Approximation and Mechanism Design

Jason D. Hartline — Northwestern University

May 15, 2010

"Game theory has a great advantage in explicitly analyzing the consequences of trading rules that presumably are really common knowledge, it is deficient to the extent it assumes other features to be common knowledge, such as one player's probability assessment about another's preferences or information.

"I forsee the progress of game theory as depending on successive reductions in the base of common knowledge required to conduct useful analysis of practical problems. Only by repeated weakening of common knowledge assumptions will the theory approximate reality."

– Robert Wilson, 1987.

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

Thesis

Informal Thesis: *approximate optimality* is

- more *descriptive* of mechanisms in practice than exact optimality,
- *prescribes* solutions to problems where exact optimality has not, and
- more *conclusive* about salient characteristics of good mechanisms.

Example 1: Gambler's Stopping Game

A Gambler's *Stopping Game*:

- *sequence* of n games,
- *prize* of game i is distributed from F_i,
- *prior-knowledge* of distributions.

On day i, gambler plays game i:

- *realizes* prize $v_i \sim F_i$,
- chooses to keep prize and *stop*, or
- discard prize and *continue*.

Example 1: Gambler's Stopping Game

A Gambler's *Stopping Game*:

- *sequence* of n games,
- *prize* of game i is distributed from F_i,
- *prior-knowledge* of distributions.

On day i, gambler plays game i:

- *realizes* prize $v_i \sim F_i$,
- chooses to keep prize and *stop*, or
- discard prize and *continue*.

Question: How should our gambler play?

Optimal Strategy

Optimal Strategy:

- threshold t_i for stopping with ith prize.
- solve with “backwards induction”.

Optimal Strategy:

- threshold t_i for stopping with ith prize.
- solve with "backwards induction".

Discussion:

- *Complicated:* n different, unrelated thresholds.
- *Inconclusive:* what are properties of good strategies?
- *Non-robust:* what if order changes? what if distribution changes?
- *Non-general:* what do we learn about variants of Stopping Game?

Threshold Strategies and Prophet Inequality

Threshold Strategy: "fix t, gambler takes first prize $v_i \geq t$".

(clearly suboptimal, may not accept prize on last day!)

Threshold Strategies and Prophet Inequality

Threshold Strategy: "fix t, gambler takes first prize $v_i \geq t$".

(clearly suboptimal, may not accept prize on last day!)

Theorem: *(Prophet Inequality)* For t such that **Pr**["no prize"] $= 1/2$,

$$\mathbf{E}[\text{prize for strategy } t] \geq \mathbf{E}[\max_i v_i]/2.$$

[Samuel-Cahn '84]

Threshold Strategies and Prophet Inequality

Threshold Strategy: "fix t, gambler takes first prize $v_i \geq t$".

(clearly suboptimal, may not accept prize on last day!)

Theorem: *(Prophet Inequality)* For t such that $\mathbf{Pr}[\text{"no prize"}] = 1/2$,

$$\mathbf{E}[\text{prize for strategy } t] \geq \mathbf{E}[\max_i v_i]/2.$$

[Samuel-Cahn '84]

Discussion:

- *Simple:* one number t.
- *Conclusive:* trade-off "stopping early" with "never stopping".
- *Robust:* change order? change distribution above or below t?
- *General:* same solution works for similar games: invariant of "tie-breaking rule"

Prophet Inequality Proof

0. Notation:

- $q_i = \textbf{Pr}[v_i < t]$.
- $x = \textbf{Pr}[\text{never stops}] = \prod_i q_i$.

1. Upper Bound on $\textbf{E}[\max]$:

2. Lower Bound on $\textbf{E}[\text{prize}]$:

3. Choose $x = 1/2$ to prove theorem.

Prophet Inequality Proof

0. Notation:

- $q_i = \mathbf{Pr}[v_i < t]$.
- $x = \mathbf{Pr}[\text{never stops}] = \prod_i q_i$.

1. Upper Bound on $\mathbf{E}[\max]$:

$$\mathbf{E}[\max] \leq t + \mathbf{E}\left[\max_i (v_i - t)^+\right]$$

2. Lower Bound on $\mathbf{E}[\text{prize}]$:

3. Choose $x = 1/2$ to prove theorem.

Prophet Inequality Proof

0. Notation:

- $q_i = \mathbf{Pr}[v_i < t]$.
- $x = \mathbf{Pr}[\text{never stops}] = \prod_i q_i$.

1. Upper Bound on $\mathbf{E}[\max]$:

$$\begin{aligned} \mathbf{E}[\max] &\leq t + \mathbf{E}\left[\max_i (v_i - t)^+\right] \\ &\leq t + \sum_i \mathbf{E}\left[(v_i - t)^+\right]. \end{aligned}$$

2. Lower Bound on $\mathbf{E}[\text{prize}]$:

3. Choose $x = 1/2$ to prove theorem.

Prophet Inequality Proof

0. Notation:

- $q_i = \mathbf{Pr}[v_i < t]$.
- $x = \mathbf{Pr}[\text{never stops}] = \prod_i q_i$.

1. Upper Bound on $\mathbf{E}[\max]$:

$$\mathbf{E}[\max] \leq t + \mathbf{E}\left[\max_i (v_i - t)^+\right]$$
$$\leq t + \sum_i \mathbf{E}\left[(v_i - t)^+\right].$$

2. Lower Bound on $\mathbf{E}[\text{prize}]$:

$\mathbf{E}[\text{prize}] \geq (1 - x)t +$

3. Choose $x = 1/2$ to prove theorem.

Prophet Inequality Proof

0. Notation:

- $q_i = \mathbf{Pr}[v_i < t]$.
- $x = \mathbf{Pr}[\text{never stops}] = \prod_i q_i$.

1. Upper Bound on $\mathbf{E}[\max]$:

$$\begin{aligned}\mathbf{E}[\max] &\le t + \mathbf{E}\left[\max_i (v_i - t)^+\right] \\ &\le t + \sum_i \mathbf{E}\left[(v_i - t)^+\right].\end{aligned}$$

2. Lower Bound on $\mathbf{E}[\text{prize}]$:

$$\mathbf{E}[\text{prize}] \ge (1-x)t + \sum_i \mathbf{E}\left[(v_i - t)^+ \mid \text{other } v_j < t\right] \mathbf{Pr}[\text{other } v_j < t]$$

3. Choose $x = 1/2$ to prove theorem.

Prophet Inequality Proof

0. Notation:

- $q_i = \mathbf{Pr}[v_i < t]$.
- $x = \mathbf{Pr}[\text{never stops}] = \prod_i q_i$.

