Auction Theory and the Internet

Jason D. Hartline
Northwestern University

Textbook Manuscript:
Mechanism Design and Approximation
`http://jasonhartline.com/MDnA/`

Mechanism Design

Basic Mechanism Design Question: How should an economic system be designed so that selfish agent behavior leads to good outcomes?

Mechanism Design

Basic Mechanism Design Question: How should an economic system be designed so that selfish agent behavior leads to good outcomes?

Internet Protocols and Applications: file sharing, reputation systems, web search, web advertising, email, Internet auctions, congestion control, etc.

Mechanism Design

Basic Mechanism Design Question: How should an economic system be designed so that selfish agent behavior leads to good outcomes?

Internet Protocols and Applications: file sharing, reputation systems, web search, web advertising, email, Internet auctions, congestion control, etc.

General Theme: resource allocation.

Case Studies

Bid sniping in eBay

- solution: activity rule (Amazon)
  (extend endtime after last bid)

Case Studies

Bid sniping in eBay

- solution: activity rule (Amazon)
  (extend endtime after last bid)

Freeloading in P2P File Sharing

- solution: tit-for-tat (BitTorrent)
  (no uploads $\Rightarrow$ no downloads)

Case Studies

Bid sniping in eBay

- solution: activity rule (Amazon)
  (extend endtime after last bid)

Freeloading in P2P File Sharing

- solution: tit-for-tat (BitTorrent)
  (no uploads ⇒ no downloads)

Email spam

- solution 1: filtering
- solution 2: micropayments
- solution 3: proofs of work

Case Studies

Bid sniping in eBay

- solution: activity rule (Amazon)
  (extend endtime after last bid)

Freeloading in P2P File Sharing

- solution: tit-for-tat (BitTorrent)
  (no uploads $\Rightarrow$ no downloads)

Email spam

- solution 1: filtering
- solution 2: micropayments
- solution 3: proofs of work

Showrooming in Search Markets

- solution 1: parity clause?
- solution 2: loyalty program?
- solution 3: paid placement?

Case Studies

Bid sniping in eBay

- solution: activity rule (Amazon)
  (extend endtime after last bid)

Freeloading in P2P File Sharing

- solution: tit-for-tat (BitTorrent)
  (no uploads ⇒ no downloads)

Email spam

- solution 1: filtering
- solution 2: micropayments
- solution 3: proofs of work

Showrooming in Search Markets

- solution 1: parity clause?
- solution 2: loyalty program?
- solution 3: paid placement?

Conclusion: incentive problems need incentive solutions.

Bid-sniping in eBay vs Amazon

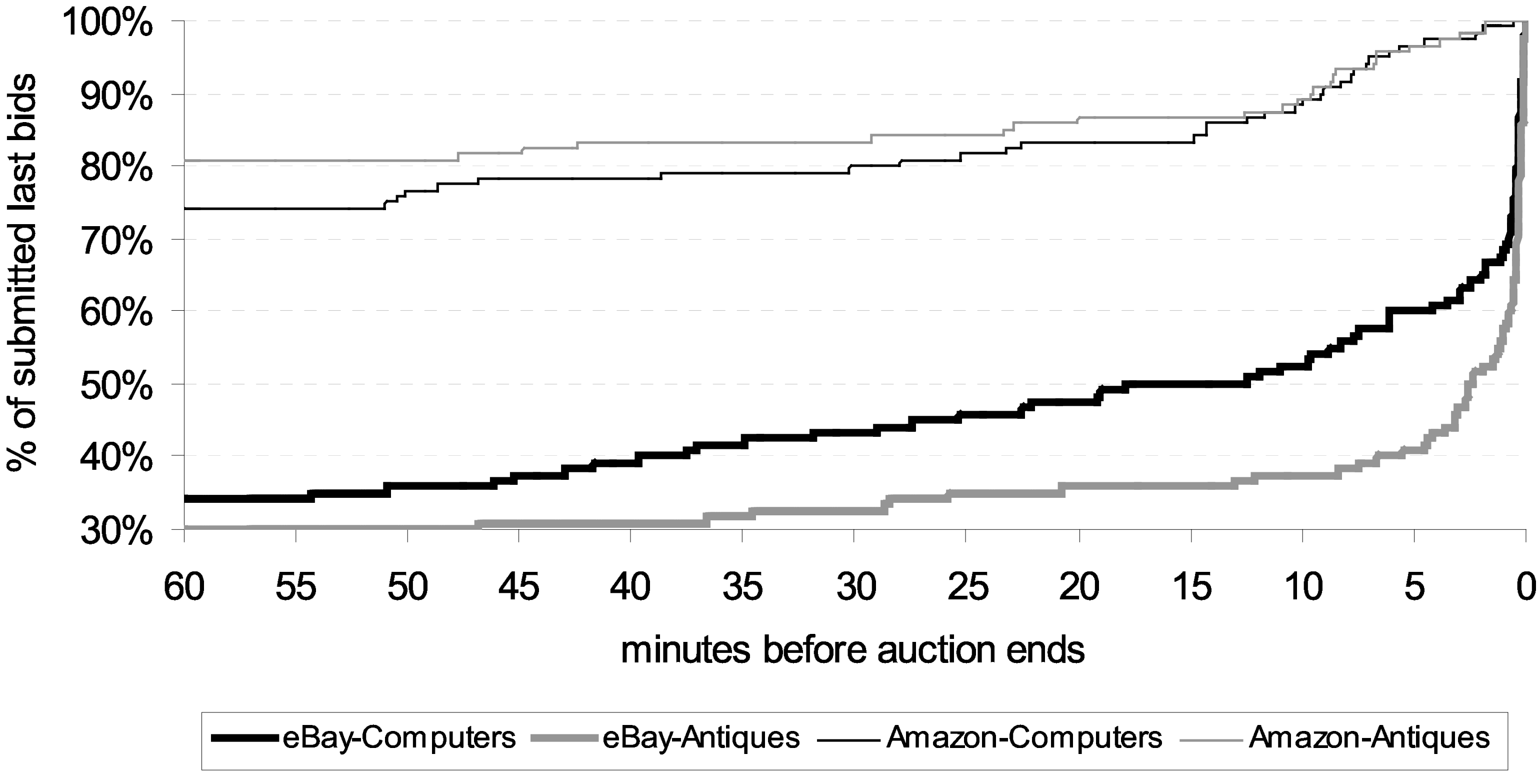

[Roth, Ockenfels, 2002]

Overview

1. single-item auction.
2. objectives: social welfare vs. seller profit.
3. applications:
    - paid search
    - retail: pricing vs. auctions,
    - intermediation: fees versus double auctions
    - competing platforms

Single-item Auction

Mechanism Design Problem: *Single-item Auction*

Given:

- one item for sale.
- n bidders (with unknown private values for item, $v_1, \ldots, v_n$)
- Bidders' objective: maximize *utility* $=$ value $-$ price paid.

Design:

- Auction to solicit bids and choose winner and payments.

Single-item Auction

Mechanism Design Problem: *Single-item Auction*

Given:

- one item for sale.
- n bidders (with unknown private values for item, $v_1, \ldots, v_n$)
- Bidders' objective: maximize *utility* $=$ value $-$ price paid.

Design:

- Auction to solicit bids and choose winner and payments.

Possible Auction Objectives:

- Maximize *social surplus*, i.e., the value of the winner.
- Maximize *seller profit*, i.e., the payment of the winner.

Objective 1: maximize social surplus

Example: The Second-price Auction

Second-price Auction

1. Solicit sealed bids.
2. Winner is highest bidder.
3. Charge winner the second-highest bid.

Example: The Second-price Auction

Second-price Auction

1. Solicit sealed bids.
2. Winner is highest bidder.
3. Charge winner the second-highest bid.

Example Input: $\mathbf{b} = (2, 6, 4, 1)$.

Example: The Second-price Auction

Second-price Auction

1. Solicit sealed bids.
2. Winner is highest bidder.
3. Charge winner the second-highest bid.

Example Input: $\mathrm{b} = (2, 6, 4, 1)$.

Questions:

- what are equilibrium strategies?
- what is equilibrium outcome?
- which has higher surplus in equilibrium?
- which has higher profit in equilibrium?

Second-price Auction Equilibrium Analysis

Second-price Auction

1. Solicit sealed bids. 2. Winner is highest bidder.
3. Charge winner the second-highest bid.

How should bidder i bid?

Second-price Auction Equilibrium Analysis

Second-price Auction

1. Solicit sealed bids. 2. Winner is highest bidder.
3. Charge winner the second-highest bid.

How should bidder i bid?

- Let $t_i = \max_{j \neq i} b_j$.
- If $b_i > t_i$, bidder i wins and pays t_i; otherwise loses.

Second-price Auction Equilibrium Analysis

Second-price Auction

1. Solicit sealed bids. 2. Winner is highest bidder.
3. Charge winner the second-highest bid.

How should bidder i bid?

- Let $t_i = \max_{j \neq i} b_j$.
- If $b_i > t_i$, bidder i wins and pays t_i; otherwise loses.

Case 1: $v_i > t_i$ **Case 2:** $v_i < t_i$

Second-price Auction Equilibrium Analysis

Second-price Auction

1. Solicit sealed bids. 2. Winner is highest bidder.
3. Charge winner the second-highest bid.

How should bidder i bid?

- Let $t_i = \max_{j \neq i} b_j$.
- If $b_i > t_i$, bidder i wins and pays t_i; otherwise loses.

Case 1: $v_i > t_i$

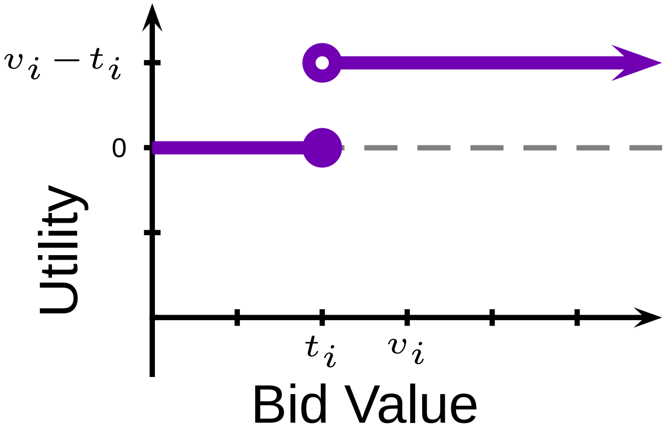

Case 2: $v_i < t_i$

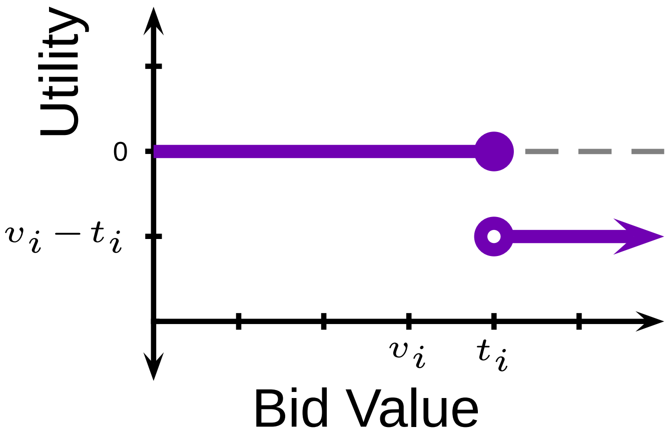

Second-price Auction Equilibrium Analysis

Second-price Auction

1. Solicit sealed bids. 2. Winner is highest bidder.
3. Charge winner the second-highest bid.

How should bidder i bid?

- Let $t_i = \max_{j \neq i} b_j$.
- If $b_i > t_i$, bidder i wins and pays t_i; otherwise loses.

Case 1: $v_i > t_i$

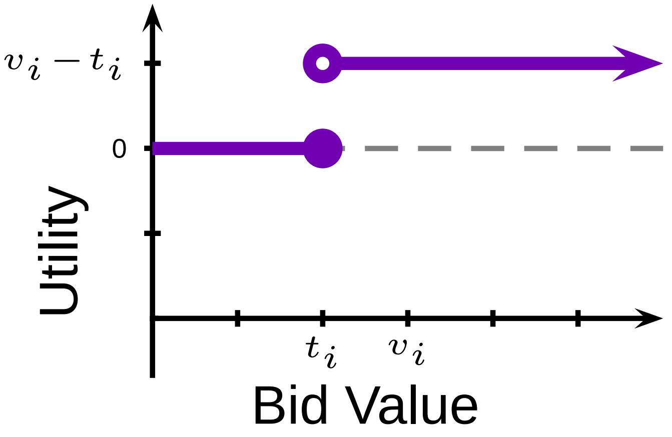

Case 2: $v_i < t_i$

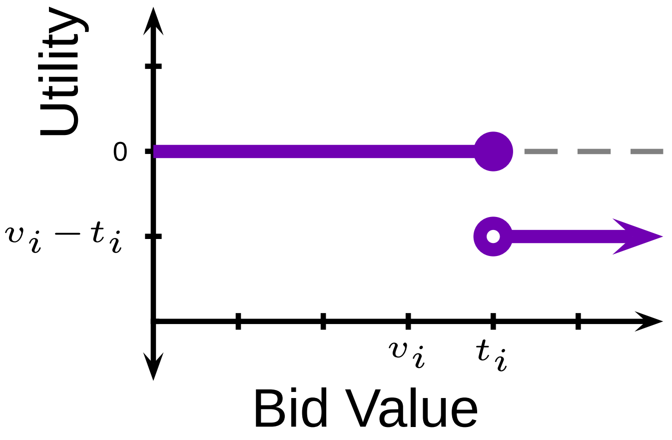

Result: Bidder i's *dominant strategy* is to bid $b_i = v_i$!

Second-price Auction Equilibrium Analysis

Second-price Auction

1. Solicit sealed bids. 2. Winner is highest bidder.
3. Charge winner the second-highest bid.

How should bidder i bid?

- Let $t_i = \max_{j \neq i} b_j$. $\Leftarrow$ *"critical value"*
- If $b_i > t_i$, bidder i wins and pays t_i; otherwise loses.

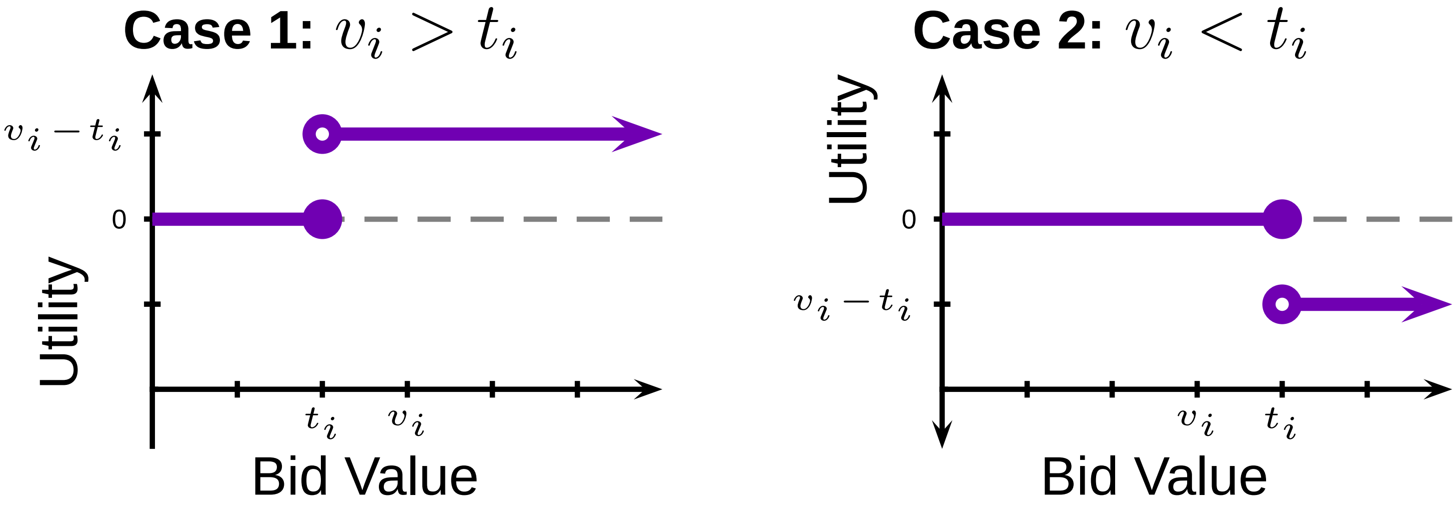

Result: Bidder i's *dominant strategy* is to bid $b_i = v_i$!

