Game Theory: introduction and applications to computer networks

Introduction

Giovanni Neglia
INRIA – EPI Maestro
3 February 2014

Part of the slides are based on a previous course with D. Figueiredo (UFRJ) and H. Zhang (Suffolk University)

Always an equilibrium with small Loss of Efficiency?

- Consider only affine cost functions,
  i.e. $c_\alpha(x) = a_\alpha + b_\alpha x$
- We will use the potential to derive a bound on the social cost of a NE
  - $P(f) <= C_S(f) <= 2\ P(f)$

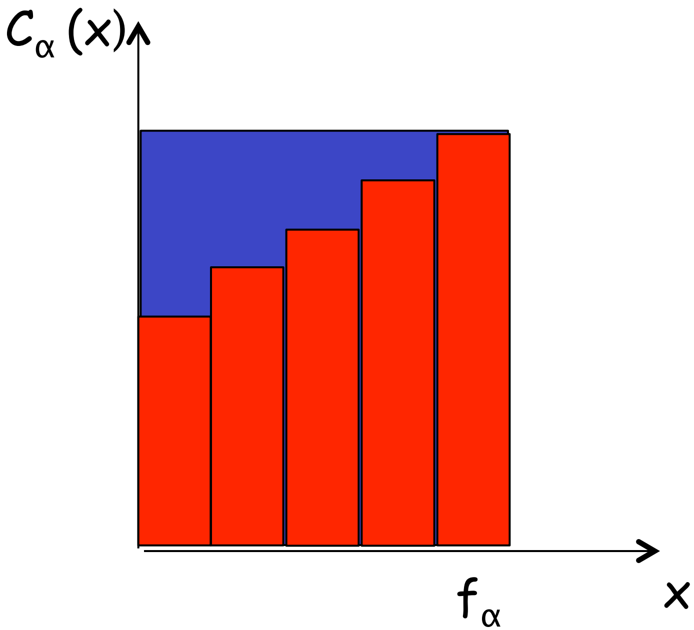

$P_\alpha(f_\alpha)$

$C_\alpha(f_\alpha)$

Always an equilibrium with small Loss of Efficiency?

- Consider only affine cost functions

  i.e. $c_\alpha(x)=a_\alpha+b_\alpha x$
- We will use the potential to derive a bound on the social cost of a NE
  - $P(f) <= C_S(f) <= 2\ P(f)$
- $P(f) = \Sigma_{\alpha\ \varepsilon\ E} P_\alpha = \Sigma_{\alpha\ \varepsilon\ E} \Sigma_{t=1,\ldots f\alpha}\ c_\alpha(t) <=$

  $<= \Sigma_{\alpha\ \varepsilon\ E} \Sigma_{t=1,\ldots f\alpha}\ c_\alpha(f_\alpha) = \Sigma_{\alpha\ \varepsilon\ E} f_\alpha c_\alpha(f_\alpha) = C_S(f)$
- $P(f) = \Sigma_{\alpha\ \varepsilon\ E} P_\alpha = \Sigma_{\alpha\ \varepsilon\ E} \Sigma_{t=1,\ldots f\alpha}(a_\alpha+b_\alpha t)=$

  $=\Sigma_{\alpha\ \varepsilon\ E} f_\alpha a_\alpha + b_\alpha f_\alpha (f_\alpha+1)/2 >= \Sigma_{\alpha\ \varepsilon\ E} f_\alpha(a_\alpha+b_\alpha f_\alpha)/2$

  $=C_S(f)/2$

Always an equilibrium with small Loss of Efficiency?

- Consider only affine cost functions
  i.e. $c_\alpha(x)=a_\alpha+b_\alpha x$
- $P(f) <= C_S(f) <= 2\ P(f)$
- Let's imagine to start from routing f_{Opt} with the optimal social cost $C_S(f_{Opt})$,
- Applying the BR dynamics we arrive to a NE with routing f_{NE} and social cost $C_S(f_{NE})$
- $C_S(f_{NE}) <= 2\ P(f_{NE}) <= 2\ P(f_{Opt}) <= 2\ C_S(f_{Opt})$
- The LoE of this equilibrium is at most 2

Same technique, different result

- Consider a network with a routing at the equilibrium
- Add some links
- Let the system converge to a new equilibrium
- The social cost of the new equilibrium can be at most 4/3 of the previous equilibrium social cost (as in the Braess Paradox)

Loss of Efficiency, Price of Anarchy, Price of Stability

- Loss of Efficiency (LoE)
  - given a NE with social cost $C_S(f_{NE})$
  - LoE = $C_S(f_{NE}) / C_S(f_{Opt})$
- Price of Anarchy (PoA) [Koutsoupias99]
  - Different settings G (a family of graph, of cost functions,...)
  - X_g =set of NEs for the setting g in G
  - PoA = $\sup_{g \in G} \sup_{NE \in X_g} \{C_S(f_{NE}) / C_S(f_{Opt})\}$ => "worst" loss of efficiency in G
- Price of Stability (PoS) [Anshelevish04]
  - PoS = $\sup_{g \in G} \inf_{NE \in X_g} \{C_S(f_{NE}) / C_S(f_{Opt})\}$ => guaranteed loss of efficiency in G

Stronger results for affine cost functions

- We have proven that for unit-traffic routing games the PoS is at most 2
- For unit-traffic routing games and single-source pairs the PoS is 4/3
- For non-atomic routing games the PoA is 4/3
  - non-atomic = infinite players each with infinitesimal traffic
- For other cost functions they can be much larger (even unbounded)