1. Upper Bound on $\mathbf{E}[\max]$:

$$\begin{aligned}\mathbf{E}[\max] &\leq t + \mathbf{E}\left[\max_i (v_i - t)^+\right] \\ &\leq t + \sum_i \mathbf{E}\left[(v_i - t)^+\right].\end{aligned}$$

2. Lower Bound on $\mathbf{E}[\text{prize}]$:

$$\mathbf{E}[\text{prize}] \geq (1-x)t + \sum_i \mathbf{E}\left[(v_i - t)^+ \mid \text{other } v_j < t\right] \overbrace{\mathbf{Pr}[\text{other } v_j < t]}^{\prod_{j \neq i} q_j}$$

3. Choose $x = 1/2$ to prove theorem.

Prophet Inequality Proof

0. Notation:

- $q_i = \mathbf{Pr}[v_i < t]$.
- $x = \mathbf{Pr}[\text{never stops}] = \prod_i q_i$.

1. Upper Bound on $\mathbf{E}[\max]$:

$$\begin{aligned}\mathbf{E}[\max] &\leq t + \mathbf{E}\left[\max_i (v_i - t)^+\right] \\ &\leq t + \sum_i \mathbf{E}\left[(v_i - t)^+\right].\end{aligned}$$

2. Lower Bound on $\mathbf{E}[\text{prize}]$:

$$\mathbf{E}[\text{prize}] \geq (1-x)t + \sum_i \mathbf{E}\left[(v_i - t)^+ \mid \text{other } v_j < t\right] \overbrace{\mathbf{Pr}[\text{other } v_j < t]}^{x \leq \prod_{j \neq i} q_j}$$

3. Choose $x = 1/2$ to prove theorem.

Prophet Inequality Proof

0. Notation:

- $q_i = \mathbf{Pr}[v_i < t]$.
- $x = \mathbf{Pr}[\text{never stops}] = \prod_i q_i$.

1. Upper Bound on $\mathbf{E}[\max]$:

$$\begin{aligned}\mathbf{E}[\max] &\leq t + \mathbf{E}\left[\max_i (v_i - t)^+\right] \\ &\leq t + \sum_i \mathbf{E}\left[(v_i - t)^+\right].\end{aligned}$$

2. Lower Bound on $\mathbf{E}[\text{prize}]$:

$$\begin{aligned}\mathbf{E}[\text{prize}] &\geq (1-x)t + \sum_i \mathbf{E}\left[(v_i - t)^+ \mid \text{other } v_j < t\right] \overbrace{\mathbf{Pr}[\text{other } v_j < t]}^{x \leq \prod_{j \neq i} q_j} \\ &\geq (1-x)t + x \sum_i \mathbf{E}\left[(v_i - t)^+ \mid \text{other } v_j < t\right]\end{aligned}$$

3. Choose $x = 1/2$ to prove theorem.

Prophet Inequality Proof

0. Notation:

- $q_i = \mathbf{Pr}[v_i < t]$.
- $x = \mathbf{Pr}[\text{never stops}] = \prod_i q_i$.

1. Upper Bound on $\mathbf{E}[\max]$:

$$\begin{aligned} \mathbf{E}[\max] &\leq t + \mathbf{E}\left[\max_i (v_i - t)^+\right] \\ &\leq t + \sum_i \mathbf{E}\left[(v_i - t)^+\right]. \end{aligned}$$

2. Lower Bound on $\mathbf{E}[\text{prize}]$:

$$\begin{aligned} \mathbf{E}[\text{prize}] &\geq (1-x)t + \sum_i \mathbf{E}\left[(v_i - t)^+ \mid \text{other } v_j < t\right] \overbrace{\mathbf{Pr}[\text{other } v_j < t]}^{x \leq \prod_{j \neq i} q_j} \\ &\geq (1-x)t + x \sum_i \mathbf{E}\left[(v_i - t)^+ \mid \text{other } v_j < t\right] \\ &= (1-x)t + x \sum_i \mathbf{E}\left[(v_i - t)^+\right]. \end{aligned}$$

3. Choose $x = 1/2$ to prove theorem.

Philosophy of Approximation

What is the point of a 2-approximation?

Philosophy of Approximation

What is the point of a 2-approximation?

- Must make tradeoff between understanding and optimality.

	$(1+\epsilon)$	constant	super-constant
Performance	great	ok	bad
Understanding	little	lots	little/none

Philosophy of Approximation

What is the point of a 2-approximation?

- Must make tradeoff between understanding and optimality.

	$(1+\epsilon)$	constant	super-constant
Performance	great	ok	bad
Understanding	little	lots	little/none

- Constant approximations identify salient features of model.

What is the point of a 2-approximation?

- Must make tradeoff between understanding and optimality.

	$(1+\epsilon)$	constant	super-constant
Performance	great	ok	bad
Understanding	little	lots	little/none

- Constant approximations identify salient features of model.
  Example: is X important in MD?
  - no, if mech without X is constant approx
  - yes, otherwise.

Philosophy of Approximation

What is the point of a 2-approximation?

- Must make tradeoff between understanding and optimality.

	$(1+\epsilon)$	constant	super-constant
Performance	great	ok	bad
Understanding	little	lots	little/none

- Constant approximations identify salient features of model.
  Example: is X important in MD? competition?
  - no, if mech without X is constant approx
  - yes, otherwise.

Philosophy of Approximation

What is the point of a 2-approximation?

- Must make tradeoff between understanding and optimality.

	$(1+\epsilon)$	constant	super-constant
Performance	great	ok	bad
Understanding	little	lots	little/none

- Constant approximations identify salient features of model.
  Example: is X important in MD? competition? transfers?
  - no, if mech without X is constant approx
  - yes, otherwise.

Philosophy of Approximation

What is the point of a 2-approximation?

- Must make tradeoff between understanding and optimality.

	$(1+\epsilon)$	constant	super-constant
Performance	great	ok	bad
Understanding	little	lots	little/none

- Constant approximations identify salient features of model.
  Example: is X important in MD? competition? transfers?
  - no, if mech without X is constant approx
  - yes, otherwise.
- Seller can always try ad hoc improvements on approximation.

