More Games and the Quality of Outcomes

Éva Tardos
Cornell University

Last two talks:

- Bounds on Quality of Nash (Price of Anarchy or Stability) via potential function or smoothness.
- Learning outcomes and smoothness bounds.
- Bounds via smoothness are tight even for pure Nash
- "Natural learning" (via Hedge algorithm) results almost always in pure Nash.

Today: Quality of Nash in other context: AdWords

Games considered are often congestion games

- Routing:
- routers choose path for packets though the Internet ✓
- Bandwidth Sharing:
- routers share limited bandwidth between processes ¼ ✓
- Facility Location:
- Decide where to host certain Web applications ✓
- Load Balancing
- Balancing load on servers (e.g. Web servers) ✓
- Network Design:
- Independent service providers building the Internet ½ ✓

Sponsored Search

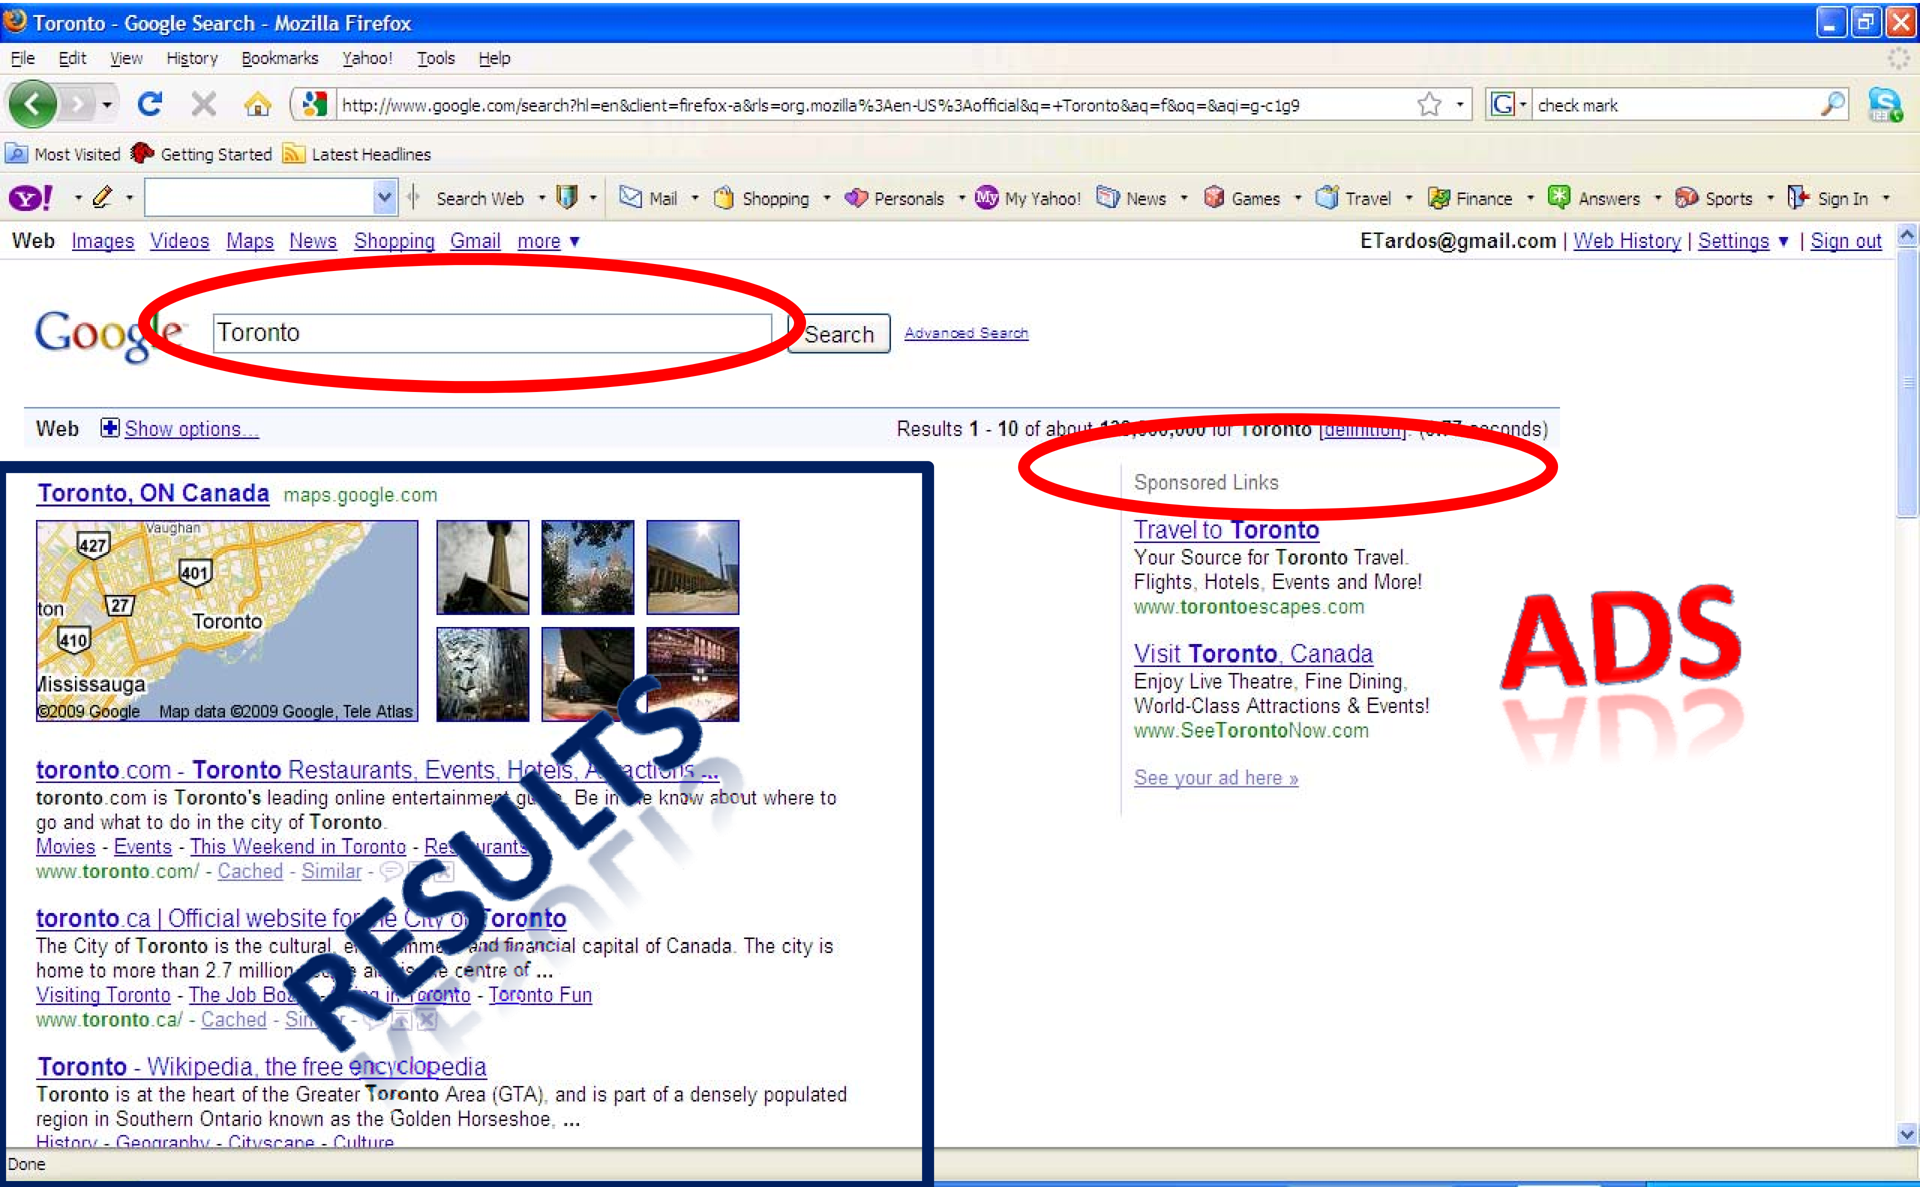

Model of Sponsored Search

Focus on ads on the side

Ordered slots, higher is better

Advertisers:

Hilton, RailEurope, CentralBudapestHotels, DestinationBudapest, RacationRentals.com, Travelzoo.com, TravelYahhoo.com, BudgetPlace.com

Model of Sponsored Search

Questions:

- which advertiser to assign to each slots, and
- how much to charge each

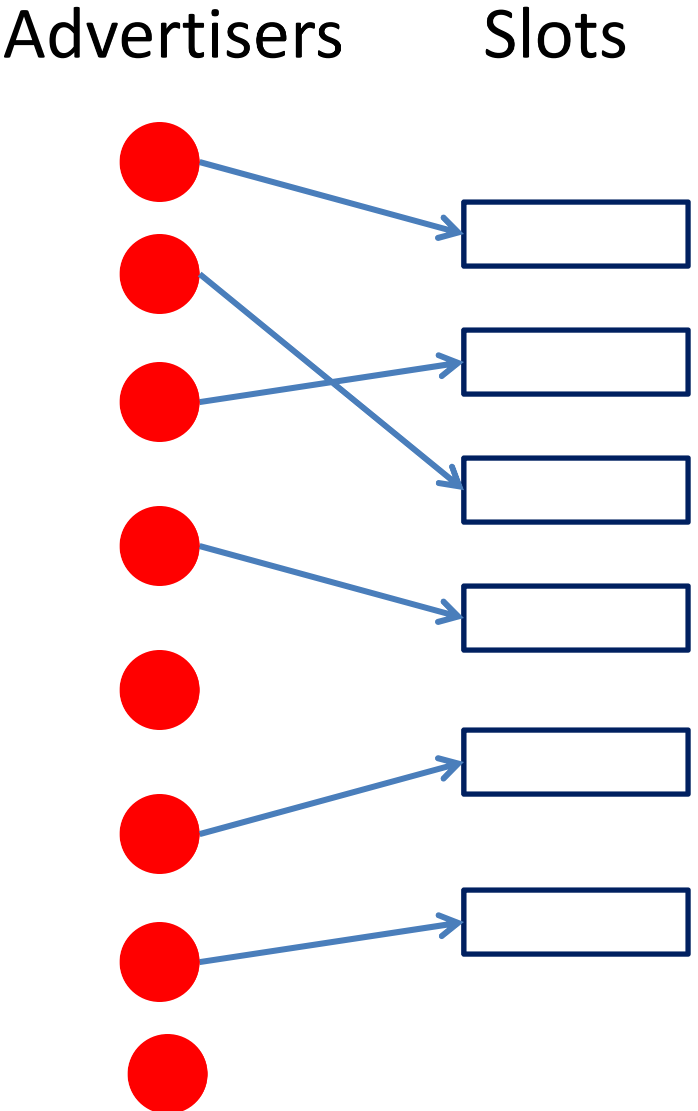

Value of Assignment

Value of a click for adv. **i** is $\mathbf{v_i}$

At per-click price of **p**, value to advertiser is

$(v_i - p) \cdot$ click-rate

Value of search engine

$p \cdot$ click-rate

Total of

$v_i \cdot$ click-rate

But what is the click rate?

