Machine Learning, Market Design, and Advertising

Jason D. Hartline
Northwestern University

December 13, 2008

Paid Search

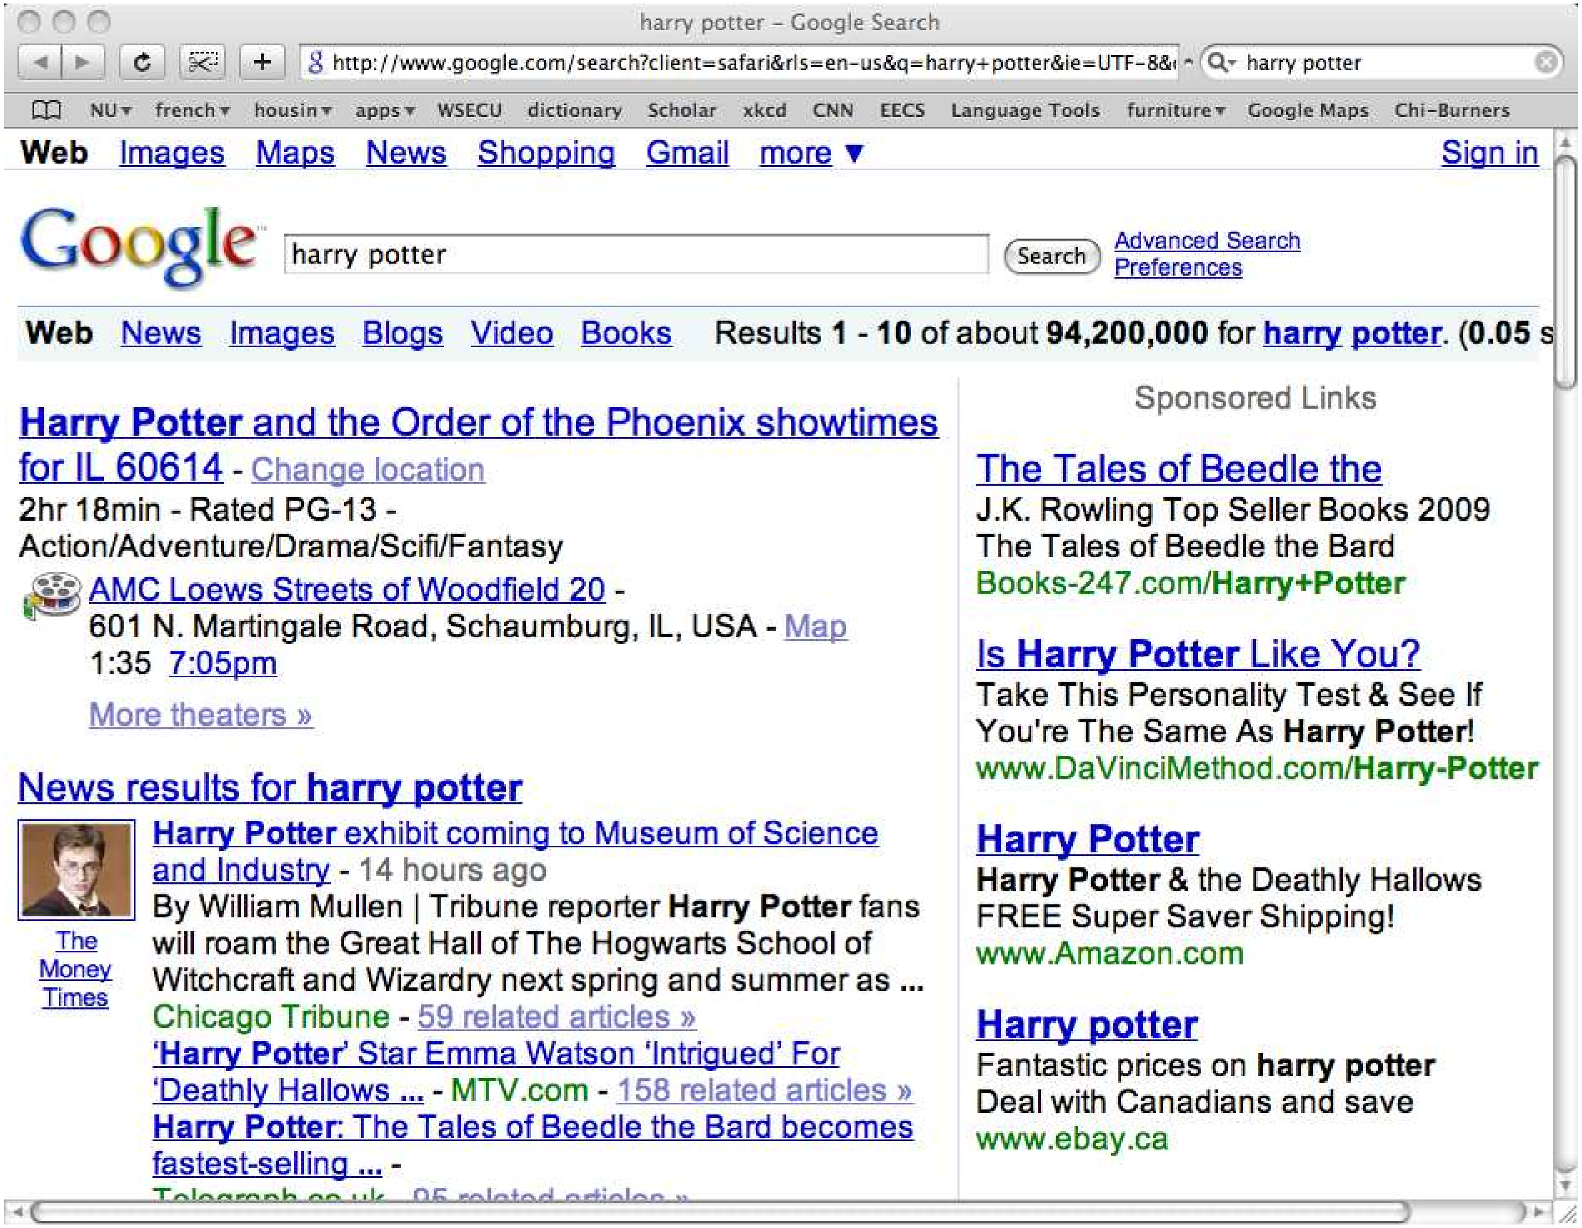

harry potter - Google Search
http://www.google.com/search?client=safari&rls=en-us&q=harry+potter&ie=UTF-8&
harry potter
NU french housin apps WSECU dictionary Scholar xkcd CNN EECS Language Tools furniture Google Maps Chi-Burners
Web Images Maps News Shopping Gmail more
Sign in
Google
harry potter
Search
Advanced Search
Preferences
Web News Images Blogs Video Books Results 1 - 10 of about 94,200,000 for harry potter. (0.05 s
Harry Potter and the Order of the Phoenix showtimes for IL 60614 - Change location
2hr 18min - Rated PG-13 - Action/Adventure/Drama/Scifi/Fantasy
AMC Loews Streets of Woodfield 20 -
601 N. Martingale Road, Schaumburg, IL, USA - Map
1:35 7:05pm
More theaters »
News results for harry potter
Harry Potter exhibit coming to Museum of Science and Industry - 14 hours ago
By William Mullen | Tribune reporter Harry Potter fans will roam the Great Hall of The Hogwarts School of Witchcraft and Wizardry next spring and summer as ...
Chicago Tribune - 59 related articles »
The Money Times
'Harry Potter' Star Emma Watson 'Intrigued' For 'Deathly Hallows ... - MTV.com - 158 related articles »
Harry Potter: The Tales of Beedle the Bard becomes fastest-selling ... -
Sponsored Links
The Tales of Beedle the
J.K. Rowling Top Seller Books 2009
The Tales of Beedle the Bard
Books-247.com/Harry+Potter
Is Harry Potter Like You?
Take This Personality Test & See If You're The Same As Harry Potter!
www.DaVinciMethod.com/Harry-Potter
Harry Potter
Harry Potter & the Deathly Hallows
FREE Super Saver Shipping!
www.Amazon.com
Harry potter
Fantastic prices on harry potter
Deal with Canadians and save
www.ebay.ca

GSP

Definition: *Generalized Second Price* (GSP) auction

- advertisers bid for keywords in advance.
- on query,
  - find all bids that *match* query.
  - rank by bid.
  - if ad clicked, charge next highest bid.

(can also scale bids by “quality” or click-through rate)

Overview

Part I: Beyond GSP.

- Advertising market overview.
- Short-comings of GSP.
- Proposal: add pre-sale market.
- Many connections to ML.

Part II: Machine learning and market design.

Part I: Beyond GSP.

Online/Search Advertising Markets

Market Participants:

Online/Search Advertising Markets

Market Participants:

- *search engine*
- *users*
- *advertisers*

Online/Search Advertising Markets

Market Participants:

- *search engine* (e.g., wants to maximize profit $=$ payments $-$ costs)
- *users*
- *advertisers*

Online/Search Advertising Markets

Market Participants:

- *search engine* (e.g., wants to maximize profit $=$ payments $-$ costs)
- *users* (e.g., want to max search/ad relevance, min search time)
- *advertisers*

Online/Search Advertising Markets

Market Participants:

- *search engine* (e.g., wants to maximize profit $=$ payments $-$ costs)
- *users* (e.g., want to max search/ad relevance, min search time)
- *advertisers* (e.g., wants max value from ads $-$ payments $-$ cost of optimizing campaign, subject to budget)

Online/Search Advertising Markets

Market Participants:

- *search engine* (e.g., wants to maximize profit $=$ payments $-$ costs)
