Approximation in Mechanism Design

Jason Hartline

Northwestern University

August 8 and 10, 2012

Manuscript available at:

http://www.eecs.northwestern.edu/~hartline/amd.pdf

Goals for Mechanism Design Theory

Mechanism Design: how can a social planner / optimizer achieve objective when participant preferences are private.

Goals for Mechanism Design Theory:

- *Descriptive:* predict/affirm mechanisms arising in practice.
- *Prescriptive:* suggest how good mechanisms can be designed.
- *Conclusive:* pinpoint salient characteristics of good mechanisms.
- *Tractable:* mechanism outcomes can be computed quickly.

Goals for Mechanism Design Theory

Mechanism Design: how can a social planner / optimizer achieve objective when participant preferences are private.

Goals for Mechanism Design Theory:

- *Descriptive:* predict/affirm mechanisms arising in practice.
- *Prescriptive:* suggest how good mechanisms can be designed.
- *Conclusive:* pinpoint salient characteristics of good mechanisms.
- *Tractable:* mechanism outcomes can be computed quickly.

Informal Thesis: *approximately optimality* is often descriptive, prescriptive, conclusive, and tractable.

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

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 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**$[\text{"no prize"}] = 1/2$,

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

[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:

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

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

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

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 $\mathbf{E}[\max]$:

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

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

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

Prophet Inequality Proof

0. Notation:

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

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

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

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

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

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

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 $\textbf{E}[\max]$:

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

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

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

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

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.

0. Notation:

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

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

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

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

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

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

Prophet Inequality Proof

0. Notation:

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

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

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

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

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

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

Prophet Inequality Proof

0. Notation:

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

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

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

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

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

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?

- Constant approximations identify details of model. [cf. Wilson '87]

Philosophy of Approximation

What is the point of a 2-approximation?

- Constant approximations identify details of model. [cf. Wilson '87]
  Example: is X a detail?
  - yes, if constant approx without X
  - no, otherwise.

Philosophy of Approximation

What is the point of a 2-approximation?

- Constant approximations identify details of model. [cf. Wilson '87]
  Example: is X a detail? competition?
  - yes, if constant approx without X
  - no, otherwise.

Philosophy of Approximation

What is the point of a 2-approximation?

- Constant approximations identify details of model. [cf. Wilson '87]
  Example: is X a detail? competition? transfers?
  - yes, if constant approx without X
  - no, otherwise.

Philosophy of Approximation

What is the point of a 2-approximation?

- Constant approximations identify details of model. [cf. Wilson '87]
  Example: is X a detail? competition? transfers?
  - yes, if constant approx without X
  - no, otherwise.

- gives relevant intuition for practice

Philosophy of Approximation

What is the point of a 2-approximation?

- Constant approximations identify details of model. [cf. Wilson '87]
  Example: is X a detail? competition? transfers?
  - yes, if constant approx without X
  - no, otherwise.

- gives relevant intuition for practice

- gives simple, robust solutions.

Philosophy of Approximation

What is the point of a 2-approximation?

- Constant approximations identify details of model. [cf. Wilson ’87]
  Example: is X a detail? competition? transfers?
  - yes, if constant approx without X
  - no, otherwise.
- gives relevant intuition for practice
- gives simple, robust solutions.
- Exact optimization is often impossible.
  (information theoretically, computationally, analytically)