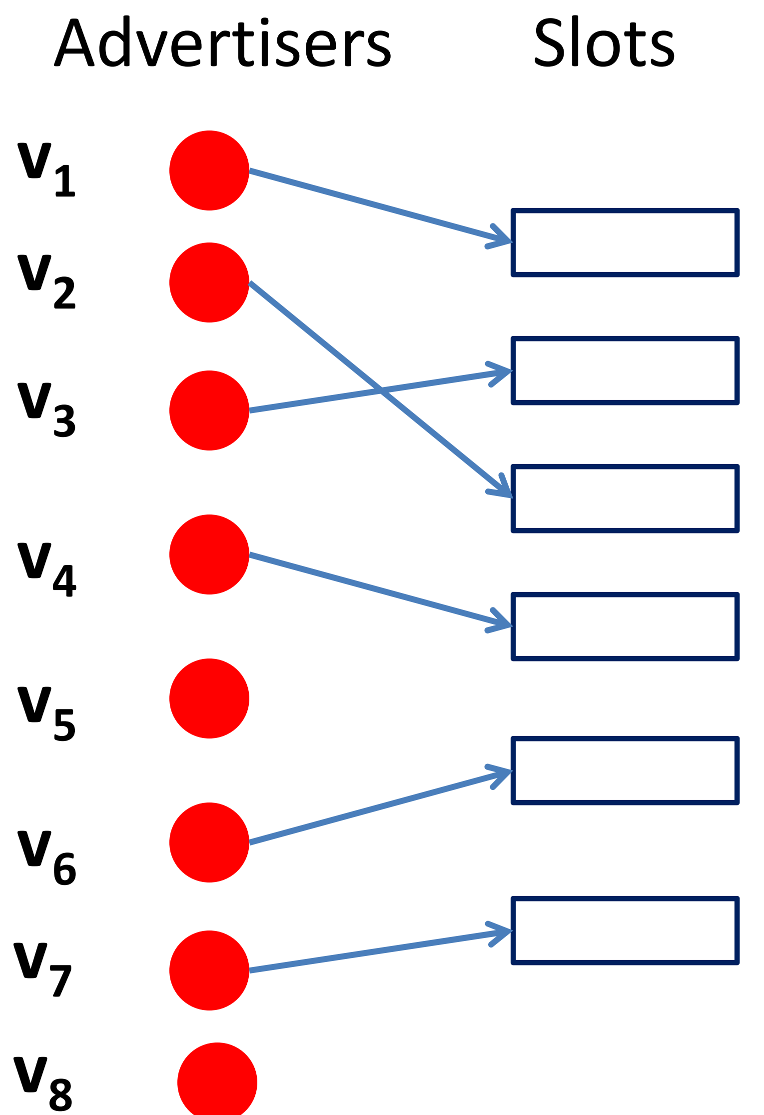

Click rate models

Simplest model:

click rate in slot ℓ is α_ℓ

Assume $\alpha_1 > \alpha_2 > \alpha_3 > \alpha_4 > \ldots$

Value of advertiser i in slot ℓ

$v_i \cdot \alpha_\ell$

Better model: click rate

also depends on advertiser i

$\alpha_\ell \cdot \gamma_i$

Advertiser i value of slot j

$v_i \cdot \gamma_i \cdot \alpha_\ell$

effective value

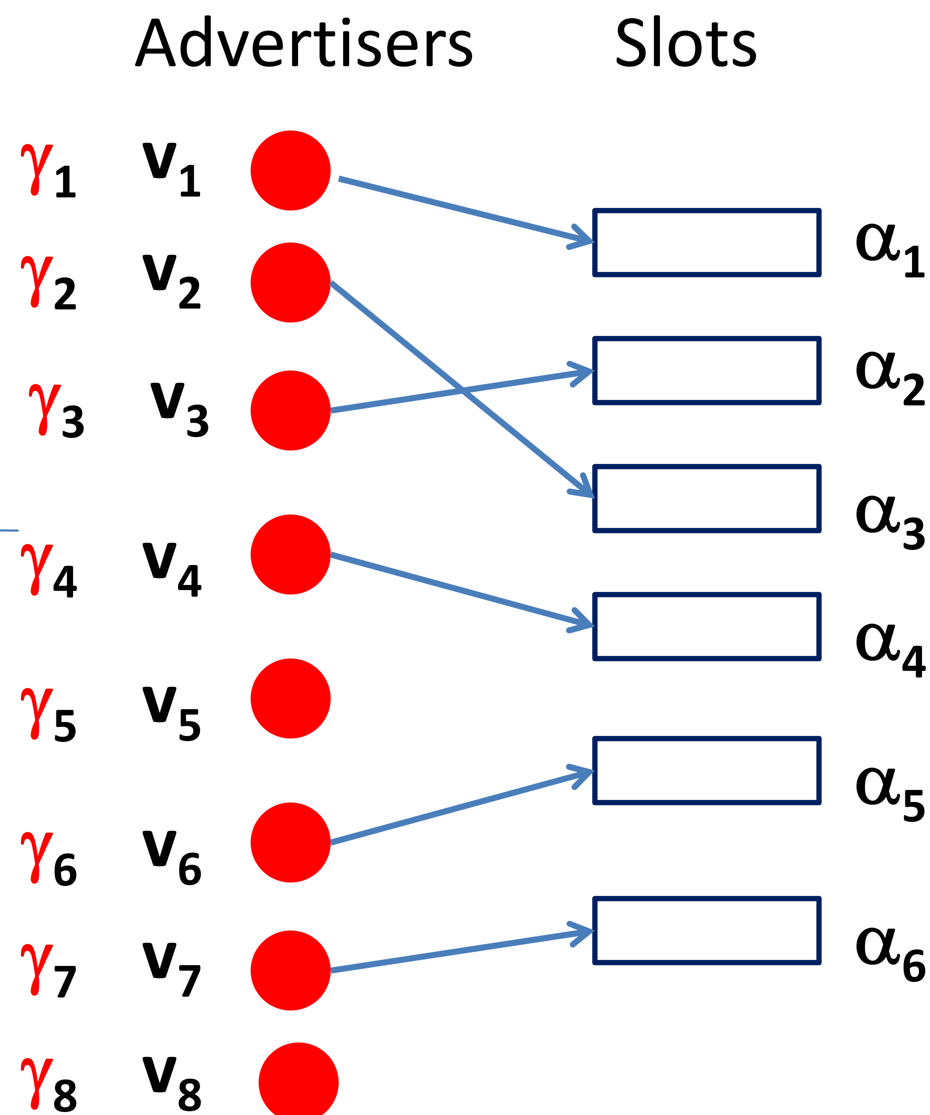

Social welfare

Talk: simplest model:

$\alpha_1 > \alpha_2 > \alpha_3 > \alpha_4 > \ldots$

Value of advertiser i in slot j

$\mathbf{v}_i \cdot \alpha_\ell$

Social welfare = total value

$\Sigma_{i \to \ell} \mathbf{v}_i \ \alpha_\ell$

Maximize social welfare:

Sort $\mathbf{v}_1 > \mathbf{v}_2 > \mathbf{v}_3 > \mathbf{v}_4 > \ldots$

Welfare $\Sigma_i \mathbf{v}_i \ \alpha_i$

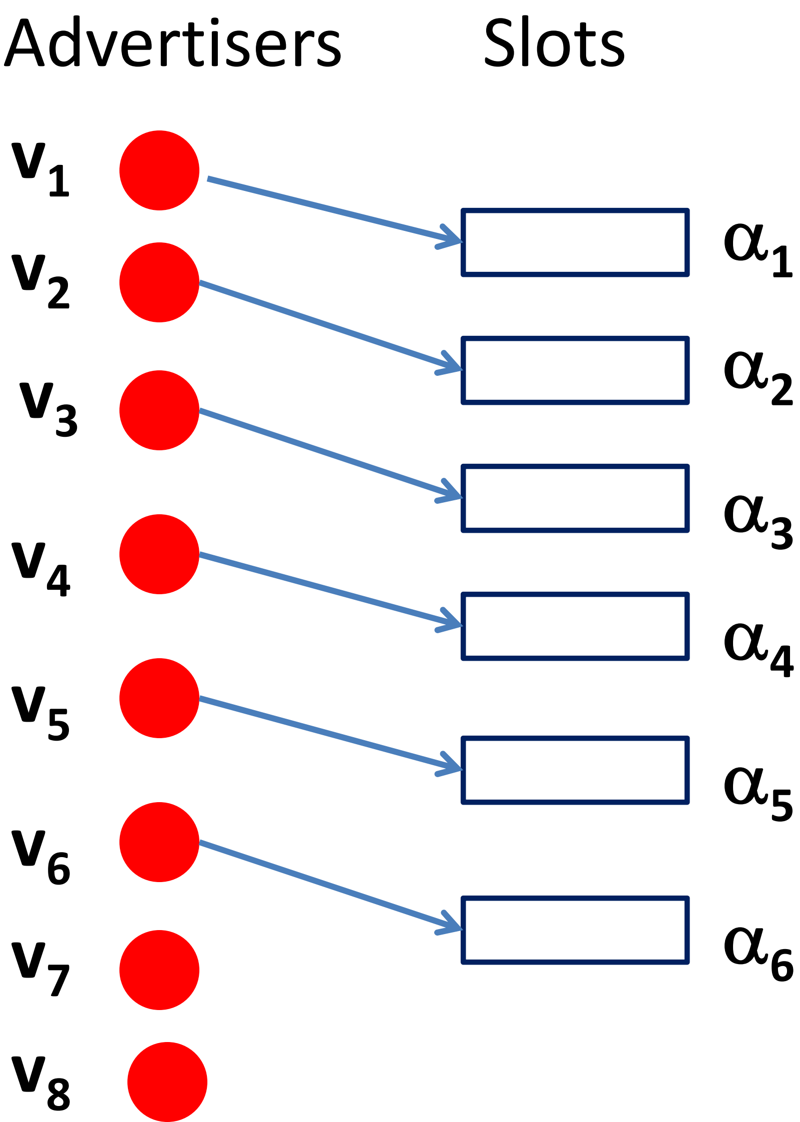

Can we elicit values?

Trouble: Value $\mathbf{v_i}$ is private, only known to user i.

Single slot: Vickrey auction (second price)

- Winner highest value $\text{argmax}_i\, \mathbf{v_i} \rightarrow$ say 1
- price $\max_{j \neq 1} \mathbf{v_j}$

Result: truthful:

Players will bid their true values $\mathbf{b_i} = \mathbf{v_i}$

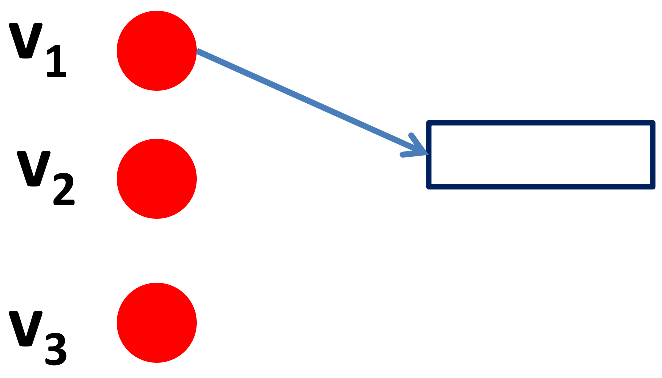

Generalized Second Prize (GSP)