- *users* (e.g., want to max search/ad relevance, min search time)
- *advertisers* (e.g., wants max value from ads $-$ payments $-$ cost of optimizing campaign, subject to budget)

Market Design Objectives:

Online/Search Advertising Markets

Market Participants:

- *search engine* (e.g., wants to maximize profit $=$ payments $-$ costs)
- *users* (e.g., want to max search/ad relevance, min search time)
- *advertisers* (e.g., wants max value from ads $-$ payments $-$ cost of optimizing campaign, subject to budget)

Market Design Objectives:

- *maximize welfare* $=$ user welfare $+$ advertiser welfare $-$ search engine costs.
- *maximize profit* $=$ payments $-$ costs.

Online/Search Advertising Markets

Market Participants:

- *search engine* (e.g., wants to maximize profit $=$ payments $-$ costs)
- *users* (e.g., want to max search/ad relevance, min search time)
- *advertisers* (e.g., wants max value from ads $-$ payments $-$ cost of optimizing campaign, subject to budget)

Market Design Objectives:

- *maximize welfare* $=$ user welfare $+$ advertiser welfare $-$ search engine costs.
- *maximize profit* $=$ payments $-$ costs.
  (short-term profit maximization is probably short-sighted)

Properties of GSP

Recall Definition: Generalized Second Price (GSP) auction

- advertisers bid for keywords in advance.
- on query,
  - find all bids that match query.
  - rank by bid.
  - if ad clicked, charge next highest bid.

Properties:

- *low-level bidding language:* bids for keywords.
- *decentralized:* advertisers are optimizers
- *local:* advertisers adapt bids to market conditions.
- *diffuse info:* advertisers know demand, engine knows supply.
- *online greedy:* allocation ignores future supply and past allocation

GSP non-optimality

Evidence of GSP Non-optimality:

- *search engine marketers* are necessary (i.e., significant bid cost).
- pervasive use of *broadmatch*.
- Many advertisers do not actively change bids.
- Budgets often *binding* (advertisers could bid less and get more).

Example: broadmatch

Broadmatch allows a single advertiser bid to match many search queries.

Example: broadmatch

Broadmatch allows a single advertiser bid to match many search queries.

Advantage: easy to specify and optimize a single bid.
(i.e., broadmatch has low bid-maintenance cost)

Example: broadmatch

Broadmatch allows a single advertiser bid to match many search queries.

Advantage: easy to specify and optimize a single bid.
(i.e., broadmatch has low bid-maintenance cost)

Disadvantage: not optimal for advertisers.
(absent bid-maintenance cost, better to submit different keyword bids)

- clicks for different keywords worth different amounts.
- demand for different keywords is different.
- supply of different keywords is different.

Example: broadmatch

Broadmatch allows a single advertiser bid to match many search queries.

Advantage: easy to specify and optimize a single bid.
(i.e., broadmatch has low bid-maintenance cost)

Disadvantage: not optimal for advertisers.
(absent bid-maintenance cost, better to submit different keyword bids)

- clicks for different keywords worth different amounts.
- demand for different keywords is different.
- supply of different keywords is different.

Note: better to have expressive bids and low bid-maintenance cost.