Picasso

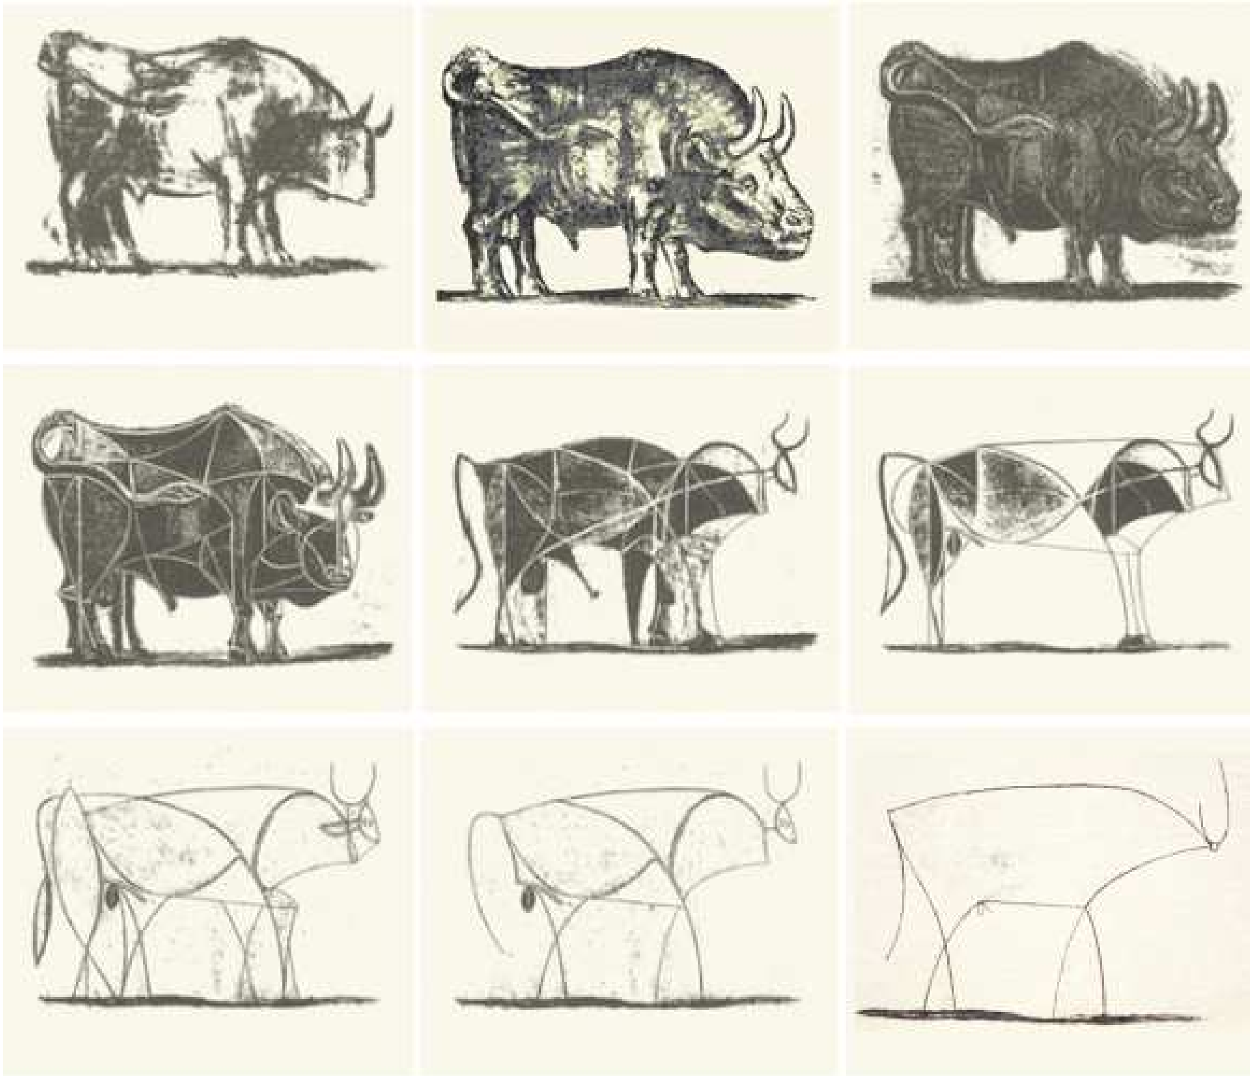

[Picasso's Bull 1945–1946 (one month)]

Questions?

Overview

Part I: Optimal Mechanism Design

- single-item auction.
- objectives: social welfare vs. seller profit.
- characterization of Bayes-Nash equilibrium.
- consequences: solving, uniqueness, and optimizing over BNE.

Part II: Approximation in Mechanism Design

- single-item auctions.
- multi-dimensional auctions.
- prior-independent auctions.
- computationally tractable mechanisms.

Overview

Part I: Optimal Mechanism Design (Chapters 2 & 3)

- single-item auction.
- objectives: social welfare vs. seller profit.
- characterization of Bayes-Nash equilibrium.
- consequences: solving, uniqueness, and optimizing over BNE.

Part II: Approximation in Mechanism Design

- single-item auctions. (Chapter 4)
- multi-dimensional auctions. (Chapter 7)
- prior-independent auctions. (Chapters 5 & 6)
- computationally tractable mechanisms. (Chapter 8)

Part IIa: 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]

1. **Thm:** BNE $\Leftrightarrow$ allocation rule is monotone.

Optimal Auction Design [Myerson '81]

1. **Thm:** BNE $\Leftrightarrow$ allocation rule is monotone.

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

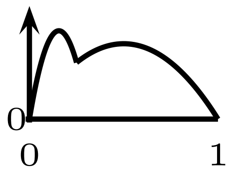

Optimal Auction Design [Myerson '81]

1. **Thm:** BNE $\Leftrightarrow$ allocation rule is monotone.

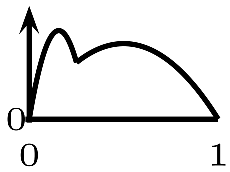

2. **Def:** *revenue curve*: $R_i(q) = q \cdot F_i^{-1}(1 - q)$.

3. **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. **Thm:** BNE $\Leftrightarrow$ allocation rule is monotone.

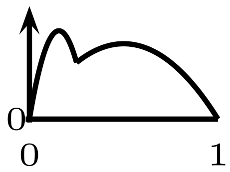

2. **Def:** *revenue curve*: $R_i(q) = q \cdot F_i^{-1}(1-q)$.
3. **Def:** *virtual value*: $\varphi_i(v_i) = v_i - \frac{1-F_i(v)}{f_i(v_i)}$ = marginal revenue.
4. **Def:** *virtual surplus*: virtual value of winner(s).