Sort by bid value

$b_1 > b_2 > b_3 > b_4 > \ldots$

Assign bidder i to slot i

Charge next price $p = b_{i+1}$

Value to bidder i

$(v_i - b_{i+1}) \cdot \alpha_i$

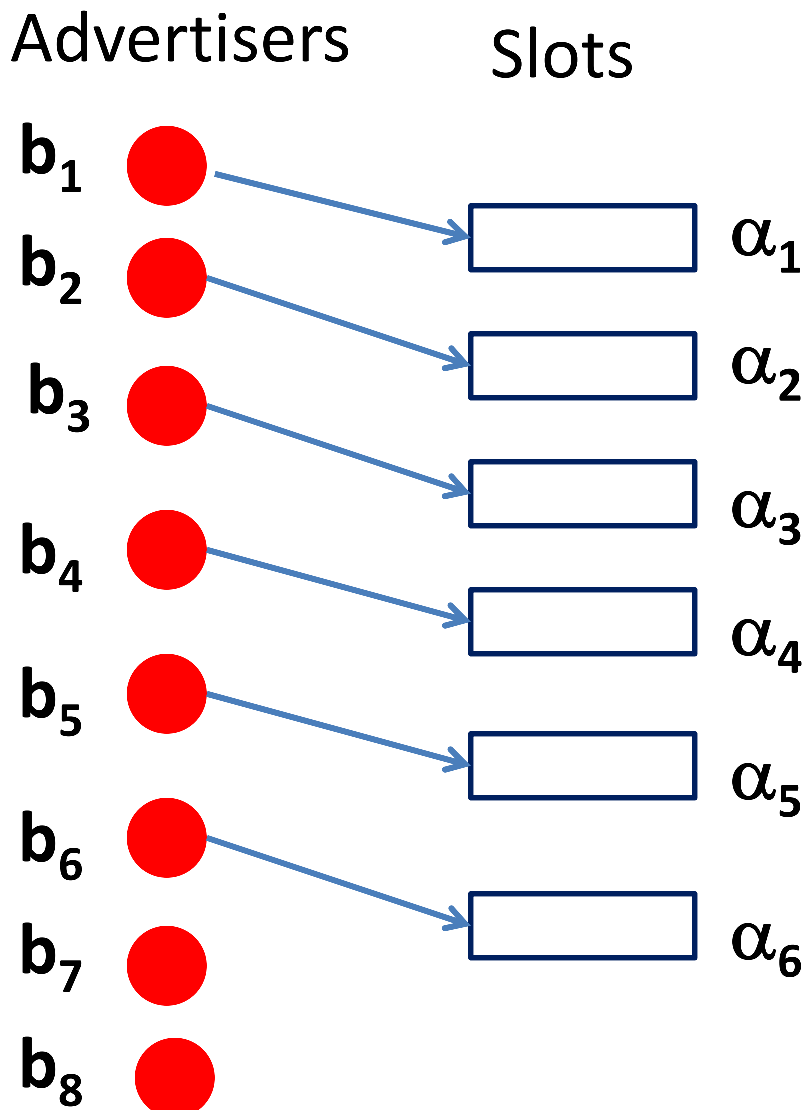

Is GSP truthful?

Is bidding $b_i = v_i$ Nash equilibrium for the bidders?

Example:

Bidder 1's value if telling the truth

$(9-8) \cdot 1 = 1$

If bidding $b_1 = 7$

$(9-1) \cdot ½ = 4$

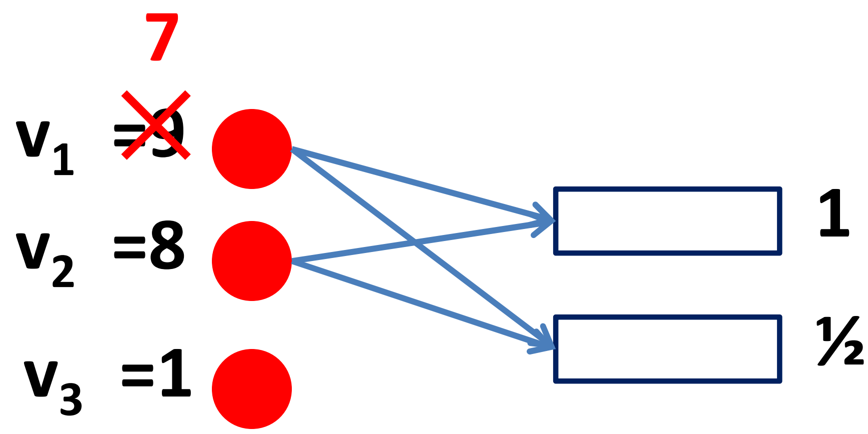

Sort by bid value

$b_1 > b_2 > b_3 > b_4 > \ldots$

Charge next price $p = b_{i+1}$

Value to bidder i

$(v_i - b_{i+1}) \cdot \alpha_i$

How good/bad are Nash equilibria?

Theorem: [Edelman & Ostrovsky & Schwarz, Varian]

There exists a Nash equilibrium set of bids that result in socially optimal assignment

bids are sorted the same order as values

Note: There is also Vickrey-Clark-Groves (VCG) mechanisms, that is always truthful. This is not VCG

∃ Optimal Nash

Maximum social welfare:

Sort $v_1 > v_2 > v_3 > v_4 > \ldots$

Welfare $\Sigma_i v_i \alpha_i$

Corresponding Nash?

user i to slot ℓ value: $v_i \alpha_\ell$

LP dual: ∃ prices p_ℓ such that

user i prefers slot i

$$i = \arg\max_\ell v_i \alpha_\ell - p_\ell$$

Nash bids $b_i = p_{i-1} / \alpha_i$

Why? Will pay exactly p_i

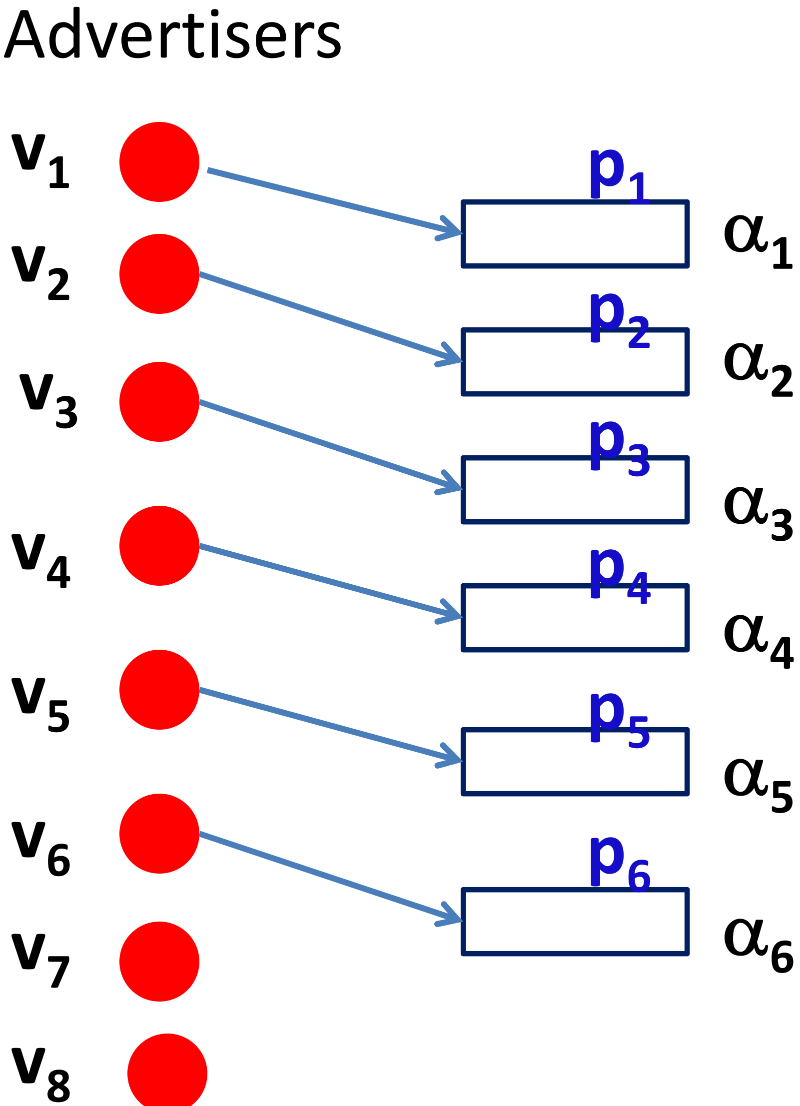

How bad can Nash equilibria be?