Example: "Harry Potter"

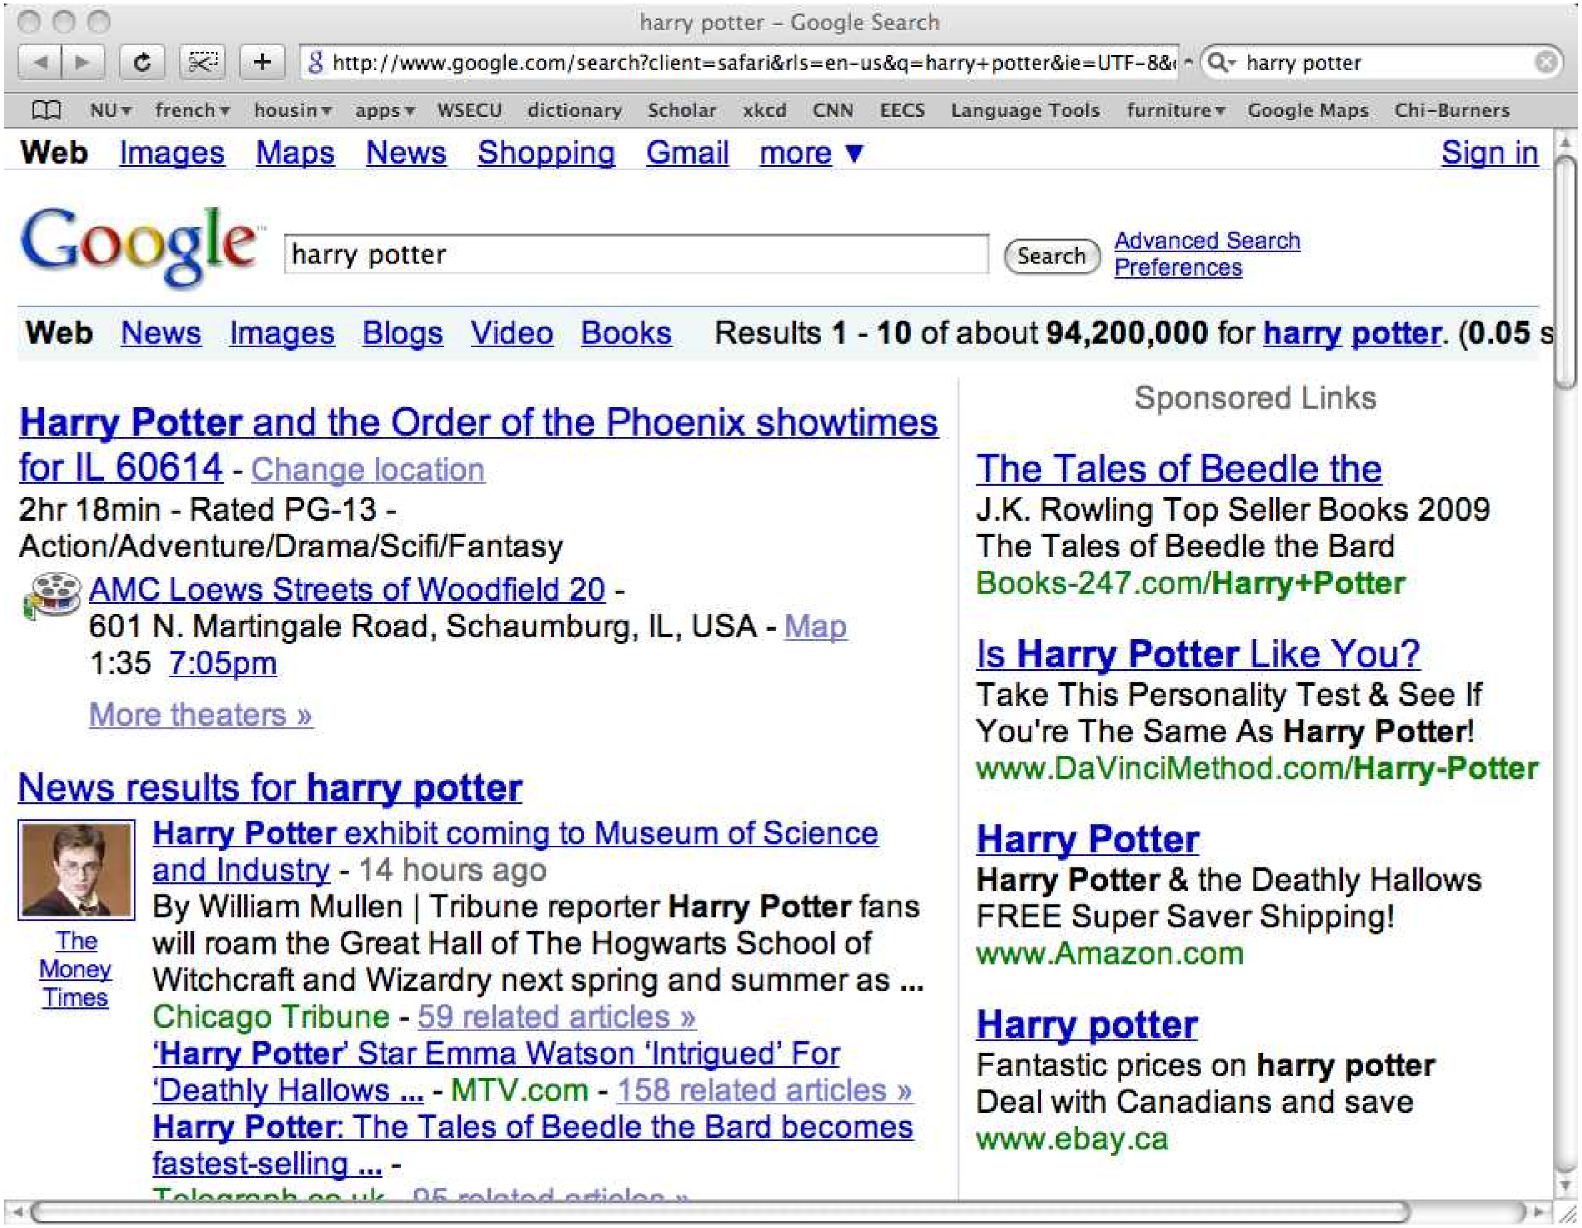

Example: “Deathly Hallows”