Performance Evaluation

Sponsored Search Markets

Giovanni Neglia

INRIA – EPI Maestro

4 February 2013

Google

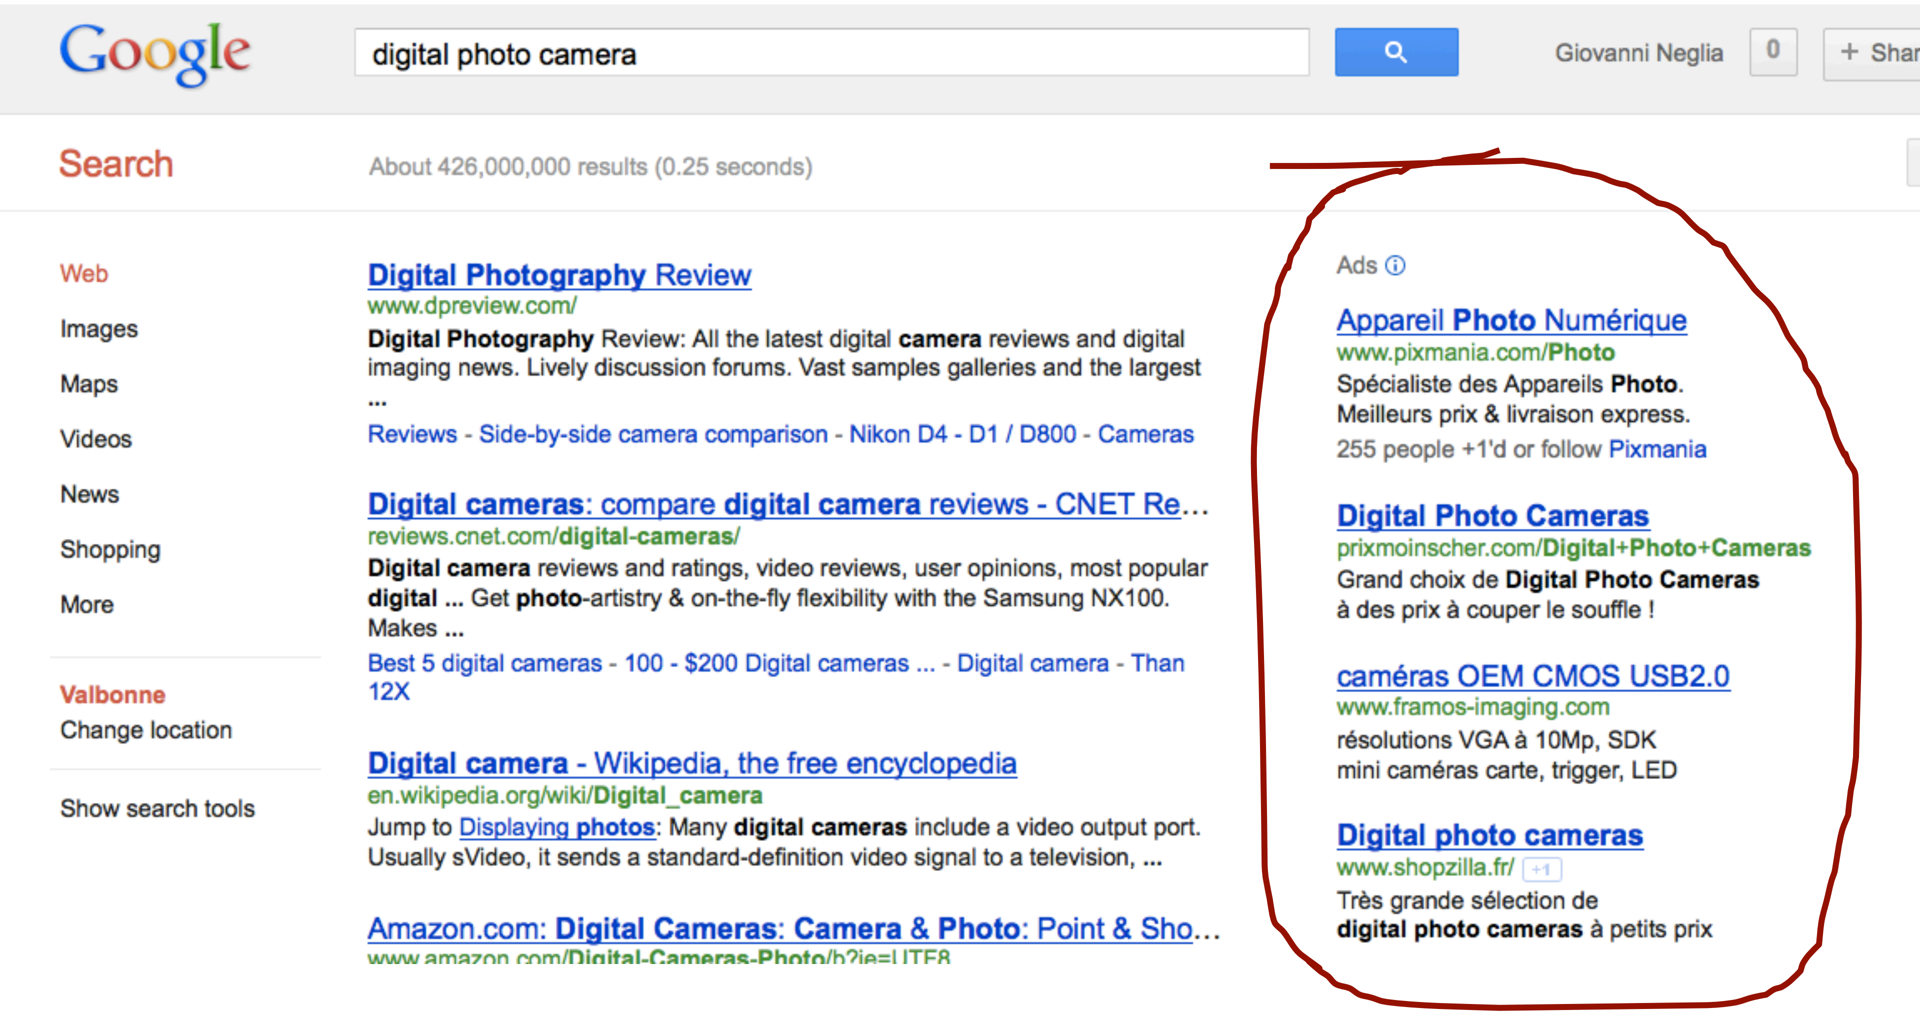
Google
digital photo camera
Giovanni Neglia
0
+ Shar
Search
About 426,000,000 results (0.25 seconds)
Web
Images
Maps
Videos
News
Shopping
More
Valbonne
Change location
Show search tools
Digital Photography Review
www.dpreview.com/
Digital Photography Review: All the latest digital camera reviews and digital imaging news. Lively discussion forums. Vast samples galleries and the largest ...
Reviews - Side-by-side camera comparison - Nikon D4 - D1 / D800 - Cameras
Digital cameras: compare digital camera reviews - CNET Re...
reviews.cnet.com/digital-cameras/
Digital camera reviews and ratings, video reviews, user opinions, most popular digital ... Get photo-artistry & on-the-fly flexibility with the Samsung NX100. Makes ...
Best 5 digital cameras - 100 - $200 Digital cameras ... - Digital camera - Than 12X
Digital camera - Wikipedia, the free encyclopedia
en.wikipedia.org/wiki/Digital_camera
Jump to Displaying photos: Many digital cameras include a video output port. Usually sVideo, it sends a standard-definition video signal to a television, ...
Amazon.com: Digital Cameras: Camera & Photo: Point & Sho...
Ads
Appareil Photo Numérique
www.pixmania.com/Photo
Spécialiste des Appareils Photo.
Meilleurs prix & livraison express.
255 people +1'd or follow Pixmania
Digital Photo Cameras
prixmoinscher.com/Digital+Photo+Cameras
Grand choix de Digital Photo Cameras
à des prix à couper le souffle !
caméras OEM CMOS USB2.0
www.framos-imaging.com
résolutions VGA à 10Mp, SDK
mini caméras carte, trigger, LED
Digital photo cameras
www.shopzilla.fr/
Très grande sélection de
digital photo cameras à petits prix

How it works

- Companies bid for keywords
- On the basis of the bids Google puts their link on a given position (first ads get more clicks)
- Companies are charged a given cost for each click (the cost depends on all the bids)

Some numbers

- ≈ 95% of Google revenues (46 billions$) from ads
  - investor.google.com/financial/tables.html
  - 87% of Google-Motorola revenues (50 billions$)
- Costs
  - "calligraphy pens" $1.70
  - "Loan consolidation" $50
  - "mesothelioma" $50 per click
- Click fraud problem

Outline

- Preliminaries
  - Auctions
  - Matching markets
- Possible approaches to ads pricing
- Google mechanism

- References
  - Easley, Kleinberg, "Networks, Crowds and Markets", ch.9,10,15

Types of auctions

- 1st price & descending bids
- 2nd price & ascending bids

Game Theoretic Model

- N players (the bidders)
- Strategies/actions: b_i is player i's bid
- For player i the good has value v_i
- p_i is player i's payment if he gets the good
- Utility:
  - v_i-p_i if player i gets the good
  - 0 otherwise
- Assumption here: values v_i are *independent* and *private*
  - i.e. very particular goods for which there is not a reference price

Game Theoretic Model

- N players (the bidders)
- Strategies: b_i is player i's bid
- Utility:
  - v_i-b_i if player i gets the good
  - 0 otherwise
- Difficulties:
  - Utilities of other players are unknown!
  - Better to model the strategy space as continuous
  - Most of the approaches we studied do not work!

2nd price auction

- Player with the highest bid gets the good and pays a price equal to the 2nd highest bid
- There is a dominant strategies
  - I.e. a strategy that is more convenient independently from what the other players do
  - **Be truthful**, i.e. bid how much you evaluate the good ($b_i=v_i$)
  - Social optimality: the bidder who value the good the most gets it!