Second-price Auction Conclusion

Second-price Auction

1. Solicit sealed bids. 2. Winner is highest bidder.
3. Charge winner the second-highest bid.

Second-price Auction Conclusion

Second-price Auction

1. Solicit sealed bids. 2. Winner is highest bidder.
3. Charge winner the second-highest bid.

Lemma: [Vickrey ’61] Truthful bidding is dominant strategy in Second-price Auction.

Second-price Auction Conclusion

Second-price Auction

1. Solicit sealed bids. 2. Winner is highest bidder.
3. Charge winner the second-highest bid.

Lemma: [Vickrey '61] Truthful bidding is dominant strategy in Second-price Auction.

Corollary: Second-price Auction maximizes social surplus.

Second-price Auction Conclusion

Second-price Auction

1. Solicit sealed bids. 2. Winner is highest bidder.
3. Charge winner the second-highest bid.

Lemma: [Vickrey '61] Truthful bidding is dominant strategy in Second-price Auction.

Corollary: Second-price Auction maximizes social surplus.

- bids $=$ values (from Lemma).
- winner is highest bidder (by definition).

$\Rightarrow$ winner is bidder with highest valuation (optimal social surplus).

Second-price Auction Conclusion

Second-price Auction

1. Solicit sealed bids. 2. Winner is highest bidder.
3. Charge winner the second-highest bid.

Lemma: [Vickrey '61] Truthful bidding is dominant strategy in Second-price Auction.

Corollary: Second-price Auction maximizes social surplus.

- bids $=$ values (from Lemma).
- winner is highest bidder (by definition).

$\Rightarrow$ winner is bidder with highest valuation (optimal social surplus).

What about revenue?

Objective 2: maximize seller profit

(other objectives are similar)

Review: Uniform Distributions

Uniform Distribution: draw value v uniformly from the interval $[0, 1]$.

Review: Uniform Distributions

Uniform Distribution: draw value v uniformly from the interval $[0, 1]$.

Cumulative Distribution Function: $F(z) = \mathbf{Pr}[v \leq z] = z$.
Probability Density Function: $f(z) = \frac{1}{dz}\mathbf{Pr}[v \leq z] = 1$.

Review: Uniform Distributions

Uniform Distribution: draw value v uniformly from the interval $[0,1]$.

Cumulative Distribution Function: $F(z) = \mathbf{Pr}[v \leq z] = z$.
Probability Density Function: $f(z) = \frac{1}{dz}\mathbf{Pr}[v \leq z] = 1$.

Order Statistics: in expectation, uniform random variables evenly divide interval.

Review: Uniform Distributions

Uniform Distribution: draw value v uniformly from the interval $[0, 1]$.

Cumulative Distribution Function: $F(z) = \mathbf{Pr}[v \leq z] = z$.
Probability Density Function: $f(z) = \frac{1}{dz}\mathbf{Pr}[v \leq z] = 1$.

Order Statistics: in expectation, uniform random variables evenly divide interval.

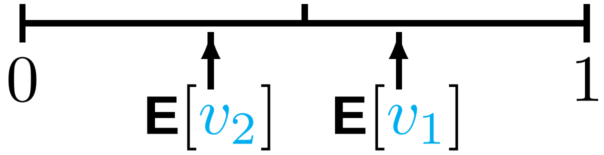

An example

Example Scenario: two bidders, uniform values

An example

Example Scenario: two bidders, uniform values

What is profit of second-price auction?

An example

Example Scenario: two bidders, uniform values

What is profit of second-price auction?

- draw values from unit interval.

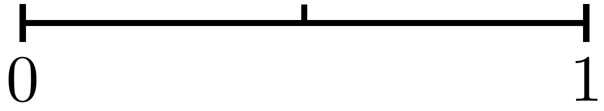

An example

Example Scenario: two bidders, uniform values

What is profit of second-price auction?

- draw values from unit interval.
- Sort values.

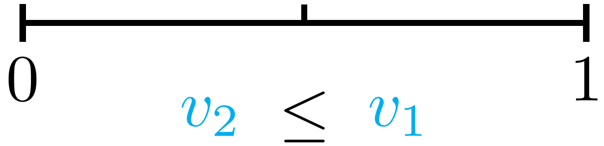

An example

Example Scenario: two bidders, uniform values

What is profit of second-price auction?

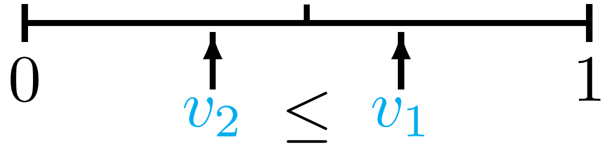

- draw values from unit interval.
- Sort values.
- In expectation, values evenly divide unit interval.

An example

Example Scenario: two bidders, uniform values

What is profit of second-price auction?

- draw values from unit interval.
- Sort values.
- In expectation, values evenly divide unit interval.
- $\mathbf{E}[\text{Profit}] = \mathbf{E}[v_2]$

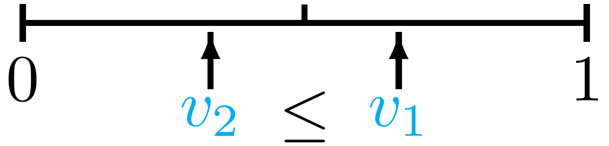

An example

Example Scenario: two bidders, uniform values

What is profit of second-price auction?

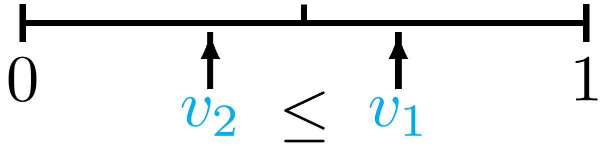

- draw values from unit interval.
- Sort values.
- In expectation, values evenly divide unit interval.
- $\mathbf{E}[\text{Profit}] = \mathbf{E}[v_2] = 1/3.$

An example

Example Scenario: two bidders, uniform values

What is profit of second-price auction?

- draw values from unit interval.
- Sort values.
- In expectation, values evenly divide unit interval.
- $\mathbf{E}[\text{Profit}] = \mathbf{E}[v_2] = 1/3.$

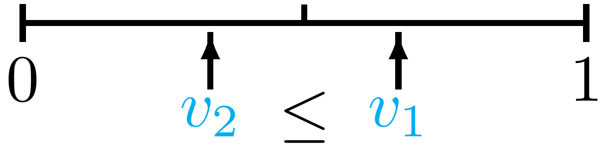

Can we get more profit?

Second-price with reserve price

Second-price Auction with reserve r

0. Insert seller-bid at r. 1. Solicit bids. 2. Winner is highest bidder. 3. Charge 2nd-highest bid.

Second-price with reserve price

Second-price Auction with reserve r

0. Insert seller-bid at r. 1. Solicit bids. 2. Winner is highest bidder. 3. Charge 2nd-highest bid.

Lemma: Second-price with reserve r has truthful DSE.

Second-price with reserve price

Second-price Auction with reserve r

0. Insert seller-bid at r. 1. Solicit bids. 2. Winner is highest bidder. 3. Charge 2nd-highest bid.

Lemma: Second-price with reserve r has truthful DSE.

What is profit of Second-price with reserve $\frac{1}{2}$ on two bidders $U[0,1]$?

Second-price with reserve price

Second-price Auction with reserve r

0. Insert seller-bid at r. 1. Solicit bids. 2. Winner is highest bidder. 3. Charge 2nd-highest bid.

Lemma: Second-price with reserve r has truthful DSE.

What is profit of Second-price with reserve $\frac{1}{2}$ on two bidders $U[0,1]$?

- draw values from unit interval.
- Sort values, $v_1 \geq v_2$

Second-price with reserve price

Second-price Auction with reserve r

0. Insert seller-bid at r. 1. Solicit bids. 2. Winner is highest bidder. 3. Charge 2nd-highest bid.

Lemma: Second-price with reserve r has truthful DSE.

What is profit of Second-price with reserve $\frac{1}{2}$ on two bidders $U[0, 1]$?

- draw values from unit interval.
- Sort values, $v_1 \geq v_2$

Case Analysis:	**Pr**$[$Case $i]$	**E**$[$Profit$]$
Case 1: $\frac{1}{2} > v_1 \geq v_2$		
Case 2: $v_1 \geq v_2 \geq \frac{1}{2}$		
Case 3: $v_1 \geq \frac{1}{2} > v_2$		

Second-price with reserve price

Second-price Auction with reserve r

0. Insert seller-bid at r. 1. Solicit bids. 2. Winner is highest bidder. 3. Charge 2nd-highest bid.

Lemma: Second-price with reserve r has truthful DSE.

What is profit of Second-price with reserve $\frac{1}{2}$ on two bidders $U[0, 1]$?

- draw values from unit interval.
- Sort values, $v_1 \geq v_2$

Case Analysis:	**Pr**$[\text{Case } i]$	**E**$[\text{Profit}]$
Case 1: $\frac{1}{2} > v_1 \geq v_2$	$1/4$	
Case 2: $v_1 \geq v_2 \geq \frac{1}{2}$	$1/4$	
Case 3: $v_1 \geq \frac{1}{2} > v_2$	$1/2$	

Second-price with reserve price

Second-price Auction with reserve r

0. Insert seller-bid at r. 1. Solicit bids. 2. Winner is highest bidder. 3. Charge 2nd-highest bid.

Lemma: Second-price with reserve r has truthful DSE.

What is profit of Second-price with reserve $\frac{1}{2}$ on two bidders $U[0,1]$?

- draw values from unit interval.
- Sort values, $v_1 \geq v_2$

Case Analysis:	**Pr**$[$Case $i]$	**E**$[$Profit$]$
Case 1: $\frac{1}{2} > v_1 \geq v_2$	$1/4$	0
Case 2: $v_1 \geq v_2 \geq \frac{1}{2}$	$1/4$	**E**$[v_2 \mid$ Case $2]$
Case 3: $v_1 \geq \frac{1}{2} > v_2$	$1/2$	$\frac{1}{2}$

Second-price with reserve price

Second-price Auction with reserve r

0. Insert seller-bid at r. 1. Solicit bids. 2. Winner is highest bidder. 3. Charge 2nd-highest bid.

Lemma: Second-price with reserve r has truthful DSE.

What is profit of Second-price with reserve $\frac{1}{2}$ on two bidders $U[0,1]$?

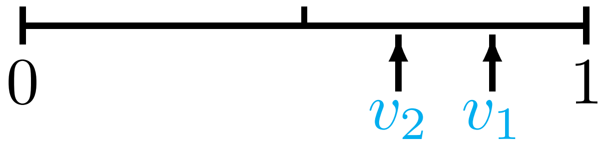

- draw values from unit interval.
- Sort values, $v_1 \geq v_2$

Case Analysis:	**Pr**[Case i]	**E**[Profit]
Case 1: $\frac{1}{2} > v_1 \geq v_2$	$1/4$	0
Case 2: $v_1 \geq v_2 \geq \frac{1}{2}$	$1/4$	**E**$[v_2 \mid \text{Case } 2] = \frac{2}{3}$
Case 3: $v_1 \geq \frac{1}{2} > v_2$	$1/2$	$\frac{1}{2}$

Second-price with reserve price

Second-price Auction with reserve r

0. Insert seller-bid at r. 1. Solicit bids. 2. Winner is highest bidder. 3. Charge 2nd-highest bid.

Lemma: Second-price with reserve r has truthful DSE.

What is profit of Second-price with reserve $\frac{1}{2}$ on two bidders $U[0,1]$?

- draw values from unit interval.

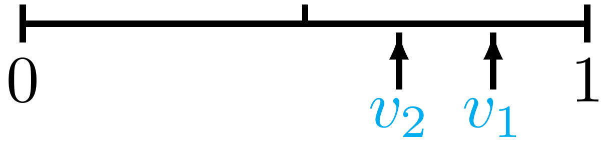

- Sort values, $v_1 \geq v_2$

Case Analysis:	**Pr**[Case i]	**E**[Profit]
Case 1: $\frac{1}{2} > v_1 \geq v_2$	$1/4$	0
Case 2: $v_1 \geq v_2 \geq \frac{1}{2}$	$1/4$	**E**$[v_2 \mid$ Case $2] = \frac{2}{3}$
Case 3: $v_1 \geq \frac{1}{2} > v_2$	$1/2$	$\frac{1}{2}$

E[profit of 2nd-price with reserve] $= \frac{1}{4} \cdot 0 + \frac{1}{4} \cdot \frac{2}{3} + \frac{1}{2} \cdot \frac{1}{2} = \frac{5}{12}$

Second-price with reserve price

Second-price Auction with reserve r

0. Insert seller-bid at r. 1. Solicit bids. 2. Winner is highest bidder. 3. Charge 2nd-highest bid.

Lemma: Second-price with reserve r has truthful DSE.

What is profit of Second-price with reserve $\frac{1}{2}$ on two bidders $U[0,1]$?

- draw values from unit interval.

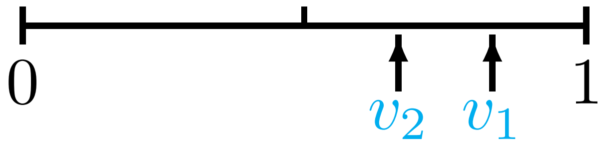

- Sort values, $v_1 \geq v_2$

Case Analysis:	**Pr**[Case i]	**E**[Profit]
Case 1: $\frac{1}{2} > v_1 \geq v_2$	$1/4$	0
Case 2: $v_1 \geq v_2 \geq \frac{1}{2}$	$1/4$	**E**$[v_2 \mid \text{Case } 2] = \frac{2}{3}$
Case 3: $v_1 \geq \frac{1}{2} > v_2$	$1/2$	$\frac{1}{2}$

$$\textbf{E}[\text{profit of 2nd-price with reserve}] = \frac{1}{4} \cdot 0 + \frac{1}{4} \cdot \frac{2}{3} + \frac{1}{2} \cdot \frac{1}{2} = \frac{5}{12}$$
$$\geq \textbf{E}[\text{profit of 2nd-price}] = \frac{1}{3}.$$

Profit Maximization

Question: What auction maximizes profit?

Profit Maximization

Question: What auction maximizes profit?

Answer: second-price with reserve (for symmetric bidders)

Microeconomics 101

Question: how should monopolist divide good across separate markets?

Microeconomics 101

Question: how should monopolist divide good across separate markets?

Demand Model: concave revenue $R(q)$ in quantity q

Microeconomics 101

Question: how should monopolist divide good across separate markets?

