Learning and Games, day 2

Price of Anarchy and Game Dynamics

Éva Tardos, Cornell

Learning and Games

Price of Anarchy and Game Dynamics

Day 2:

- Price of Anarchy on learning outcomes
- Can we really learn this well?

Next: what can learning do that Nash cannot?

Summary from yesterday

simple games and variants:

- matching pennies,
- coordination,
- prisoner’s dilemma,
- Rock-paper-scissor

Congestion games, such as traffic routing

Auction games

Summary from yesterday (2)

- Fictitious play, and no-regret learning. Leaning algorithms that get

$$\sum_\tau c_i(s^\tau) \le (1+\epsilon) \min_x \sum_\tau c_i(x, s_i) + O(\frac{\log n}{\epsilon})$$

$$\sum_\tau u_i(s^\tau) \ge (1-\epsilon) \max_x \sum_\tau u_i(x, s_i) - O(\frac{\log n}{\epsilon})$$

n=# strategies for player

Comments: Given time T, the best possible $\epsilon = \sqrt{\log n/T}$

- Without knowing T, use variable $\epsilon = \sqrt{\log n/\boldsymbol{t}}$

choose new random r_x each step!

Summary from yesterday (3)

Outcome for learning in games

Coarse correlated equilibrium: a convex set of probability distributions on strategy vectors p_s probability that strategy vector s used

Comment: convergence to the set, but may not be to a point

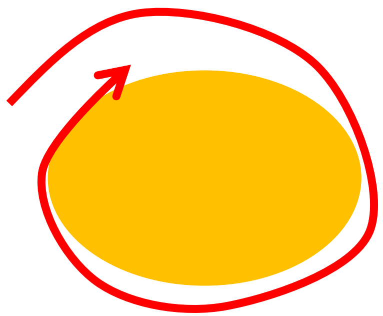

Outcomes in games:

- Fictitious play: can be a mess (such as coordination game)
- No-regret learning in 2 person 0-sum games: converges to Nash both in value and in marginal distribution (but not in actual play, see RPS)
- Leaning outcome in congestion games to be continued

What can we say about learning outcome?

Limit distribution σ of play (strategy vectors s=$(s_1, s_2, \ldots, s_n)$)

- all players i have no regret for all strategies x

$$E_{s\sim\sigma}\left(c_i(s)\right) \leq E_{s\sim\sigma}(c_i(x, s_{-i}))$$

Hart & Mas-Colell: Long term average play is (coarse) correlated equilibrium

How good are coarse correlated equilibria??

Dynamics of rock-paper-scissor

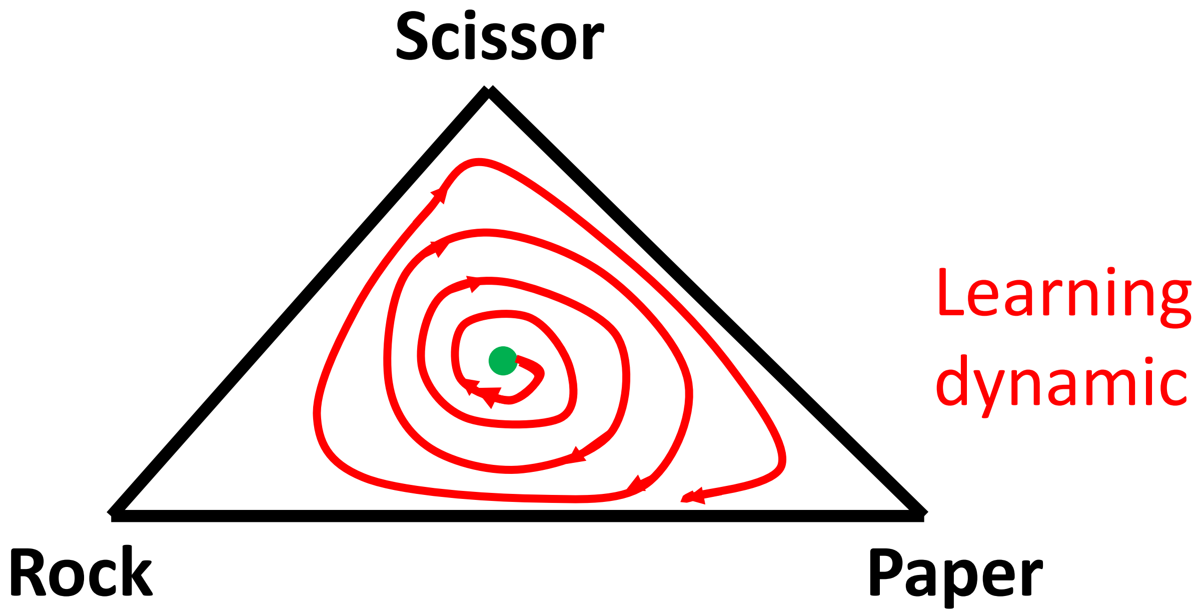

Nash:

	$\frac{1}{3}$	$\frac{1}{3}$	$\frac{1}{3}$
	R	**P**	**S**
R	0 **0**	1 -1	-1 1
P	-1 1	 **0**	1 -1
S	1 -1	-1 1	**0** **0**

Payoffs/utility

- Doesn't converge
- correlates on shared history

Coarse Correlated Equilibria: Prob x on diagonal, and prob (1-3x)/6 off diagonal, with $0 \leq 1/3 \leq 1$

Outcome of learning in games: cost minimization

- Finite set of players 1,…,n
- strategy sets S_i for player i:
- Resulting in strategy vector: s=$(s_1, \dots, s_n)$ for each $s_i \in S_i$
- Cost of player i: $c_i(s)$ or $c_i(s_i, s_{-i})$

  Pure Nash equilibrium if $c_i(s) \leq c_i(s_i', s_{-i})$ for all players and all alternate strategies $s'_i \in S_i$

- Social welfare: $\text{cost}(s) = \sum_i c_i(s)$

  Optimum: $OPT = \min_s \sum_i c_i(s)$

Quality of Learning Outcome

Price of Anarchy [Koutsoupias-Papadimitriou'99]

$$PoA = \max_{s\ Nash} \frac{cost(s)}{Opt}$$

Assuming **no-regret learners** in fixed game: [Blum, Hajiaghayi, Ligett, Roth'08, Roughgarden'09]

$$PoA = \lim_{T\to\infty} \frac{\sum_{t=1}^{T} cost(s^t)}{T\ Opt}$$

Example: Model of Routing Game

- A directed graph $G = (V,E)$
- source–sink pairs s_i,t_i for $i=1,..,k$

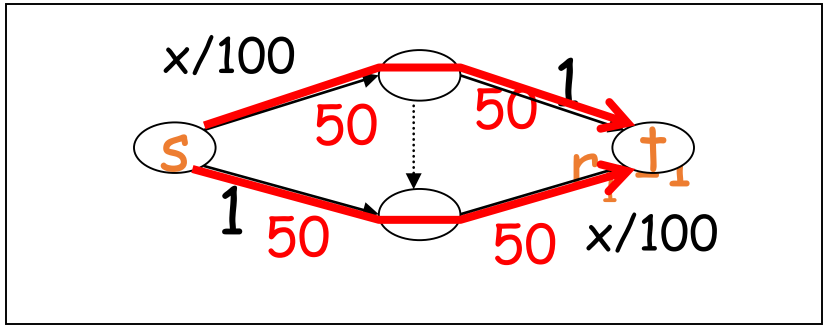

- Goal minimum delay:

  delay adds along path

  edge-cost/delay is a function $c_e(\cdot)$ of the load on the edge e

Delay Functions

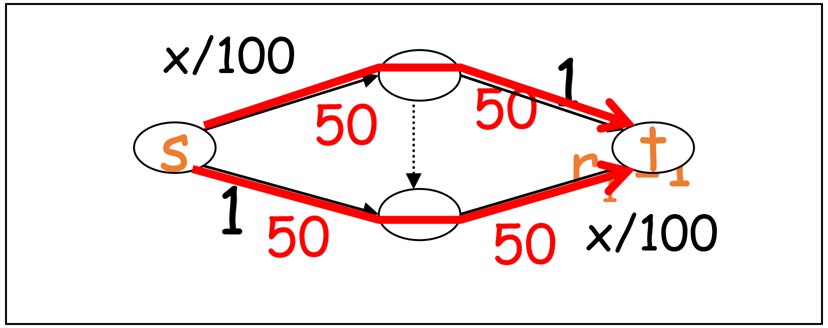

Assume $c_e(x)$ continuous and monotone increasing in load x on edge

No capacity of edges for now

Example to model capacity u:

$c_e(x) = a/(u-x)$

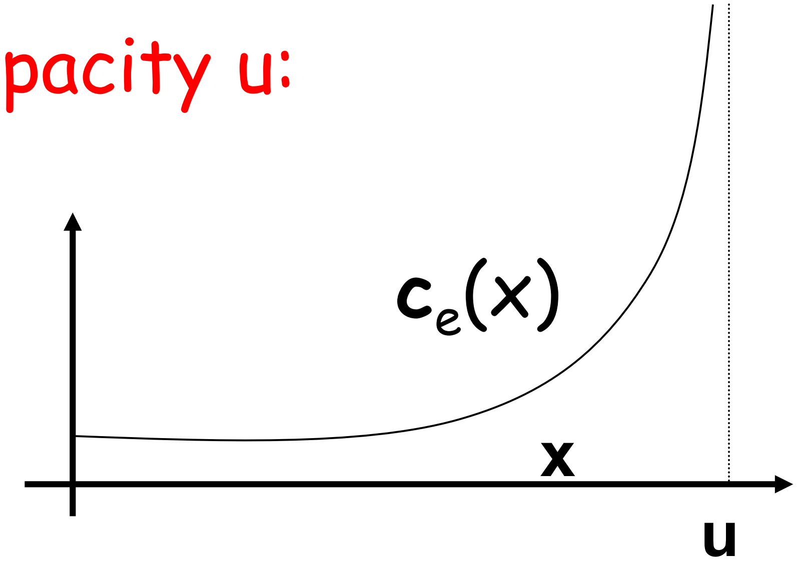

Goal's of the Game: min delay

Personal objective: minimize

Costs

$c_P(f)$ = sum of ~~delays~~ of edges along **P** (wrt. flow **f**)

$$\mathbf{c_P(f)} = \sum_{e \in P} c_e(f_e)$$

Overall objective:

Cost -

C(f) = total ~~delay~~ of a flow **f**: $= \Sigma_P\, f_P \cdot c_P(f)$

= - social welfare

or total/average delay

Also:

C(f) = $\Sigma_e\, f_e \cdot c_e(f_e)$

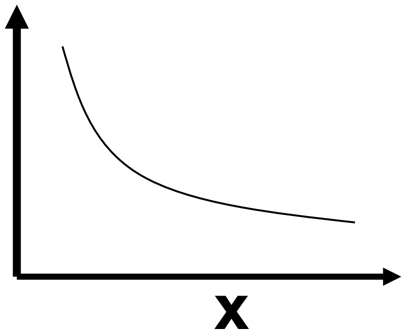

Equilibrium

Each use must not regret not following her optimal path

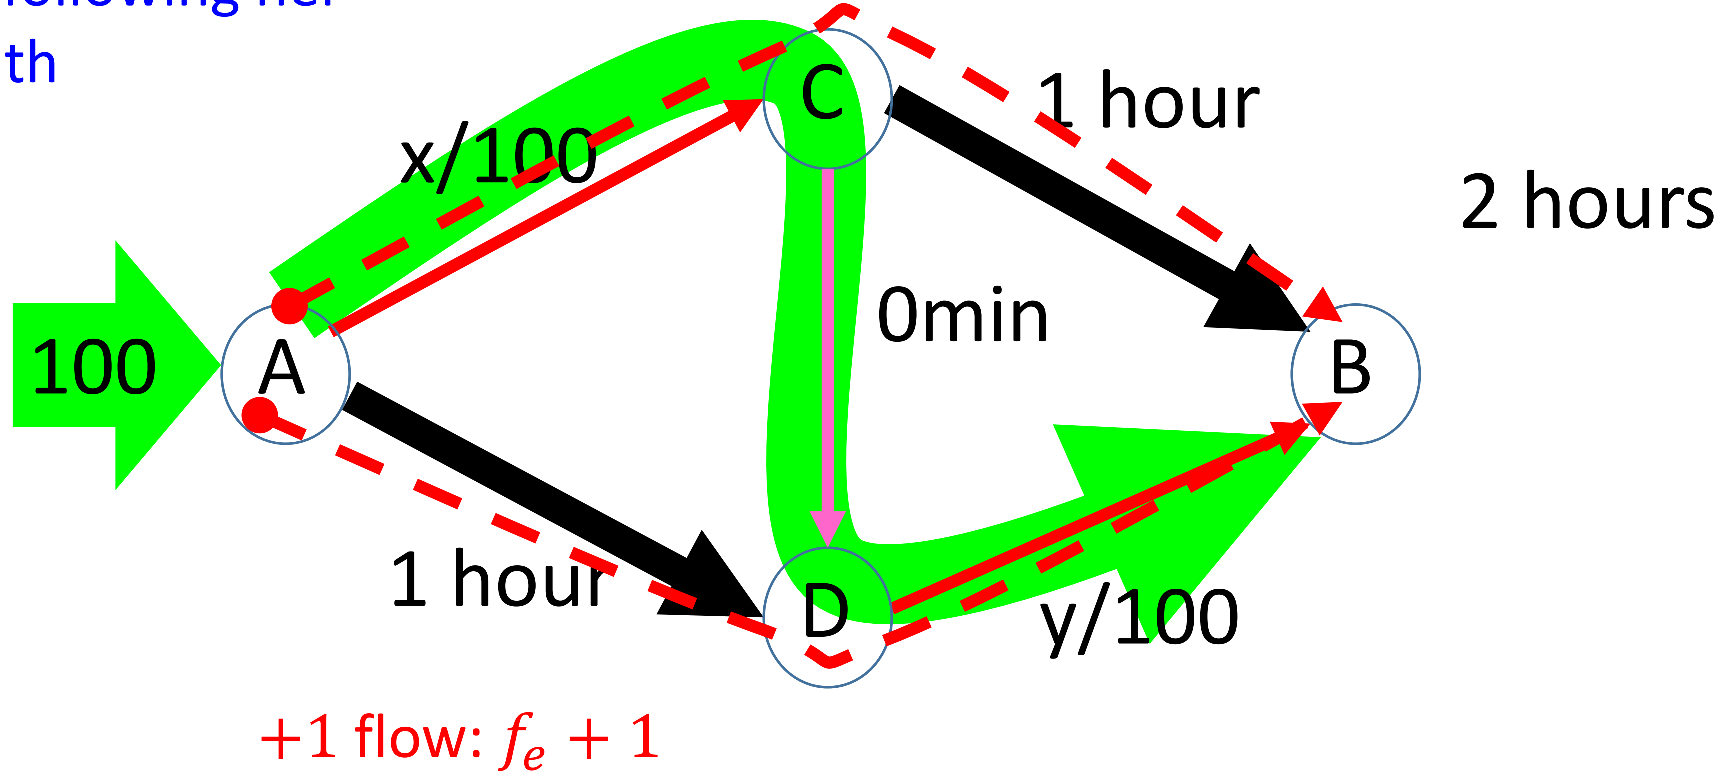

No regret inequality for flow

- f_e Nash flow on edge e, P path used by Nash, Q path used by opt

No regret =

$$\sum_{e \in P} c_e(f_e) \leq \sum_{e \in P \cap Q} c_e(f_e) + \sum_{e \in Q \setminus P} c_e(f_e + 1)$$

- Without the +1 nonatomic flow: assumes +1 is too small to really make a difference

No-regret inequality with small flow unit

For flow using path Q and alternate Q we have

$$\sum_{e\in P} c_e(f_e) \le \sum_{e\in P\cap Q} c_e(f_e) + \sum_{e\in Q\setminus P} c_e(f_e+\delta)$$