$b_i=v_i$ is the highest bid

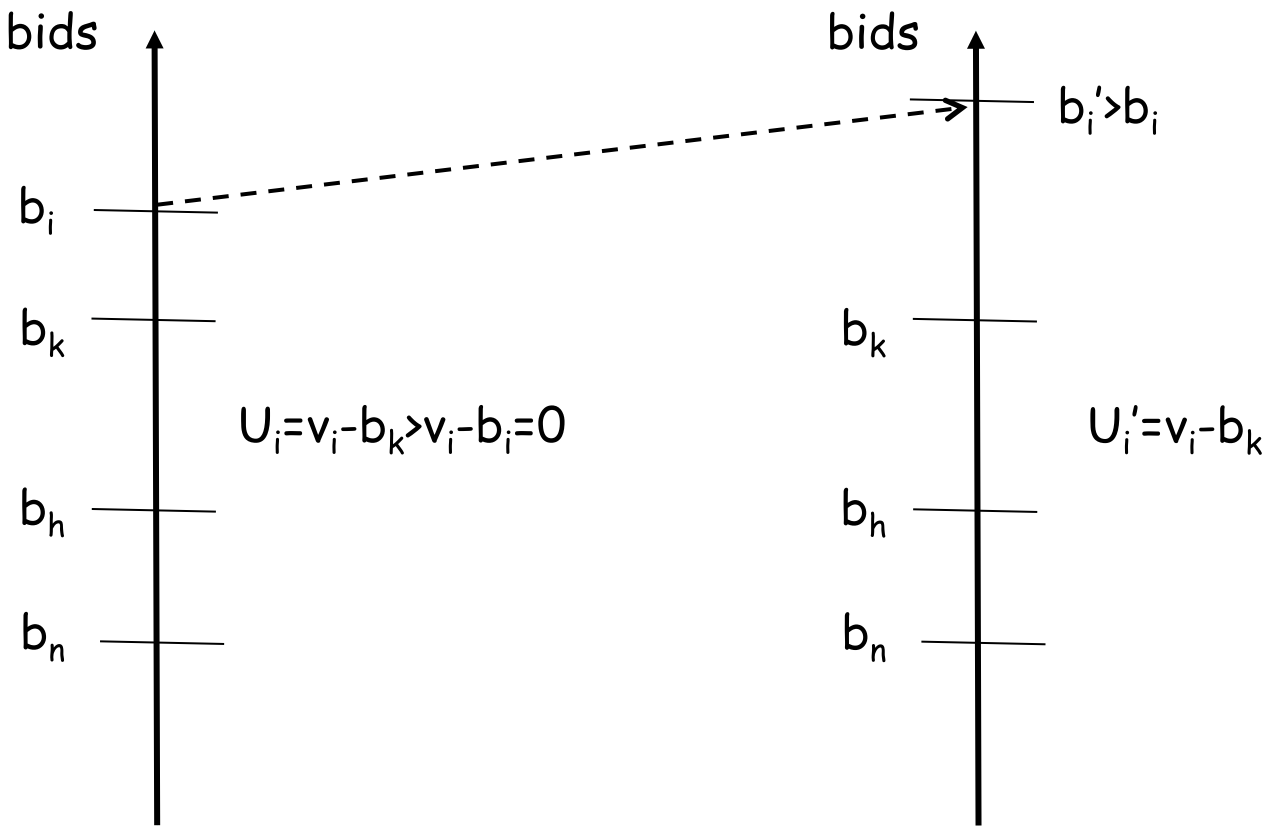

Bidding more than v_i is not convenient

$b_i=v_i$ is the highest bid

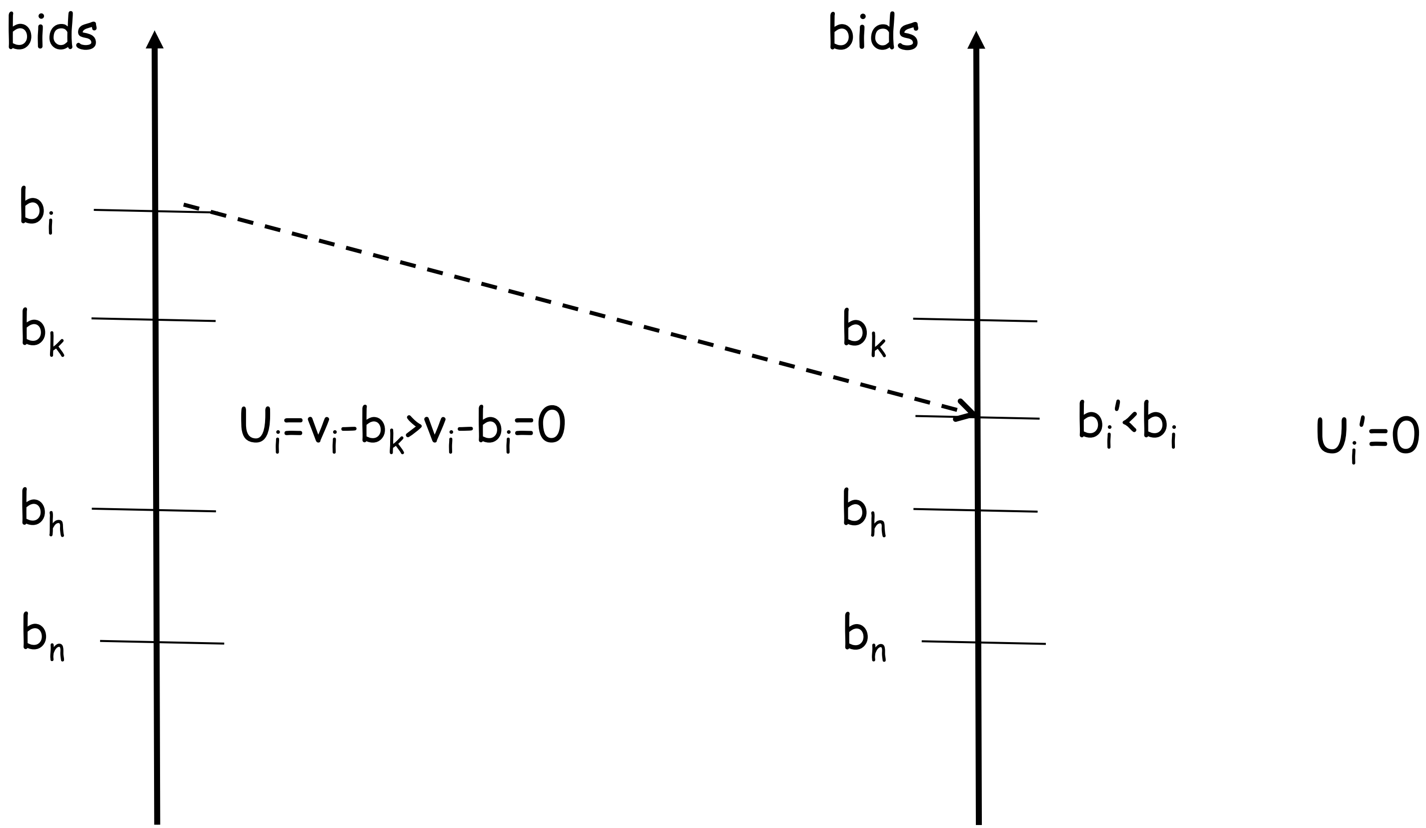

Bidding less than v_i is not convenient (may be unconvenient)