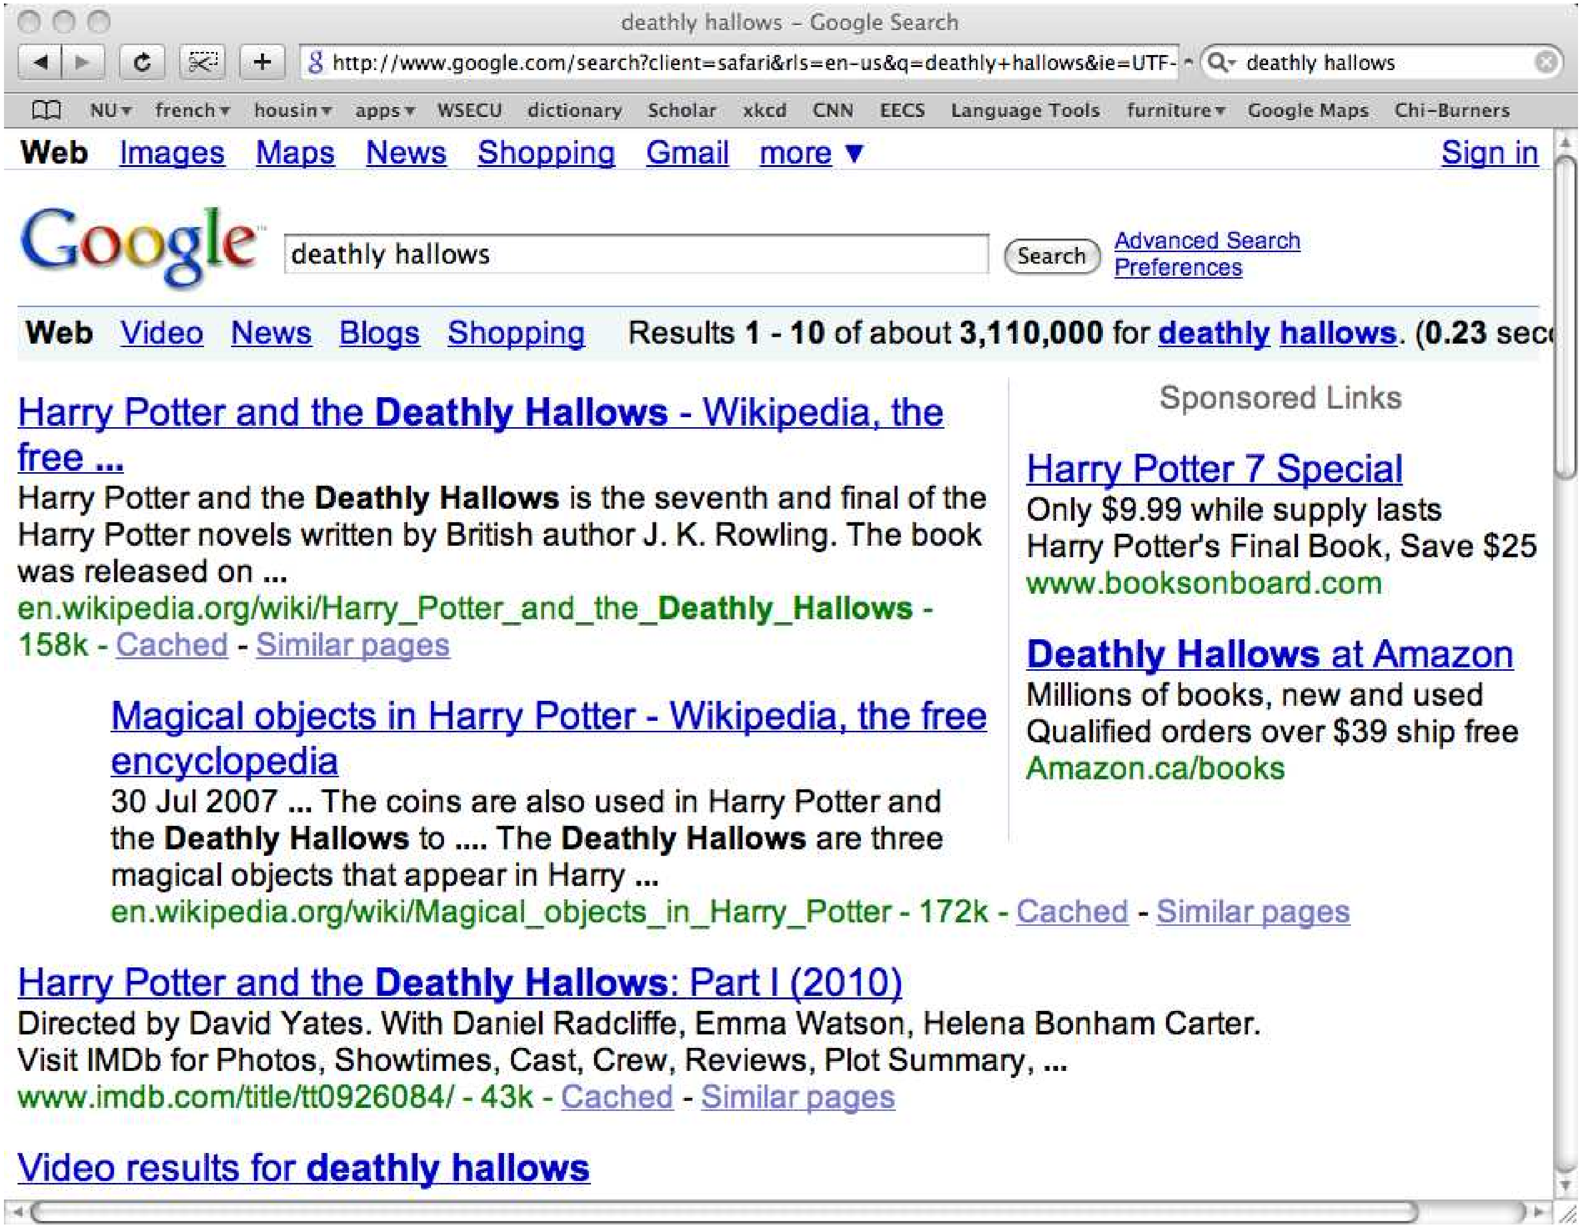

Broadmatch Discussion

Discussion:

- Compare Amazon's value-per-click:
  Probably "Harry Potter" $<$ "Deathly Hallows"
- Compare advertiser competition:
  Probably "Harry Potter" $>$ "Deathly Hallows"
- Compare keyword supply:
  Probably "Harry Potter" $>$ "Deathly Hallows"

Broadmatch Discussion

Discussion:

- Compare Amazon's value-per-click:
  Probably "Harry Potter" $<$ "Deathly Hallows"
- Compare advertiser competition:
  Probably "Harry Potter" $>$ "Deathly Hallows"
- Compare keyword supply:
  Probably "Harry Potter" $>$ "Deathly Hallows"

Conclusion: Amazon should bid differently for "H.P." vs "D.H."

Broadmatch Discussion

Discussion:

- Compare Amazon's value-per-click:
  Probably "Harry Potter" $<$ "Deathly Hallows"
- Compare advertiser competition:
  Probably "Harry Potter" $>$ "Deathly Hallows"
- Compare keyword supply:
  Probably "Harry Potter" $>$ "Deathly Hallows"

Conclusion: Amazon should bid differently for "H.P." vs "D.H."

Suggestion:

- Use "conversion tracking" to learn *conversion rates*.
  (compatible with GSP)

Broadmatch Discussion

Discussion:

- Compare Amazon's value-per-click:
  Probably "Harry Potter" $<$ "Deathly Hallows"

- Compare advertiser competition:
  Probably "Harry Potter" $>$ "Deathly Hallows"

- Compare keyword supply:
  Probably "Harry Potter" $>$ "Deathly Hallows"

Conclusion: Amazon should bid differently for "H.P." vs "D.H."

