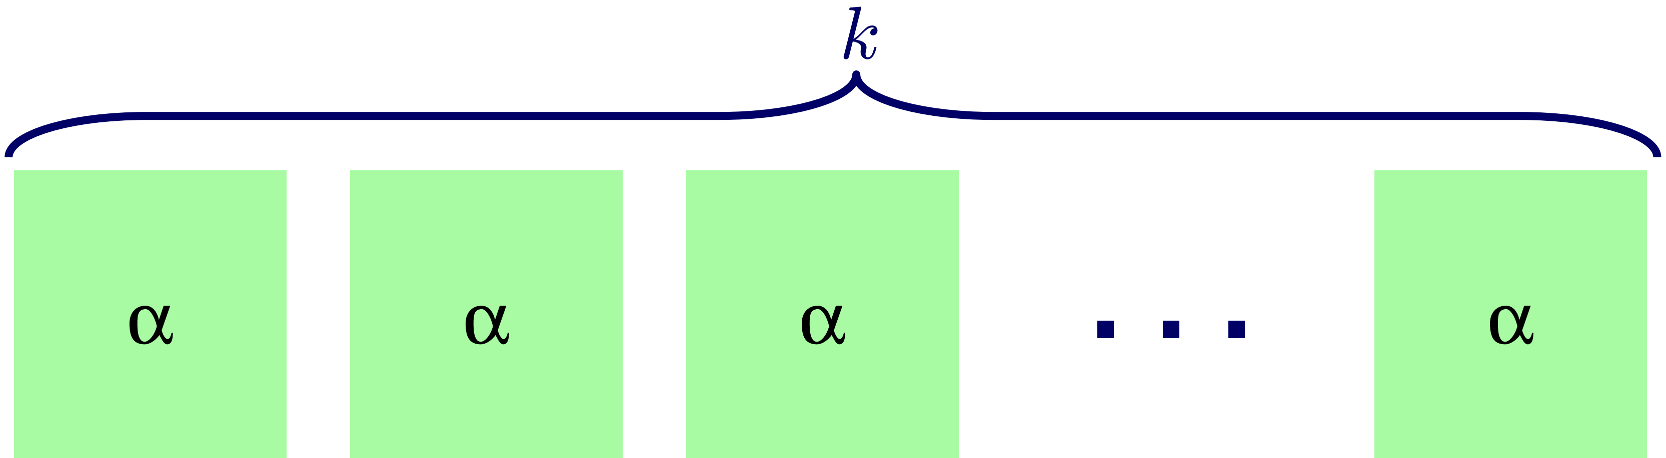

[Second price auction: review]

In a second price auction:

Each bidder submits a bid.

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

Easy result:
It is a *dominant strategy* for each bidder to bid their true valuation.

Dynamic auctions with learning

- Therefore a one period model is "easy":
  Every bidder i bidding E[v_i] is an equilibrium.
- But a dynamic model is much harder:

  (1) Bidders will want to *overbid,* to learn about their valuation.

  (2) Rational bidders will also learn about *their competitors,* so an optimal strategy will be structurally complex.

No insight into structure of equilibria in a game with finitely many players

Dynamic auctions with learning: mean field

Dynamic auctions: mean field

We consider the limit where n, $k \to \infty$, and study (stationary) mean field equilibrium.

x_t = (w_t, ℓ_t) : posterior of an agent at time t

w_t = number of periods with win and success

ℓ_t = number of periods with win and failure

a_t : bid of an agent at time t

g : population bid distribution in mean field limit

π(x_t, a_t, g) = expected payoff given current posterior

P($\cdot$ | x_t, a_t, g) = posterior update

Dynamic auctions: mean field

Write $\pi(x,a,g) = q(a \mid g)\, \mu(x) - p(a \mid g)$

- $q(a \mid g) = g(a)^{\alpha - 1}$ = probability of winning
- $\mu(x) = w/(w + \ell)$ = conditional mean valuation
- $p(a \mid g) = a\, q(a \mid g) - \int_0^a q(z \mid g)\, dz$ = expected payment