$b_i=v_i$ is not the highest bid

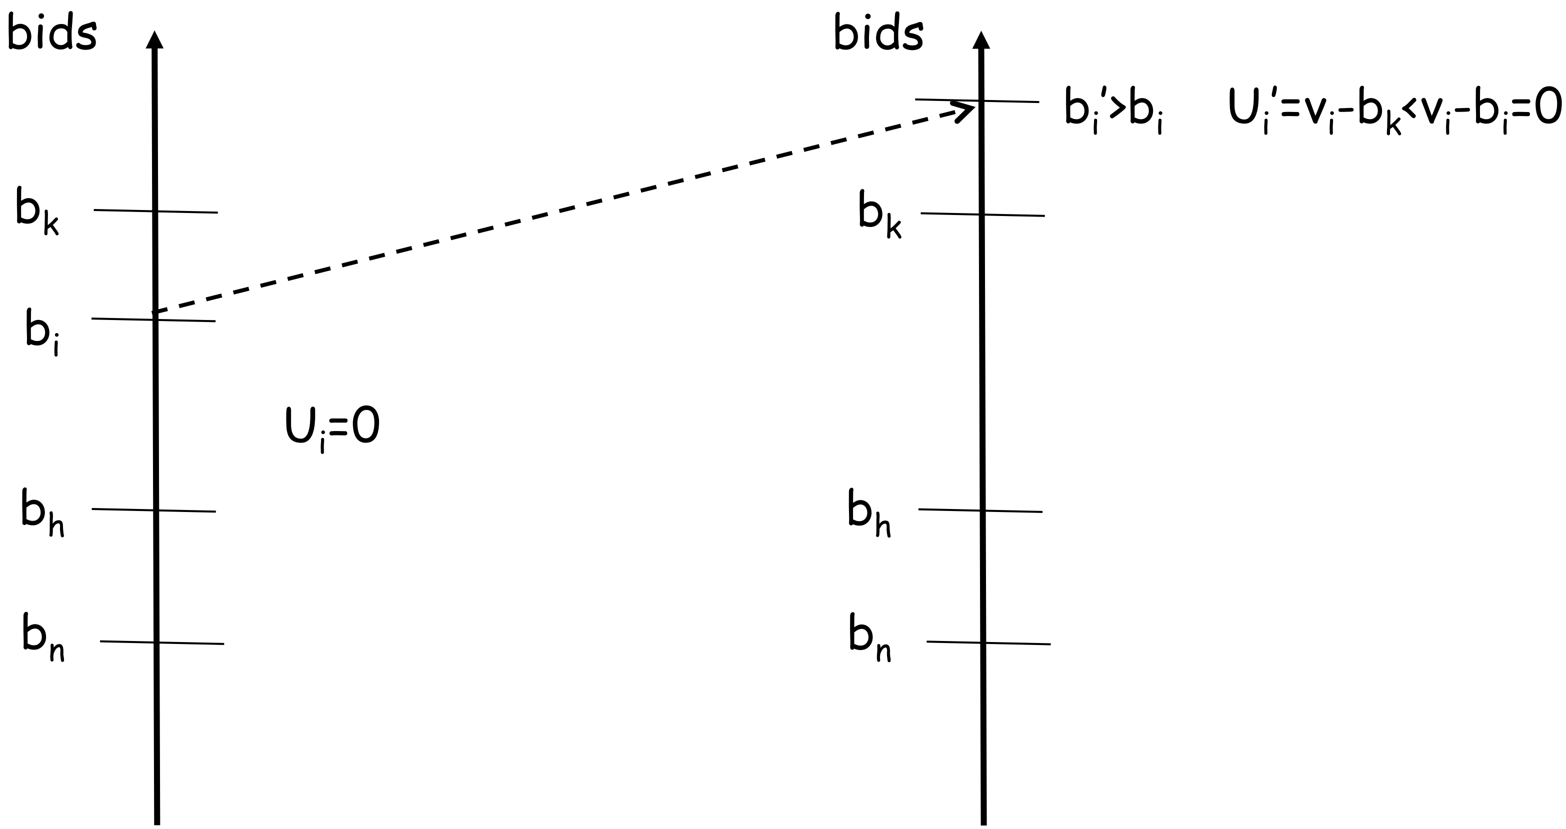

Bidding more than v_i is not convenient (may be unconvenient)

$b_i = v_i$ is not the highest bid

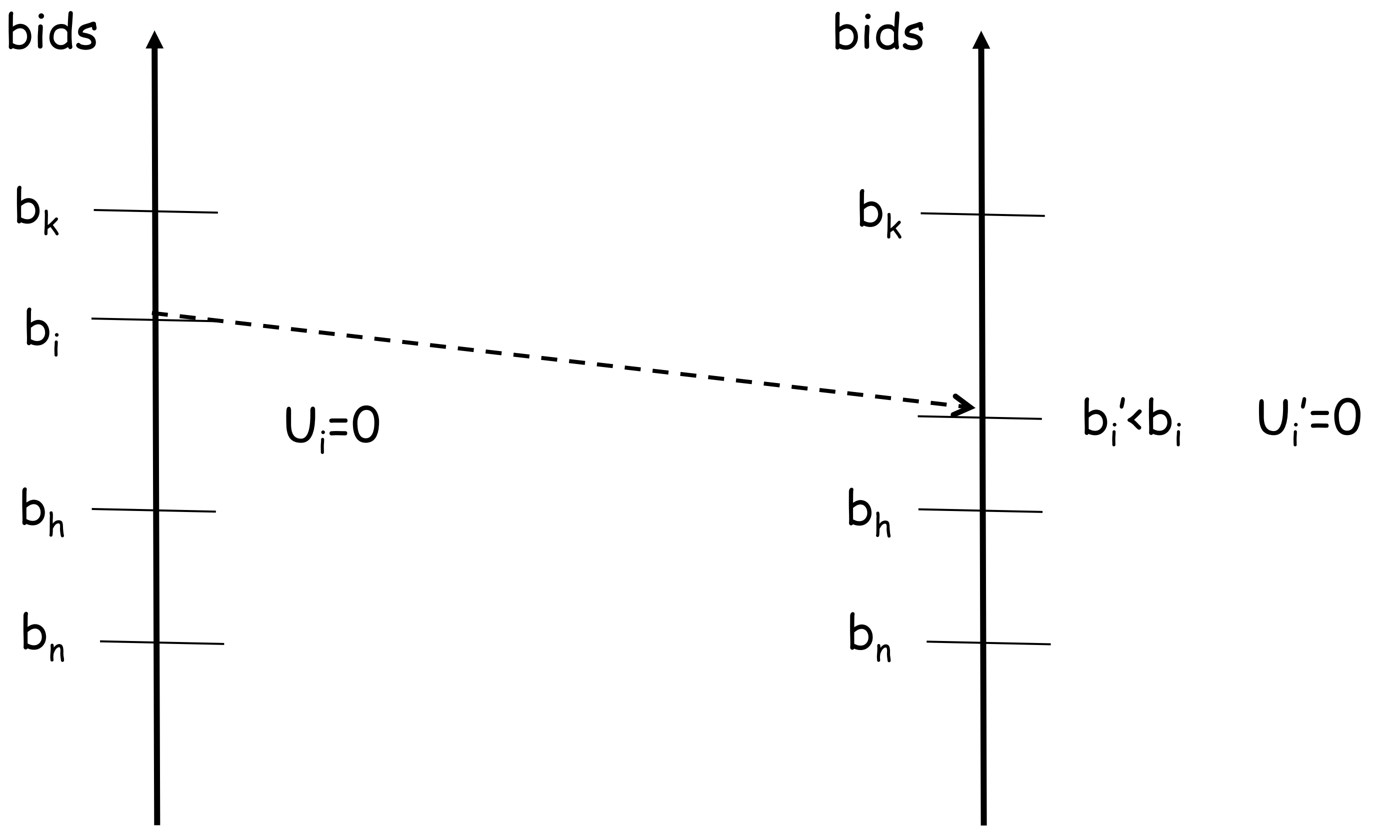

Bidding less than v_i is not convenient

Seller revenue

- N bidders
- Values are independent random values between 0 and 1
- Expected i^{th} largest utility is $(N+1-i)/(N+1)$
- Expected seller revenue is $(N-1)/(N+1)$

1^{st} price auction

- Player with the highest bid gets the good and pays a price equal to her/his bid
- Being truthful is not a dominant strategy anymore!
- How to study it?