Overview

1. Single-dimensional Bayesian settings.

   (e.g., single-item auctions)

2. Multi-dimensional Bayesian settings.

   (e.g., multi-item auctions)

3. Prior-free settings.

Part I: Approximation for single-dimensional Bayesian mechanism design

(where agent preferences are given by a private value for service, zero for no service; preferences are drawn from a distribution)

Example 2: Single-item auction

Problem: Bayesian Single-item Auction Problem

- a single item for sale,
- n buyers, and
- a dist. $\mathbf{F} = F_1 \times \cdots \times F_n$ from which the consumers' values for the item are drawn.

Goal: seller opt. auction for $\mathbf{F}$.

Example 2: Single-item auction

Problem: Bayesian Single-item Auction Problem

- a single item for sale,
- n buyers, and
- a dist. $\mathbf{F} = F_1 \times \cdots \times F_n$ from which the consumers' values for the item are drawn.

Goal: seller opt. auction for $\mathbf{F}$.

Question: What is optimal auction?

Optimal Auction Design [Myerson '81]

Optimal Auction Design [Myerson '81]

1. **Def:** *revenue curve*: $R_i(q) = q \cdot F_i^{-1}(1-q)$.

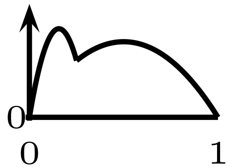

Optimal Auction Design [Myerson '81]

1. **Def:** *revenue curve*: $R_i(q) = q \cdot F_i^{-1}(1-q)$.

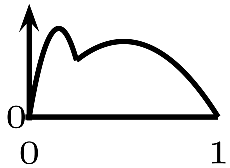

2. **Def:** *virtual value*: $\varphi_i(v_i) = v_i - \frac{1-F_i(v)}{f_i(v_i)}$ = marginal revenue.

Optimal Auction Design [Myerson '81]

1. **Def:** *revenue curve*: $R_i(q) = q \cdot F_i^{-1}(1-q)$.

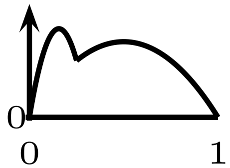

2. **Def:** *virtual value*: $\varphi_i(v_i) = v_i - \frac{1-F_i(v)}{f_i(v_i)}$ = marginal revenue.

3. **Def:** *virtual surplus*: virtual value of winner(s).

Optimal Auction Design [Myerson '81]

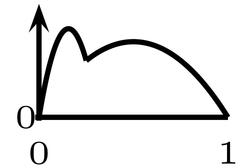

1. **Def:** *revenue curve*: $R_i(q) = q \cdot F_i^{-1}(1-q)$.
2. **Def:** *virtual value*: $\varphi_i(v_i) = v_i - \frac{1-F_i(v)}{f_i(v_i)} =$ marginal revenue.
3. **Def:** *virtual surplus*: virtual value of winner(s).
4. **Thm:** **E**[revenue] = **E**[virtual surplus].

Optimal Auction Design [Myerson '81]

1. **Def:** *revenue curve*: $R_i(q) = q \cdot F_i^{-1}(1-q)$.

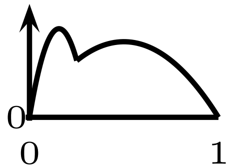

2. **Def:** *virtual value*: $\varphi_i(v_i) = v_i - \frac{1-F_i(v)}{f_i(v_i)}$ = marginal revenue.

3. **Def:** *virtual surplus*: virtual value of winner(s).

4. **Thm:** $\mathbf{E}[\text{revenue}] = \mathbf{E}[\text{virtual surplus}]$.

5. **Def:** F_i is *regular* iff revenue curve concave iff virtual values monotone.

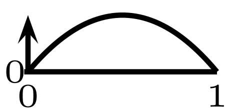

Optimal Auction Design [Myerson '81]

1. **Def:** *revenue curve*: $R_i(q) = q \cdot F_i^{-1}(1-q)$.

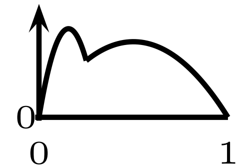

2. **Def:** *virtual value*: $\varphi_i(v_i) = v_i - \frac{1-F_i(v)}{f_i(v_i)}$ = marginal revenue.

3. **Def:** *virtual surplus*: virtual value of winner(s).

4. **Thm:** $\mathbf{E}$[revenue] $=$ $\mathbf{E}$[virtual surplus].

5. **Def:** F_i is *regular* iff revenue curve concave iff virtual values monotone.

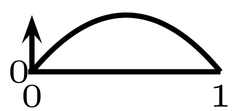

6. **Thm:** for regular dists, optimal auction sells to bidder with highest positive virtual value.

Optimal Auction Design [Myerson '81]

1. **Def:** *revenue curve*: $R_i(q) = q \cdot F_i^{-1}(1-q)$.

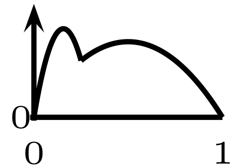

2. **Def:** *virtual value*: $\varphi_i(v_i) = v_i - \frac{1-F_i(v)}{f_i(v_i)}$ = marginal revenue.

3. **Def:** *virtual surplus*: virtual value of winner(s).

4. **Thm:** **E**[revenue] = **E**[virtual surplus].

5. **Def:** F_i is *regular* iff revenue curve concave iff virtual values monotone.

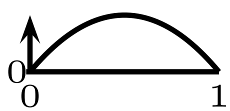

6. **Thm:** for regular dists, optimal auction sells to bidder with highest positive virtual value.

7. **Cor:** for iid, regular dists, optimal auction is *Vickrey with monopoly reserve price* $\varphi^{-1}(0)$.

Optimal Auctions

Optimal Auctions:

- *iid, regular distributions*: Vickrey with monopoly reserve price.
- *general*: sell to bidder with highest positive virtual value.

Optimal Auctions:

- *iid, regular distributions*: Vickrey with monopoly reserve price.
- *general*: sell to bidder with highest positive virtual value.

Discussion:

- iid, regular case: seems unlikely in practice.
- general case: nobody runs optimal auction (too complicated?).

Approximation with reserve prices

Question: when is reserve pricing a good approximation?

Approximation with reserve prices

Question: when is reserve pricing a good approximation?