But also: There exists Nash equilibrium of very bad social value

$\mathbf{b_1 = 0}$ $\mathbf{v_1 = 1}$

is Nash: Why? $\mathbf{b_2 = 1}$ $\mathbf{v_2 = 0}$

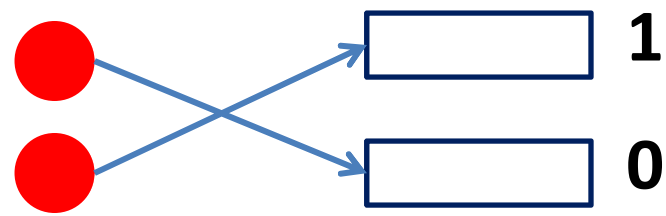

Player 1 current value: $(1\text{-}0) \cdot 0 = \mathbf{0}$

alternate option: bid $b>1$ and pay $p=1$

for a value of $(1\text{-}p) \cdot 1 = \mathbf{0}$

Social welfare **0**

Optimum social welfare 1

Note: this is Vickrey auction…

Is bad Nash realistic?

$b_1 = 0$ $v_1 = 1$

$b_2 = 1$ $v_2 = 0$

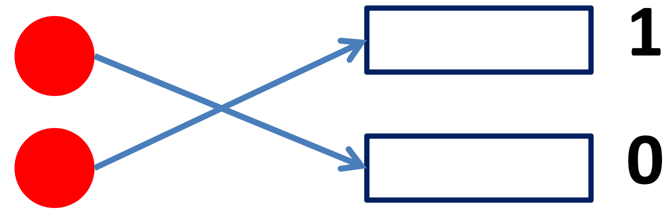

Notice: **b_2=1 extremely dangerous… $b_2 >> v_2$**

- if a new bidder shows up between $b_i > b > v_i$ can be forced to pay above value
- bidding … $b_2 = v_2$ is just as good and not dangerous

Assumption: Conservative bidders

$b_i \leq v_i$ for all bidders i

How good are Nash equilibria with conservative bidders?

Theorem: [Renato Paes Leme & Tardos'09]

All Nash equilibria with conservative bids has social welfare at least 1.618 fraction of the maximum possible