Suggestion:

- Use "conversion tracking" to learn *conversion rates*.
  (compatible with GSP)

- Use auction where advertisers bid *true value-per-click*.
  (incompatible with GSP)

Challenges and Tasks

Challenges:

1. complex advertiser and user preferences.
2. online supply.
3. large tail.
4. incentives (esp. with budgets)

Challenges and Tasks

Challenges:	**Tasks:**
1. complex advertiser and user preferences.	1. learn preferences.
2. online supply.	2. predict future supply
3. large tail.	3. cluster tail.
4. incentives (esp. with budgets)	4. pricing based mech. design.

Challenges and Tasks

Challenges:	**Tasks:**
1. complex advertiser and user preferences.	1. learn preferences.
2. online supply.	2. predict future supply
3. large tail.	3. cluster tail.
4. incentives (esp. with budgets)	4. pricing based mech. design.

Note: These do not fit into GSP model.

Challenges and Tasks

Challenges:	**Tasks:**
1. complex advertiser and user preferences.	1. learn preferences.
2. online supply.	2. predict future supply
3. large tail.	3. cluster tail.
4. incentives (esp. with budgets)	4. pricing based mech. design.

Note: These do not fit into GSP model.

What would be a better mechanism?

Beyond GSP

Rethinking the Ad Market Mechanism:

Combine pre-sale (offline) mechanism with spot (online) mech.

Beyond GSP

Rethinking the Ad Market Mechanism:

Combine pre-sale (offline) mechanism with spot (online) mech.

Almost all mature markets have pre-sales!

Rethinking the Ad Market Mechanism:

Combine pre-sale (offline) mechanism with spot (online) mech.

Almost all mature markets have pre-sales!

Related Examples:

- *timber*: 20% spot auction, 80% pre-sale (prices from spot)
- *pollution allowance*: short and medium-term markets.
- *electricity markets*: short ($\leq$ 1 day), medium (1–3 years), long-term (4–20 years) markets.

Beyond GSP

Rethinking the Ad Market Mechanism:

Combine pre-sale (offline) mechanism with spot (online) mech.

Almost all mature markets have pre-sales!

Related Examples:

- *timber*: 20% spot auction, 80% pre-sale (prices from spot)
- *pollution allowance*: short and medium-term markets.
- *electricity markets*: short ($\leq$ 1 day), medium (1–3 years), long-term (4–20 years) markets.

How should we design the advertising pre-sale market?

Part II: Machine learning and market design.

Setting

Setting:

- can estimate supply.
- can estimate preferences.
  (if advertisers provide automated reports)
- can cluster tail.

Market Design Goal:

- incentivize advertisers to provide automated reports.
- optimize objective.

Pricing-based mechanisms

Definition: an *offer* is a "menu" that maps *bundles of goods* to *prices*.

Pricing-based mechanisms

Definition: an *offer* is a “menu” that maps *bundles of goods* to *prices*.

Note: *advertiser preference* and *offer* induce a *demand* and *payment*.

- *demand*: sell advertiser their most preferred bundle.
  (at given prices)

- *payment*: charge bundle’s price.

Pricing-based mechanisms

Definition: an *offer* is a “menu” that maps *bundles of goods* to *prices*.

Note: *advertiser preference* and *offer* induce a *demand* and *payment*.

- *demand*: sell advertiser their most preferred bundle.
  (at given prices)

- *payment*: charge bundle’s price.

Note: for advertiser to get most preferred bundle, search engine needs to have accurate model of advertiser preferences.

Pricing-based mechanisms

Definition: an *offer* is a “menu” that maps *bundles of goods* to *prices*.

Note: *advertiser preference* and *offer* induce a *demand* and *payment*.

- *demand*: sell advertiser their most preferred bundle.
  (at given prices)

- *payment*: charge bundle’s price.

Note: for advertiser to get most preferred bundle, search engine needs to have accurate model of advertiser preferences.

Claim: For any fixed offer, reporting true preferences is optimal.

Pricing-based mechanisms

Definition: an *offer* is a "menu" that maps *bundles of goods* to *prices*.

Note: *advertiser preference* and *offer* induce a *demand* and *payment*.

- *demand*: sell advertiser their most preferred bundle.
  (at given prices)

- *payment*: charge bundle's price.

Note: for advertiser to get most preferred bundle, search engine needs to have accurate model of advertiser preferences.

Claim: For any fixed offer, reporting true preferences is optimal.

Advertiser may as well opt-in to automated resports.

Pricing-based mechanisms

Definition: an *offer* is a "menu" that maps *bundles of goods* to *prices*.

Note: *advertiser preference* and *offer* induce a *demand* and *payment*.

- *demand*: sell advertiser their most preferred bundle.
  (at given prices)

- *payment*: charge bundle's price.

Note: for advertiser to get most preferred bundle, search engine needs to have accurate model of advertiser preferences.

Claim: For any fixed offer, reporting true preferences is optimal.

Advertiser may as well opt-in to automated resports.

Claim: many justifications for pricing-based approach.

Limited Supply

Consider:

- make the same offer to all advertisers,

Limited Supply

Consider:

- make the same offer to all advertisers,
- but supply of keyword impressions is limited,

Limited Supply

Consider:

- make the same offer to all advertisers,
- but supply of keyword impressions is limited,
- so offer may result in over-demanded keywords.

Limited Supply

Consider:

- make the same offer to all advertisers,
- but supply of keyword impressions is limited,
- so offer may result in over-demanded keywords.

Solution: *random priority*: order advertisers at random, make offer "while supplies last".

Limited Supply

Consider:

- make the same offer to all advertisers,
- but supply of keyword impressions is limited,
- so offer may result in over-demanded keywords.

Solution: *random priority*: order advertisers at random, make offer "while supplies last".

Result: Well defined expected performance of any offer.

Limited Supply

Consider:

- make the same offer to all advertisers,
- but supply of keyword impressions is limited,
- so offer may result in over-demanded keywords.

Solution: *random priority*: order advertisers at random, make offer "while supplies last".

Result: Well defined expected performance of any offer.

Natural Objective: for class of offers $\mathcal{G}$, find offer that maximizes objective payoff. (e.g., social welfare, profit, etc.)

Optimizing Offers

Optimization Challenge: given preferences and supplies, compute offer with highest performance.

Optimizing Offers

Optimization Challenge: given preferences and supplies, compute offer with highest performance.

- intractable for general preferences.
- focus on properties of advertising enable tractability.

Optimizing Offers

Optimization Challenge: given preferences and supplies, compute offer with highest performance.