Non-atomic flow, when each user is small, but total flow remains the same: Limit as size as $\delta \to 0$

- Nash inequality limit for path P used and alternate !

$$\sum_{e\in P} c_e(f_e) \le \sum_{e\in Q} c_e(f_e)$$

Exercise

a. An alternate definition for equilibrium of a non-atomic flow would be for each path P carrying flow and each alternate path Q if more a δ amount of P to Q to get a new flow $\tilde{f}$ then $\sum_{e \in P} c_e(f_e) \leq \sum_{e \in Q} c_e(\widetilde{f_e})$

   Under what conditions is this equivalent to the definition given.

b. Nash equilibrium of non-atomic flow is the true optimum of

$$\Phi(f) = \sum_e \int_0^{f_e} c_e(\xi) d\xi.$$

   Note that this is convex if c_e are monotone increasing

Price of Anarchy: proof technique [Roughgarden'09]

- What we can work with:

  Optimum $s^* = (s_1^*, s_2^*, \dots, s_n^*)$

  Nash: $s = (s_1, s_2, \dots, s_n)$

- What we know:

$$c_i(s) \le c_i(s_i', s_{-i}) \text{ for all i and all } s_i' \in S_i$$

Use it for all players and sum

$$c(s) = \sum_i c_i(s) \le \sum_i c_i(s_i^*, s_{-i})$$

Proof smooth games

Nash property gave us (s is Nash, s* optimum)

$$c(s) = \sum_i c_i(s) \le \sum_i c_i(s_i^*, s_{-i})$$

Game is smooth if for some μ<1 and λ>0 and all s and s*

$$\sum_i c_i(s_i^*, s_{-i}) \le \lambda \boxed{c(s^*)} + \mu \boxed{c(s)}$$

(λ,μ)-smooth

If Opt <<cost(s), some player will want to deviate to s_i^*

Theorem: (λ,μ)-smooth game ⇒

Price of anarchy at most $\lambda/(1-\mu)$

Learning and price of anarchy (in smooth games)

Use approx no-regret learning:

$\sum_t c_i(s^t) \leq (1+\epsilon)\sum_t c_i(s_i^*, s_{-i}^t) + R$ for all players

A cost minimization game is (λ,μ)-smooth (λ > 0; μ< 1):

$\sum_t \sum_i c_i(s_i^*, s_{-i}^t) \leq \lambda \sum_t Opt + \mu \sum_t c(s^t)$

A approx. no-regret sequence s^t has

$$\frac{1}{T}\sum_t c(s^t) \leq \frac{(1+\epsilon)\lambda}{1-(1+\epsilon)\mu}\mathrm{Opt} + \frac{\mathrm{n}}{\mathrm{T}(1-(1+\epsilon)\mu)}\mathrm{R}$$

Note the convergence speed! $R = \frac{\log d}{\epsilon}$, so error

$$\boxed{\frac{\mathrm{n}}{\mathrm{T}} \cdot \frac{\log d}{\epsilon(1-(1+\epsilon)\mu)}}$$

Foster, Li, Lykouris, Sridharan, T, NIPS’16

Proving smoothness for flows

What we need $\sum_i c_i(s_i^*, s_{-i}) \leq \lambda c(s^*) + \mu\, c(s)$

Nash inequality for s to t user using path P with alternate path Q

$$\sum_{e\in P} c_e(f_e) \leq \sum_{e\in Q} c_e(f_e)$$

Sum over paths $Q_i = P_i^*$ in opt with f Nash and f^* optimal

$$\sum_P f_p \sum_{e\in P} c_e(f_e) \leq \sum_Q f_Q^* \sum_{e\in Q} c_e(f_e)$$

and rearranging sums

$$\sum_e f_e c_e(f_e) \leq \sum_e f_e^* c_e(f_e))$$

We need

$$\sum_e f_e^* c_e(f_e) \leq \lambda \sum_e f_e^* c_e(f_e^*) + \mu \sum_e f_e c_e(f_e)$$

Claim: true edge by edge

Linear delay is smooth

Claim: $f^* \bullet c(f) \leq f^* \bullet c(f^*) + \frac{1}{4} f \bullet c(f)$

assuming $c(f)$ linear: $\lambda = 1$; $\mu = \frac{1}{4}$

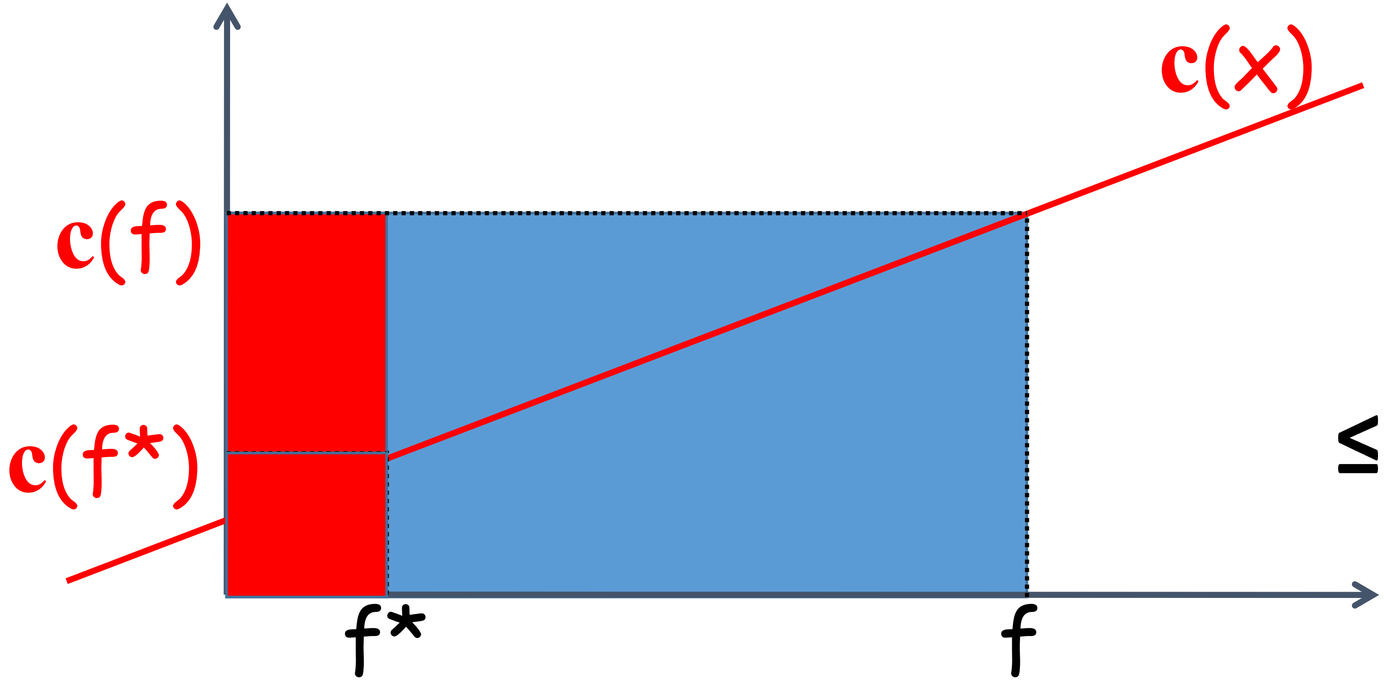

$f^* c(f)$
$- f^* c(f^*)$
$\leq \frac{1}{4} f\, c(f)$

Sharper results for non-atomic games

Theorem (Roughgarden-& '02):

In any network with linear cost functions the worst price of anarchy (in non-atomic games) is at most 4/3

Proof: $\left(1, \frac{1}{4}\right)$-smooth implies price of anarchy of $\lambda/(1-\mu)$= $1/(1-1/4)$=4/3