Demand Model: concave revenue $R(q)$ in quantity q

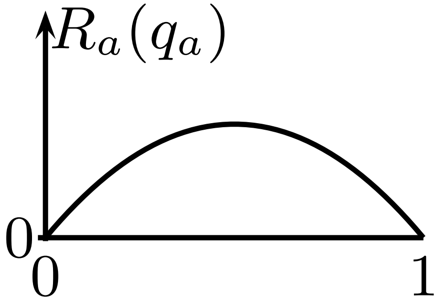

Microeconomics 101

Question: how should monopolist divide good across separate markets?

Demand Model: concave revenue $R(q)$ in quantity q

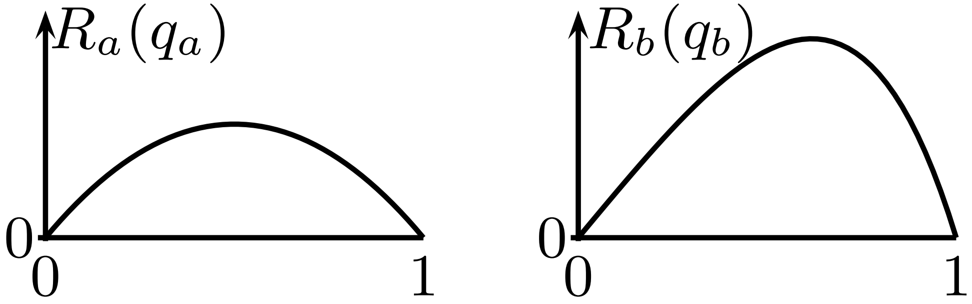

Microeconomics 101

Question: how should monopolist divide good across separate markets?

Demand Model: concave revenue $R(q)$ in quantity q

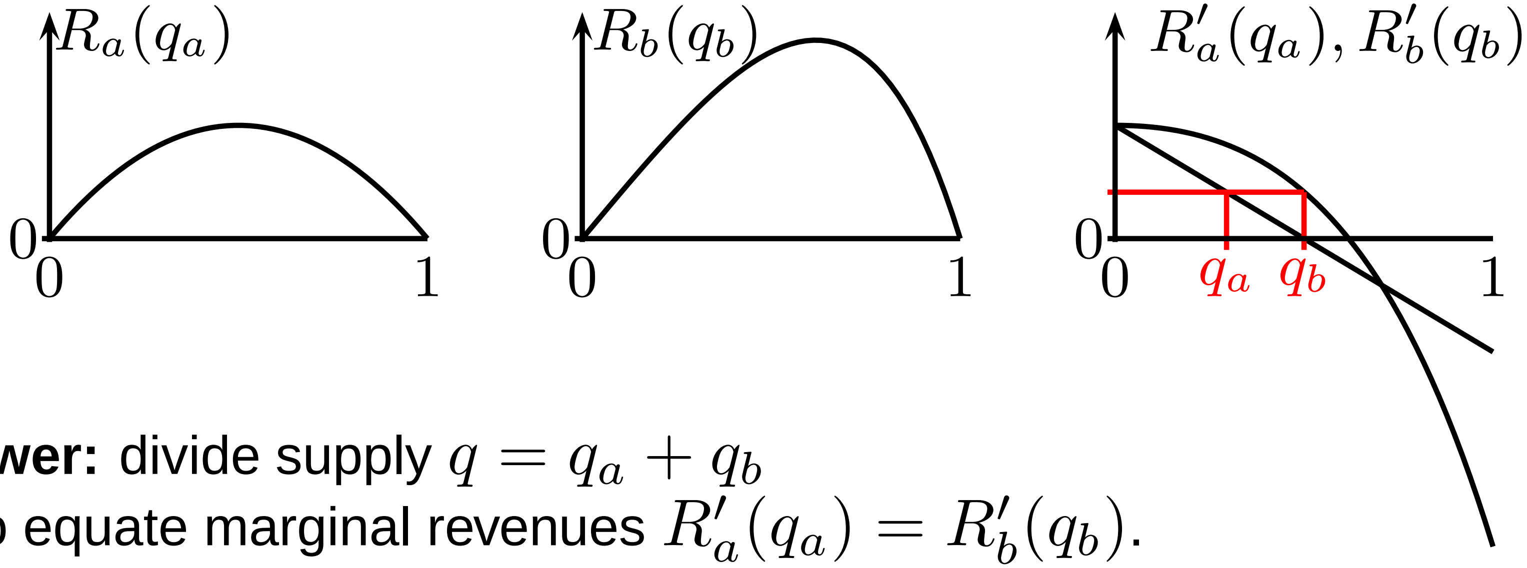

Answer: divide supply $q = q_a + q_b$
to equate marginal revenues $R'_a(q_a) = R'_b(q_b)$.

Example: two bidders, uniform values

Theorem: optimal auction maximizes “marginal revenue”.

[Myerson ’81; Bulow, Roberts ’89]

Example: two bidders, uniform values

Theorem: optimal auction maximizes "marginal revenue".

[Myerson '81; Bulow, Roberts '89]

Example: two bidders, uniform values

- sorted value: $V(q) = 1 - q$.

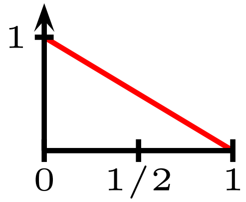

Example: two bidders, uniform values

Theorem: optimal auction maximizes "marginal revenue".

[Myerson '81; Bulow, Roberts '89]

Example: two bidders, uniform values

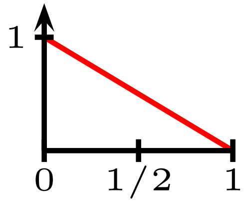

- sorted value: $V(q) = 1 - q$.
- *revenue curve*: from offering price $V(q)$.

Example: two bidders, uniform values

Theorem: optimal auction maximizes “marginal revenue”.

[Myerson ’81; Bulow, Roberts ’89]

Example: two bidders, uniform values

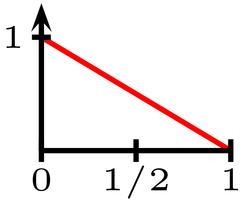

- sorted value: $V(q) = 1 - q$.
- *revenue curve*: from offering price $V(q)$.

$\Rightarrow$ buys and pays $V(q)$ with probability q.

$\Rightarrow R(q) = q \cdot V(q) = q - q^2$.

Example: two bidders, uniform values

Theorem: optimal auction maximizes "marginal revenue".

[Myerson '81; Bulow, Roberts '89]

Example: two bidders, uniform values

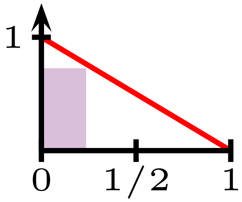

- sorted value: $V(q) = 1 - q$.

- *revenue curve*: from offering price $V(q)$.

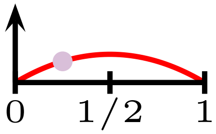

$\Rightarrow$ buys and pays $V(q)$ with probability q.

$\Rightarrow$ $R(q) = q \cdot V(q) = q - q^2$.

Example: two bidders, uniform values

Theorem: optimal auction maximizes "marginal revenue".

[Myerson '81; Bulow, Roberts '89]

Example: two bidders, uniform values

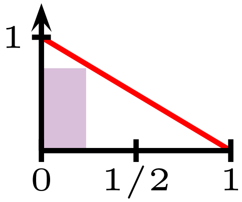

- sorted value: $V(q) = 1 - q$.
- *revenue curve*: from offering price $V(q)$.

  $\Rightarrow$ buys and pays $V(q)$ with probability q.

  $\Rightarrow$ $R(q) = q \cdot V(q) = q - q^2$.

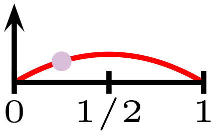

- *marginal revenue*: $R'(q) = 1 - 2q$.

Example: two bidders, uniform values

Theorem: optimal auction maximizes “marginal revenue”.

[Myerson ’81; Bulow, Roberts ’89]

Example: two bidders, uniform values

- sorted value: $V(q) = 1 - q$.

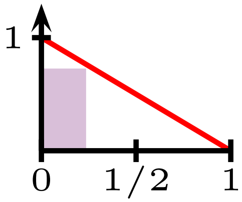

- *revenue curve*: from offering price $V(q)$.

  $\Rightarrow$ buys and pays $V(q)$ with probability q.

  $\Rightarrow$ $R(q) = q \cdot V(q) = q - q^2$.

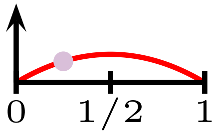

- *marginal revenue*: $R'(q) = 1 - 2q$.

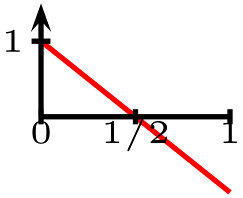

Example: two bidders, uniform values

Theorem: optimal auction maximizes “marginal revenue”.

[Myerson ’81; Bulow, Roberts ’89]

Example: two bidders, uniform values

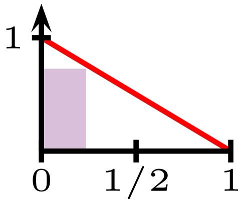

- sorted value: $V(q) = 1 - q$.

- *revenue curve*: from offering price $V(q)$.

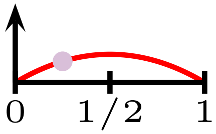

$\Rightarrow$ buys and pays $V(q)$ with probability q.

$\Rightarrow$ $R(q) = q \cdot V(q) = q - q^2$.

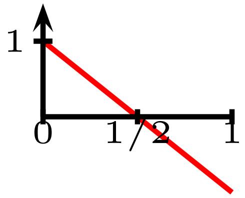

- *marginal revenue*: $R'(q) = 1 - 2q$.

- Maximize “maginal revenue”

$\Rightarrow$ sell to bidder i with highest positive $R'(q_i)$.

$\Rightarrow$ sell to bidder i with highest value at least $1/2$.

Example: two bidders, uniform values

Theorem: optimal auction maximizes “marginal revenue”.

[Myerson ’81; Bulow, Roberts ’89]

Example: two bidders, uniform values

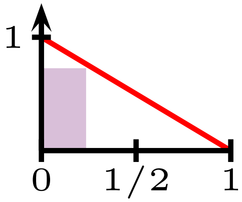

- sorted value: $V(q) = 1 - q$.
- *revenue curve*: from offering price $V(q)$.

$\Rightarrow$ buys and pays $V(q)$ with probability q.

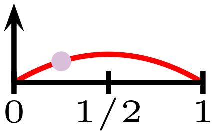

$\Rightarrow$ $R(q) = q \cdot V(q) = q - q^2$.

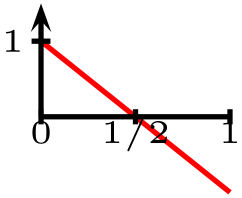

- *marginal revenue*: $R'(q) = 1 - 2q$.
- Maximize “maginal revenue”

$\Rightarrow$ sell to bidder i with highest positive $R'(q_i)$.

$\Rightarrow$ sell to bidder i with highest value at least $1/2$.

$\Rightarrow$ second-price auction with reserve $1/2$.

Example: two bidders, uniform values

Theorem: optimal auction maximizes “marginal revenue”.

[Myerson ’81; Bulow, Roberts ’89]

Example: two bidders, uniform values

- sorted value: $V(q) = 1 - q$.

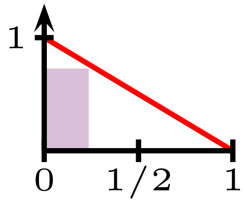

- *revenue curve*: from offering price $V(q)$.

$\Rightarrow$ buys and pays $V(q)$ with probability q.

$\Rightarrow$ $R(q) = q \cdot V(q) = q - q^2$.

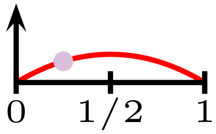

- *marginal revenue*: $R'(q) = 1 - 2q$.

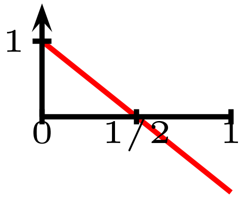

- Maximize “maginal revenue”

$\Rightarrow$ sell to bidder i with highest positive $R'(q_i)$.

$\Rightarrow$ sell to bidder i with highest value at least $1/2$.

$\Rightarrow$ second-price auction with reserve $1/2$.

Corollary: for symmetric bidders, second-price w. reserve is optimal.

[Myerson ’81]

Optimal Auction Observations

Observations:

- single auction maximizes surplus (for any distribution).
- pretending to value the good increases seller profit.
- which mechanism has better profit depends on distribution.

Optimal Auction Observations

Observations:

- single auction maximizes surplus (for any distribution).
- pretending to value the good increases seller profit.
- which mechanism has better profit depends on distribution.

Questions?

Applications

1. paid search
   (e.g., Google ads)

2. retail: auctions vs. pricing
   (e.g., eBay Auctions vs. Buy it Now)

3. intermediation: double auctions vs. fee on sale.
   (e.g., real estate, eBay, Booking.com)

4. competing platforms (e.g., Google ads vs. Bing ads)