- intractable for general preferences.
- focus on properties of advertising enable tractability.

Incentive Challenge: advertisers can manipulate this optimal offer.

Can we design mech. where it is optimal to report true preferences?

Approach 1: Random Sampling Auction

Random Sampling Optimal Offer Auction, RSOO$_{\mathcal{G}}$

1. Randomly partition bidders into two sets, S_1 and S_2.
2. compute optimal offers, g_1 and g_2, for each set.
3. Offer g_1 to S_2 and g_2 to S_1.

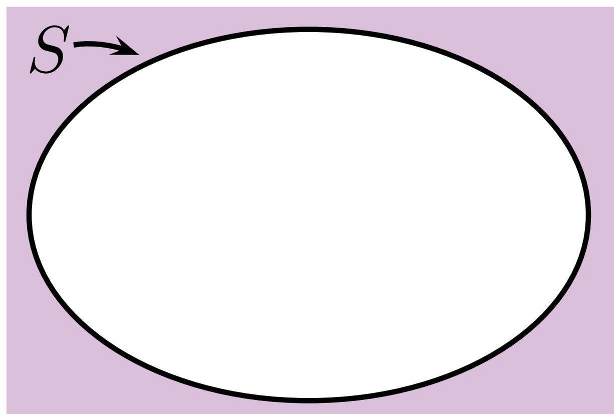

Approach 1: Random Sampling Auction

Random Sampling Optimal Offer Auction, RSOO$_{\mathcal{G}}$

1. Randomly partition bidders into two sets, S_1 and S_2.
2. compute optimal offers, g_1 and g_2, for each set.
3. Offer g_1 to S_2 and g_2 to S_1.

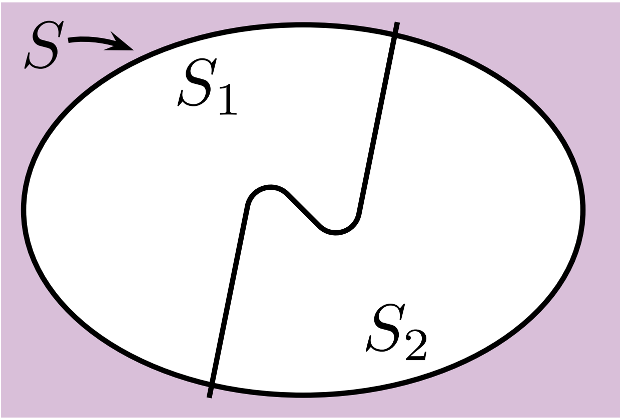

Approach 1: Random Sampling Auction

Random Sampling Optimal Offer Auction, $\text{RSOO}_{\mathcal{G}}$

1. Randomly partition bidders into two sets, S_1 and S_2.
2. compute optimal offers, g_1 and g_2, for each set.
3. Offer g_1 to S_2 and g_2 to S_1.

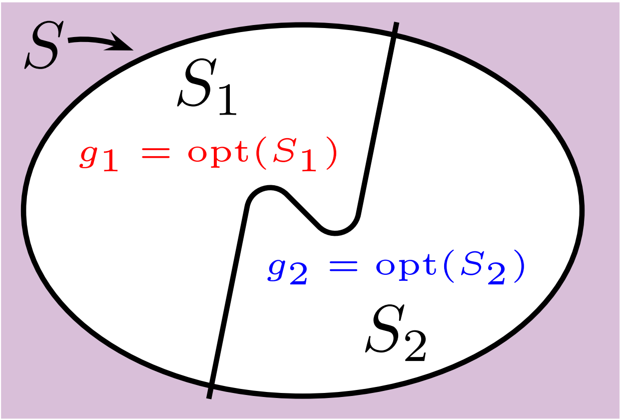

Approach 1: Random Sampling Auction

Random Sampling Optimal Offer Auction, RSOO$_{\mathcal{G}}$

1. Randomly partition bidders into two sets, S_1 and S_2.
2. compute optimal offers, g_1 and g_2, for each set.
3. Offer g_1 to S_2 and g_2 to S_1.

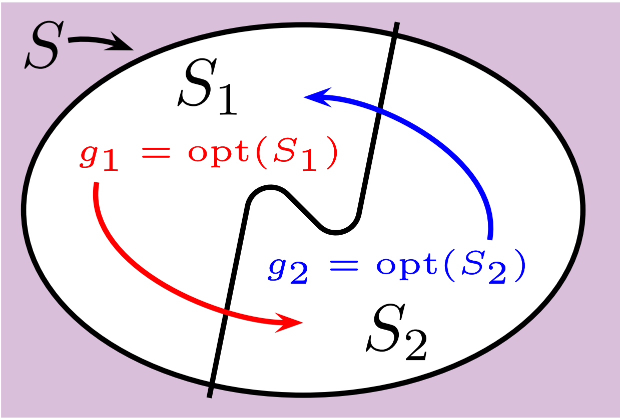

Approach 1: Random Sampling Auction

Random Sampling Optimal Offer Auction, RSOO$_{\mathcal{G}}$

1. Randomly partition bidders into two sets, S_1 and S_2.
2. compute optimal offers, g_1 and g_2, for each set.
3. Offer g_1 to S_2 and g_2 to S_1.

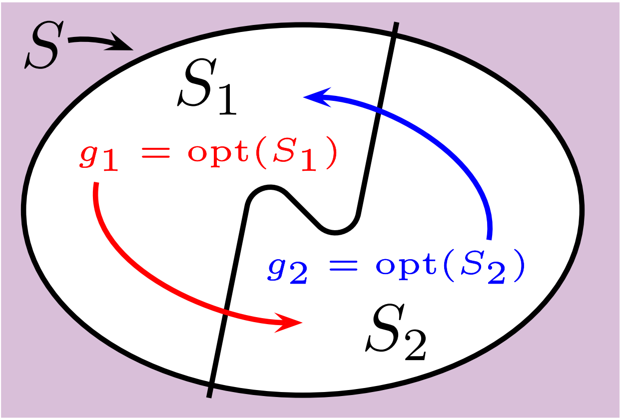

Claim: In RSOO$_{\mathcal{G}}$, reporting true preferences is optimal.

Approach 1: Random Sampling Auction

Random Sampling Optimal Offer Auction, RSOO$_{\mathcal{G}}$

1. Randomly partition bidders into two sets, S_1 and S_2.
2. compute optimal offers, g_1 and g_2, for each set.
3. Offer g_1 to S_2 and g_2 to S_1.