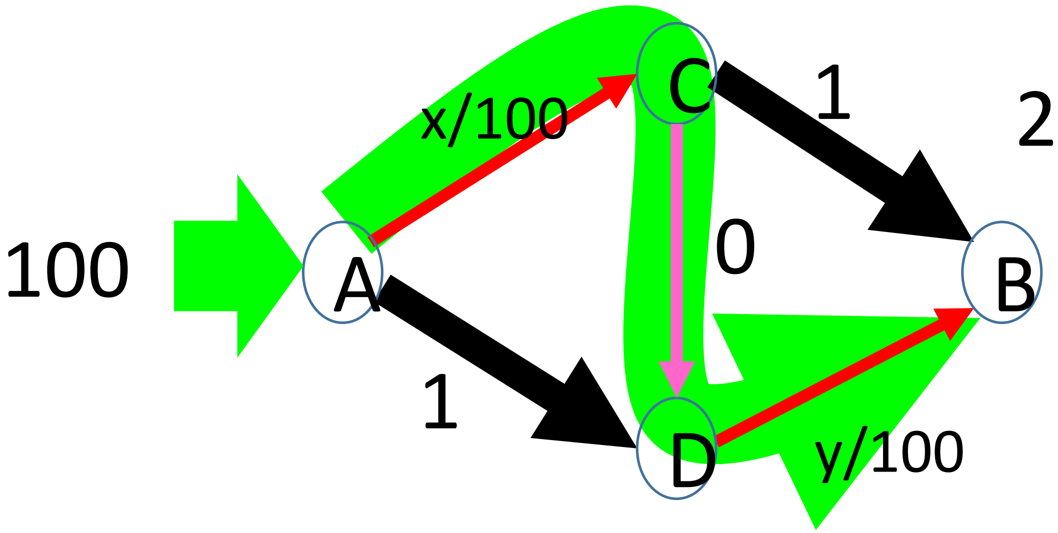

2 vs 1.5

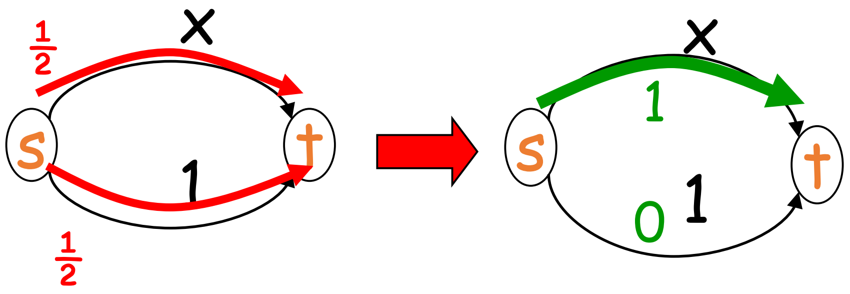

Sharper results for non-atomic games

Theorem (Roughgarden'03):

In any network with any class of convex continuous latency functions the worst price of anarchy (in non-atomic games) is always on two edge network

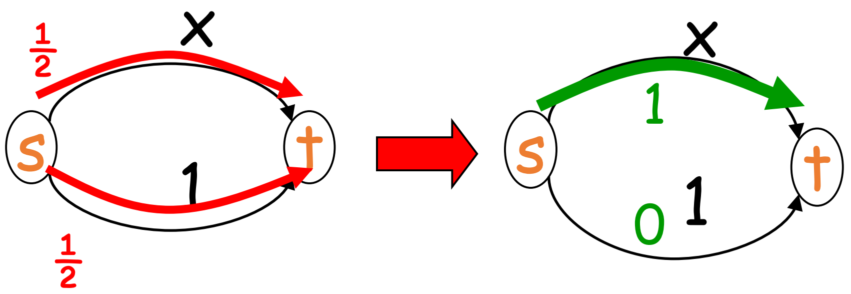

Corollary:
price of anarchy for degree d polynomials is O(d/log d).

Monotone delay is (1,1)-smooth

Claim: $\mathbf{f}^{*}\boldsymbol{c}(\boldsymbol{f}) \leq \boldsymbol{max}\left(\boldsymbol{f}^{*}\boldsymbol{c}(\boldsymbol{f}^{*}), \boldsymbol{f}\boldsymbol{c}(\boldsymbol{f})\right) \leq \boldsymbol{f}^{*}\boldsymbol{c}(\boldsymbol{f}^{*}) + \boldsymbol{f}\ \boldsymbol{c}(\boldsymbol{f})$

assuming $\mathbf{c}$ (f) monotone: λ =1; μ=1

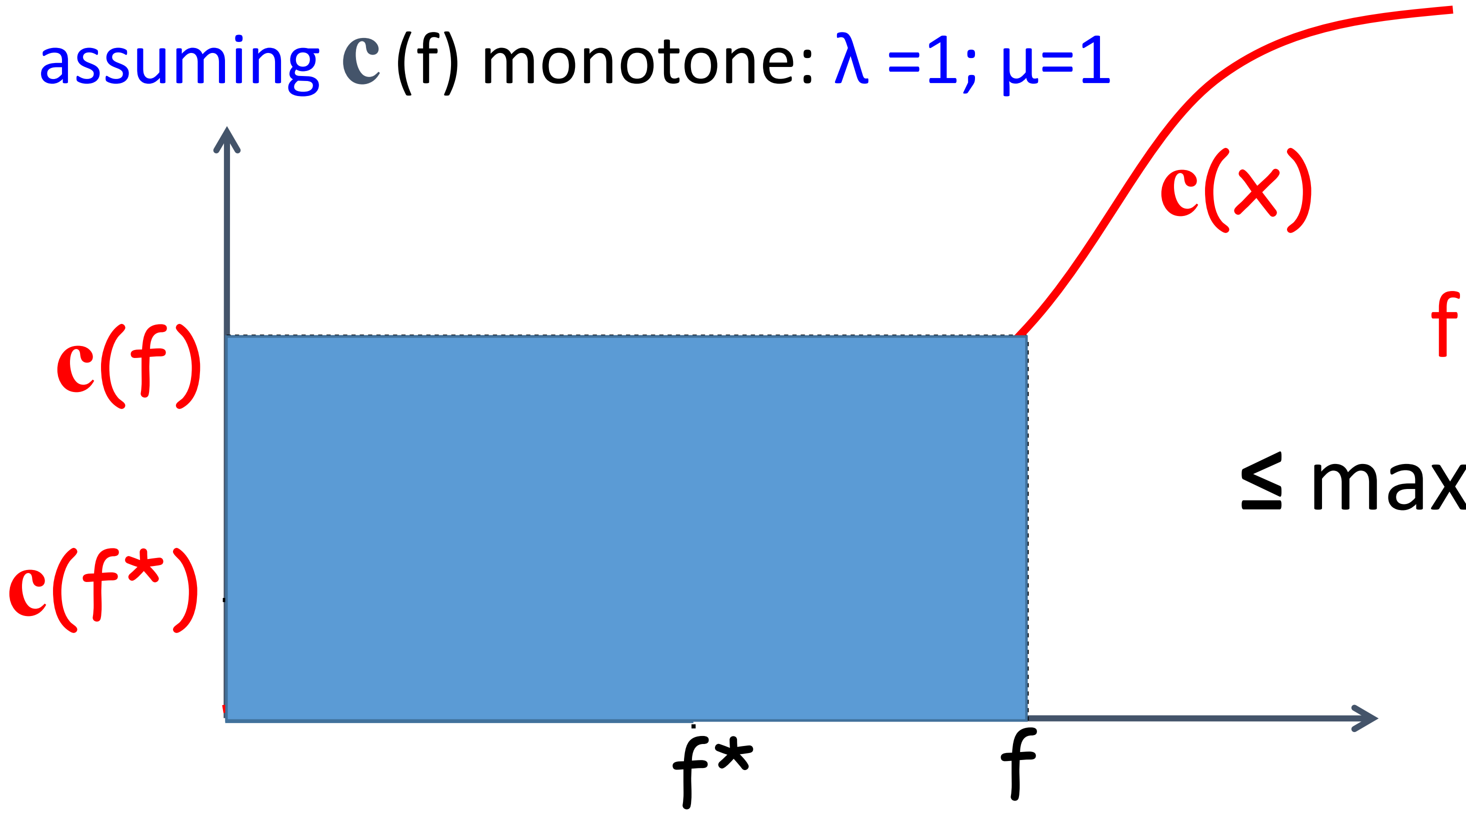

f*c(f)

≤ max(f c(f),f*c(f*))

High Social Welfare: Price of Anarchy in Routing

Theorem (Roughgarden-T'02):

In any network with continuous, non-decreasing cost and very small users

cost of Nash with rates r_i for all i	$\leq$	cost of opt with rates $2r_i$ for all i

Proof if Nash carries ½ of the flow

$$\sum_e f_e c_e(f_e) \underset{\text{Nash}}{\leq} \sum_e \frac{1}{2} f_e^* c_e(f_e) \underset{\text{smooth}}{\leq} \frac{1}{2}[\lambda \sum_e f_e^* c_e(f_e^*) + \mu \sum_e f_e c_e(f_e)]$$

Implying $c(f) \leq \frac{\frac{1}{2}\lambda}{(1-\frac{1}{2}\mu)} c(f^*)$, so (1,1)- smooth implies the theorem!