Paid Search

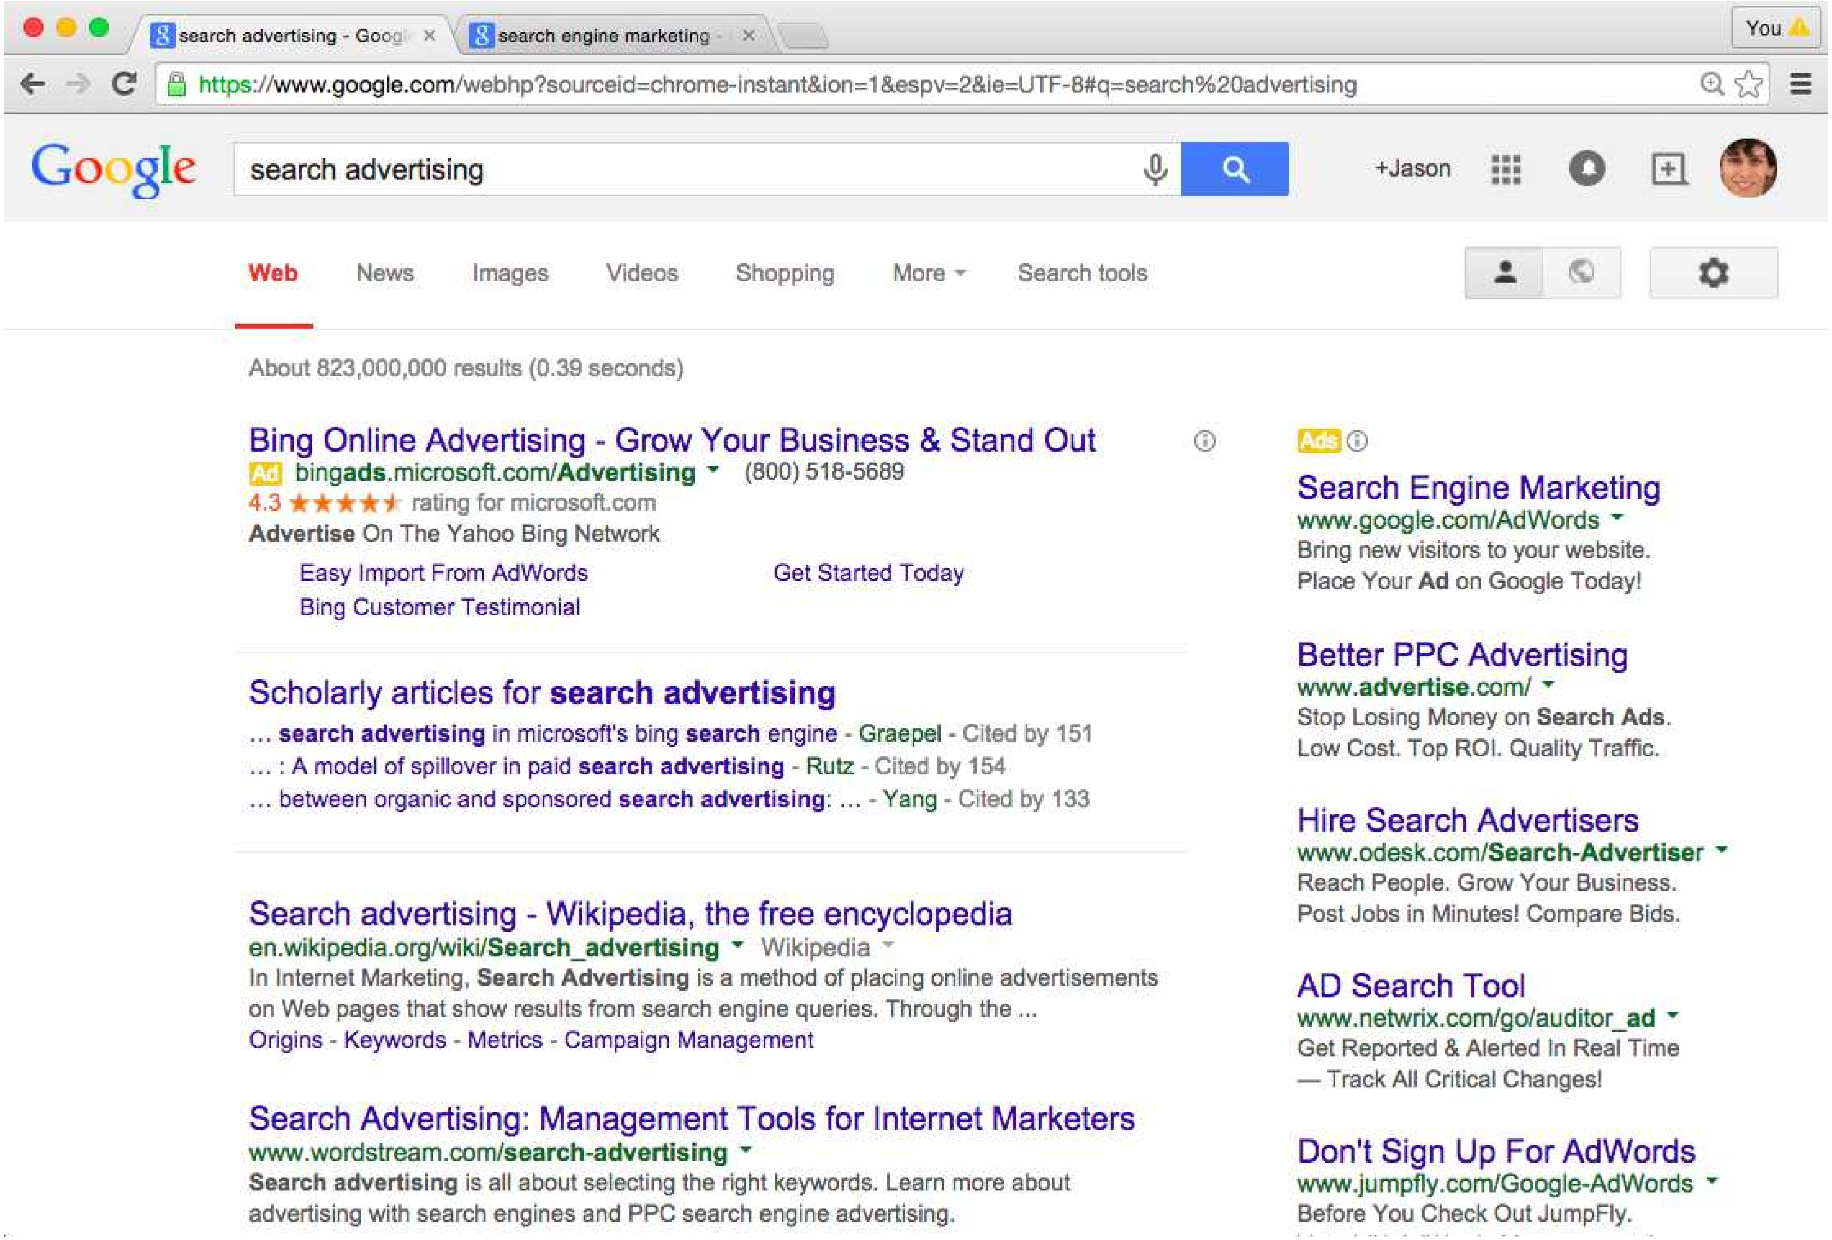

About 823,000,000 results (0.39 seconds)

Bing Online Advertising - Grow Your Business & Stand Out
Ad bingads.microsoft.com/**Advertising** (800) 518-5689
4.3 ★★★★★ rating for microsoft.com
Advertise On The Yahoo Bing Network
Easy Import From AdWords
Get Started Today
Bing Customer Testimonial

Scholarly articles for **search advertising**
… **search advertising** in microsoft's bing **search** engine - Graepel - Cited by 151
… : A model of spillover in paid **search advertising** - Rutz - Cited by 154
… between organic and sponsored **search advertising**: … - Yang - Cited by 133

Search advertising - Wikipedia, the free encyclopedia
en.wikipedia.org/wiki/**Search_advertising** Wikipedia
In Internet Marketing, **Search Advertising** is a method of placing online advertisements on Web pages that show results from search engine queries. Through the ...
Origins - Keywords - Metrics - Campaign Management

Search Advertising: Management Tools for Internet Marketers
www.wordstream.com/**search-advertising**
Search advertising is all about selecting the right keywords. Learn more about advertising with search engines and PPC search engine advertising.

Ads

Search Engine Marketing
www.google.com/AdWords
Bring new visitors to your website.
Place Your **Ad** on Google Today!

Better PPC Advertising
www.**advertise**.com/
Stop Losing Money on **Search Ads**.
Low Cost. Top ROI. Quality Traffic.

Hire Search Advertisers
www.odesk.com/**Search-Advertiser**
Reach People. Grow Your Business.
Post Jobs in Minutes! Compare Bids.

AD Search Tool
www.netwrix.com/go/auditor_**ad**
Get Reported & Alerted In Real Time
— Track All Critical Changes!

Don't Sign Up For AdWords
www.jumpfly.com/Google-AdWords
Before You Check Out JumpFly.

Generalized Second Price Auction

Generalized Second Price Auction

Generalized Second Price Auction [Google '02]

1. A user issues a query.
2. Find all ads matching query terms and exceed reserve.
3. Rank ads by bid × *click-through-rate*.
4. Charge advertiser *minimum bid to maintain position* if user clicks on ad.

Generalized Second Price Auction

Generalized Second Price Auction [Google '02]

1. A user issues a query.
2. Find all ads matching query terms and exceed reserve.
3. Rank ads by bid × *click-through-rate*.
4. Charge advertiser *minimum bid to maintain position* if user clicks on ad.

Note: GSP optimal if reserve is chosen for distribution.

Generalized Second Price Auction

Generalized Second Price Auction [Google '02]

1. A user issues a query.
2. Find all ads matching query terms and exceed reserve.
3. Rank ads by bid × *click-through-rate*.
4. Charge advertiser *minimum bid to maintain position* if user clicks on ad.

Note: GSP optimal if reserve is chosen for distribution.

Reserve Optimization: [Ostrovsky, Schwarz '11]

Generalized Second Price Auction

Generalized Second Price Auction [Google ’02]

1. A user issues a query.
2. Find all ads matching query terms and exceed reserve.
3. Rank ads by bid × *click-through-rate*.
4. Charge advertiser *minimum bid to maintain position* if user clicks on ad.

Note: GSP optimal if reserve is chosen for distribution.

Reserve Optimization: [Ostrovsky, Schwarz ’11]

- learn distribution family: log-normal.

Generalized Second Price Auction

Generalized Second Price Auction [Google ’02]

1. A user issues a query.
2. Find all ads matching query terms and exceed reserve.
3. Rank ads by bid × *click-through-rate*.
4. Charge advertiser *minimum bid to maintain position* if user clicks on ad.

Note: GSP optimal if reserve is chosen for distribution.

Reserve Optimization: [Ostrovsky, Schwarz ’11]

- learn distribution family: log-normal.
  (**Note:** bids $\neq$ values, must use *econometrics*)

Generalized Second Price Auction

Generalized Second Price Auction [Google '02]

1. A user issues a query.
2. Find all ads matching query terms and exceed reserve.
3. Rank ads by bid × *click-through-rate*.
4. Charge advertiser *minimum bid to maintain position* if user clicks on ad.

Note: GSP optimal if reserve is chosen for distribution.

Reserve Optimization: [Ostrovsky, Schwarz '11]

- learn distribution family: log-normal.
  (**Note:** bids $\neq$ values, must use *econometrics*)
- estimate distribution parameters for each keyword.

Generalized Second Price Auction

Generalized Second Price Auction [Google '02]

1. A user issues a query.
2. Find all ads matching query terms and exceed reserve.
3. Rank ads by bid $\times$ *click-through-rate*.
4. Charge advertiser *minimum bid to maintain position* if user clicks on ad.

Note: GSP optimal if reserve is chosen for distribution.

Reserve Optimization: [Ostrovsky, Schwarz '11]

- learn distribution family: log-normal.
  (**Note:** bids $\neq$ values, must use *econometrics*)
- estimate distribution parameters for each keyword.
- compute optimal price from estimated distribution.

Generalized Second Price Auction

Generalized Second Price Auction [Google '02]

1. A user issues a query.
2. Find all ads matching query terms and exceed reserve.
3. Rank ads by bid × *click-through-rate*.
4. Charge advertiser *minimum bid to maintain position* if user clicks on ad.

Note: GSP optimal if reserve is chosen for distribution.

Reserve Optimization: [Ostrovsky, Schwarz '11]

- learn distribution family: log-normal.
  (**Note:** bids $\neq$ values, must use *econometrics*)
- estimate distribution parameters for each keyword.
- compute optimal price from estimated distribution.

Conclusion: improved Yahoo!'s revenue by 5-10 percent (billions!)

Retail: Auctions vs. Pricing

Auctions	**Pricing**

Retail: Auctions vs. Pricing

Auctions	**Pricing**
Pros: ● optimal.	

Retail: Auctions vs. Pricing

Auctions

Pros:

- optimal.

Cons:

- centralized.
- complex for buyers.
- slow, delayed.

Pricing

Retail: Auctions vs. Pricing

Auctions	**Pricing**
Pros: ● optimal. Cons: ● centralized. ● complex for buyers. ● slow, delayed.	Cons: ● non-optimal.

Retail: Auctions vs. Pricing

Auctions	**Pricing**
Pros:	Cons:
● optimal.	● non-optimal.
Cons:	Pros:
● centralized.	● decentralized.
● complex for buyers.	● simple for buyers.
● slow, delayed.	● fast, immediate.

Retail: Auctions vs. Pricing

Auctions

Pros:

- optimal.

Cons:

- centralized.
- complex for buyers.
- slow, delayed.

Pricing

Cons:

- non-optimal.

Pros:

- decentralized.
- simple for buyers.
- fast, immediate.

Theorem: for pricing k units: loss at most $1/\sqrt{2\pi k}$ of optimal.
(e.g., $k = 1$: 37%; $k = 10$: 13%; $k = 100$: 4% in the worst case!)

Intermediation (w. revenue maximization)

Double Auctions	Fee on sale

Intermediation (w. revenue maximization)

Double Auctions

- buyer and seller bid.
- trade if “marg. rev. $>$ marg. loss”
- buyer pays “critical value”; seller receives “critical cost”; broker keeps difference.

Fee on sale

Intermediation (w. revenue maximization)

Double Auctions

- buyer and seller bid.
- trade if “marg. rev. $>$ marg. loss”
- buyer pays “critical value”; seller receives “critical cost”; broker keeps difference.

Fee on sale

- seller posts price
- buyer takes it or leaves it.
- seller pays fee to broker.

Intermediation (w. revenue maximization)

Double Auctions

- buyer and seller bid.
- trade if "marg. rev. $>$ marg. loss"
- buyer pays "critical value"; seller receives "critical cost"; broker keeps difference.

Pros:

- optimal.

Cons:

- centralized.
- complex for buyers.
- slow, delayed.
- budget imbalance.

Fee on sale

- seller posts price
- buyer takes it or leaves it.
- seller pays fee to broker.

Intermediation (w. revenue maximization)

Double Auctions

- buyer and seller bid.
- trade if "marg. rev. $>$ marg. loss"
- buyer pays "critical value"; seller receives "critical cost"; broker keeps difference.

Pros:

- optimal.

Cons:

- centralized.
- complex for buyers.
- slow, delayed.
- budget imbalance.

Fee on sale

- seller posts price
- buyer takes it or leaves it.
- seller pays fee to broker.

Cons:

- usually non-optimal.

Pros:

- decentralized.
- simple for buyers.
- fast, immediate.
- budget balance.

Intermediation (w. revenue maximization)

Double Auctions

- buyer and seller bid.
- trade if “marg. rev. $>$ marg. loss”
- buyer pays “critical value”; seller receives “critical cost”; broker keeps difference.

Pros:

- optimal.

Cons:

- centralized.
- complex for buyers.
- slow, delayed.
- budget imbalance.

Fee on sale

- seller posts price
- buyer takes it or leaves it.
- seller pays fee to broker.

Cons:

- usually non-optimal.

Pros:

- decentralized.
- simple for buyers.
- fast, immediate.
- budget balance.

Theorem: fee on sale is sometimes optimal; usually close to optimal.

[Loertscher, Niedermayer 2011]

Duopoly Platform Design

Recall: optimal auction is second-price with reserve (for monopolist).

Question: what about competition between two platforms?
(bidders may choose to go to competing platform)

Duopoly Platform Design

Recall: optimal auction is second-price with reserve (for monopolist).

Question: what about competition between two platforms?
(bidders may choose to go to competing platform)

Related Results:

Duopoly Platform Design

Recall: optimal auction is second-price with reserve (for monopolist).

Question: what about competition between two platforms?
(bidders may choose to go to competing platform)

Related Results: Revenue of second-price auction:

- no reserve and n bidder $>$ optimal reserve and $n-1$ bidders.
[Bulow, Klemperer '96]

Duopoly Platform Design

Recall: optimal auction is second-price with reserve (for monopolist).

Question: what about competition between two platforms?
(bidders may choose to go to competing platform)

Related Results: Revenue of second-price auction:

- no reserve and n bidder $>$ optimal reserve and $n-1$ bidders.
[Bulow, Klemperer '96]
- with entry cost, no reserve is optimal.[McAfee, McMillan '87]

Example: Google ads vs. Bing ads

Recall: how should monopolist divide good across separate markets?