1^{st} price auction

- Assumption: for each player the other values are i.i.d. random variables between 0 and 1
  - to overcome the fact that utilities are unknown
- Player i's strategy is a function s() mapping value v_i to a bid b_i
  - s() strictly increasing, differentiable function
  - $0 \leq s(v) \leq v \rightarrow s(0)=0$
- We investigate if there is a strategy s() common to all the players that leads to a Nash equilibrium

1^{st} price auction

- Assumption: for each player the other values are i.i.d. random variables between 0 and 1
- Player i's strategy is a function s() mapping value v_i to a bid b_i
- Expected payoff of player i if all the players plays s():
  - $U_i(s,\ldots s,\ldots s) = v_i^{N-1} (v_i - s(v_i))$

prob. i wins

i's payoff if he/she wins

1^{st} price auction

- Expected payoff of player i if all the players play s():
  - $U_i(s,\ldots s,\ldots s) = v_i^{N-1} (v_i - s(v_i))$
- What if i plays a different strategy t()?
  - If all players playing s() is a NE, then :
  - $U_i(s,\ldots s,\ldots s) = v_i^{N-1} (v_i - s(v_i)) \geq v_i^{N-1} (v_i - t(v_i)) =$

    $= U_i(s,\ldots t,\ldots s)$
- Difficult to check for all the possible functions t() different from s()
- Help from the **revelation principle**

The Revelation Principle

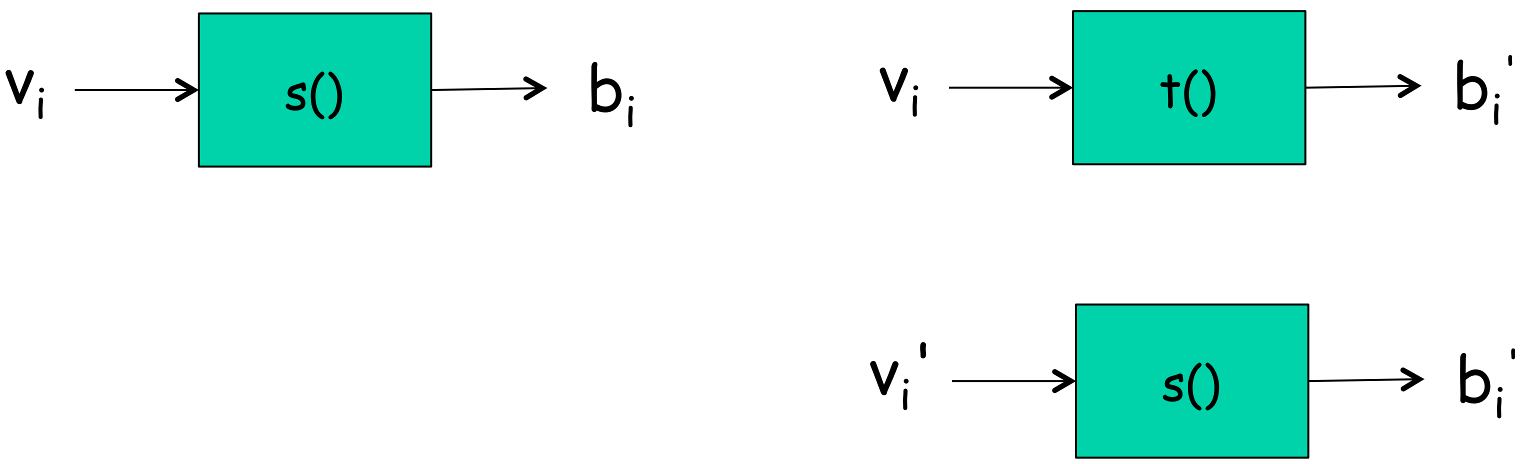

- All the strategies are equivalent to bidder i supplying to s() a different value of v_i

1^{st} price auction

- Expected payoff of player i if all the players plays s():
  - $U_i(v_1,\ldots v_i,\ldots v_N) = U_i(s,\ldots s,\ldots s) = v_i^{N-1} (v_i - s(v_i))$
- What if i plays a different strategy t()?
- By the revelation principle:
  - $U_i(s,\ldots t,\ldots s) = U_i(v_1,\ldots v,\ldots v_N) = v^{N-1} (v_i - s(v))$
- If $v_i^{N-1} (v_i - s(v_i)) \geq v^{N-1} (v_i - s(v))$ for each v (and for each v_i)
  - Then all players playing s() is a NE

1st price auction

- If $v_i^{N-1}(v_i - s(v_i)) \geq v^{N-1}(v_i - s(v))$ for each v (and for each v_i)
  - Then all players playing s() is a NE
- $f(v) = v_i^{N-1}(v_i - s(v_i)) - v^{N-1}(v_i - s(v))$ is minimized for $v = v_i$
- $f'(v) = 0$ for $v = v_i$,
  - i.e. $(N-1)\, v_i^{N-2}(v_i - s(v)) + v_i^{N-1} s'(v_i) = 0$ for each v_i
  - $s'(v_i) = (N-1)(1 - s(v_i)/v_i)$, $s(0)=0$
  - Solution: $s(v_i) = (N-1)/N\, v_i$

1^{st} price auction

- All players bidding according to
  $s(v) = (N-1)/N \, v$ is a NE
- Remarks
  - They are not truthful
  - The more they are, the higher they should bid
- Expected seller revenue
  - $((N-1)/N)\, E[v_{max}] = ((N-1)/N)\, (N/(N+1)) = (N-1)/(N+1)$
  - Identical to 2^{nd} price auction!
  - A general revenue equivalence principle

Outline

- Preliminaries
  - Auctions
  - Matching markets
- Possible approaches to ads pricing
- Google mechanism

- References
  - Easley, Kleinberg, "Networks, Crowds and Markets", ch.9,**10**,15

Matching Markets

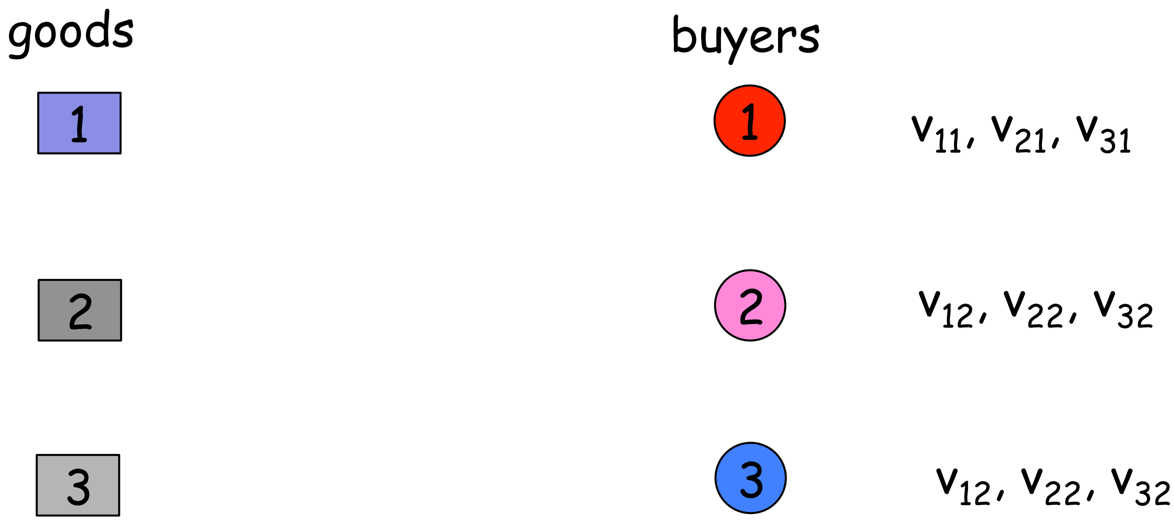

v_{ij}: value that buyer j gives to good i

How to match a set of different goods to a set of buyers with different evaluations