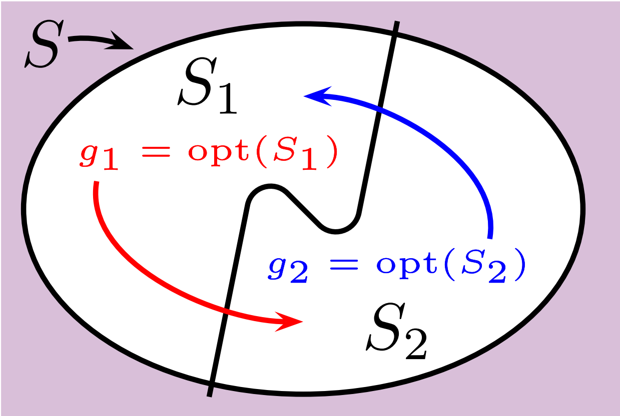

Claim: In RSOO$_{\mathcal{G}}$, reporting true preferences is optimal.

Question: when does RSOO$_{\mathcal{G}}$ perform well?

Approach 1: Random Sampling Auction

Random Sampling Optimal Offer Auction, $\mathrm{RSOO}_{\mathcal{G}}$

1. Randomly partition bidders into two sets, S_1 and S_2.
2. compute optimal offers, g_1 and g_2, for each set.
3. Offer g_1 to S_2 and g_2 to S_1.

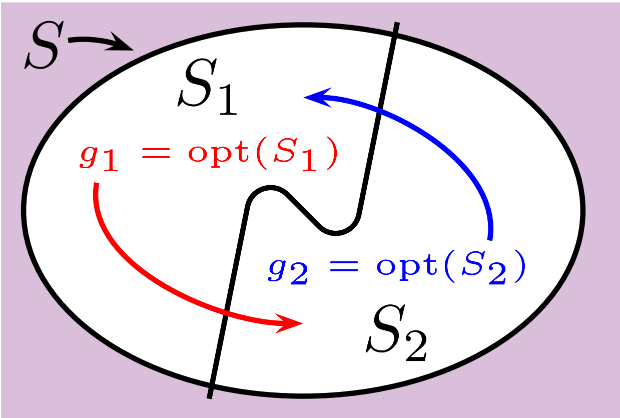

Claim: In $\mathrm{RSOO}_{\mathcal{G}}$, reporting true preferences is optimal.

Question: when does $\mathrm{RSOO}_{\mathcal{G}}$ perform well?

Note: close connection to sample complexity and machine learning.

RSOO Performance

Theorem: (Approximately) For any linear objective (e.g., welfare or profit), class of offers $\mathcal{G}$, and ϵ;

$$\mathbf{E}\big[\text{RSOO}_{\mathcal{G}}\big] \geq (1-\epsilon)\,\text{OPT}_{\mathcal{G}}$$

as long as

$$\text{OPT}_{\mathcal{G}} \geq \frac{h}{\epsilon^2} \log \frac{|\mathcal{G}|}{\epsilon}$$

and h is upper bound on payoff from any agent.

RSOO Performance

Theorem: (Approximately) For any linear objective (e.g., welfare or profit), class of offers $\mathcal{G}$, and ϵ;

$$\textbf{E}\big[\text{RSOO}_{\mathcal{G}}\big] \geq (1-\epsilon)\,\mathrm{OPT}_{\mathcal{G}}$$

as long as

$$\mathrm{OPT}_{\mathcal{G}} \geq \tfrac{h}{\epsilon^2} \log \tfrac{|\mathcal{G}|}{\epsilon}$$

and h is upper bound on payoff from any agent.

Interpretation: convergence rate is $O(h \log |\mathcal{G}|)$.

Example: tee shirts

Example: Selling tee shirts.

- Bidders with valuations in $[1, h]$ for a tee shirt.
- Reasonable offers: $\mathcal{G} = \{\text{price } 2^i$ for $i \in \{1, \ldots, \log h\}\}$.
- Convergence Rate: $O(h \log |\mathcal{G}|) = O(h \log \log h)$

Extensions

Recall Interpretation: convergence rate is $O(h \log |\mathcal{G}|)$.

Extensions:

- use *covering* arguments to improve bounds.
- use *structural-risk-minimization* to penalize for “complex” offers.

Selected References:

- Pricing Algorithms: E.g., [Gurusuami et al., 2005]
- Unlimited Supply: [Balcan et al., 2005]
- Limited Supply: [Balcan et al., unpublished]

Approach 2: Differential Privacy

Definition: A function f satisfies *ϵ-differential privacy* if for S and S' differing in one coordinate and set R in range of f,

$$\mathbf{Pr}[f(S) \in R] \leq e^{\epsilon} \times \mathbf{Pr}[f(S') \in R]$$

Approach 2: Differential Privacy

Definition: A function f satisfies *ϵ-differential privacy* if for S and S' differing in one coordinate and set R in range of f,

$$\mathbf{Pr}[f(S) \in R] \leq e^{\epsilon} \times \mathbf{Pr}[f(S') \in R]$$

Note: if near optimal offer can be computed with ϵ-diff. privacy, advertisers cannot manipulate it.

Approach 2: Differential Privacy

Definition: A function f satisfies *ϵ-differential privacy* if for S and S' differing in one coordinate and set R in range of f,

$$\mathbf{Pr}[f(S) \in R] \leq e^{\epsilon} \times \mathbf{Pr}[f(S') \in R]$$

Note: if near optimal offer can be computed with ϵ-diff. privacy, advertisers cannot manipulate it.

Comment: in fact, perhaps all services that use private data should satisfy ϵ-differential privacy.

Approach 2: Differential Privacy

Definition: A function f satisfies *ϵ-differential privacy* if for S and S' differing in one coordinate and set R in range of f,

$$\textbf{Pr}[f(S) \in R] \leq e^{\epsilon} \times \textbf{Pr}[f(S') \in R]$$

Note: if near optimal offer can be computed with ϵ-diff. privacy, advertisers cannot manipulate it.

Comment: in fact, perhaps all services that use private data should satisfy ϵ-differential privacy.