Demand Model: concave revenue $R(q)$ in quantity q

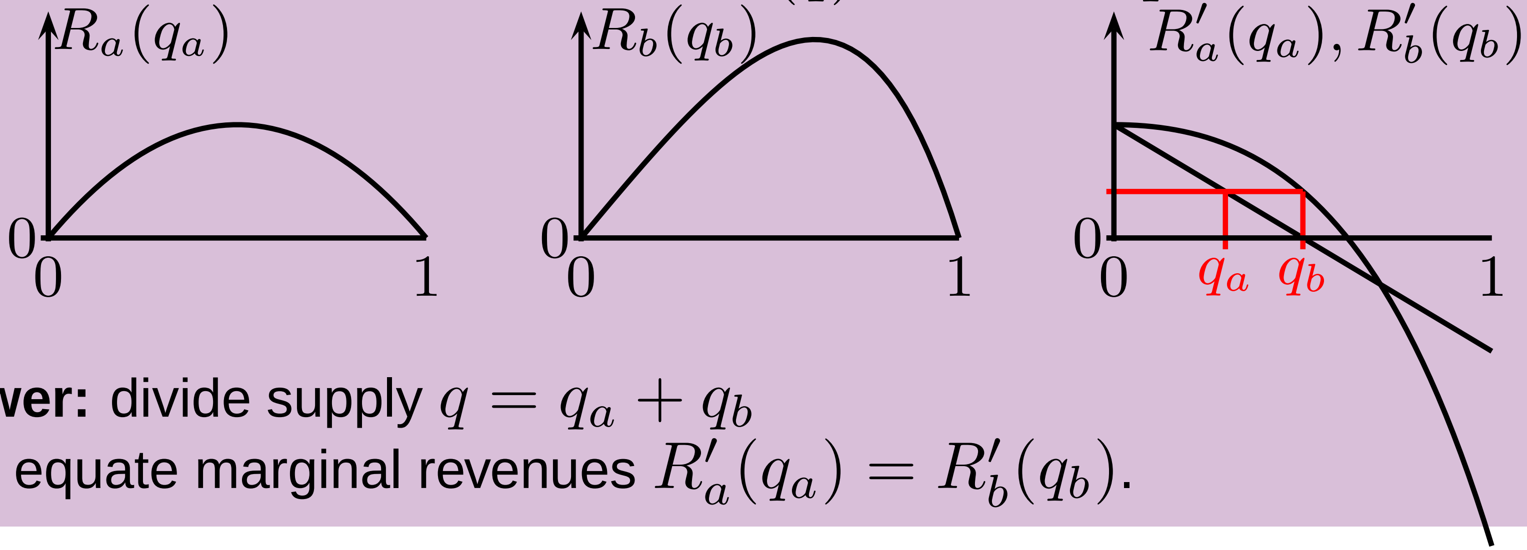

Answer: divide supply $q = q_a + q_b$

to equate marginal revenues $R'_a(q_a) = R'_b(q_b)$.

Example: Google ads vs. Bing ads

Recall: how should monopolist divide good across separate markets?

Demand Model: concave revenue $R(q)$ in quantity q

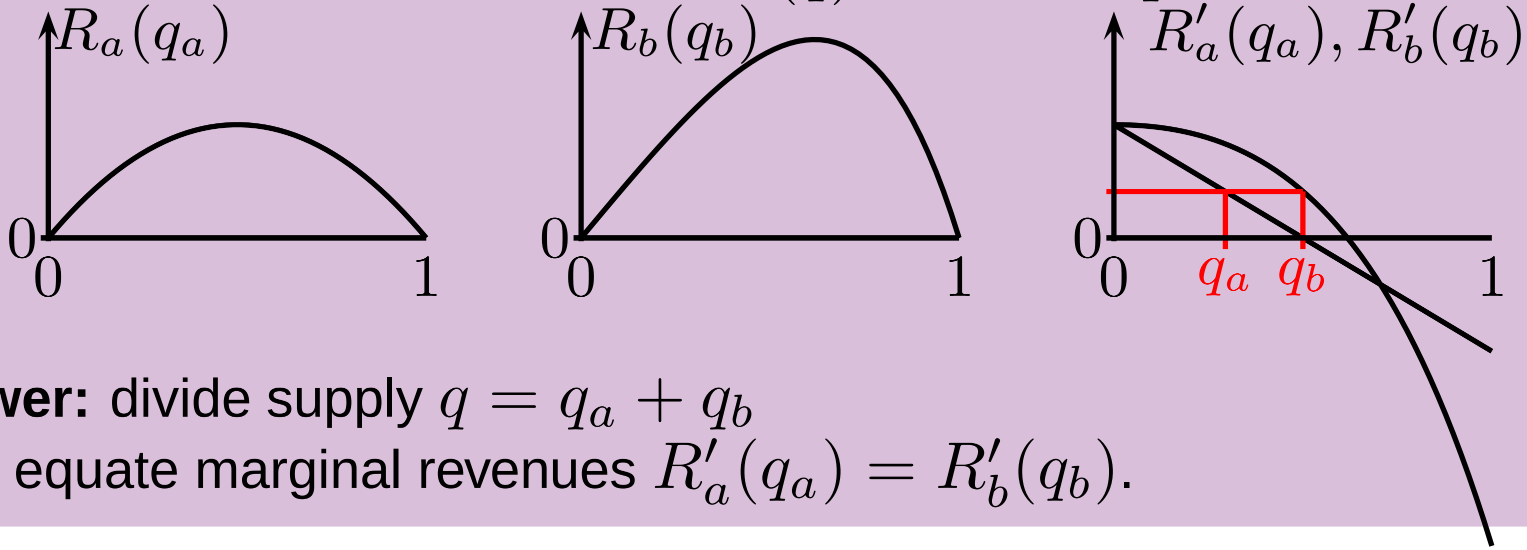

Answer: divide supply $q = q_a + q_b$

to equate marginal revenues $R'_a(q_a) = R'_b(q_b)$.

Related Question: How should advertiser divide budget across Bing and Google?

Example: Google ads vs. Bing ads

Recall: how should monopolist divide good across separate markets?

Demand Model: concave revenue $R(q)$ in quantity q

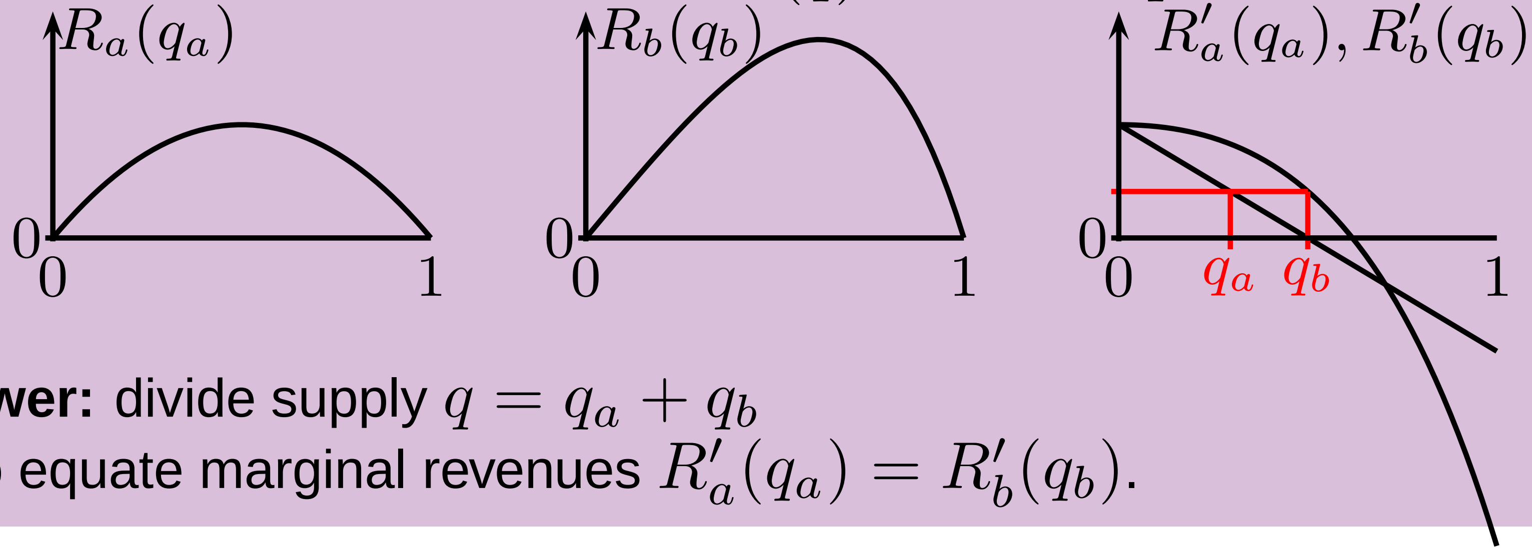

Answer: divide supply $q = q_a + q_b$

to equate marginal revenues $R'_a(q_a) = R'_b(q_b)$.

Related Question: How should advertiser divide budget across Bing and Google?

Answer: The same.

Should Bing Raise Prices?

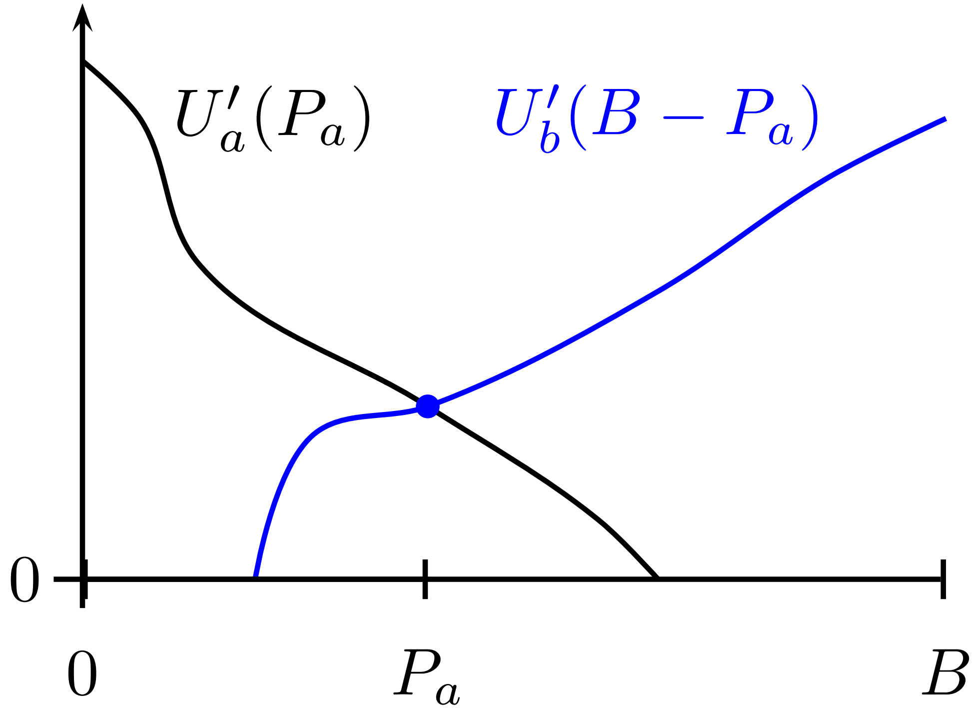

Should Bing Raise Prices?

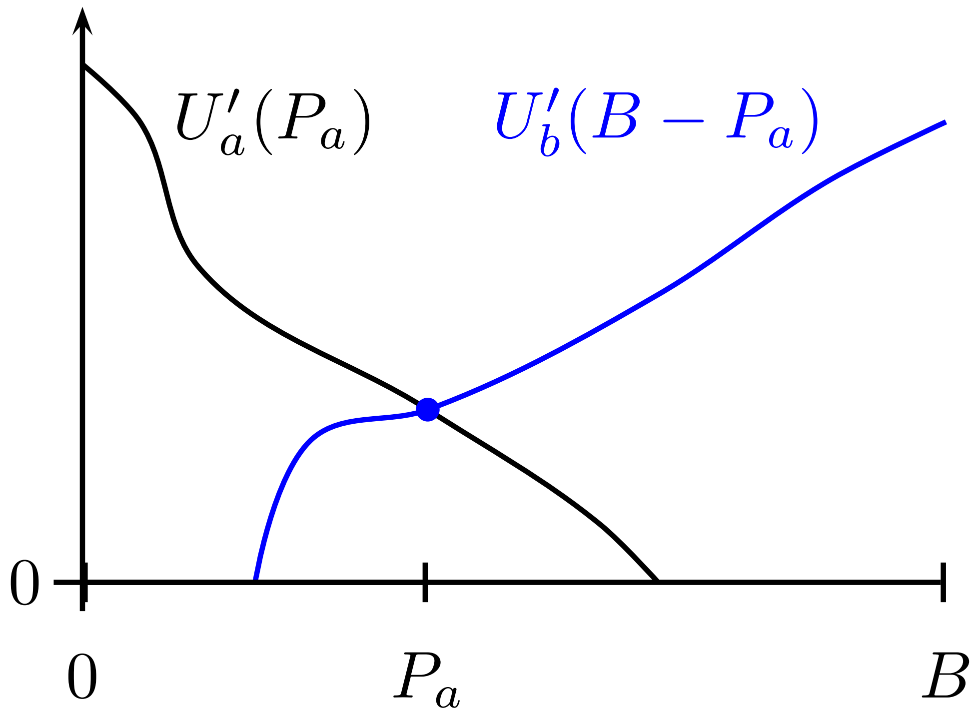

Question: what if a increases ad prices?

Should Bing Raise Prices?

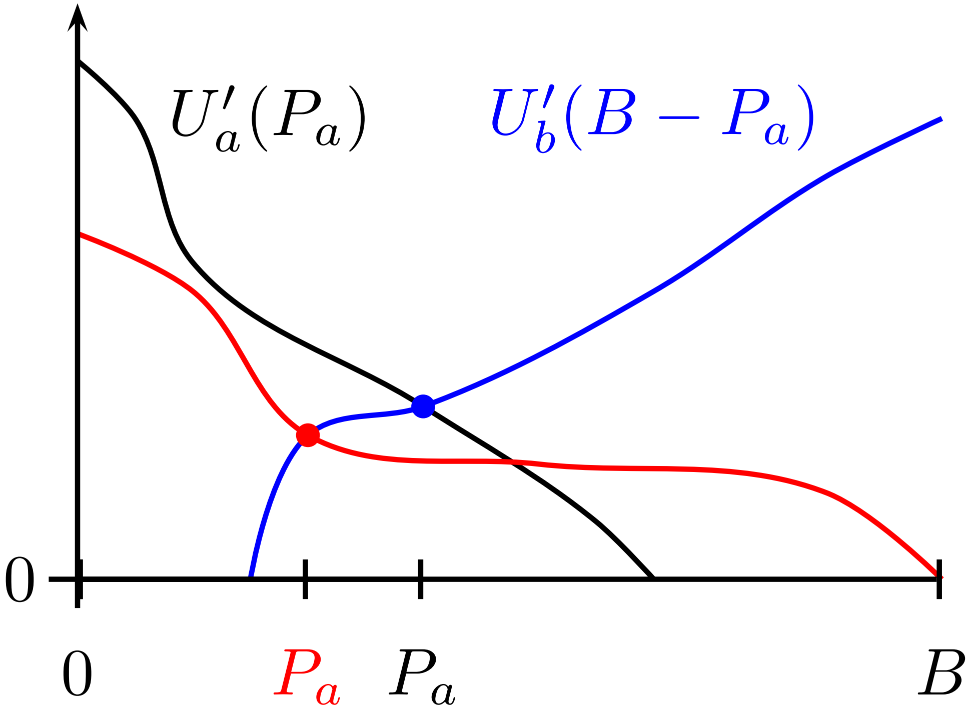

Question: what if a increases ad prices?

Should Bing Raise Prices?