Bellman equation given g:

$$\underline{V}(x|g) = \max_{a \in [0,1]} \Big\{ q(a|g)\mu(x) - p(a|g) + \beta q(a|g)\mu(x)\underline{V}(x + e_1|g) + \\ \beta q(a|g)(1-\mu(x))\underline{V}(x+e_2|g) + \beta(1-q(a|g))\underline{V}(x|g) \Big\}$$

Dynamic auctions: mean field

Rewrite:

$$\underline{V}(x|g) = \frac{1}{1-\beta} \max_{a \in [0,1]} \{q(a|g)\xi(x|g) - p(a|g)\}$$

where:

$$\begin{aligned}\xi(x|g) = &\mu(x) + \beta\mu(x)(\underline{V}(x+e_1|g) - \underline{V}(x|g)) \\ &+ \beta(1-\mu(x))(\underline{V}(x+e_2|g) - \underline{V}(x|g))\end{aligned}$$

Dynamic auctions: mean field

Key observation in mean field model:

At state x, a bidder's payoff is proportional to her payoff in a *standard second price auction*, against α -1 i.i.d. bidders drawn from g each period, where she has "valuation" $\xi(x\,|\,g)$.

We show: $0 \leq \xi(x\,|\,g) \leq 1$ for all x

$\Rightarrow$ *bidding* $\xi(x\,|\,g)$ *is optimal at state* x*!*

Mean field equilibrium

The strategy ξ and bid distribution g constitute a *mean field equilibrium (MFE)* if:

(1) $\xi(\cdot \mid g)$ is an **optimal strategy** given g and

(2) g is the **steady state bid distribution** given ξ

Distribution g → Strategy $\xi(\cdot \mid g)$ → New bid distribution $\Phi(g)$

A MFE bid distribution g is a fixed point of Φ: $g = \Phi(g)$

Dynamic auctions: MFE

Theorem:

There exists a MFE of the dynamic auction with learning where at time t every bidder i bids their *virtual valuation* given posterior:

$\mathbf{E}_t[v_i] + \beta \times \mathbf{E}_t$[future marginal benefit from one additional observation]

[Iyer, Johari, Sundararajan]

A simple structural description of equilibrium!

Proof technique

We use Brouwer's fixed point theorem:

- Given g, find optimal strategy $\xi(x\,|\,g)$.
- Given g and $\xi(x\,|\,g)$, as well as initial distribution over valuations and states, find stationary distribution of resulting state Markov process
- Find new induced bid distribution g' = F(g)
- Show: F is continuous if we endow continuous cdfs on [0,1] with the sup norm
- Show: Can restrict attention to a compact set

Approximation

Asymptotic equilibrium

Is MFE a good approximation to equilibrium behavior in a finite system?

A MFE (ξ, g) has the ***AE property*** if as number of players $\to \infty$,

Profit under *any* strategy, given others play ξ	–	Profit under ξ, given others play ξ

$\longrightarrow$ 0, almost surely

Asymptotic equilibrium

In the dynamic auction model, the same continuity properties used to establish existence also serve to establish a version of the AE property.

Conclusion

Big picture

Model of multiple interacting agents

Traditional game theory makes *the model* so complex...