Vickrey: Nash equilibria with conservative bids has optimal social welfare

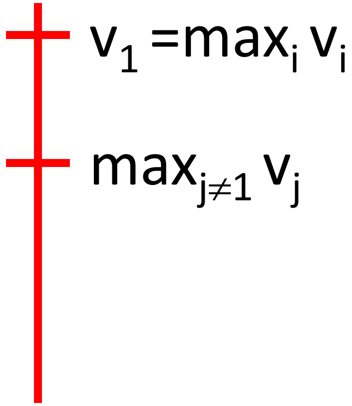

Proof: induction on number of slots

If top value assigned to top slot: use induction

Else let β and λ be the ratio in values as shown

Lemma $1/\beta + 1/\lambda \geq 1$

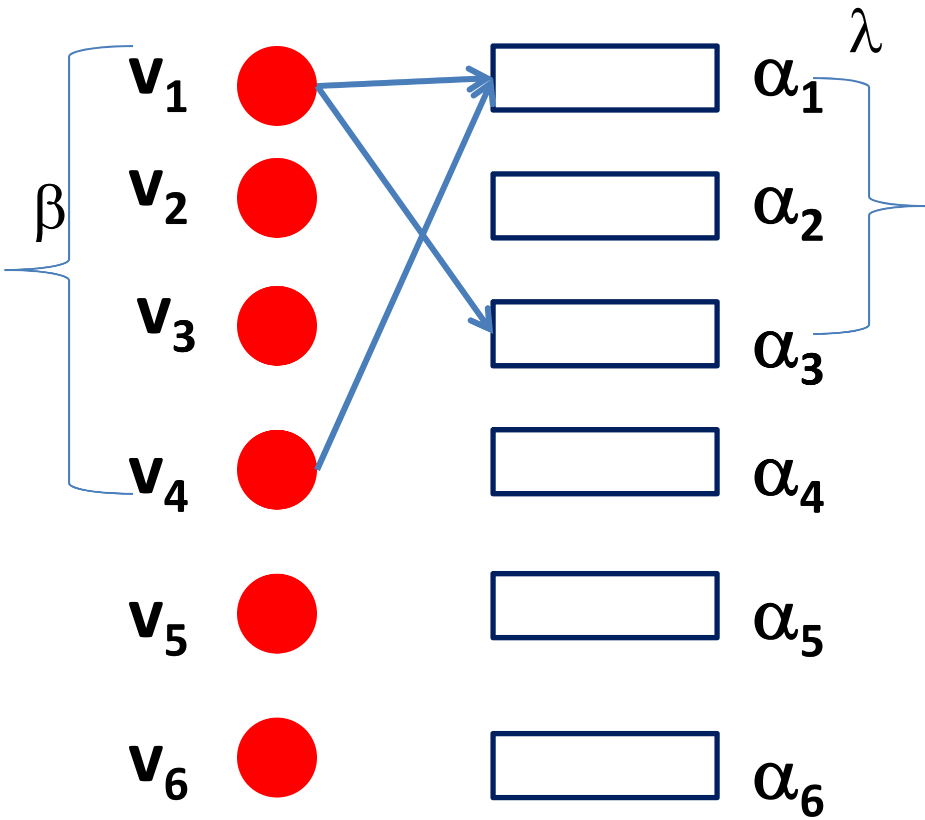

Worst case for two slots

Else let β and λ be the ratio in values as shown

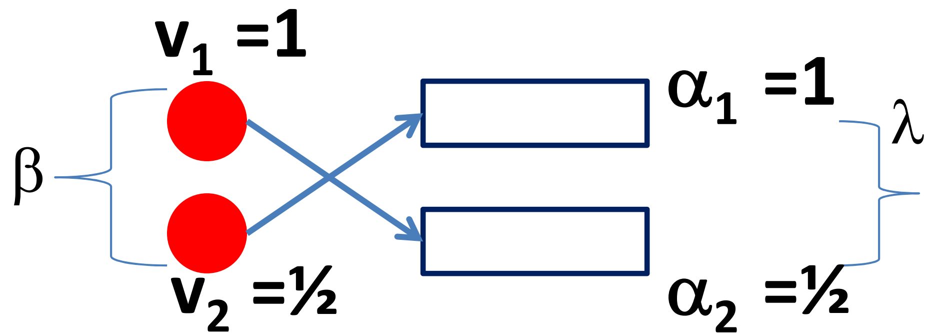

Lemma $1/\beta+1/\lambda \geq 1$

Worst case for two slots

Optimal value: 1+¼

Conservative Nash solution bids b_1 =0 and b_2=½

total value: 1

Value for bidders: $(v_1 - 0)\alpha_2$ =½ and: $(v_2-0)\ \alpha_1$=½

Alternative for bidder 1: b_1 >½ for value $(v_1 - ½)\alpha_1$ =½

Optimum $1+ 1/\beta\lambda$ Nash value $1/\beta+1/\lambda$

Proof: simple bound of 2

induction on number of slots

If top value assigned to top slot: use induction

Else let β and λ be the ratio in values as shown

Lemma $1/\beta + 1/\lambda \geq 1$

hence either β or $\lambda \leq 2$

We get total value

at least ½ $\mathbf{v_1} \cdot \alpha_1$ (as $\lambda \leq 2$)

+ ½ **rest** by induction

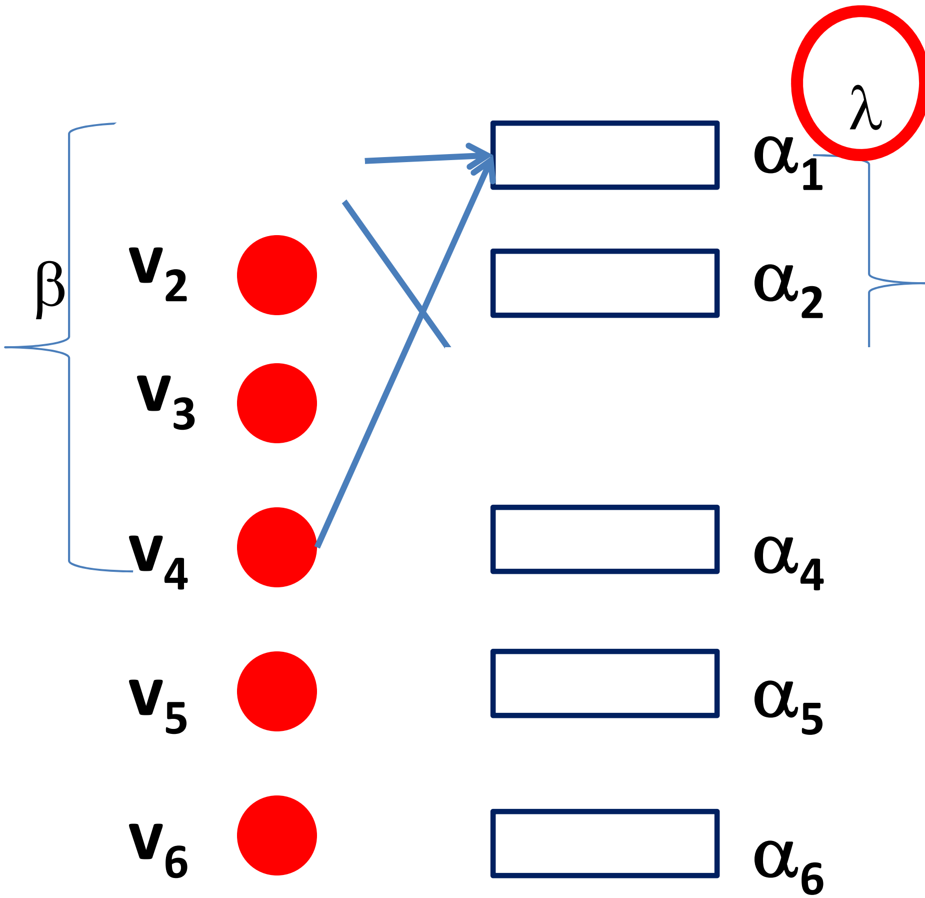

Proof: Improved bound: $\gamma=(1+\sqrt{5})/2\approx 1.618$

If top value $\rightarrow$ to top slot: use induction

Else β and λ: gap as shown

Lemma $1/\beta+1/\lambda \geq 1$

We get total value at least

$1/\lambda\, \mathbf{v_1} \cdot \alpha_1 + \gamma$ **rest**

$\geq\; 1/\lambda\, \mathbf{v_1} \cdot \alpha_1 + \gamma[(\alpha_1 - \alpha_\ell)\, \mathbf{v}_\ell + \Sigma_{j\geq 2}\, \mathbf{v_j}\; \alpha_j\,]$

$\geq\; 1/\lambda\, v_1 \cdot \alpha_1 + \gamma(1/\beta)\, (\alpha_1 - \alpha_\ell)\, v_1 + \Sigma_{j\geq 2} v_j\, \alpha_j$

$\geq\; 1/\lambda\, v_1 \cdot \alpha_1 + \gamma(1/\beta)(1-1/\lambda)\alpha_1 v_1 + \Sigma_{j\geq 2} v_j \alpha_j$

$\geq\; (1/\lambda + \gamma(1-1/\lambda)^2\,)\alpha_1\, v_1 + \gamma\, \Sigma_{j\geq 2}\, v_j\, \alpha_j$