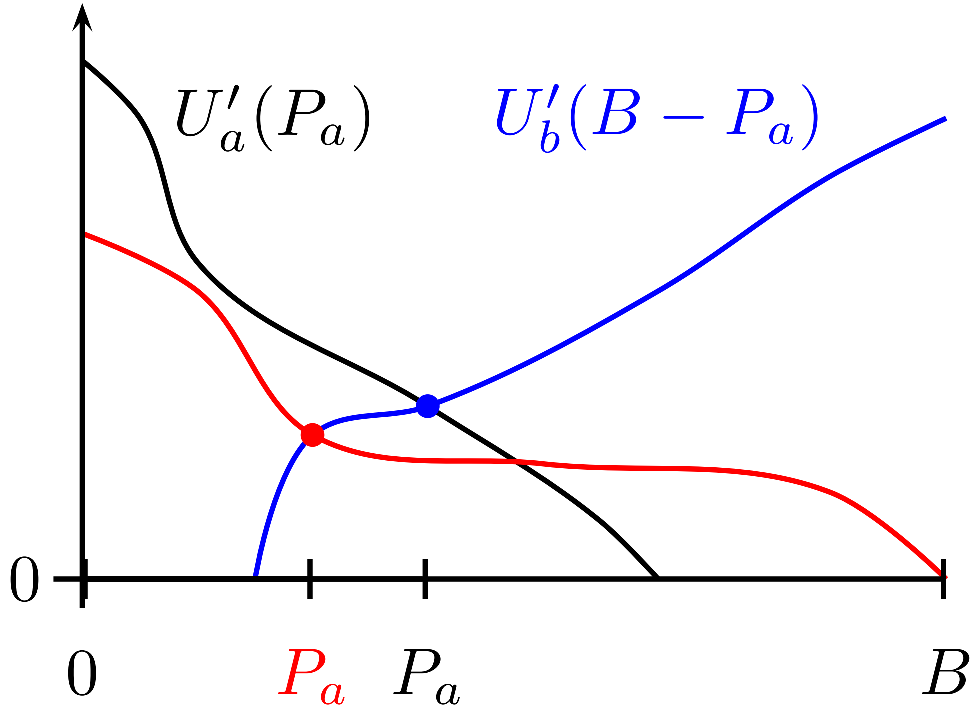

Question: what if a increases ad prices?

Answer: advertisers moves spend from a to b.

Conclusions

1. single-item auction.
2. objectives: social welfare vs. seller profit.
3. applications:
    - paid search
    - retail: pricing vs. auctions,
    - intermediation: fees versus double auctions
    - competing platforms

Conclusions

1. single-item auction.
2. objectives: social welfare vs. seller profit.
3. applications:
    - paid search
    - retail: pricing vs. auctions,
    - intermediation: fees versus double auctions
    - competing platforms

Questions?

Mechanism Design for the Classroom (Optimization of Scoring Rules)

Jason Hartline
ML from Human Preferences – November 13,

Northwestern University (visiting Stanford 2023–2024)
hartline@northwestern.edu

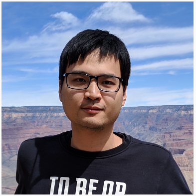

Yingkai Li

Liren Shan

Yifan Wu

Mechanism Design for the Classroom

The classroom as a "computer":

- students: local optimizers
- grader/instructor: imprecise operators
- syllabus: rules that map actions to grades
- student incentives: minimize work, maximize grade
- goal: minimize work, maximize learning, fairly assess

Mechanism Design for the Classroom

The classroom as a "computer":

- students: local optimizers
- grader/instructor: imprecise operators
- syllabus: rules that map actions to grades
- student incentives: minimize work, maximize grade
- goal: minimize work, maximize learning, fairly assess

Basic Questions: What is best syllabus?

Mechanism Design for the Classroom

The classroom as a "computer":

- students: local optimizers
- grader/instructor: imprecise operators
- syllabus: rules that map actions to grades
- student incentives: minimize work, maximize grade
- goal: minimize work, maximize learning, fairly assess

Basic Questions: What is best syllabus?

Examples:

Mechanism Design for the Classroom

The classroom as a "computer":

- students: local optimizers
- grader/instructor: imprecise operators
- syllabus: rules that map actions to grades
- student incentives: minimize work, maximize grade
- goal: minimize work, maximize learning, fairly assess

Basic Questions: What is best syllabus?

Examples:

- grading randomized exams: ex post fairness?

Mechanism Design for the Classroom

The classroom as a "computer":

- students: local optimizers
- grader/instructor: imprecise operators
- syllabus: rules that map actions to grades
- student incentives: minimize work, maximize grade
- goal: minimize work, maximize learning, fairly assess

Basic Questions: What is best syllabus?

Examples:

- grading randomized exams: ex post fairness?
- grading with partial credit: incentivizing precise answers?

Mechanism Design for the Classroom

The classroom as a “computer”:

- students: local optimizers
- grader/instructor: imprecise operators
- syllabus: rules that map actions to grades
- student incentives: minimize work, maximize grade
- goal: minimize work, maximize learning, fairly assess

Basic Questions: What is best syllabus?

Examples:

- grading randomized exams: ex post fairness?
- grading with partial credit: incentivizing precise answers?
- group projects: incentivizing teamwork?

Mechanism Design for the Classroom

The classroom as a "computer":

- students: local optimizers
- grader/instructor: imprecise operators
- syllabus: rules that map actions to grades
- student incentives: minimize work, maximize grade
- goal: minimize work, maximize learning, fairly assess

Basic Questions: What is best syllabus?

Examples:

- grading randomized exams: ex post fairness?
- grading with partial credit: incentivizing precise answers?
- group projects: incentivizing teamwork?
- peer grading: incentives for accurate peer reviews?

Mechanism Design for the Classroom

The classroom as a "computer":

- students: local optimizers
- grader/instructor: imprecise operators
- syllabus: rules that map actions to grades
- student incentives: minimize work, maximize grade
- goal: minimize work, maximize learning, fairly assess

Basic Questions: What is best syllabus?

Examples:

- grading randomized exams: ex post fairness? [Chen, Hartline, Zoeter FORC'23]
- grading with partial credit: incentivizing precise answers? [Chen, Hartline, Zoeter]
- group projects: incentivizing teamwork?
- peer grading: incentives for accurate peer reviews? [Li, Hartline, Shan, Wu EC'22]

A Peer Grading Story

1. A peer grading platform (PeerPal).
2. Grading peer reviews with proper scoring rules is horrible!
3. (Quick fix: Manually grade the peer reviews.)
4. Optimization of scoring rules.
5. Fundamental Role of Scoring Rules

PeerPal

Peer grading system:

- Canvas plugin (`https://www.peerpal.io/`)
- used in “Intro to CS”, “Intro to Algorithms”, “Intro to Online Markets”, “Mechanism Design”, etc.

PeerPal

Peer grading system:

- Canvas plugin (`https://www.peerpal.io/`)
- used in “Intro to CS”, “Intro to Algorithms”, “Intro to Online Markets”, “Mechanism Design”, etc.

Main Algorithms:

- matching peers and TAs to submissions
- grading submissions from peer reviews
- grading peer reviews from TA reviews

Advantages of Peer Grading

Advantages of Peer Grading: (observations from Intro to Algs)

- learning by reviewing.

- reduces teacher grading.

- promptness of feedback.

Advantages of Peer Grading

Advantages of Peer Grading: (observations from Intro to Algs)

- learning by reviewing.
  (learn material: 60% agree; learn to write better: 55% agree
  (worse students agree more: A: 52%; B: 54%; C: 75%; D: 80%)
- reduces teacher grading.
- promptness of feedback.

Advantages of Peer Grading

Advantages of Peer Grading: (observations from Intro to Algs)

- learning by reviewing.
  (learn material: 60% agree; learn to write better: 55% agree
  (worse students agree more: A: 52%; B: 54%; C: 75%; D: 80%)
- reduces teacher grading.
  (TAs graded 1/5 of student work.)
- promptness of feedback.

Advantages of Peer Grading

Advantages of Peer Grading: (observations from Intro to Algs)

- learning by reviewing.
  (learn material: 60% agree; learn to write better: 55% agree
  (worse students agree more: A: 52%; B: 54%; C: 75%; D: 80%)
- reduces teacher grading.
  (TAs graded 1/5 of student work.)
- promptness of feedback.
  (peer review feedback in 3 days, grades in 5 days; versus 2 weeks)

Advantages of Peer Grading

Advantages of Peer Grading: (observations from Intro to Algs)

- learning by reviewing.
  (learn material: 60% agree; learn to write better: 55% agree
  (worse students agree more: A: 52%; B: 54%; C: 75%; D: 80%)
- reduces teacher grading.
  (TAs graded 1/5 of student work.)
- promptness of feedback.
  (peer review feedback in 3 days, grades in 5 days; versus 2 weeks)

Potential Disadvantages: Inaccurate grades, student unrest, . . .

Advantages of Peer Grading

Advantages of Peer Grading: (observations from Intro to Algs)

- learning by reviewing.
  (learn material: 60% agree; learn to write better: 55% agree
  (worse students agree more: A: 52%; B: 54%; C: 75%; D: 80%)
- reduces teacher grading.
  (TAs graded 1/5 of student work.)
- promptness of feedback.
  (peer review feedback in 3 days, grades in 5 days; versus 2 weeks)

Potential Disadvantages: Inaccurate grades, student unrest, . . .

(3.7% appeal rate; 1-6% strongly disagree with survey questions)

Advantages of Peer Grading

Advantages of Peer Grading: (observations from Intro to Algs)

- learning by reviewing.
  (learn material: 60% agree; learn to write better: 55% agree
  (worse students agree more: A: 52%; B: 54%; C: 75%; D: 80%)
- reduces teacher grading.
  (TAs graded 1/5 of student work.)
- promptness of feedback.
  (peer review feedback in 3 days, grades in 5 days; versus 2 weeks)

Potential Disadvantages: Inaccurate grades, student unrest, . . .

(3.7% appeal rate; 1-6% strongly disagree with survey questions)

Main Challenge: incentivizing accurate peer reviews.

(i.e., "grading the grading")

Grading Peer Reviews

Example Scenario:

- 100 students
- submit homeworks in pairs $\Rightarrow$ 50 submissions.
- each review three submissions $\Rightarrow$ 300 peer reviews.
- need to grade: 50 submissions, 300 peer reviews.

Grading Peer Reviews

Example Scenario:

- 100 students
- submit homeworks in pairs ⇒ 50 submissions.
- each review three submissions ⇒ 300 peer reviews.
- need to grade: 50 submissions, 300 peer reviews.

Approach:

1. pick 10 submissions for TA to review.
2. assign each peer 1 of these 10 submissions at random to review.
3. assign each peer 2 of remaining 40 submissions at random to review.
4. grade Step 2 peer reviews against TA review. (don't grade other reviews)

Grading Peer Reviews

Example Scenario:

- 100 students
- submit homeworks in pairs ⇒ 50 submissions.
- each review three submissions ⇒ 300 peer reviews.
- need to grade: 50 submissions, 300 peer reviews.

Approach:

1. pick 10 submissions for TA to review.
2. assign each peer 1 of these 10 submissions at random to review.
3. assign each peer 2 of remaining 40 submissions at random to review.
4. grade Step 2 peer reviews against TA review. (don't grade other reviews)

Remaining challenge: grading peer reviews from TA review.

Grading Peer Reviews

Example Scenario:

- 100 students
- submit homeworks in pairs $\Rightarrow$ 50 submissions.
- each review three submissions $\Rightarrow$ 300 peer reviews.
- need to grade: 50 submissions, 300 peer reviews.

Approach:

1. pick 10 submissions for TA to review.
2. assign each peer 1 of these 10 submissions at random to review.
3. assign each peer 2 of remaining 40 submissions at random to review.
4. grade Step 2 peer reviews against TA review. (don't grade other reviews)

Remaining challenge: grading peer reviews from TA review.

Idea: use proper scoring rule! [McCarthy PNAS'56] [Savage JASA'71] [Gneiting, Raftery JASA'07] [. . .].

A Peer Grading Story

1. A peer grading platform (PeerPal).
2. Grading peer reviews with proper scoring rules is horrible!
3. (Quick fix: Manually grade the peer reviews.)
4. Optimization of scoring rules.
5. Fundamental Role of Scoring Rules

Proper Scoring Rules

Grading Review with Proper Scoring Rule

Proper Scoring Rules

Grading Review with Proper Scoring Rule

- TA score $\theta \in [0, 1]$ (truth)

Proper Scoring Rules

Grading Review with Proper Scoring Rule

- TA score $\theta \in [0, 1]$ (truth)
- Peer belief $\mu \in [0, 1]$; peer report $r \in [0, 1]$

Proper Scoring Rules

Grading Review with Proper Scoring Rule

- TA score $\theta \in [0,1]$ (truth)
- Peer belief $\mu \in [0,1]$; peer report $r \in [0,1]$
- quadratic scoring rule: $S(r,\theta) = 1 - (\theta - r)^2$

Proper Scoring Rules

Grading Review with Proper Scoring Rule

- TA score $\theta \in [0, 1]$ (truth)
- Peer belief $\mu \in [0, 1]$; peer report $r \in [0, 1]$
- quadratic scoring rule: $S(r, \theta) = 1 - (\theta - r)^2$
- (if multiple rubric elements, average across rubric)

Proper Scoring Rules

Grading Review with Proper Scoring Rule

- TA score $\theta \in [0, 1]$ (truth)
- Peer belief $\mu \in [0, 1]$; peer report $r \in [0, 1]$
- quadratic scoring rule: $S(r, \theta) = 1 - (\theta - r)^2$
- (if multiple rubric elements, average across rubric)

Example

Proper Scoring Rules

Grading Review with Proper Scoring Rule

- TA score $\theta \in [0,1]$ (truth)
- Peer belief $\mu \in [0,1]$; peer report $r \in [0,1]$
- quadratic scoring rule: $S(r,\theta) = 1 - (\theta - r)^2$
- (if multiple rubric elements, average across rubric)

Example

- peer report $r = 0.8$

Proper Scoring Rules

Grading Review with Proper Scoring Rule

- TA score $\theta \in [0, 1]$ (truth)
- Peer belief $\mu \in [0, 1]$; peer report $r \in [0, 1]$
- quadratic scoring rule: $S(r, \theta) = 1 - (\theta - r)^2$
- (if multiple rubric elements, average across rubric)

Example

- peer report $r = 0.8$
- TA report $\theta = 0.3$

Proper Scoring Rules

Grading Review with Proper Scoring Rule

- TA score $\theta \in [0, 1]$ (truth)
- Peer belief $\mu \in [0, 1]$; peer report $r \in [0, 1]$
- quadratic scoring rule: $S(r, \theta) = 1 - (\theta - r)^2$
- (if multiple rubric elements, average across rubric)

Example

- peer report $r = 0.8$
- TA report $\theta = 0.3$
- score $S(r, \theta) = 1 - 0.5^2 = 0.75$

Proper Scoring Rules

Grading Review with Proper Scoring Rule

- TA score $\theta \in [0, 1]$ (truth)
- Peer belief $\mu \in [0, 1]$; peer report $r \in [0, 1]$
- quadratic scoring rule: $S(r, \theta) = 1 - (\theta - r)^2$
- (if multiple rubric elements, average across rubric)

Theorem
Reporting $r = \mu$ is optimal for peer.

Example

- peer report $r = 0.8$
- TA report $\theta = 0.3$
- score $S(r, \theta) = 1 - 0.5^2 = 0.75$

Proper Scoring Rules

Grading Review with Proper Scoring Rule

- TA score $\theta \in [0, 1]$ (truth)
- Peer belief $\mu \in [0, 1]$; peer report $r \in [0, 1]$
- quadratic scoring rule: $S(r, \theta) = 1 - (\theta - r)^2$
- (if multiple rubric elements, average across rubric)

Theorem
Reporting $r = \mu$ is optimal for peer.