Matching Markets

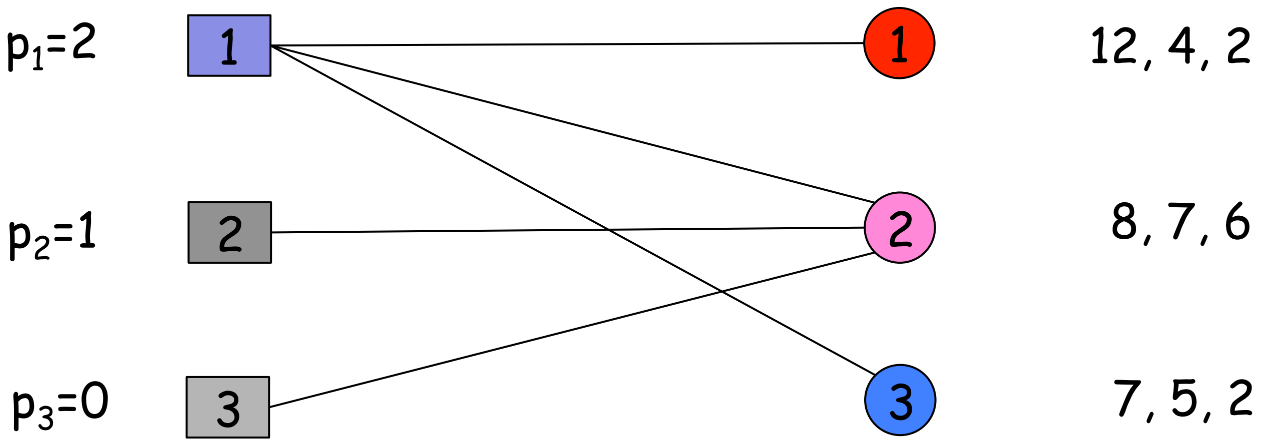

Which goods buyers like most? Preferred seller graph

How to match a set of different goods to a set of buyers with different evaluations

Matching Markets

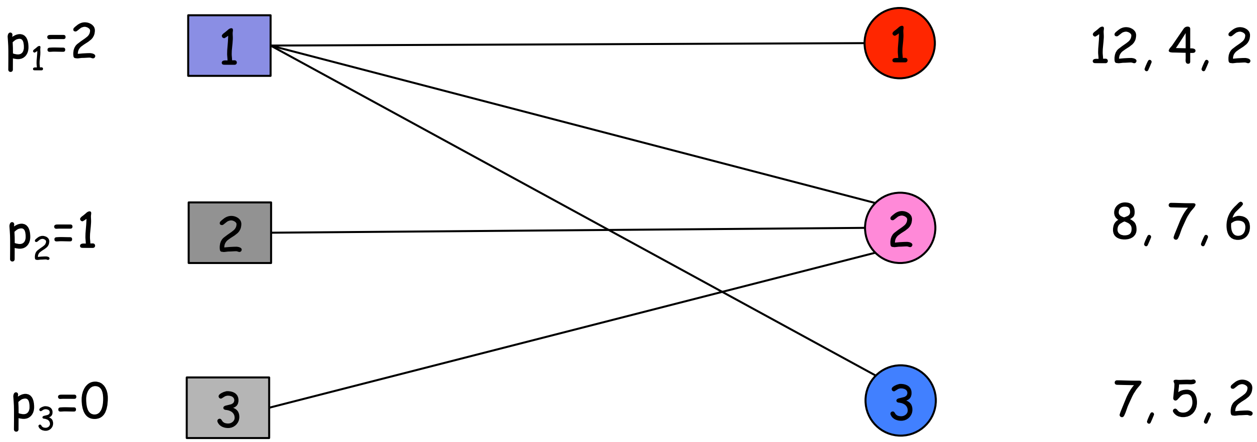

Which goods buyers like most? Preferred seller graph

- Given the prices, look for a perfect matching on the preferred seller graph
- There is no such matching for this graph

Matching Markets

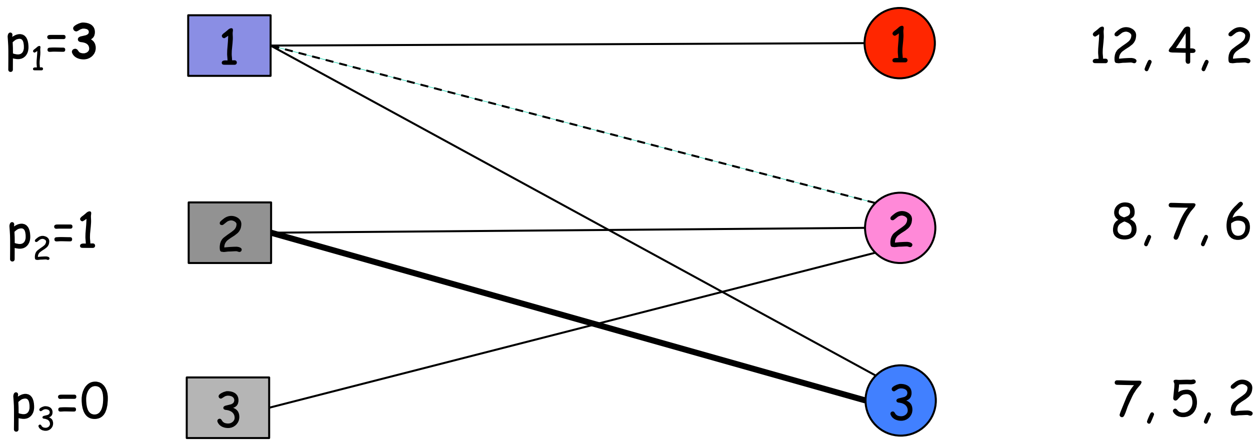

Which goods buyers like most? Preferred seller graph

- But with different prices, there is

Matching Markets

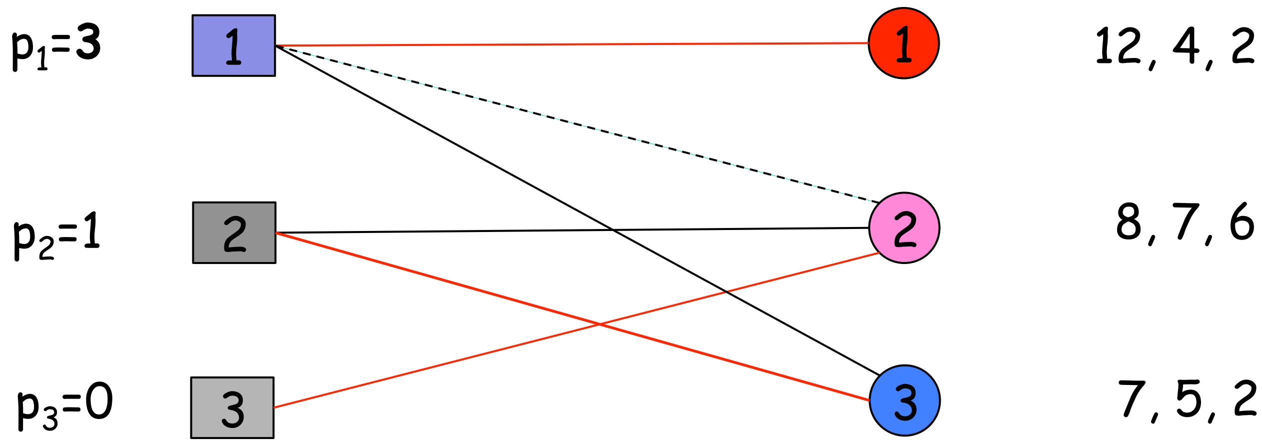

Which goods buyers like most? Preferred seller graph

- But with different prices, there is
- Such prices are **market clearing prices**

Market Clearing Prices

- They always exist
  - And can be easily calculated if valuations are known
- They are socially optimal in the sense that they maximize the sum of all the payoffs in the network (both sellers and buyers)

Outline

- Preliminaries
  - Auctions
  - Matching markets
- Possible approaches to ads pricing
- Google mechanism

- References
  - Easley, Kleinberg, "Networks, Crowds and Markets", ch.9,10,**15**