Thm: Vickrey with reserve $=$ *constant virtual price* with **Pr**$[\text{no sale}] = 1/2$ is a 2-approximation. [Chawla, H, Malec, Sivan '10]

Approximation with reserve prices

Question: when is reserve pricing a good approximation?

Thm: Vickrey with reserve $=$ *constant virtual price* with **Pr**$[\text{no sale}] = 1/2$ is a 2-approximation. [Chawla, H, Malec, Sivan '10]

Proof: apply prophet inequality (tie-breaking by value) to virtual values.

Approximation with reserve prices

Question: when is reserve pricing a good approximation?

Thm: Vickrey with reserve $=$ *constant virtual price* with **Pr**$[$no sale$] = 1/2$ is a 2-approximation. [Chawla, H, Malec, Sivan '10]

Proof: apply prophet inequality (tie-breaking by value) to virtual values.

prophet inequality	Vickrey with reserves
prizes	virtual values
threshold t	virtual price
E$[$max prize$]$	**E**$[$optimal revenue$]$
E$[$prize for $t]$	**E**$[$Vickrey revenue$]$

Approximation with reserve prices

Question: when is reserve pricing a good approximation?

Thm: Vickrey with reserve $=$ *constant virtual price* with **Pr**$[$no sale$] = 1/2$ is a 2-approximation. [Chawla, H, Malec, Sivan '10]

Proof: apply prophet inequality (tie-breaking by value) to virtual values.

prophet inequality	Vickrey with reserves
prizes	virtual values
threshold t	virtual price
E[max prize]	**E**[optimal revenue]
E[prize for t]	**E**[Vickrey revenue]

Discussion:

- constant virtual price $\Rightarrow$ bidder-specific reserves.
- *simple* reserve prices natural, practical, and easy to find.
- *robust* posted pricing with arbitrary tie-breaking works fine too.

Anonymous Reserves

Question: for non-identical distributions, is *anonymous reserve* approximately optimal?

(e.g., eBay)

Anonymous Reserves

Question: for non-identical distributions, is *anonymous reserve* approximately optimal?

(e.g., eBay)

Thm: non-identical, regular distributions, Vickrey with *anonymous reserve price* is 3-approximation. [H, Roughgarden '09]

Anonymous Reserves

Question: for non-identical distributions, is *anonymous reserve* approximately optimal?

(e.g., eBay)

Thm: non-identical, regular distributions, Vickrey with *anonymous reserve price* is 3-approximation. [H, Roughgarden '09]

Proof: more complicated extension of prophet inequalities.

Anonymous Reserves

Question: for non-identical distributions, is *anonymous reserve* approximately optimal?

(e.g., eBay)

Thm: non-identical, regular distributions, Vickrey with *anonymous reserve price* is 3-approximation. [H, Roughgarden '09]

Proof: more complicated extension of prophet inequalities.

Discussion:

- theorem is not tight, actual bound is in $[2, 3]$.
- justifies wide prevalence.
- approximation good for *platform design*.

Extensions

Beyond single-item auctions: *general feasibility constraints*.

Extensions

Beyond single-item auctions: *general feasibility constraints*.

Thm: for non-identical regular distributions, VCG with monopoly reserves is often a 2-approximation. [H, Roughgarden '09]

Thm: non-identical (possibly irregular) distributions, *posted pricing mechanisms* are often constant approximations.
[Chawla, H, Malec, Sivan '10]

Extensions

Beyond single-item auctions: *general feasibility constraints*.

Thm: for non-identical regular distributions, VCG with monopoly reserves is often a 2-approximation. [H, Roughgarden '09]

Thm: non-identical (possibly irregular) distributions, *posted pricing mechanisms* are often constant approximations.
[Chawla, H, Malec, Sivan '10]

Proof technique:

- optimal mechanism is a virtual surplus maximizer.
- reserve-price mechanisms are virtual surplus approximators.

Extensions

Beyond single-item auctions: *general feasibility constraints*.

Thm: for non-identical regular distributions, VCG with monopoly reserves is often a 2-approximation. [H, Roughgarden '09]

Thm: non-identical (possibly irregular) distributions, *posted pricing mechanisms* are often constant approximations. [Chawla, H, Malec, Sivan '10]

Proof technique:

- optimal mechanism is a virtual surplus maximizer.
- reserve-price mechanisms are virtual surplus approximators.

Basic Open Question: to what extent to simple mechanisms approximate (well understood but complex) optimal ones?

Challenges: non-downward-closed settings, negative virtual values.

Part II: Approximation for multi-dimensional Bayesian mechanism design

(where agent preferences are given by values for each available service, zero for no service; preferences drawn from distribution)

Example 3: unit-demand pricing

Problem: Bayesian Unit-Demand Pricing

- a single, unit-demand consumer.
- n items for sale.
- a dist. $\mathbf{F} = F_1 \times \cdots \times F_n$ from which the consumer's values for each item are drawn.

Goal: seller optimal *item-pricing* for $\mathbf{F}$.

Example 3: unit-demand pricing

Problem: Bayesian Unit-Demand Pricing

- a single, unit-demand consumer.
- n items for sale.
- a dist. $\mathbf{F} = F_1 \times \cdots \times F_n$ from which the consumer's values for each item are drawn.

Goal: seller optimal *item-pricing* for $\mathbf{F}$.

Question: What is optimal pricing?

Optimal Pricing

Optimal Pricing: consider distribution, feasibility constraints, incentive constraints, and solve!

Optimal Pricing

Optimal Pricing: consider distribution, feasibility constraints, incentive constraints, and solve!

Discussion:

- little conceptual insight and
- not generally tractable.

Analogy

Challenge: approximate optimal but we do not understand it?

Analogy

Challenge: approximate optimal but we do not understand it?

Problem: Bayesian Unit-demand Pricing (a.k.a., MD-PRICING)

- a single, *unit-demand* buyer,
- n items for sale, and
- a dist. $\mathbf{F}$ from which the consumer's value for each item is drawn.

Goal: seller opt. item-pricing for $\mathbf{F}$.

Problem: Bayesian Single-item Auction (a.k.a., SD-AUCTION)

- a single item for sale,
- n buyers, and
- a dist. $\mathbf{F}$ from which the consumers' values for the item are drawn.

Goal: seller opt. auction for $\mathbf{F}$.

Analogy

Challenge: approximate optimal but we do not understand it?