Proof.

Example

- peer report $r = 0.8$
- TA report $\theta = 0.3$
- score $S(r, \theta) = 1 - 0.5^2 = 0.75$

Proper Scoring Rules

Grading Review with Proper Scoring Rule

- TA score $\theta \in [0, 1]$ (truth)
- Peer belief $\mu \in [0, 1]$; peer report $r \in [0, 1]$
- quadratic scoring rule: $S(r, \theta) = 1 - (\theta - r)^2$
- (if multiple rubric elements, average across rubric)

Theorem
Reporting $r = \mu$ is optimal for peer.

Proof.

- let $u(r) = 1 - r + r^2$

Example

- peer report $r = 0.8$
- TA report $\theta = 0.3$
- score $S(r, \theta) = 1 - 0.5^2 = 0.75$

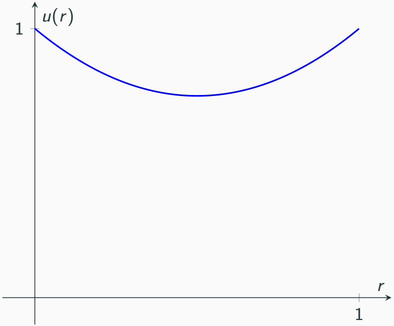

Proper Scoring Rules

Grading Review with Proper Scoring Rule

- TA score $\theta \in [0, 1]$ (truth)
- Peer belief $\mu \in [0, 1]$; peer report $r \in [0, 1]$
- quadratic scoring rule: $S(r, \theta) = 1 - (\theta - r)^2$
- (if multiple rubric elements, average across rubric)

Theorem
Reporting $r = \mu$ is optimal for peer.

Proof.

- let $u(r) = 1 - r + r^2$
- algebra ⇒ can rewrite as:

$$S(r, \theta) = u(r) + u'(r)\,(\theta - r) + \kappa(\theta).$$

Example

- peer report $r = 0.8$
- TA report $\theta = 0.3$
- score $S(r, \theta) = 1 - 0.5^2 = 0.75$

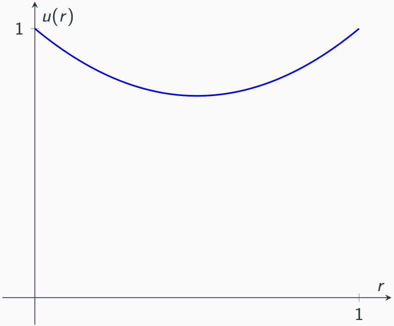

Proper Scoring Rules

Grading Review with Proper Scoring Rule

- TA score $\theta \in [0,1]$ (truth)
- Peer belief $\mu \in [0,1]$; peer report $r \in [0,1]$
- quadratic scoring rule: $S(r,\theta) = 1 - (\theta - r)^2$
- (if multiple rubric elements, average across rubric)

Theorem
Reporting $r = \mu$ is optimal for peer.

Proof.

- let $u(r) = 1 - r + r^2$
- algebra $\Rightarrow$ can rewrite as:
$$S(r,\theta) = u(r) + u'(r)\,(\theta - r) + \kappa(\theta).$$
- report cannot affect κ (so ignore it)

Example

- peer report $r = 0.8$
- TA report $\theta = 0.3$
- score $S(r,\theta) = 1 - 0.5^2 = 0.75$

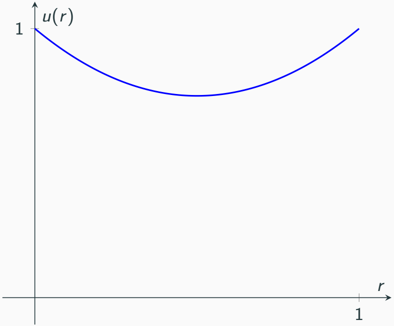

Proper Scoring Rules

Grading Review with Proper Scoring Rule

- TA score $\theta \in [0,1]$ (truth)
- Peer belief $\mu \in [0,1]$; peer report $r \in [0,1]$
- quadratic scoring rule: $S(r,\theta) = 1 - (\theta - r)^2$
- (if multiple rubric elements, average across rubric)

Theorem
Reporting $r = \mu$ is optimal for peer.

Proof.

- let $u(r) = 1 - r + r^2$
- algebra $\Rightarrow$ can rewrite as:

$$S(r,\theta) = u(r) + u'(r)\,(\theta - r) + \kappa(\theta).$$

- report cannot affect κ (so ignore it)
- let supporting tangent at r be:

$$h_r(\theta) = u(r) + u'(r)(\theta - r)$$

Example

- peer report $r = 0.8$
- TA report $\theta = 0.3$
- score $S(r,\theta) = 1 - 0.5^2 = 0.75$

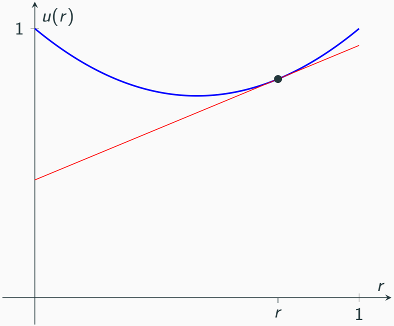

Proper Scoring Rules

Grading Review with Proper Scoring Rule

- TA score $\theta \in [0,1]$ (truth)
- Peer belief $\mu \in [0,1]$; peer report $r \in [0,1]$
- quadratic scoring rule: $S(r,\theta) = 1 - (\theta - r)^2$
- (if multiple rubric elements, average across rubric)

Theorem
Reporting $r = \mu$ is optimal for peer.

Proof.

- let $u(r) = 1 - r + r^2$
- algebra $\Rightarrow$ can rewrite as:
$$S(r,\theta) = u(r) + u'(r)\,(\theta - r) + \kappa(\theta).$$
- report cannot affect κ (so ignore it)
- let supporting tangent at r be:
$$h_r(\theta) = u(r) + u'(r)(\theta - r)$$
- loss from report r at belief μ: $u(\mu) - h_r(\mu)$. $\square$

Example

- peer report $r = 0.8$
- TA report $\theta = 0.3$
- score $S(r,\theta) = 1 - 0.5^2 = 0.75$

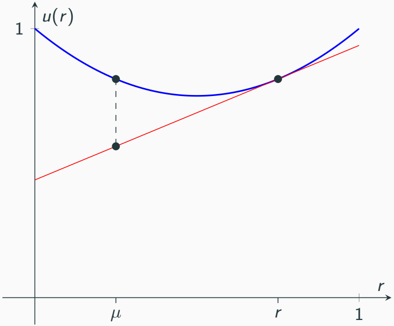

Standard Scoring Rules are Horrible

The Lazy Peer Strategy

Standard Scoring Rules are Horrible

The Lazy Peer Strategy

- assume: TA grade $\theta \in [0.6, 1]$

Standard Scoring Rules are Horrible

The Lazy Peer Strategy

- assume: TA grade $\theta \in [0.6, 1]$
- strategy: always report $r = 0.8$

Standard Scoring Rules are Horrible

The Lazy Peer Strategy

- assume: TA grade $\theta \in [0.6, 1]$
- strategy: always report $r = 0.8$
- $S(r, \theta) \geq 1 - (0.2)^2 = 0.96$

Standard Scoring Rules are Horrible

The Lazy Peer Strategy

- assume: TA grade $\theta \in [0.6, 1]$
- strategy: always report $r = 0.8$
- $S(r, \theta) \geq 1 - (0.2)^2 = 0.96$

Result
Very little incentive for effort!

A Peer Grading Story

1. A peer grading platform (PeerPal).
2. Grading peer reviews with proper scoring rules is horrible!
3. (Quick fix: Manually grade the peer reviews.)
4. Optimization of scoring rules.
5. Fundamental Role of Scoring Rules

Review Grading By Hand

Submission 42

⋮

contents of submission

⋮

	Peer 1	Peer 2	Peer 3	TA Score	TA Comment
Algorithm	8*	9*	10	9	good solution . . .
Correctness	5*	7*	10	6	missing base case . . .
Clarity	8*	8*	10	8	easy to follow . . .
Quantitative	9	10	5		
Qualitative	8	8	0		
Feedback	see TA review	see TA review	must provide detailed review		

A Peer Grading Story

1. A peer grading platform (PeerPal).
2. Grading peer reviews with proper scoring rules is horrible!
3. (Quick fix: Manually grade the peer reviews.)
4. Optimization of scoring rules.
5. Fundamental Role of Scoring Rules

Summary: Optimization of Scoring Rules

Optimal Scoring Rule for Incentivizing Binary Effort

- peers choose effort or no effort
- maximize: difference in score for effort vs no effort
- subject to: proper and bounded scoring rule.

Summary: Optimization of Scoring Rules

Optimal Scoring Rule for Incentivizing Binary Effort

$$\begin{aligned} \max_{\text{scoring rule}} \quad & \mathbf{E}_{\text{state, belief with effort}}[\text{score with effort} - \text{score without effort}] \\ \text{s.t.} \quad & \text{scoring rule is proper (optimal to truthfully report belief)} \\ & \text{scoring rule is bounded} \end{aligned}$$

Summary: Optimization of Scoring Rules

Optimal Scoring Rule for Incentivizing Binary Effort

$$\max_{\text{scoring rule}} \mathbf{E}_{\text{state, belief with effort}}[\text{score with effort} - \text{score without effort}]$$

$$\text{s.t. scoring rule is proper (optimal to truthfully report belief)}$$

$$\text{scoring rule is bounded}$$

Theorem
optimal single-dimensional scoring rule:

choose side of prior mean, score linear in state

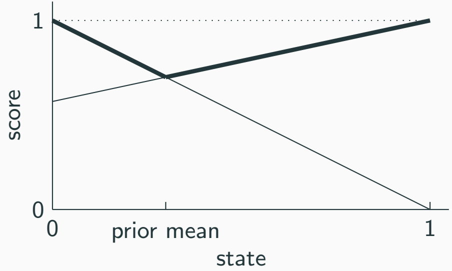

Summary: Optimization of Scoring Rules

Optimal Scoring Rule for Incentivizing Binary Effort

$$\max_{\text{scoring rule}} \mathbf{E}_{\text{state, belief with effort}}[\text{score with effort} - \text{score without effort}]$$

$$\text{s.t. scoring rule is proper (optimal to truthfully report belief)}$$

$$\text{scoring rule is bounded}$$

Theorem
optimal single-dimensional scoring rule:

choose side of prior mean, score linear in state

(standard scoring rules like quadratic not approx optimal)

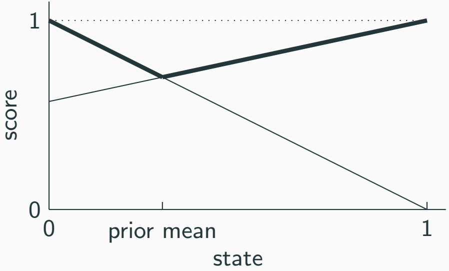

Summary: Optimization of Scoring Rules

Optimal Scoring Rule for Incentivizing Binary Effort

$$\max_{\text{scoring rule}} \mathbf{E}_{\text{state, belief with effort}}[\text{score with effort} - \text{score without effort}]$$

$$\text{s.t. scoring rule is proper (optimal to truthfully report belief)}$$

$$\text{scoring rule is bounded}$$

Theorem
optimal single-dimensional scoring rule:

choose side of prior mean, score linear in state

(standard scoring rules like quadratic not approx optimal)

Theorem
approximately optimal multi-dimensional scoring rule:

maximum over optimal separate scoring rules

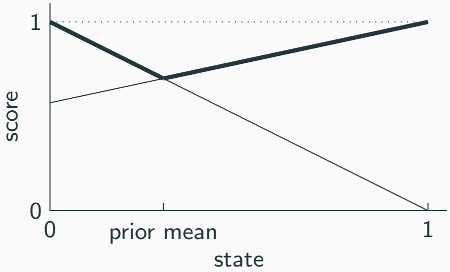

Summary: Optimization of Scoring Rules

Optimal Scoring Rule for Incentivizing Binary Effort

$$\max_{\text{scoring rule}} \mathbf{E}_{\text{state, belief with effort}}[\text{score with effort} - \text{score without effort}]$$

$$\text{s.t. scoring rule is proper (optimal to truthfully report belief)}$$

$$\text{scoring rule is bounded}$$

Theorem
optimal single-dimensional scoring rule:

choose side of prior mean, score linear in state

(standard scoring rules like quadratic not approx optimal)

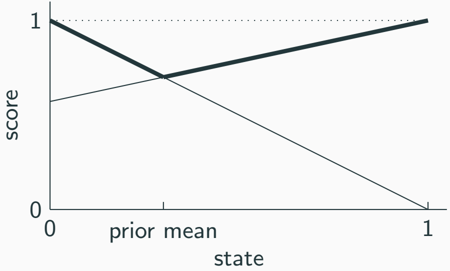

Theorem
approximately optimal multi-dimensional scoring rule:

maximum over optimal separate scoring rules

(average of separate scoring rules not approx optimal)

Single-dimensional Optimal Scoring Rules

Theorem
optimal single-dimensional scoring rule:

choose side of prior mean, score linear in state

Single-dimensional Optimal Scoring Rules

Theorem
optimal single-dimensional scoring rule:

choose side of prior mean, score linear in state

Proof.

Single-dimensional Optimal Scoring Rules

Theorem
optimal single-dimensional scoring rule:

choose side of prior mean, score linear in state

Proof.

- consider ex post bounded scoring rule defined by convex u

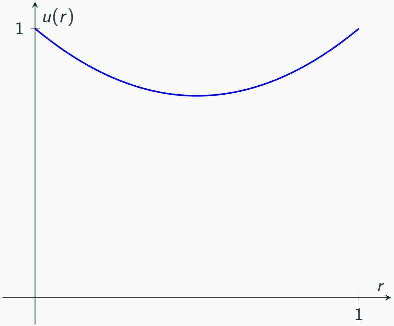

Single-dimensional Optimal Scoring Rules

Theorem
optimal single-dimensional scoring rule:

choose side of prior mean, score linear in state

Proof.

- consider ex post bounded scoring rule defined by convex u
- replace $u(r)$ with V-shape at μ_{prior}

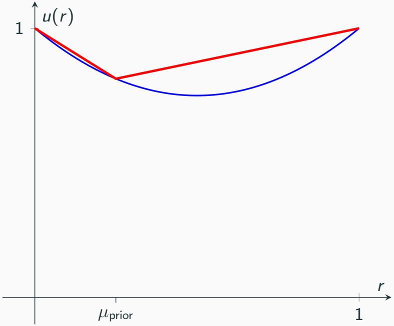

Single-dimensional Optimal Scoring Rules

Theorem
optimal single-dimensional scoring rule:

choose side of prior mean, score linear in state

Proof.

- consider ex post bounded scoring rule defined by convex u
- replace $u(r)$ with V-shape at μ_{prior}
- objective $\mathbf{E}[u(\mu_{\text{posterior}}) - u(\mu_{\text{prior}})]$ weakly increased:

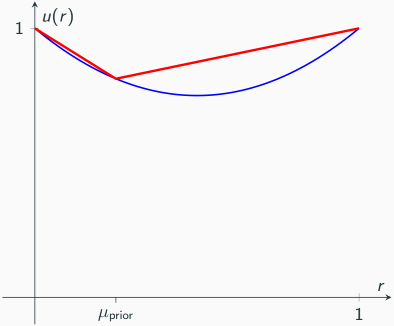

Single-dimensional Optimal Scoring Rules

Theorem
optimal single-dimensional scoring rule:

choose side of prior mean, score linear in state

Proof.