Exercise

A popular delay model is $c_e(x) = \frac{a_e}{u_e - x}$, modeling

capacity u_e and delay on empty road $\frac{a_e}{u_e}$

a. Show that for any rates and any capacities, optimal flow has total cost $\geq$ Cost of Nash with double capacities $u'_e = 2u_e$

b. Anything useful follows if capacities $u'_e = \alpha \cdot u_e$ for some other $\alpha > 1$

Linear delay atomic flow

Atomic game (players with >0 traffic) with linear delay (5/3,1/3)-smooth (Awerbuch-Azar-Epstein'05 & Christodoulou-Koutsoupias'05)

$\Rightarrow$ 2.5 price of anarchy

- Need to prove: for all nonnegative integers $x = f^*(e)$ and $y = f(e)$

$$x(y+1) \leq \frac{5}{3}x^2 + \frac{1}{3}y^2$$

That is: 3xy +3x $\leq 5x^2 + y^2$

HW ??

Theorem: Price of anarchy for polynomials of degree at most d at most exponential in d: $O(2^d\, d^{d+1})$

Suri-Toth-Zhou SPAA'04 (special case)
Awerbuch-Azar-Epstein STOC'05
Christodoulou-Koutsoupias STOC'05

Homework

Smoothess for value maximization games

- Utility of player i: $u_i(s)$ or $u_i(s_i, s_{-i})$

  Pure Nash equilibrium if $u_i(s) \geq u_i(s_i', s_{-i})$ for all players and all alternate strategies $s'_i \in S_i$

- Suppose $\sum_i u_i(s_i^*, s_{-i}) \geq \lambda \sum_i u_i(s^*) - \mu \sum_i u_i(s)$ for some $\lambda, \mu > 0$, an optimal solution vector s^* and any solution s. What does this imply about the price of anarchy?

[Roughgarden'09]

$\sum_i u_i(s) \geq \frac{\boldsymbol{\lambda}}{\boldsymbol{\mu+1}} \sum_i u_i(s^*)$

A utility game: Auctions as (Bayesian) game

First Example: Single item first price

- Auction sets a price p (full info, pure Nash).

Players

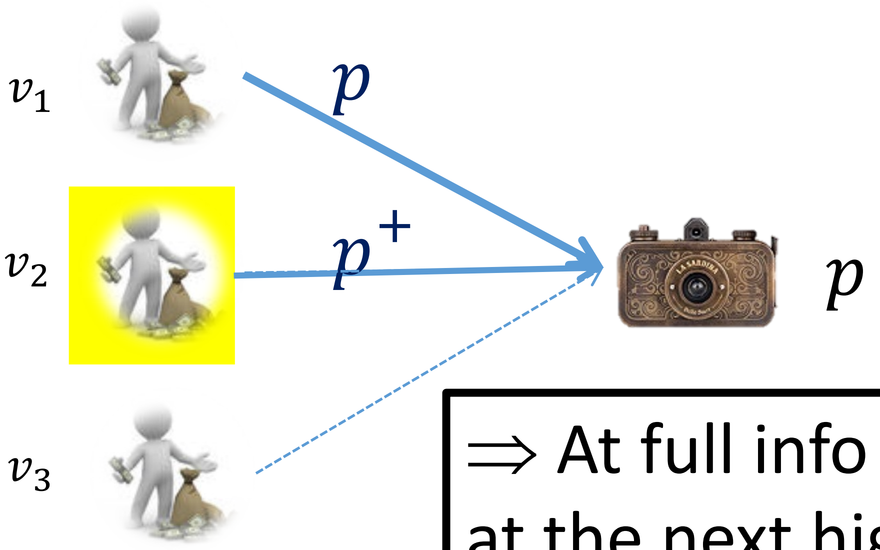

Can win by bidding p^+

$\Rightarrow$ At full info Nash player with max v_i wins at the next highest price

First price auction with uncertainty?

- Bayesian game
- Randomized bid

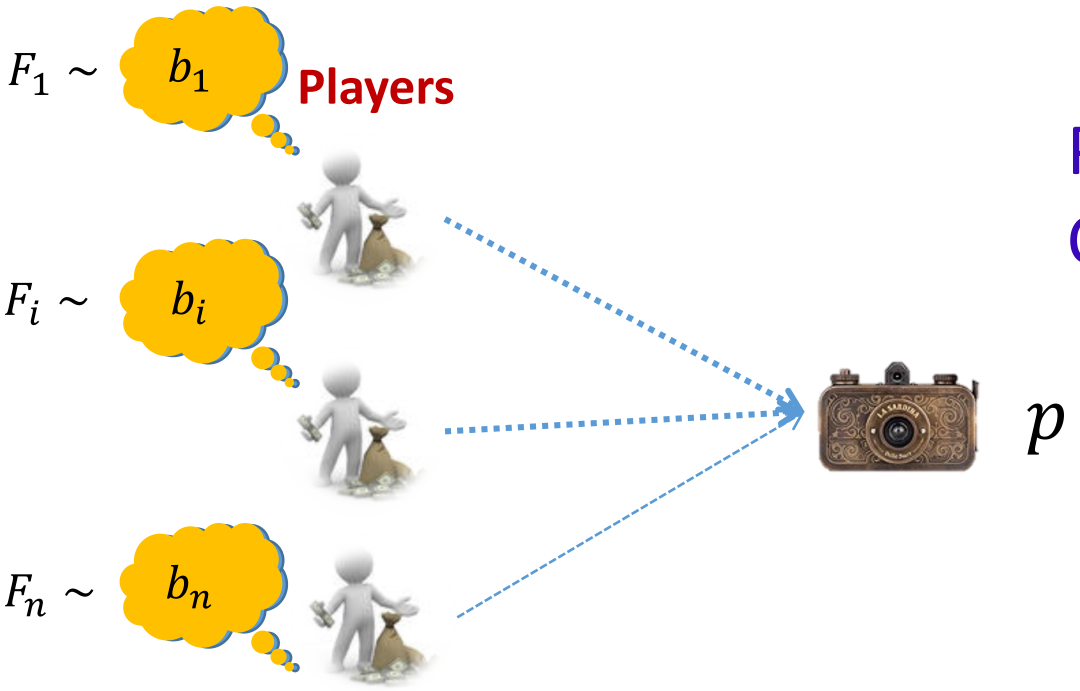

Price p random
Cannot bid p^+

Auction games:

- Finite set of players 1,…,n
- strategy sets S_i for player i: bid on some items (not a finite set)
- Resulting in strategy vector: s=$(s_1, \dots, s_n)$ for each $s_i \in S_i$
- Utility player i: $u_i(s)$ or $u_i(s_i, s_{-i})$
  - We assume quasi-linear utility, and no externalities:
  - If player wins set if items A_i and pays p_i her value is $u_i(A_i, p_i) = v_i(A_i) - p_i$

- Social welfare? (include auctioneer): $\sum_i v_i(A_i) = \sum_i u_i(A_i) + \underbrace{\sum_i p_i}_{\text{Revenue}}$