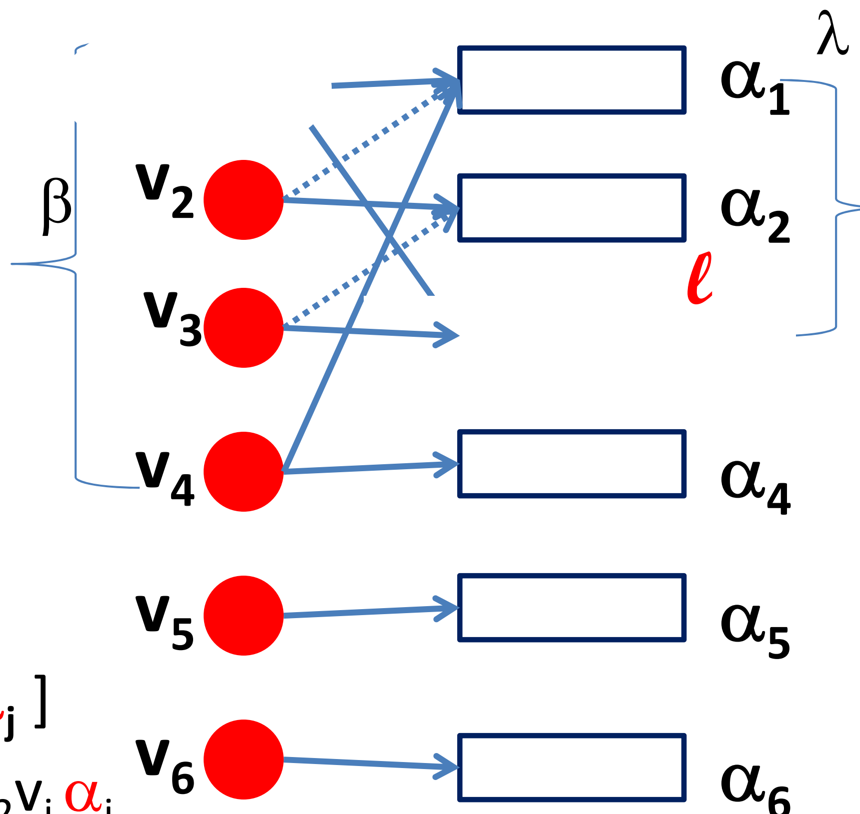

Main Lemma

Lemma $1/\beta + 1/\lambda \geq 1$

Proof: bidder 1 is happy

Current value

$$(\mathbf{v_1} - \mathbf{b}_{\ell+1}) \cdot \alpha_\ell \leq \mathbf{v_1} \cdot \alpha_\ell$$

Alternate option: $> \mathbf{b_i}$

$$(\mathbf{v_1} - \mathbf{b_i}) \cdot \alpha_1 \geq (\mathbf{v_1} - \mathbf{v_i}) \cdot \alpha_1$$

conservative

Nash $\Rightarrow \mathbf{v_1} \cdot \alpha_\ell \geq (\mathbf{v_1} - \mathbf{v_i}) \cdot \alpha_1$

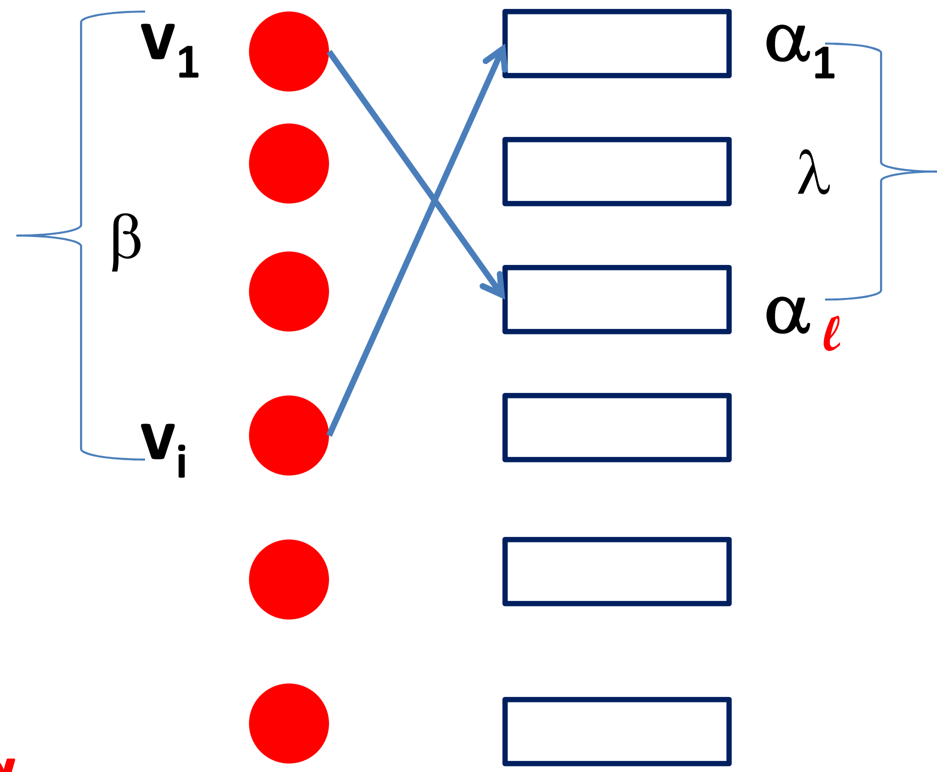

Conclusion

On line advertisement: fun assignment problem

Mechanism used GSP (generalized second price). Simple, but not always good.

- There exists Nash with optimal value
- Resulting social value not too bad (assuming conservative bidders)

Questions: which Nash is realistic?

Learning in context with many identical bid options

Other models of click-rate? (externalities?)

Summary

Network Games and Quality of Nash

- Examples routing and cost-sharing
- Outcome : Nash & Price of Anarchy /Stability
- Smoothness: a common and powerful proof technique

Learning in Network Games

- No-regret learning
- (coarse) correlated equilibrium & Total Price of Anarchy
- Smoothness bounds also valid
- Natural learning leads to good outcome

Quality in other games: Ad-Auctions

Many natural questions:

- Other classes of games and other learning methods?