- consider ex post bounded scoring rule defined by convex u
- replace $u(r)$ with V-shape at μ_{prior}
- objective $\mathbf{E}[u(\mu_{\text{posterior}}) - u(\mu_{\text{prior}})]$ weakly increased:
  - $u(\mu_{\text{posterior}})$ weakly increased.

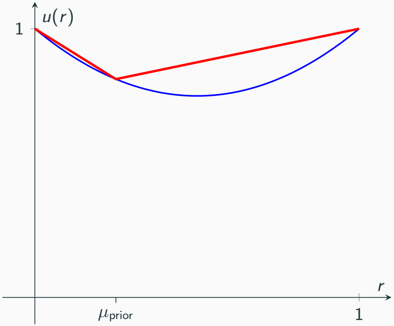

Single-dimensional Optimal Scoring Rules

Theorem
optimal single-dimensional scoring rule:

choose side of prior mean, score linear in state

Proof.

- consider ex post bounded scoring rule defined by convex u
- replace $u(r)$ with V-shape at μ_{prior}
- objective $\mathbf{E}[u(\mu_{\text{posterior}}) - u(\mu_{\text{prior}})]$ weakly increased:
  - $u(\mu_{\text{posterior}})$ weakly increased.
  - $u(\mu_{\text{prior}})$ is unchanged.

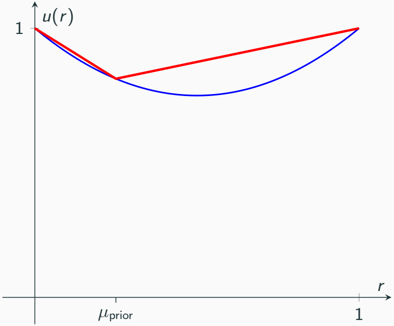

Single-dimensional Optimal Scoring Rules

Theorem
optimal single-dimensional scoring rule:

choose side of prior mean, score linear in state

Proof.

- consider ex post bounded scoring rule defined by convex u
- replace $u(r)$ with V-shape at μ_{prior}
- objective $\mathbf{E}[u(\mu_{\text{posterior}}) - u(\mu_{\text{prior}})]$ weakly increased:
  - $u(\mu_{\text{posterior}})$ weakly increased.
  - $u(\mu_{\text{prior}})$ is unchanged.
- score for extremal reports weakly less extreme

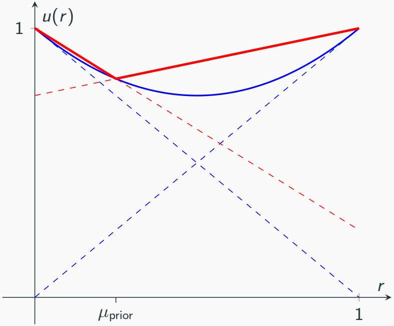

Single-dimensional Optimal Scoring Rules

Theorem
optimal single-dimensional scoring rule:

choose side of prior mean, score linear in state

Proof.

- consider ex post bounded scoring rule defined by convex u
- replace $u(r)$ with V-shape at μ_{prior}
- objective $\mathbf{E}[u(\mu_{\text{posterior}}) - u(\mu_{\text{prior}})]$ weakly increased:
  - $u(\mu_{\text{posterior}})$ weakly increased.
  - $u(\mu_{\text{prior}})$ is unchanged.
- score for extremal reports weakly less extreme
  - still ex post bounded.

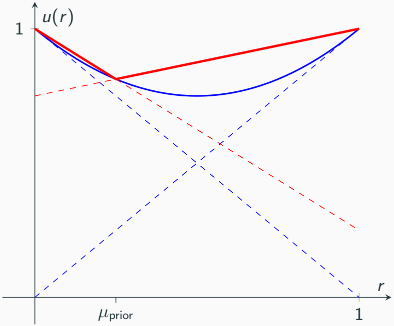

Single-dimensional Optimal Scoring Rules

Theorem
optimal single-dimensional scoring rule:

choose side of prior mean, score linear in state

Proof.

- consider ex post bounded scoring rule defined by convex u
- replace $u(r)$ with V-shape at μ_{prior}
- objective $\mathbf{E}[u(\mu_{\text{posterior}}) - u(\mu_{\text{prior}})]$ weakly increased:
  - $u(\mu_{\text{posterior}})$ weakly increased.
  - $u(\mu_{\text{prior}})$ is unchanged.
- score for extremal reports weakly less extreme
  - still ex post bounded. □

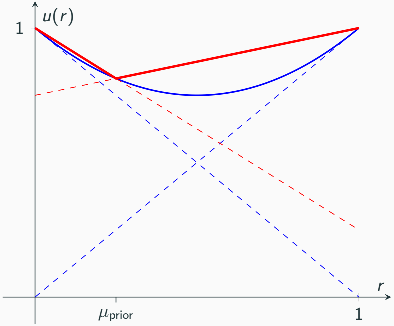

A Peer Grading Story

1. A peer grading platform (PeerPal).
2. Grading peer reviews with proper scoring rules is horrible!
3. (Quick fix: Manually grade the peer reviews.)
4. Optimization of scoring rules.
5. Fundamental Role of Scoring Rules

A Value of Data (via “Revelation Principle”)

analysis of dataset ⟶ decision optimization ⟶ payoff from decision

A Value of Data (via "Revelation Principle")

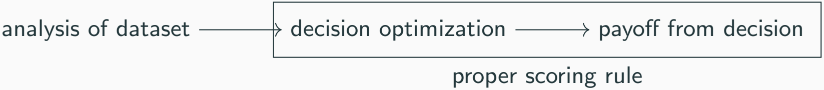

analysis of dataset → decision optimization → payoff from decision

proper scoring rule

A Value of Data (via “Revelation Principle”)

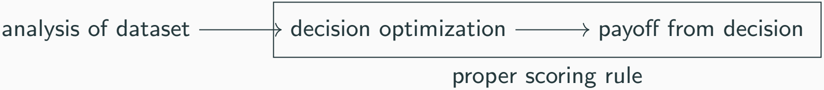

analysis of dataset → decision optimization → payoff from decision

proper scoring rule

Interpretations

A Value of Data (via "Revelation Principle")

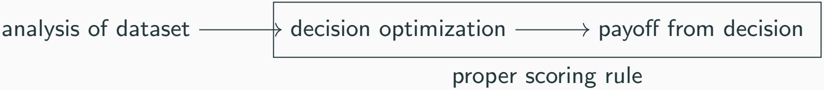

analysis of dataset → decision optimization → payoff from decision

proper scoring rule

Interpretations

- scoring rules are fundamental for understanding good data analyses

A Value of Data (via “Revelation Principle”)

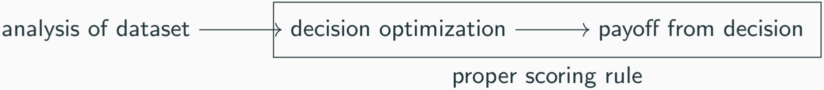

analysis of dataset → decision optimization → payoff from decision

proper scoring rule

Interpretations

- scoring rules are fundamental for understanding good data analyses
- optimal scoring rules for binary effort ⇒ setting-independent value of dataset

A Value of Data (via “Revelation Principle”)

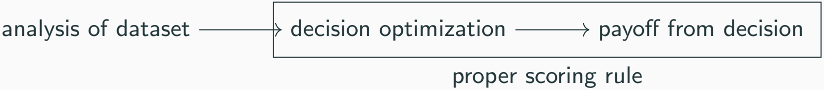

analysis of dataset → decision optimization → payoff from decision

proper scoring rule

Interpretations

- scoring rules are fundamental for understanding good data analyses
- optimal scoring rules for binary effort ⇒ setting-independent value of dataset

Example (Rational Agent Framework for Data Visualization)

[Wu, Guo, Mamakos, Hartline, Hullman VIS’23]

A Value of Data (via “Revelation Principle”)

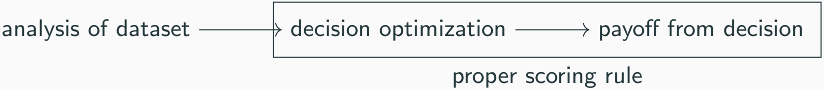

analysis of dataset → decision optimization → payoff from decision

proper scoring rule

Interpretations

- scoring rules are fundamental for understanding good data analyses
- optimal scoring rules for binary effort ⇒ setting-independent value of dataset

Example (Rational Agent Framework for Data Visualization)

- researcher shows behavioral subjects different visual stimuli.

[Wu, Guo, Mamakos, Hartline, Hullman VIS’23]

A Value of Data (via "Revelation Principle")

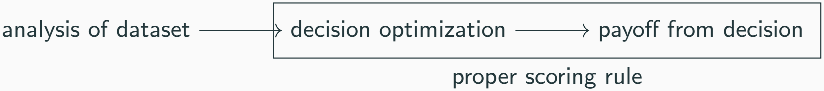

analysis of dataset → decision optimization → payoff from decision

proper scoring rule

Interpretations

- scoring rules are fundamental for understanding good data analyses
- optimal scoring rules for binary effort $\Rightarrow$ setting-independent value of dataset

Example (Rational Agent Framework for Data Visualization)

- researcher shows behavioral subjects different visual stimuli.
- measure performance in decision problem (a.k.a., scoring rule).

[Wu, Guo, Mamakos, Hartline, Hullman VIS'23]

A Value of Data (via "Revelation Principle")

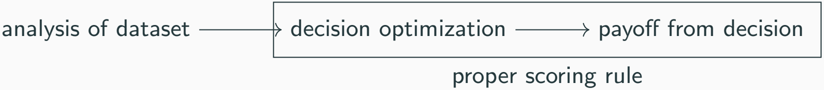

analysis of dataset → decision optimization → payoff from decision

proper scoring rule

Interpretations

- scoring rules are fundamental for understanding good data analyses
- optimal scoring rules for binary effort ⇒ setting-independent value of dataset

Example (Rational Agent Framework for Data Visualization)

- researcher shows behavioral subjects different visual stimuli.
- measure performance in decision problem (a.k.a., scoring rule).
- benchmark against rational agent with and without stimuli.

[Wu, Guo, Mamakos, Hartline, Hullman VIS'23]

A Value of Data (via “Revelation Principle”)

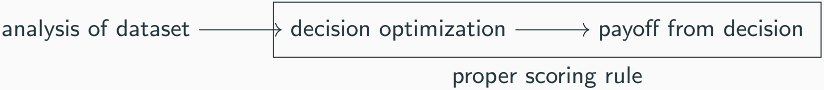

analysis of dataset ⟶ decision optimization ⟶ payoff from decision

proper scoring rule

Interpretations

- scoring rules are fundamental for understanding good data analyses
- optimal scoring rules for binary effort ⇒ setting-independent value of dataset

Example (Rational Agent Framework for Data Visualization)

- researcher shows behavioral subjects different visual stimuli.
- measure performance in decision problem (a.k.a., scoring rule).
- benchmark against rational agent with and without stimuli.

[Wu, Guo, Mamakos, Hartline, Hullman VIS’23]

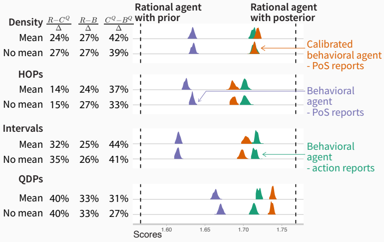

	$\frac{R-C^Q}{\Delta}$	$\frac{R-B}{\Delta}$	$\frac{C^Q-B^Q}{\Delta}$
Density			
Mean	24%	27%	42%
No mean	27%	27%	39%
HOPs			
Mean	14%	24%	37%
No mean	15%	27%	33%
Intervals			
Mean	32%	25%	44%
No mean	35%	26%	41%
QDPs			
Mean	40%	33%	31%
No mean	40%	33%	27%

Mechanism Design for the Classroom

The classroom as a "computer":

- students: local optimizers
- grader/instructor: imprecise operators
- syllabus: rules that map actions to grades
- student incentives: minimize work, maximize grade
- goal: minimize work, maximize learning, fairly assess

Basic Questions: What is best syllabus?

Examples:

- grading randomized exams: ex post fairness? [Chen, Hartline, Zoeter FORC'23]
- grading with partial credit: incentivizing precise answers? [Chen, Hartline, Zoeter]
- group projects: incentivizing teamwork?
- peer grading: incentives for accurate peer reviews? [Li, Hartline, Shan, Wu EC'22]

Related Work

Related Work:

- characterizing scoring rules:
  - eliciting full distribution [McCarthy '56; Gneiting, Raftery '07]
  - eliciting the mean [Abernethy, Frongillo '12]
  - set of elicitable properties (e.g., variance is not directly elicitable) [Lambert '11]

Related Work

Related Work:

- characterizing scoring rules:
  - eliciting full distribution [McCarthy '56; Gneiting, Raftery '07]
  - eliciting the mean [Abernethy, Frongillo '12]
  - set of elicitable properties (e.g., variance is not directly elicitable) [Lambert '11]
- maximize effort with quadratic scoring rules [Osband '89]

Related Work

Related Work:

- characterizing scoring rules:
  - eliciting full distribution [McCarthy '56; Gneiting, Raftery '07]
  - eliciting the mean [Abernethy, Frongillo '12]
  - set of elicitable properties (e.g., variance is not directly elicitable) [Lambert '11]
- maximize effort with quadratic scoring rules [Osband '89]
- maximize effort in a binary state model with costly samples [Neyman, Noarov, Weinberg '21]

Related Work

Related Work:

- characterizing scoring rules:
  - eliciting full distribution [McCarthy '56; Gneiting, Raftery '07]
  - eliciting the mean [Abernethy, Frongillo '12]
  - set of elicitable properties (e.g., variance is not directly elicitable) [Lambert '11]
- maximize effort with quadratic scoring rules [Osband '89]
- maximize effort in a binary state model with costly samples [Neyman, Noarov, Weinberg '21]
- framework adopted by follow-up works:
  - optimizing max-min objective without knowledge about prior and signal [Chen and Yu '21]
  - optimization of peer prediction mechanisms [Kong '21]
  - bounded expected score [Papireddygari, Waggoner '22]
  - maximizing effort under multi-dimensional effort model [Hartline, Li, Shan, Wu '23]
  - benchmark for visualization experiments [Wu, Guo, Mamakos, Hartline, Hullman '23]

Related Work

Related Work:

- characterizing scoring rules:
  - eliciting full distribution [McCarthy '56; Gneiting, Raftery '07]
  - eliciting the mean [Abernethy, Frongillo '12]
  - set of elicitable properties (e.g., variance is not directly elicitable) [Lambert '11]
- maximize effort with quadratic scoring rules [Osband '89]
- maximize effort in a binary state model with costly samples [Neyman, Noarov, Weinberg '21]
- framework adopted by follow-up works:
  - optimizing max-min objective without knowledge about prior and signal [Chen and Yu '21]
  - optimization of peer prediction mechanisms [Kong '21]
  - bounded expected score [Papireddygari, Waggoner '22]
  - maximizing effort under multi-dimensional effort model [Hartline, Li, Shan, Wu '23]
  - benchmark for visualization experiments [Wu, Guo, Mamakos, Hartline, Hullman '23]
- ex post value of information [Frankel, Kamenica '19]