Ads pricing

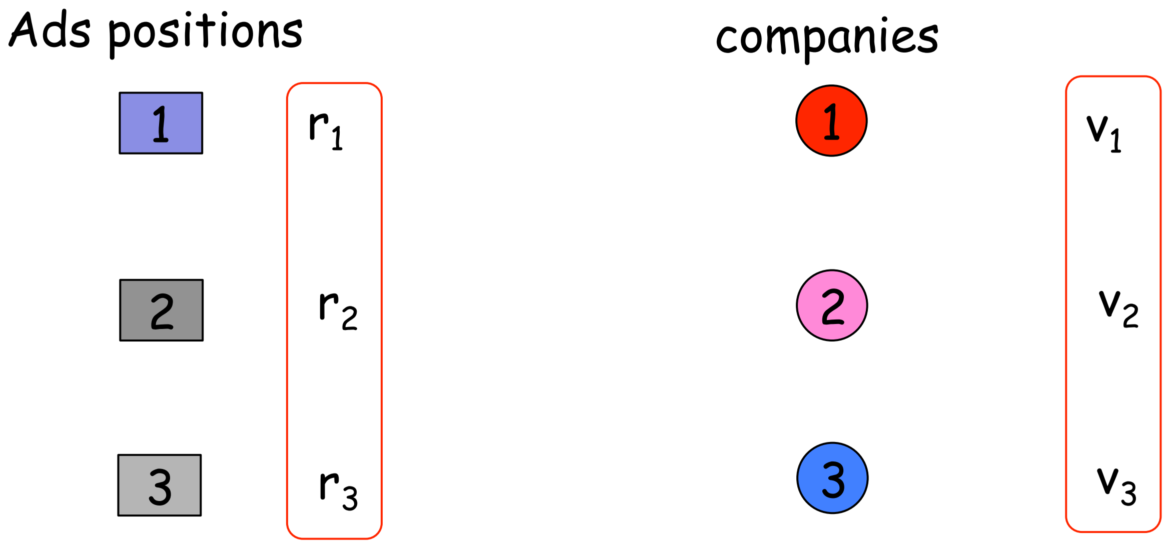

r_i: click rate for an ad in position i (assumed to be independent from the ad and known a priori)

v_i: value that company i gives to a click

How to rank ads from different companies

Ads pricing as a matching market

Ads positions		companies	
1	r_1	1	v_1r_1, v_1r_2, v_1r_3
2	r_2	2	v_2r_1, v_2r_2, v_2r_3
3	r_3	3	v_3r_1, v_3r_2, v_3r_3

r_i: click rate for an ad in position i (assumed to be independent from the ad and known a priori)

v_i: value that company i gives to a click

- Problem: Valuations are not known!
- … but we could look for something as 2nd price auctions

The VCG mechanism

- The correct way to generalize 2nd price auctions to multiple goods
- Vickrey-Clarke-Groves
- Every buyers should pay a price equal to the social value loss for the others buyers
  - Example: consider a 2nd price auction with $v_1 > v_2 > \dots v_N$
    - With 1 present the others buyers get 0
    - Without 1, 2 would have got the good with a value v_2
    - then the social value loss for the others is v_2

The VCG mechanism

- The correct way to generalize 2nd price auctions to multiple goods
- Vickrey-Clarke-Groves
- Every buyers should pay a price equal to the social value loss for the others buyers
  - If V_B^S is the maximum total valuation over all the possible perfect matchings of the set of sellers S and the set of buyers B,
  - If buyer j gets good i, he/she should be charged $V_{B-j}^{S} - V_{B-j}^{S-i}$

VCG example

Ads positions		companies	
1	$r_1=10$	1	$v_1=3$
2	$r_2=5$	2	$v_2=2$
3	$r_3=2$	3	$v_3=1$

r_i: click rate for an ad in position i (assumed to be independent from the ad and known a priori)