Problem: Bayesian Unit-demand Pricing (a.k.a., MD-PRICING)

- a single, *unit-demand* buyer,
- n items for sale, and
- a dist. $\mathbf{F}$ from which the consumer's value for each item is drawn.

Goal: seller opt. item-pricing for $\mathbf{F}$.

Problem: Bayesian Single-item Auction (a.k.a., SD-AUCTION)

- a single item for sale,
- n buyers, and
- a dist. $\mathbf{F}$ from which the consumers' values for the item are drawn.

Goal: seller opt. auction for $\mathbf{F}$.

Note: Same informational structure.

Analogy

Challenge: approximate optimal but we do not understand it?

Problem: Bayesian Unit-demand Pricing (a.k.a., MD-PRICING)

- a single, *unit-demand* buyer,
- n items for sale, and
- a dist. $\mathbf{F}$ from which the consumer's value for each item is drawn.

Goal: seller opt. item-pricing for $\mathbf{F}$.

Problem: Bayesian Single-item Auction (a.k.a., SD-AUCTION)

- a single item for sale,
- n buyers, and
- a dist. $\mathbf{F}$ from which the consumers' values for the item are drawn.

Goal: seller opt. auction for $\mathbf{F}$.

Note: Same informational structure.
Thm: for any indep. distributions, MD-PRICING $\leq$ SD-AUCTION.

Analogy

Challenge: approximate optimal but we do not understand it?

Problem: Bayesian Unit-demand Pricing (a.k.a., MD-PRICING) • a single, *unit-demand* buyer, • n items for sale, and • a dist. $\mathbf{F}$ from which the consumer's value for each item is drawn. **Goal:** seller opt. item-pricing for $\mathbf{F}$.	**Problem:** Bayesian Single-item Auction (a.k.a., SD-AUCTION) • a single item for sale, • n buyers, and • a dist. $\mathbf{F}$ from which the consumers' values for the item are drawn. **Goal:** seller opt. auction for $\mathbf{F}$.

Note: Same informational structure.

Thm: for any indep. distributions, MD-PRICING $\leq$ SD-AUCTION.

Thm: a constant virtual price for MD-PRICING is 2-approx.

[Chawla, H, Malec, Sivan '10]

Analogy

Challenge: approximate optimal but we do not understand it?

Problem: Bayesian Unit-demand Pricing (a.k.a., MD-PRICING)

- a single, *unit-demand* buyer,
- n items for sale, and
- a dist. $\mathbf{F}$ from which the consumer's value for each item is drawn.

Goal: seller opt. item-pricing for $\mathbf{F}$.

Problem: Bayesian Single-item Auction (a.k.a., SD-AUCTION)

- a single item for sale,
- n buyers, and
- a dist. $\mathbf{F}$ from which the consumers' values for the item are drawn.

Goal: seller opt. auction for $\mathbf{F}$.

Note: Same informational structure.

Thm: for any indep. distributions, MD-PRICING $\leq$ SD-AUCTION.

Thm: a constant virtual price for MD-PRICING is 2-approx.

[Chawla, H, Malec, Sivan '10]

Proof: prophet inequality.

Multi-item Auctions

Sequential Posted Pricing: agents arrive in sequence, offer posted prices.

Multi-item Auctions

Sequential Posted Pricing: agents arrive in sequence, offer posted prices.

Thm: in many unit-demand settings, sequential posted pricings are a constant approximation to the optimal mechanism.

[Chawla, H, Malec, Sivan '10]

Multi-item Auctions

Sequential Posted Pricing: agents arrive in sequence, offer posted prices.

Thm: in many unit-demand settings, sequential posted pricings are a constant approximation to the optimal mechanism.

[Chawla, H, Malec, Sivan '10]

Approach:

1. *Analogy:* "single-dimensional analog"

   (replace unit-demand agent with many single-dimensional agents)

Multi-item Auctions

Sequential Posted Pricing: agents arrive in sequence, offer posted prices.

Thm: in many unit-demand settings, sequential posted pricings are a constant approximation to the optimal mechanism.

[Chawla, H, Malec, Sivan '10]

Approach:

1. *Analogy:* "single-dimensional analog"

   (replace unit-demand agent with many single-dimensional agents)

2. *Upper bound:* SD-AUCTION $\geq$ MD-PRICING

   (competition increases revenue)

Multi-item Auctions

Sequential Posted Pricing: agents arrive in sequence, offer posted prices.

Thm: in many unit-demand settings, sequential posted pricings are a constant approximation to the optimal mechanism.

[Chawla, H, Malec, Sivan '10]

Approach:

1. *Analogy:* "single-dimensional analog"

   (replace unit-demand agent with many single-dimensional agents)

2. *Upper bound:* SD-AUCTION $\geq$ MD-PRICING

   (competition increases revenue)

3. *Reduction:* MD-PRICING $\geq$ SD-PRICING

   (pricings don't use competition)

Multi-item Auctions

Sequential Posted Pricing: agents arrive in sequence, offer posted prices.

Thm: in many unit-demand settings, sequential posted pricings are a constant approximation to the optimal mechanism.

[Chawla, H, Malec, Sivan '10]

Approach:

1. *Analogy:* "single-dimensional analog"

   (replace unit-demand agent with many single-dimensional agents)

2. *Upper bound:* SD-AUCTION $\geq$ MD-PRICING

   (competition increases revenue)

3. *Reduction:* MD-PRICING $\geq$ SD-PRICING

   (pricings don't use competition)

4. *Instantiation:* SD-PRICING $\geq \frac{1}{\beta}$SD-AUCTION

   (virtual surplus approximation)

Sequential Posted Pricing Discussion

Sequential Posted Pricing: agents arrive in sequence, offer posted prices.

Thm: in many unit-demand settings, sequential posted pricings are a constant approximation to the optimal mechanism.

[Chawla, H, Malec, Sivan ’10]

Sequential Posted Pricing Discussion

Sequential Posted Pricing: agents arrive in sequence, offer posted prices.

Thm: in many unit-demand settings, sequential posted pricings are a constant approximation to the optimal mechanism.

[Chawla, H, Malec, Sivan '10]

Discussion:

- *robust* to agent ordering, collusion, etc.
- *conclusive*: competition not important for approximation.
- *practical*: posted pricings widely prevalent. (e.g., eBay)
- role of randomization is crucial. [Chawla, Malec, Sivan '10]

Sequential Posted Pricing Discussion

Sequential Posted Pricing: agents arrive in sequence, offer posted prices.