Bayes Nash analysis

Strategy: bid as a function of value $b_i(v)$

Nash: $E_{v_{-i}b}\left[u_i\left(b(v)\right)|v_i\right] \geq E_{v_{-i}b_{-i}}\left[u_i\left(b_i', b_{-i}(v_{-i})\right)|v_i\right]$

for all b_i'

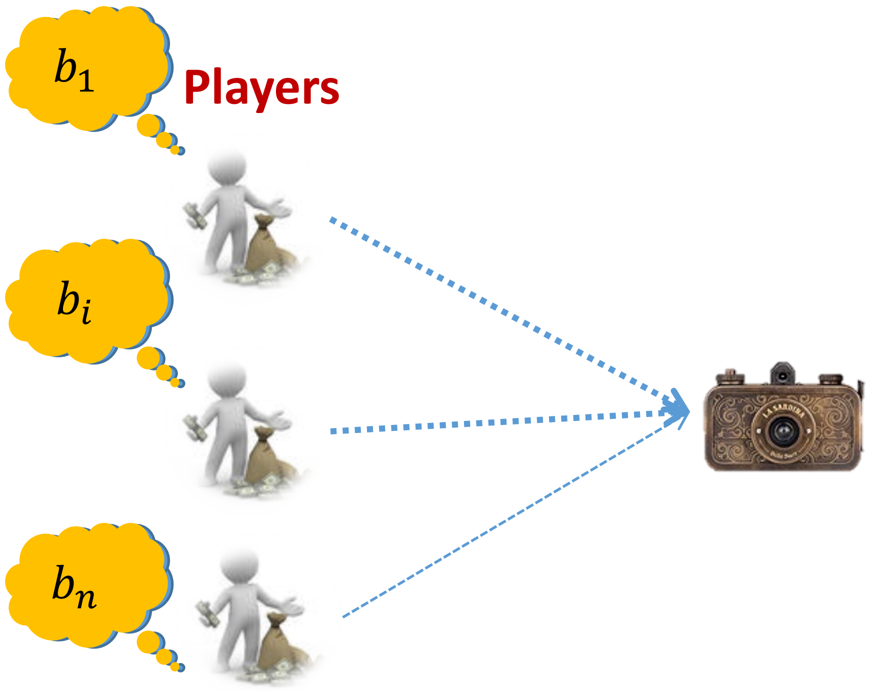

Example: [0,1] uniform value independent

- Two players:
- Assume both use deterministic, monotone, and identical bidding functions b(v)
  - Person with larger value wins
  - Bid must maximize utility:

    alternate bid for a player with value v: bid b(z) (pretend to have value z)

$$\mathrm{v} = \underset{z}{\operatorname{argmax}}\, z\,(v - b(z))$$

Prob of winning   value   price

$\rightarrow$ v-b(v)-vb'(v)=0

Solved by b(v)=v/2

First price single item auction

- Uniform independent [0,1] value n players:

  bid b(v)=$\frac{n-1}{n}v$ (more competition bid more aggressively)
- Independent identical distributions $\mathcal{F}$ and n players:

  bid b(v)= E(max of n-1 draws from $\mathcal{F}$|each $\leq v$)

BTW, Second price auction: bid your value,

first price bid = expected payment

revenue equivalence (Meyerson)

If distribution not identical and independent: big mess!!!

Smoothness for auctions

Auction game is λ-smooth if for some λ>0 and some strategy s* and all s we have

$$\sum_i u_i(s_i^*, s_{-i}) \geq \lambda opt - Rev(s)$$

R(s) = revenue at bid vector s

Theorem: [Syrgkanis-T'13] λ-smooth auction game ⇒

Price of anarchy for any $\leq \frac{1}{\lambda}$

Social welfare: $\sum_i u_i(s) + R(s)$

Robust Analysis: first price auction

No regret: $u_i(b) \geq \quad u_i\left(\frac{1}{2}v_i, b_{-i}\right) \geq \frac{1}{2}v_i - p\ ,0$

either i wins or price above p $\geq \frac{1}{2}v_i$

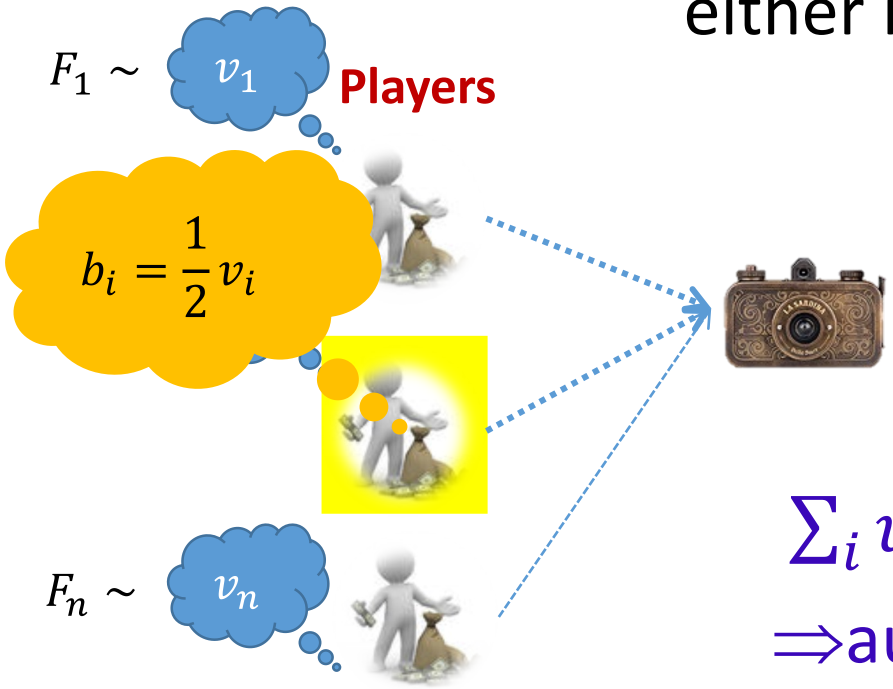

- Apply this to the top value

\+ winner doesn't regret paying

$\sum_i u_i\left(\frac{v_i}{2}, b_{-i}\right) \geq \left(\max\left(\frac{v_i}{2}\right) - p\right) + \sum_i 0$

$\Rightarrow$auction is 1/2- smooth

$\Rightarrow$a price of anarchy of 2

(actually… $(e-1)/e \approx 0.63$)

Bayes Nash analysis: Bayesian extension (I)

Strategy: bid as a function of value $b_i(v)$

Nash: $E_{v_{-i}b}\left[u_i\big(b(v)\big)|v_i\right] \geq E_{v_{-i}b_{-i}}\left[u_i\big(b_i', b_{-i}(v_{-i})\big)|v_i\right]$

for all b_i'

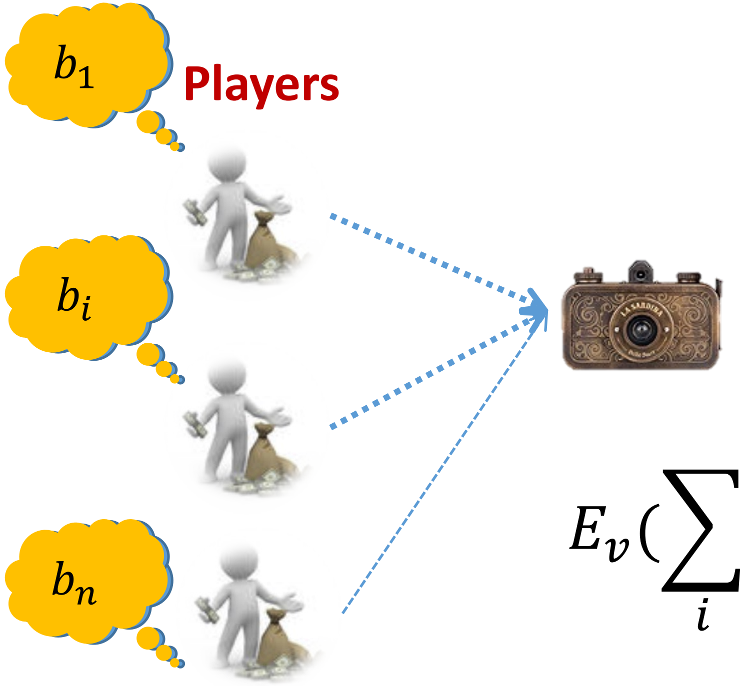

Same bound on price of anarchy,
same prof (take expectation)

$$E_v\Big(\sum_i u_i(b)\Big) \geq \sum_i E_v\left(u_i\left(\frac{v_i}{2}, b_i\right)\right) \geq \lambda E_v(Opt(v)) - \mu E_v(Rev(b)$$

No need to bid $\frac{v_i}{2}$ just don't regret it!