v_i: value that company i gives to a click

VCG example

Ads positions	companies	
1	1	30, 15, 6
2	2	20, 10, 4
3	3	10, 5, 2

VCG example

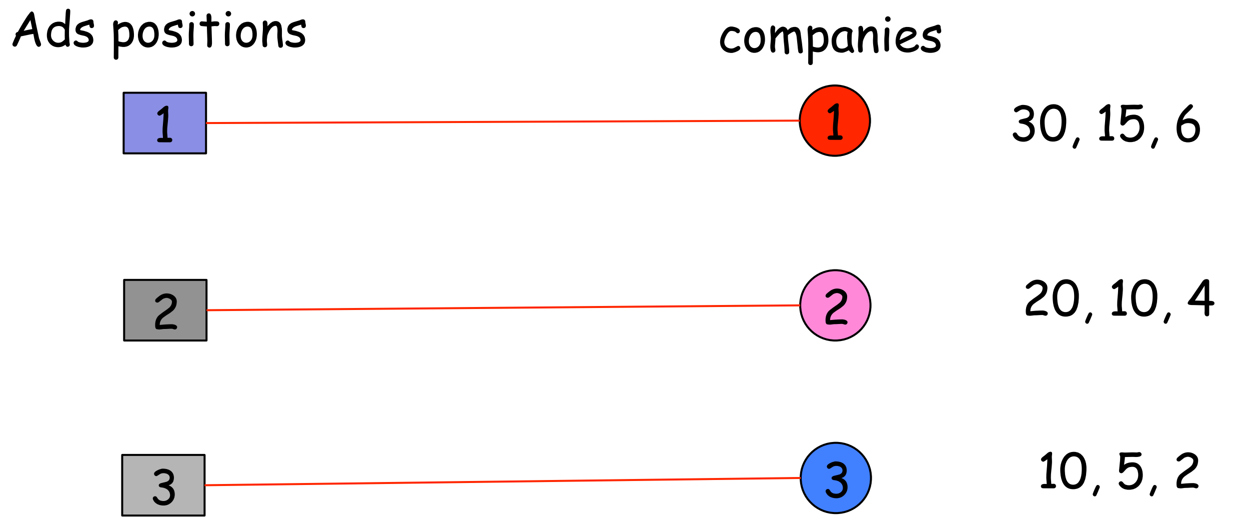

- This is the maximum weight matching
- 1 gets 30, 2 gets 10 and 3 gets 2

VCG example

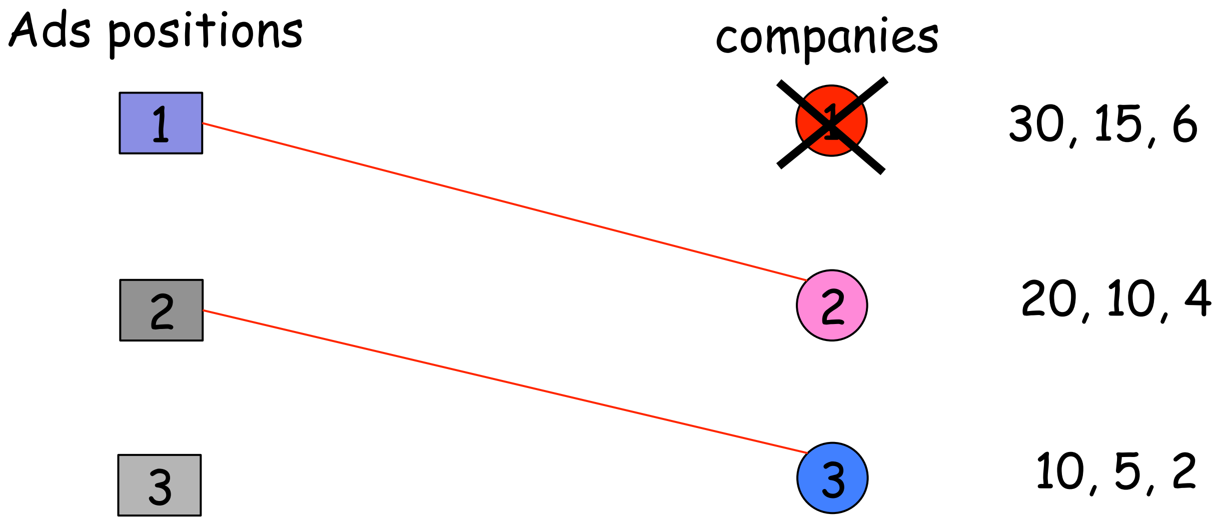

- If 1 weren't there, 2 and 3 would get 25 instead of 12,
- Then 1 should pay 13

VCG example

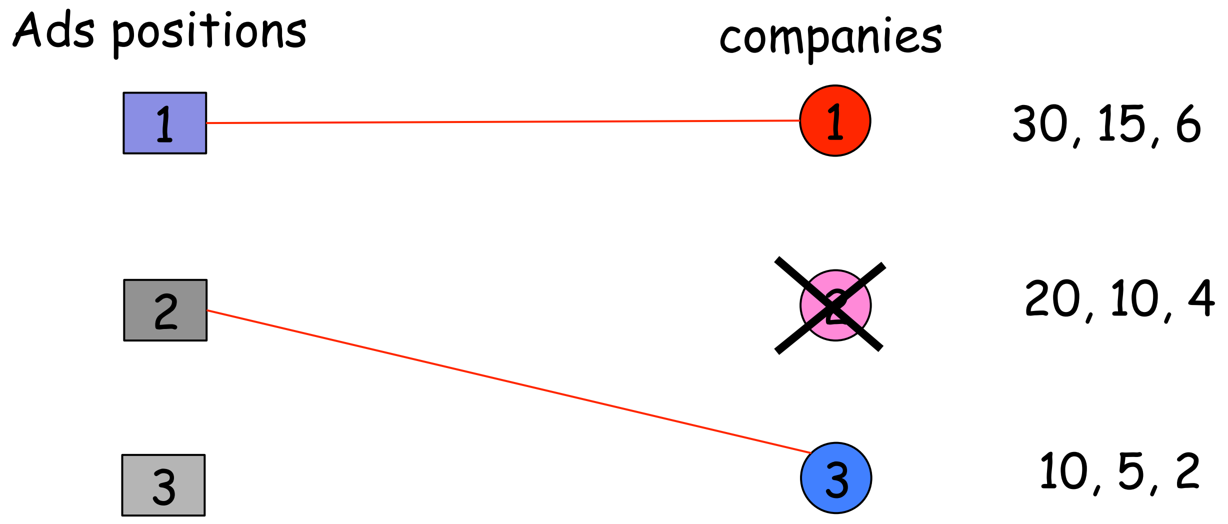

- If 2 weren't there, 1 and 3 would get 35 instead of 32,
- Then 2 should pay 3

VCG example

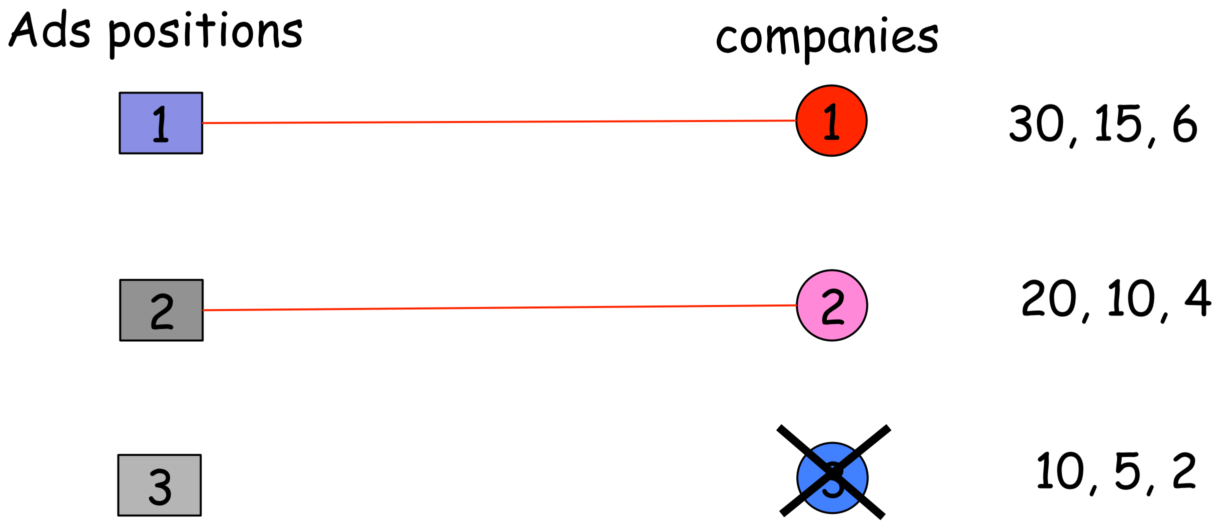

- If 3 weren't there, nothing would change for 1 and 2,
- Then 3 should pay 0

The VCG mechanism

- Every buyers should pay a price equal to the social value loss for the others buyers
  - If V_B^S is the maximum total valuation over all the possible perfect matchings of the set of sellers S and the set of buyers B,
  - If buyer j gets good i, he/she should be charged $V_{B-j}^{S} - V_{B-j}^{S-i}$
- Under this price mechanism, truth-telling is a dominant strategy

Outline

- Preliminaries
  - Auctions
  - Matching markets
- Possible approaches to ads pricing
- Google mechanism

- References
  - Easley, Kleinberg, "Networks, Crowds and Markets", ch.9,10,**15**

Google's GSP auction

- Generalized Second Price
- Once all the bids are collected $b_1 > b_2 > \ldots b_N$
- Company i pays b_{i+1}
- In the case of a single good (position), GSP is equivalent to a 2nd price auction, and also to VCG
- But why Google wanted to implement something different???

GSP properties

- Truth-telling may not be an equilibrium

GSP example

Ads positions		companies	
1	$r_1=10$	1	$v_1=7$
2	$r_2=4$	2	$v_2=6$
3	$r_3=0$	3	$v_3=1$

r_i: click rate for an ad in position i (assumed to be independent from the ad and known a priori)

v_i: value that company i gives to a click

- If each player bids its true evaluation, 1 gets a payoff equal to 10
- If 1 bids 5, 1 gets a payoff equal to 24

GSP properties

- Truth-telling may not be an equilibrium
- There is always at least 1 NE maximizing total advertiser valuation

GSP example

Ads positions

1 $r_1=10$

2 $r_2=4$

3 $r_3=0$

r_i: click rate for an ad in position i (assumed to be independent from the ad and known a priori)

companies

1 $v_1=7$

2 $v_2=6$

3 $v_3=1$

v_i: value that company i gives to a click

- Multiple NE
  - 1 bids 5, 2 bids 4 and 3 bids 2
  - 1 bids 3, 2 bids 5 and 3 bids 1

GSP properties

- Truth-telling may not be an equilibrium
- There is always at least 1 NE maximizing total advertiser valuation
- Revenues can be higher or lower than VCG
  - Attention: the revenue equivalence principle does not hold for auctions with multiple goods!
  - Google was targeting higher revenues…
  - … not clear if they did the right choice.

GSP example

Ads positions		companies	
1	$r_1=10$	1	$v_1=7$
2	$r_2=4$	2	$v_2=6$
3	$r_3=0$	3	$v_3=1$

- Multiple NE
  - 1 bids 5, 2 bids 4, 3 bids 2 → google's revenue=48
  - 1 bids 3, 2 bids 5, 3 bids 1 → google's revenue=34
- With VCG, google's revenue=44

Other issues

- Click rates are unknown and depend on the ad!
  - Concrete risk: low-quality advertiser bidding high may reduce the search engine's revenue
  - Google's solution: introduce and ad-quality factor taking into account actual click rate, relevance of the page and its ranking
    - Google is very secretive about how to calculate it => the market is more opaque
- Complex queries, nobody paid for
  - Usually engines extrapolate from simpler bids