Thm: in many unit-demand settings, sequential posted pricings are a constant approximation to the optimal mechanism.

[Chawla, H, Malec, Sivan '10]

Discussion:

- *robust* to agent ordering, collusion, etc.
- *conclusive*: competition not important for approximation.
- *practical*: posted pricings widely prevalent. (e.g., eBay)
- role of randomization is crucial. [Chawla, Malec, Sivan '10]

Open Question: identify upper bounds beyond unit-demand settings that are

- conceptually tractable and
- approximable.

Part III: Approximation for prior-free mechanism design.

(mechanisms should be good for any set of agent preferences, not just given distributional assumptions)

The problem with priors

Prior assumption: the mechanism designer knows the distribution of agent preferences.

The problem with priors

Prior assumption: the mechanism designer knows the distribution of agent preferences.

Where does prior come from:

- historical data

  then using prior affects incentives of earlier transactions.
  (e.g. Coase Conjecture)

- market analysis

  accuracy depends on market size, auctions are for small markets.

The problem with priors

Prior assumption: the mechanism designer knows the distribution of agent preferences.

Where does prior come from:

- historical data

  then using prior affects incentives of earlier transactions.
  (e.g. Coase Conjecture)

- market analysis

  accuracy depends on market size, auctions are for small markets.

Question: can we design good auctions without knowledge of prior-distribution?

Resource augmentation

Approach 1: “resource” augmentation.

Resource augmentation

Approach 1: "resource" augmentation.

Thm: for iid, regular, single-item auctions, the Vickrey auction on $n + 1$ bidders has more revenue than the optimal auction on n bidders.

[Bulow, Klemperer '96]

Resource augmentation

Approach 1: “resource” augmentation.

Thm: for iid, regular, single-item auctions, the Vickrey auction on $n + 1$ bidders has more revenue than the optimal auction on n bidders.

[Bulow, Klemperer ’96]

Discussion: [Dhangwatnotai, Roughgarden, Yan ’10]

- “recruit one more bidder” is prior-free strategy.
- “bicriteria” approximation result.
- *conclusive:* competition more important than optimization.

Resource augmentation

Approach 1: "resource" augmentation.

Thm: for iid, regular, single-item auctions, the Vickrey auction on $n + 1$ bidders has more revenue than the optimal auction on n bidders.

[Bulow, Klemperer '96]

Discussion: [Dhangwatnotai, Roughgarden, Yan '10]

- "recruit one more bidder" is prior-free strategy.
- "bicriteria" approximation result.
- *conclusive:* competition more important than optimization.
- *non-generic*: e.g., for k-unit auctions, need k additional bidders.

Special Case: $n = 1$

Special Case: for regular distribution, the Vickrey revenue from two bidders is at least the optimal revenue from one bidder.

Special Case: $n = 1$

Special Case: for regular distribution, the Vickrey revenue from two bidders is at least the optimal revenue from one bidder.

Geometric Proof: [Dhangwatnotai, Roughgarden, Yan ’10]

Special Case: $n = 1$

Special Case: for regular distribution, the Vickrey revenue from two bidders is at least the optimal revenue from one bidder.

Geometric Proof: [Dhangwatnotai, Roughgarden, Yan '10]

- each bidder in Vickrey views other bid as "random reserve".

Special Case: $n = 1$

Special Case: for regular distribution, the Vickrey revenue from two bidders is at least the optimal revenue from one bidder.

Geometric Proof: [Dhangwatnotai, Roughgarden, Yan '10]

- each bidder in Vickrey views other bid as "random reserve".
- Vickrey revenue $= 2\times$ random reserve revenue.

Special Case: $n = 1$

Special Case: for regular distribution, the Vickrey revenue from two bidders is at least the optimal revenue from one bidder.

Geometric Proof: [Dhangwatnotai, Roughgarden, Yan '10]

- each bidder in Vickrey views other bid as "random reserve".
- Vickrey revenue $= 2\times$ random reserve revenue.
- random reserve revenue $\geq \frac{1}{2}\times$ optimal reserve revenue:

Special Case: $n = 1$

Special Case: for regular distribution, the Vickrey revenue from two bidders is at least the optimal revenue from one bidder.

Geometric Proof: [Dhangwatnotai, Roughgarden, Yan '10]

- each bidder in Vickrey views other bid as "random reserve".
- Vickrey revenue $= 2\times$ random reserve revenue.
- random reserve revenue $\geq \frac{1}{2}\times$ optimal reserve revenue:

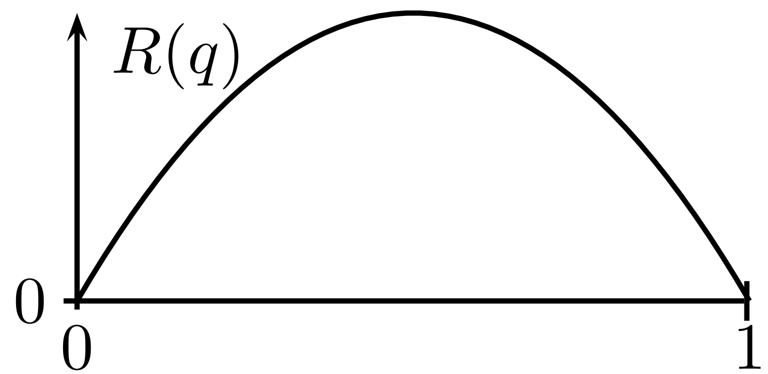

Special Case: $n = 1$

Special Case: for regular distribution, the Vickrey revenue from two bidders is at least the optimal revenue from one bidder.