Smoothness and Bayesian games

We had $b_i^*(v) = \frac{v_i}{2}$, 0.5-smooth: Bid depends only on the players own value!

Theorem: Auction is λ-smooth and b_i^* is a function of v_i only, then price of anarchy bounded by $1/\lambda$ for arbitrary (private value) type distributions

Proof: just take expectations!

Price of anarchy in multi-item

- First price is auction Hassidim, Kaplan, Mansour, Nisan EC'11)
  Price of anarchy 1.58…
- All pay auction price of anarchy 2
- First position auction (GFP) is price of anarchy 2
- Variants with second price (see also Christodoulou, Kovacs, Schapira ICALP'08) price of anarchy 2

Other applications include:

- public goods
- Fair sharing (Kelly, Johari-Tsitsiklis) price of anarchy 1.33
- Walrasian Mechanism (Babaioff, Lucier, Nisan, and Paes Leme EC'13)

All pay auction (example)

Trouble: b_i*(v) depends on other player's valuation!

Claim: all pay auction is 1/2-smooth

Max value player: b_i*(v) uniform random [0,v].

All others: bid b_i*(v)=0

i not the top value: $u_i(b_i^*, b_{-i}) = 0$

i is the top value, and suppose max other bid is b.

If b>v_i we are set: $\sum_i u_i(b_i^*, b_{-i}) \geq -\frac{v_i}{2} \geq \frac{1}{2} Opt \; - b$

Else expected value for player i

$$\mathrm{E}(\mathrm{u_i}(\mathrm{b_i^*}, b_{-i})) = -\frac{v_i}{2} + v_i \frac{v_i - b}{v_i} \geq \frac{1}{2} v_i - b$$

Bayesian extension theorem

Theorem [Syrgkanis-T'13] Auction game is λ-auction smooth, and values are drawn from independent distribution, than the Price of anarchy in the Bayesian game is at most $1/\lambda$

Extension theorem: OK to only think about the full information game!

Proof idea: bid b*(v)….

Trouble: depends on other players and hence we don't know……

Bayesian extension theorem

Notation v=$(v_1, \dots v_n)$ value vector and use $b_i^*(v) = b_i^*(v_i, v_{-i})$

Idea: random sample opponent w_{-i}, and bid $b_i^*(v_i, w_{-i})$

Any fixed value v_i, and any player i we get

$$E_{w_{-i}b_{-i}}(u_i(b_i^*(v_i, w_{-i}), b_{-i}|v_i) \leq E_{b_{-i}}(u_i(b)|v_i)$$

Rename $w_{-i} = v_{-i}$, and also take expectation over v_i

$$E_{vb}(u_i(b_i^*(v), b_{-i}) \leq E_{vb}(u_i(b))$$

Bayesian extension theorem (cont.)

$$E_{vb}(u_i(b_i^*(v), b_{-i}) \leq E_{vb}(u_i(b))$$

Recall smoothness: for all fixed v and b

$$\sum_i u_i(b_i^*(v), b_{-i} \middle| v_i) \geq \lambda\, Opt(v) - \mu Rev(b)$$

Combine and take expectation over b and v (these are independent in the above!!!)

$$E_{vb}(\sum_i u_i(b)) \geq E_{vb}(\sum_i u_i(b_i^*(v), b_{-i})) \geq \lambda E_v(Opt(v) - \mu\, E_b(R(b))$$

Multiple items (e.g. unit demand bidders)

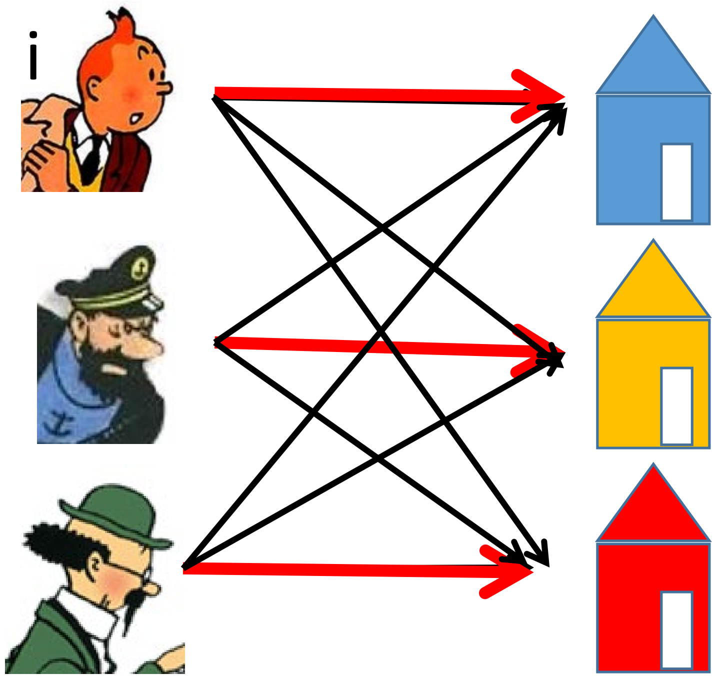

Value if i gets subset S is $v_i(S)$

for example: $v_i(S) = \max_{j \in S} v_{ij}$

Optimum is max value matching!

$$\max_{M^*} \sum_{ij \in M^*} v_{ij}$$

Extension also if $v_i(A)$ submodular function of set A

Also for diminishing value of added items:

$$A \subset B \Rightarrow \; v_i(A+x) - v_i(A) \geq v_i(B+x) - v_i(B)$$

Multi-item first prize auction with unit demand bidders

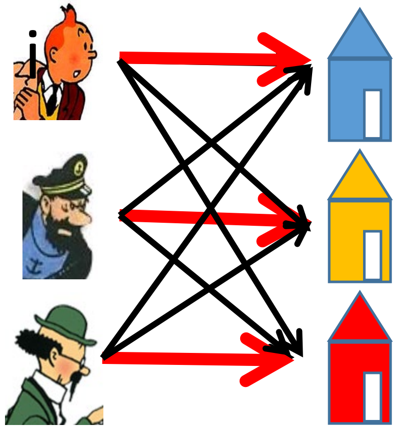

- Optimal solution $\max_{M^*} \sum_{ij \in M^*} v_{ij}$
- A bid vector b^* inducing optimal solution i bids $v_{ij}/2$ on item j_i^* assigned in i in opt ($(i, j_i^*) \in M^*$)

- Smoothness?
- $\sum_i u_i(b_i^*, b_{-i}) \geq 1/2 \sum_i v_{ij_i^*} - \sum_j \max_i b_{ij} = \frac{1}{2} OPT - Rev$
- True item by item!

Trouble: bidding is very hard!

- Special case: unit demand buyer, all items has the vale value $v \gg 0$
- There are n items
- Opponents bid 1 on some items, and $h > v$ on all others
- Possible set that they bid 1 on: $S_1, S_2, \ldots, S_k$ uniformly likely
- Use 2nd price or assume you can bid 1, and will win (and pay 1) if max other bid is 1

$v > \mathrm{nk}$ implies,

- optimal bid always wins some item
- Wanted: T s.t. $S_i \cap T \neq \emptyset$ for all i and $\sum_i |S_i \cap T|$ as small as possible

Fining optimal strategy NP-complete

- Given sets $S_1, S_2, \ldots, S_k$
- Wanted: T s.t. $S_i \cap T \neq \emptyset$ for all i and $\sum_i |S_i \cap T|$ as small as possible

Assume: every s the number $\#\{i: s \in S_i\} = \mathrm{r}$ is the same

Then $\sum_i |S_i \cap T| = r|T|$

Wanted: T s.t. $S_i \cap T \neq \emptyset$ for all i and $|T|$ minimal

This is hitting set

Still NP-complete (= set-cover with equal size sets)

in fact, hard to approximate within $\approx \log r$

What is possible to do?