Selected References:

- Differential Privacy: [Dwork, 2006]
- Differential Privacy Auction: [McSherry and Talwar, 2007]

Approach 2: Differential Privacy Auction

Privacy Preserving Optimal Offer Auction, DPOO$_{\mathcal{G}}$

1. Compute approximately optimal offer g with ϵ-diff. privacy.
2. Offer g to all advertisers.

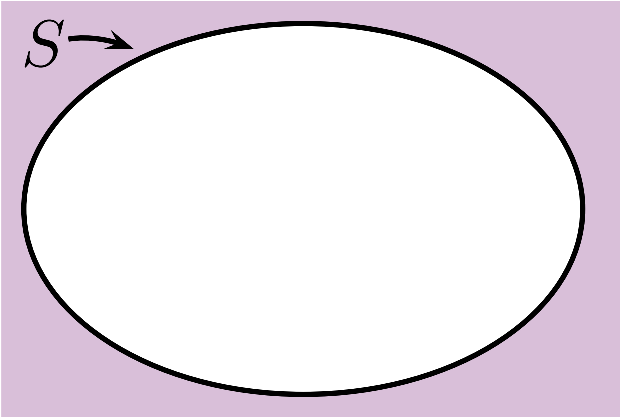

Approach 2: Differential Privacy Auction

Privacy Preserving Optimal Offer Auction, $\text{DPOO}_{\mathcal{G}}$

1. Compute approximately optimal offer g with ϵ-diff. privacy.
2. Offer g to all advertisers.

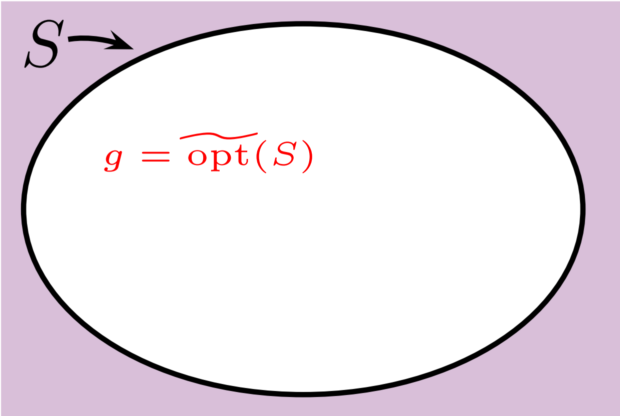

Approach 2: Differential Privacy Auction

Privacy Preserving Optimal Offer Auction, DPOO$_{\mathcal{G}}$

1. Compute approximately optimal offer g with ϵ-diff. privacy.
2. Offer g to all advertisers.

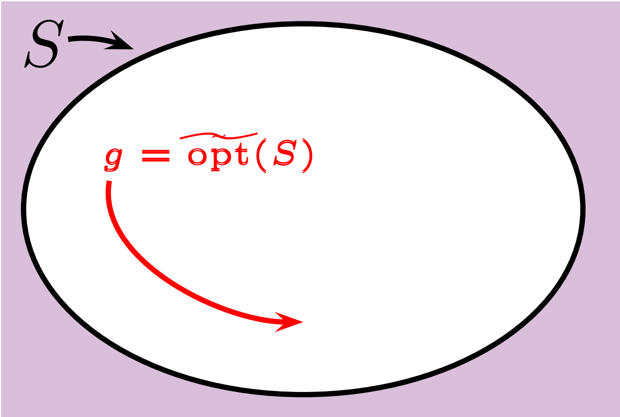

Approach 2: Differential Privacy Auction

Privacy Preserving Optimal Offer Auction, DPOO$_{\mathcal{G}}$

1. Compute approximately optimal offer g with ϵ-diff. privacy.
2. Offer g to all advertisers.

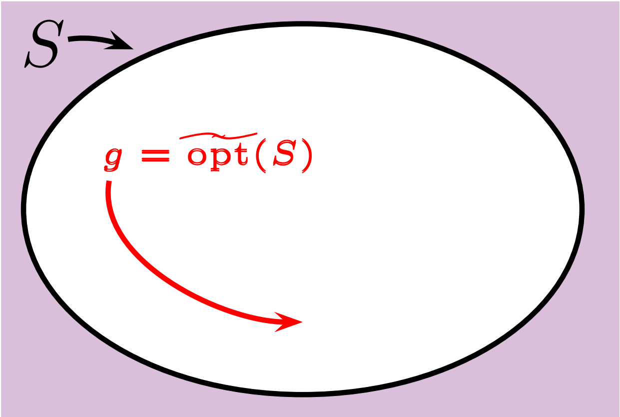

Claim: DPOO$_{\mathcal{G}}$ is has near optimal performance.

Approach 2: Differential Privacy Auction

Privacy Preserving Optimal Offer Auction, DPOO$_{\mathcal{G}}$

1. Compute approximately optimal offer g with ϵ-diff. privacy.
2. Offer g to all advertisers.

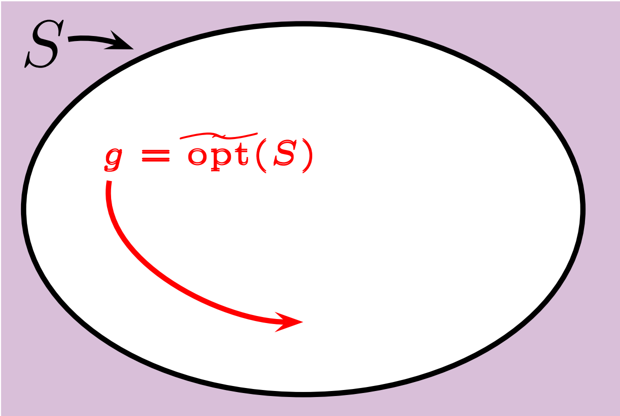

Claim: DPOO$_{\mathcal{G}}$ is has near optimal performance.

Claim: With high probability in DPOO$_{\mathcal{G}}$, reporting true preferences is optimal.

Approach 2: Differential Privacy Auction

Privacy Preserving Optimal Offer Auction, DPOO$_{\mathcal{G}}$

1. Compute approximately optimal offer g with ϵ-diff. privacy.
2. Offer g to all advertisers.

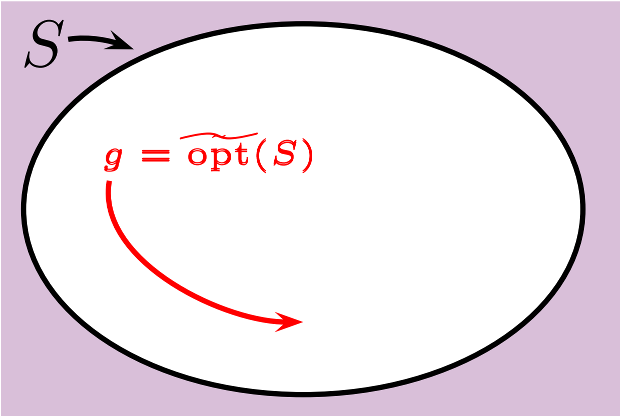

Claim: DPOO$_{\mathcal{G}}$ is has near optimal performance.

Claim: With high probability in DPOO$_{\mathcal{G}}$, reporting true preferences is optimal.

Note: "high probability" is as $\mathrm{OPT} \gg h \log |\mathcal{G}|$.

Conclusions

1. GSP unlikely to optimize desired objectives.
2. ML can significantly help advertising market design.
   - predict supply.
   - learn preferences.
   - cluster tail.
   - pricing-based mechanisms.
3. advertising markets need pre-sale market.
4. pricing-based mechanisms may be right way to go.