Geometric Proof: [Dhangwatnotai, Roughgarden, Yan '10]

- each bidder in Vickrey views other bid as "random reserve".
- Vickrey revenue $= 2\times$ random reserve revenue.
- random reserve revenue $\geq \frac{1}{2}\times$ optimal reserve revenue:

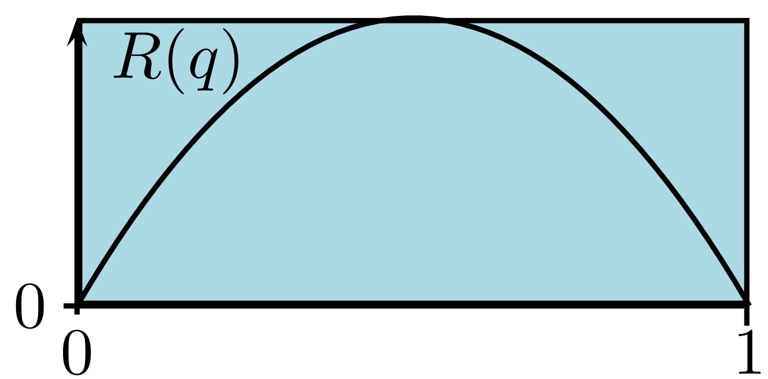

Special Case: $n = 1$

Special Case: for regular distribution, the Vickrey revenue from two bidders is at least the optimal revenue from one bidder.

Geometric Proof: [Dhangwatnotai, Roughgarden, Yan '10]

- each bidder in Vickrey views other bid as "random reserve".
- Vickrey revenue $= 2\times$ random reserve revenue.
- random reserve revenue $\geq \frac{1}{2}\times$ optimal reserve revenue:

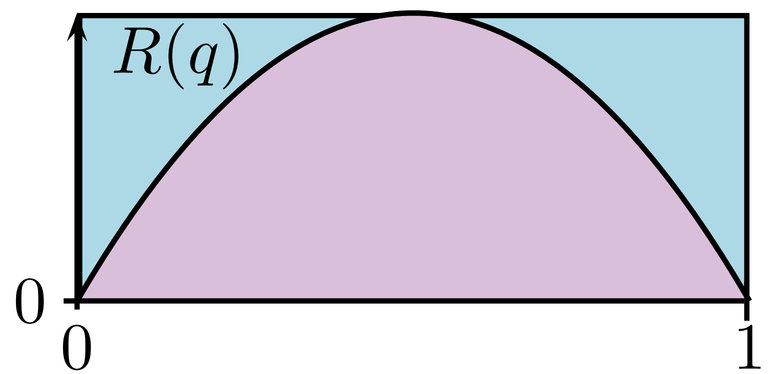

Special Case: $n = 1$

Special Case: for regular distribution, the Vickrey revenue from two bidders is at least the optimal revenue from one bidder.

Geometric Proof: [Dhangwatnotai, Roughgarden, Yan '10]

- each bidder in Vickrey views other bid as "random reserve".
- Vickrey revenue $= 2\times$ random reserve revenue.
- random reserve revenue $\geq \frac{1}{2}\times$ optimal reserve revenue:

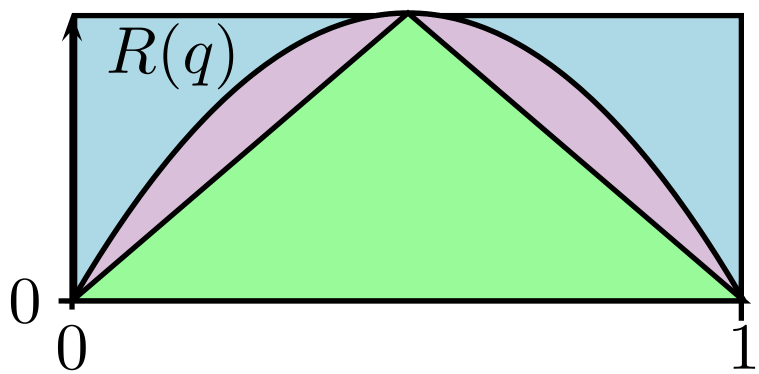

Special Case: $n = 1$

Special Case: for regular distribution, the Vickrey revenue from two bidders is at least the optimal revenue from one bidder.

Geometric Proof: [Dhangwatnotai, Roughgarden, Yan '10]

- each bidder in Vickrey views other bid as "random reserve".
- Vickrey revenue $= 2\times$ random reserve revenue.
- random reserve revenue $\geq \frac{1}{2}\times$ optimal reserve revenue:

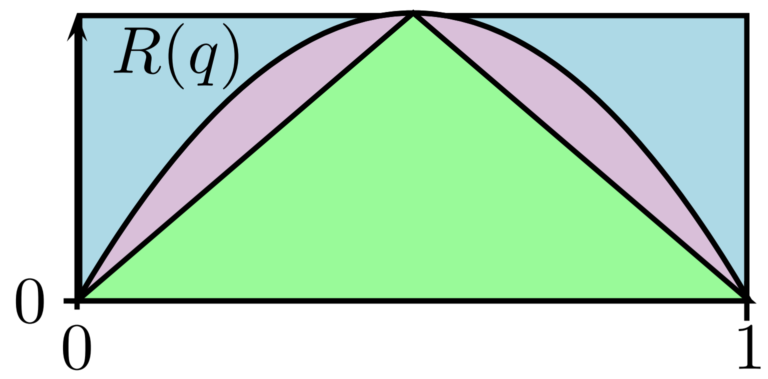

- So Vickrey with two bidders $\geq$ optimal revenue from one bidder.

Example 4: digital goods

Question: how should a profit-maximizing seller sell a *digital good* (n bidder, n copies of item)?

Example 4: digital goods

Question: how should a profit-maximizing seller sell a *digital good* (n bidder, n copies of item)?

Bayesian Optimal Solution: if values are iid from known distribution, post the monopoly price $\varphi^{-1}(0)$. [Myerson '81]

Example 4: digital goods

Question: how should a profit-maximizing seller sell a *digital good* (n bidder, n copies of item)?

Bayesian Optimal Solution: if values are iid from known distribution, post the monopoly price $\varphi^{-1}(0)$. [Myerson '81]

Discussion:

- optimal,
- simple, but
- not prior-free

Approximation via Single Sample

Single-Sample Auction: (for digital goods)

[Dhangwatnotai, Roughgarden, Yan ’10]

1. pick random agent i as sample.

2. offer all other agents price v_i.

3. reject i.

Approximation via Single Sample

Single-Sample Auction: (for digital goods)
[Dhangwatnotai, Roughgarden, Yan '10]

1. pick random agent i as sample.

2. offer all other agents price v_i.

3. reject i.