Why is no-regret so hard? So many bids to consider $(b_1, b_2, \ldots, b_n)$ all possible bids on all items

Simplifications:

- Do not bid $b_j > v_j$, still bid space is $\prod_j [0, v_j]$
- Discretize, only bid multiples of ϵ, being off my an ϵ can only cause ϵ regret! Only $\prod_j v_j/\epsilon$ options
  - Assume (k-1)$\epsilon < b < k\epsilon$
  - If b wins: so does kϵ and pays too much by ϵ
  - If kϵ wins and b looses kϵ is better off.
- Bid on a single item only? Regret can be huge!

Bidding options that are possible to not regret [Daskalakis-Syrgkanis'16]

- Idea: strategy space names set S of items to buy, regardless of price

- If no regret:

$$\sum_\tau v_i(s^\tau) - p_i(s^\tau) \geq (1-\epsilon)\max_{S_i} \sum_\tau v_i(S_i, s_{-i}^\tau) - \sum_\tau p(S_i, s_{-i}^\tau) - Regret$$

Items in j ∈ S_i are evaluated against their average price! $|\mathrm{T}| v_j - \sum_\tau p^\tau(j)$

Choosing sets, versus bidding for a set

- Second price:

  selected set S: bid v_j for $j \in S$ and 0 elsewhere. This is strictly better!

- Is no regret for this good enough for social welfare?

  Let S_i^* be set awarded to i in optimum. We get

  $\sum_\tau u_i(S^\tau) \geq Tv_i(S_i^*) - \sum_\tau Rev^\tau\,(S_i^*)$

  Sum over all players

  $\sum_\tau \sum_i u_i(s^\tau) \geq T \sum_i v_i(S_i^*) - \sum_\tau \sum_i Rev^\tau(S_i^*) = T\ OPT - \sum_\tau Rev^\tau$

Choosing sets, versus bidding for a set

First price:

selected set S: bid $\frac{1}{2}v_j$ for $j \in S$ and 0 elsewhere.

If no regret:

$$\sum_\tau v_i(s^\tau) - p_i(s^\tau) \geq \frac{1}{2}\max_{S_i} \sum_\tau v_i(S_i, s_{-i}^\tau) - \sum_\tau p(S_i, s_{-i}^\tau) - Regret$$

Is this no regret for this good enough for social welfare?

Let S_i^* be set awarded to i in optimum. We get

$$\sum_\tau u_i(S^\tau) \geq \frac{1}{2}Tv_i(S_i^*) - \sum_\tau Rev^\tau\,(S_i^*)$$

Sum over all players

$$\sum_\tau \sum_i u_i(s^\tau) \geq \frac{1}{2}T\sum_i v_i(S_i^*) - \sum_\tau \sum_i Rev^\tau(S_i^*) = \frac{1}{2}T\,OPT - \sum_\tau Rev^\tau$$

Magic Fictitious play and no regret

Fictitious play = best respond to past history of other players

$$s_i^t = argmax_x \sum_{\tau=1}^{t-1} u_i(x, s_{-i}^\tau)$$

Magic enhancement of Fictitious play with response included

$$s_i^t = argmin_x \sum_{\tau=1}^{t} u_i(x, s_{-i}^\tau)$$

Theorem 1: Magic fictitious play has no regret.

Proof: by induction we claim that

By choice of $\boldsymbol{s_i^t}$

$$\sum_{\tau=1}^{t} u_i(s^\tau) \geq \sum_{\tau=1}^{t} u_i\left(\boldsymbol{s_i^t}, s_{-i}^\tau\right) = \max_x \sum_{\tau=1}^{t} u_i(x, s_{-i}^\tau)$$

IH $\Downarrow$ with $x = s_i^t$

$$\sum_{\tau=1}^{t} u_i(s^\tau) = \sum_{\tau=1}^{t-1} u_i(s^\tau) + u_i(s^t) \geq \sum_{\tau=1}^{t-1} u_i\left(s_i^t, s_{-i}^\tau\right) + u_i(s^t)$$

QED

Follow the perturbed leader has small regret (Theorem)

Follow the perturbed leader: chose a random r_j, for all items j

$$\text{select } argmin_x [\textstyle\sum_{j\in x} r_j + \sum_{\tau=1}^{t-1} c_i(x, s_{-i}^{\tau})]$$

Step 1: Magic Follow the perturbed leader has regret at most $\max_x \sum_{j\in x} r_j$

$$\text{select } argmin_x [\textstyle\sum_{j\in x} r_j + \sum_{\tau=1}^{t} c_i(x, s_{-i}^{\tau})]$$

Proof: as before

$$\sum_{\tau=1}^{t} c_i(s^\tau) - r_{s_i^1} \le \sum_{\tau=1}^{t} c_i\left(\boldsymbol{s_i^t}, s_{-i}^{\tau}\right) - r_{s_i^t} \le \min_x \sum_{\tau=1}^{t} c_i(x, s_{-i}^{\tau}) - r_x$$

IH

$$\sum_{\tau=1}^{t} c_i(s^\tau) - r_{s_i^1} = \sum_{\tau=1}^{t-1} c_i(s^\tau) - r_{s_i^1} + c_i(s^t) \le \sum_{\tau=1}^{t-1} c_i\left(s_i^t, s_{-i}^{\tau}\right) - r_{s_i^t} + c_i(s^t)$$

QED

Real follow the perturbed leader

Let r_j random: number of coins till you get H, if probability of H is ϵ

So $E(r_x) = \frac{|x|}{\epsilon}$ Also, for n strategies $E(\max_x \sum_{j\in x} r_j) = O(\frac{n}{\epsilon})$

Step 2: if $\max u_i(s) \leq 1$, then in any one step, the probability that magic perturbed follow the leader makes a different choice than real $\leq \epsilon$

Alternate way to flip the coins.

Start with r_x=1 all x

While more than one x possible

Take largest x, and flip a coin for a j in x.
If all coins already H: x eliminated

When one x left: flip coins for x till H

If ≠H, then adding $u_i(x, s_{-i}^t)$ or not makes no difference, prob=$1-\epsilon$

Follow perturbed leader: small regret

Assuming we always follow magic version: regret at most $\max_x r_x$

- Expected value $E(\max_x r_x) = \frac{n}{\epsilon}$
- expected total utility loss when not following the magic leader is at most an ϵ fraction
- Total regret at most

$$\sum_\tau^t u_i(s^t) \leq (1-\epsilon) \max_x \sum_\tau^t c_i(x, s_i^t) - \frac{n}{\epsilon}$$

Theorem: Select $\epsilon = \sqrt{\frac{n}{T}}$ then resulting regret at most $O(\sqrt{Tn})$

Valuations beyond unit demand

- Unit demand $v_i(S) = \max_{j \in S} v_{ij}$
- Additive $v_i(S) = \sum_{j \in S} v_{ij}$

XOS = mix of the two $v_i(S) = \max_k \sum_{j \in S} v_{ij}^k$

Fact: unit demand is XOS: $v_{ij}^k = v_{ij}$ if $k = j$, and 0 otherwise

Submodular: $A \subset B$ we have $v_i(A + j) - v_i(A) \geq v_i(B + j) - b(B)$

Lemma: Submodular is XOS: for any order π we have v_{ij} = marginal value of j in this order

Plans for next two lectures: things that learning can do beyond getting to CCE

So far we had: learning outcome is as good as Price of Anarchy proven via smoothness arguments (and almost all PoA proofs are smoothness arguments)

Things we hope learning can do:

- Adjust to changing environments (churn)
- Do better than the worst case Nash (or better than any Nash?)