Big picture

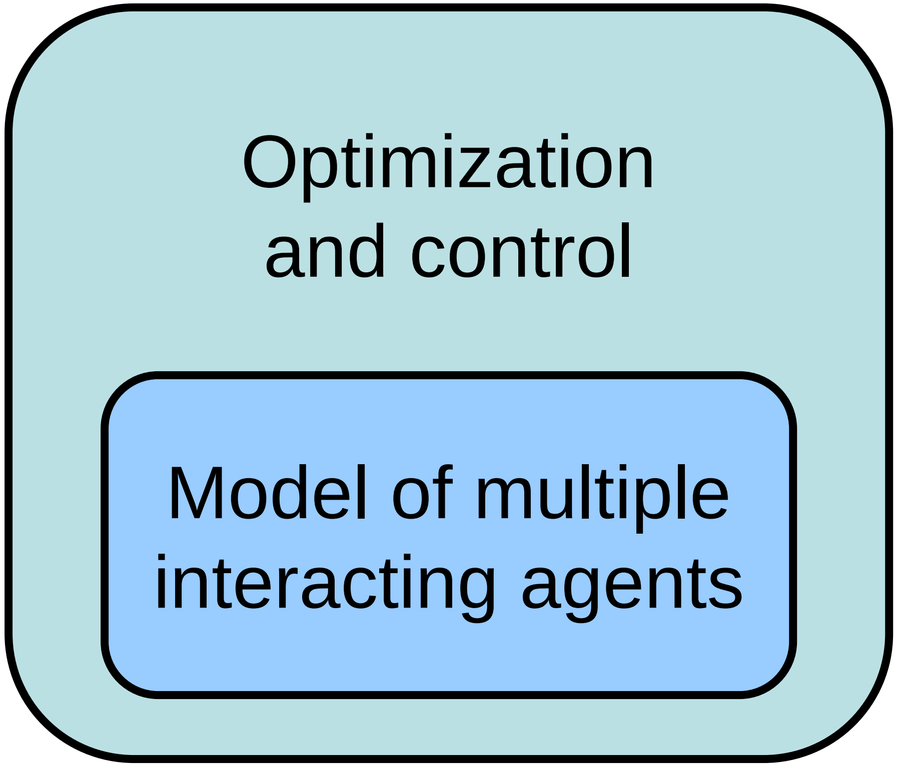

Traditional game theory makes *the model* so complex…

…that *optimization and control* are intractable.

Big picture

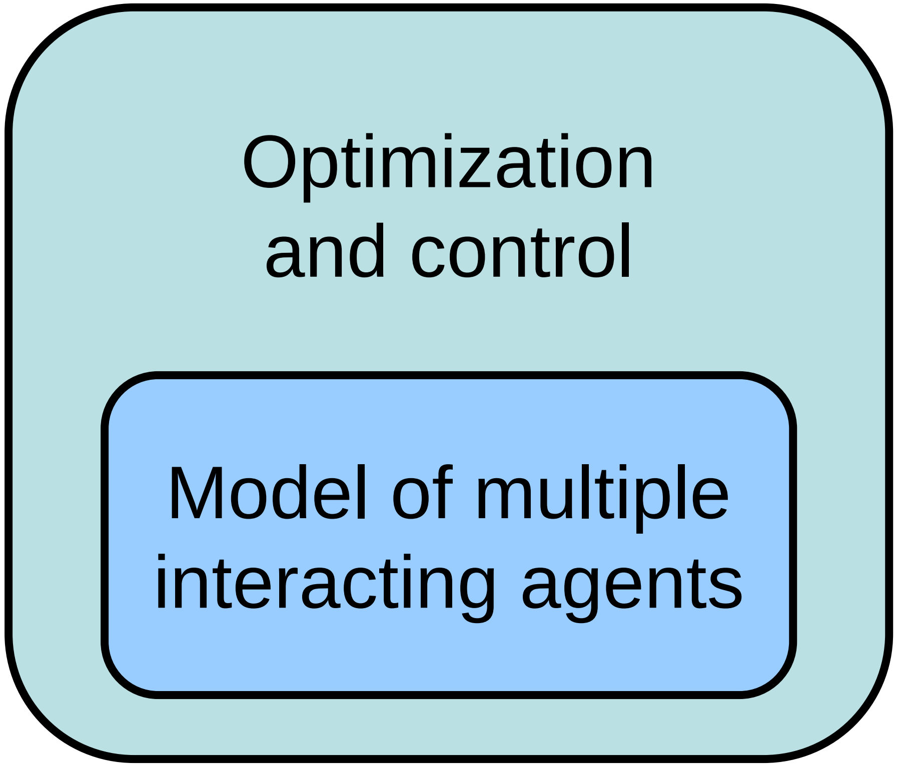

MFE simplifies the model, so we can gain structural insight into equilibria.

Conclusion

Ongoing work:

- Multiarmed bandit games
- Queueing games
- Other markets

Other issues of interest:

- Uniqueness of equilibrium
- Other objectives: average cost, regret, etc.
- Learning in other contexts
- Efficiency and distributed control