Thm: for iid, regular distributions, single sample auction on $(n+1)$-agents is 2-approx to optimal on n agents.
[Dhangwatnotai, Roughgarden, Yan '10]

Approximation via Single Sample

Single-Sample Auction: (for digital goods)
[Dhangwatnotai, Roughgarden, Yan '10]

1. pick random agent i as sample.
2. offer all other agents price v_i.
3. reject i.

Thm: for iid, regular distributions, single sample auction on $(n+1)$-agents is 2-approx to optimal on n agents.
[Dhangwatnotai, Roughgarden, Yan '10]

Proof: from geometric argument.

Approximation via Single Sample

Single-Sample Auction: (for digital goods)

[Dhangwatnotai, Roughgarden, Yan '10]

1. pick random agent i as sample.

2. offer all other agents price v_i.

3. reject i.

Thm: for iid, regular distributions, single sample auction on $(n+1)$-agents is 2-approx to optimal on n agents.

[Dhangwatnotai, Roughgarden, Yan '10]

Proof: from geometric argument.

Discussion:

- *prior-free*.
- *conclusive*, don't need precise distribution, only need single sample for approximation. more samples can improve approximation factor.
- *generic*, applies to general settings.

Average-case vs Worst-case

Note: prior-free auction cannot be optimal in every setting.

Average-case vs Worst-case

Note: prior-free auction cannot be optimal in every setting.

Average Case Approximation: $\exists \mathcal{A}, \forall \mathbf{F} \in$ IID,

$$\mathsf{E}_{\mathbf{v}\sim\mathbf{F}}[\mathcal{A}(\mathbf{v})] \geq \frac{\mathsf{E}_{\mathbf{v}\sim\mathbf{F}}[\mathrm{OPT}_{\mathbf{F}}(\mathbf{v})]}{\beta}.$$

Average-case vs Worst-case

Note: prior-free auction cannot be optimal in every setting.

Average Case Approximation: $\exists \mathcal{A}, \forall \mathbf{F} \in \text{IID},$

$$\mathbf{E}_{\mathbf{v} \sim \mathbf{F}}[\mathcal{A}(\mathbf{v})] \geq \frac{\mathbf{E}_{\mathbf{v} \sim \mathbf{F}}[\mathrm{OPT}_{\mathbf{F}}(\mathbf{v})]}{\beta}.$$

Worst Case Approximation: $\exists \mathcal{A}, \forall \mathbf{v},$

$$\mathcal{A}(\mathbf{v}) \geq \frac{\sup_{\mathbf{F} \in \text{IID}} \mathrm{OPT}_{\mathbf{F}}(\mathbf{v})}{\beta}.$$

Average-case vs Worst-case

Note: prior-free auction cannot be optimal in every setting.

Average Case Approximation: $\exists \mathcal{A}, \forall \mathbf{F} \in \text{IID}$,

$$\mathbf{E}_{\mathbf{v}\sim\mathbf{F}}[\mathcal{A}(\mathbf{v})] \geq \frac{\mathbf{E}_{\mathbf{v}\sim\mathbf{F}}[\text{OPT}_{\mathbf{F}}(\mathbf{v})]}{\beta}.$$

Worst Case Approximation: $\exists \mathcal{A}, \forall \mathbf{v}$,

$$\mathcal{A}(\mathbf{v}) \geq \frac{\sup_{\mathbf{F}\in\text{IID}} \text{OPT}_{\mathbf{F}}(\mathbf{v})}{\beta}.$$

Notes:

- worst-case approximation implies average-case approximation.
- $\sup_{\mathbf{F}\in\text{IID}} \text{OPT}_{\mathbf{F}}(\mathbf{v})$ is *prior-free performance benchmark*.
- for digital goods, prior-free benchmark $=$ optimal posted price revenue.

Approximation via Random Sampling

Random Sampling Auction: (for digital goods)

[Goldberg, H, Wright ’01]

1. Randomly partition agents into two sets.

2. Compute optimal posted prices for each set.

3. Offer prices to opposite set.

Approximation via Random Sampling

Random Sampling Auction: (for digital goods)

[Goldberg, H, Wright '01]

1. Randomly partition agents into two sets.

2. Compute optimal posted prices for each set.

3. Offer prices to opposite set.

Thm: Random sampling auction is worst-case 4.68-approximation.*

[Aleai, Malekian, Srinivasan '09]

Approximation via Random Sampling

Random Sampling Auction: (for digital goods)

[Goldberg, H, Wright '01]

1. Randomly partition agents into two sets.
2. Compute optimal posted prices for each set.
3. Offer prices to opposite set.

Thm: Random sampling auction is worst-case 4.68-approximation.*

[Aleai, Malekian, Srinivasan '09]

Conjecture: Random sampling auction is worst-case 4-approximation.

Approximation via Random Sampling

Random Sampling Auction: (for digital goods)

[Goldberg, H, Wright '01]

1. Randomly partition agents into two sets.
2. Compute optimal posted prices for each set.
3. Offer prices to opposite set.

Thm: Random sampling auction is worst-case 4.68-approximation.*

[Aleai, Malekian, Srinivasan '09]

Conjecture: Random sampling auction is worst-case 4-approximation.

Discussion:

- *conclusive*, market analysis can be done "on the fly"
- *worst-case* is for $n = 2$.
- *practical*, bounds approach 1 in limit with n.
- *generic*, analysis extends beyond digital goods.

Extensions

Prior-free results extend to limited supply, downward-closed settings, non-identical distributions, other objectives, etc.

[citations omitted]

Extensions

Prior-free results extend to limited supply, downward-closed settings, non-identical distributions, other objectives, etc.

[citations omitted]

Open Questions:

- non-downward-closed settings?
- multi-dimensional settings?
- beyond the *revelation principle*?

Conclusions

Conclusions:

1. Approximation is more predictive, descriptive, and conclusive than exact optimality.

Conclusions

Conclusions:

1. Approximation is more predictive, descriptive, and conclusive than exact optimality.

2. Key step for approximation: concise description of upper bound.

Conclusions

Conclusions:

1. Approximation is more predictive, descriptive, and conclusive than exact optimality.

2. Key step for approximation: concise description of upper bound.

3. Approximation mechanisms for multi-dimensional and prior-free settings.

Conclusions

Conclusions:

1. Approximation is more predictive, descriptive, and conclusive than exact optimality.

2. Key step for approximation: concise description of upper bound.

3. Approximation mechanisms for multi-dimensional and prior-free settings.

Basic Open Question: attack economic impossibility w. approximation.