Optimal Auction Design [Myerson '81]

1. **Thm:** BNE $\Leftrightarrow$ allocation rule is monotone.

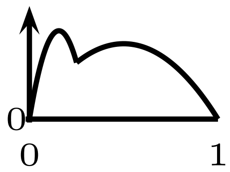

2. **Def:** *revenue curve*: $R_i(q) = q \cdot F_i^{-1}(1-q)$.

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

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

5. **Thm:** **E**[revenue] $=$ **E**[virtual surplus]. (via "revenue equivalence")

Optimal Auction Design [Myerson '81]

1. **Thm:** BNE $\Leftrightarrow$ allocation rule is monotone.

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

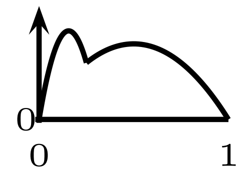

3. **Def:** *virtual value*: $\varphi_i(v_i) = v_i - \frac{1-F_i(v)}{f_i(v_i)}$ = marginal revenue.

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

5. **Thm:** **E**[revenue] = **E**[virtual surplus]. (via "revenue equivalence")

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

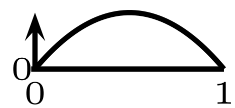

1. **Thm:** BNE $\Leftrightarrow$ allocation rule is monotone.

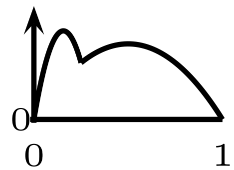

2. **Def:** *revenue curve*: $R_i(q) = q \cdot F_i^{-1}(1-q)$.

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

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

5. **Thm:** $\mathbf{E}$[revenue] $=$ $\mathbf{E}$[virtual surplus]. (via "revenue equivalence")

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

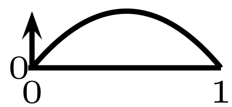

7. **Thm:** for regular dists, optimal auction sells to bidder with highest positive virtual value.

Optimal Auction Design [Myerson '81]

1. **Thm:** BNE $\Leftrightarrow$ allocation rule is monotone.

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

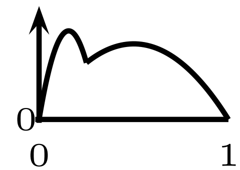

3. **Def:** *virtual value*: $\varphi_i(v_i) = v_i - \frac{1-F_i(v)}{f_i(v_i)}$ = marginal revenue.

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

5. **Thm:** $\mathbf{E}[\text{revenue}] = \mathbf{E}[\text{virtual surplus}]$. (via "revenue equivalence")

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

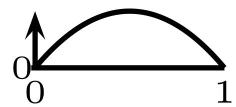

7. **Thm:** for regular dists, optimal auction sells to bidder with highest positive virtual value.

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

Optimal Auctions

Optimal Auctions:

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

Optimal Auctions:

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

Discussion:

- iid, regular case: seems very special.
- general case: optimal auction rarely used. (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: second-price with reserve $=$ *constant virtual price* with **Pr**$[\text{no sale}] = 1/2$ is a 2-approximation.

[Chawla, Hartline, Malec, Sivan '10]

Approximation with reserve prices

Question: when is reserve pricing a good approximation?

Thm: second-price with reserve $=$ *constant virtual price* with **Pr**$[\text{no sale}] = 1/2$ is a 2-approximation.

[Chawla, Hartline, Malec, Sivan '10]

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

Approximation with reserve prices

Question: when is reserve pricing a good approximation?

Thm: second-price with reserve $=$ *constant virtual price* with **Pr**$[$no sale$] = 1/2$ is a 2-approximation.

[Chawla, Hartline, Malec, Sivan '10]

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

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

Approximation with reserve prices

Question: when is reserve pricing a good approximation?

Thm: second-price with reserve $=$ *constant virtual price* with **Pr**$[$no sale$] = 1/2$ is a 2-approximation.

[Chawla, Hartline, Malec, Sivan '10]

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

prophet inequality	second-price with reserves
prizes	virtual values
threshold t	virtual price
E$[$max prize$]$	**E**$[$optimal revenue$]$
E$[$prize for $t]$	**E**$[$second-price 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, collusion fine, etc.

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, second-price with *anonymous reserve price* is 4-approximation. [Hartline, Roughgarden '09]

Anonymous Reserves

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

(e.g., eBay)

Thm: non-identical, regular distributions, second-price with *anonymous reserve price* is 4-approximation. [Hartline, 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, second-price with *anonymous reserve price* is 4-approximation. [Hartline, Roughgarden '09]

Proof: more complicated extension of prophet inequalities.

Discussion:

- theorem is not tight, actual bound is in $[2, 4]$.
- justifies wide prevalence.

Extensions

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

Extensions

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

Thm: non-identical (possibly irregular) distributions, *posted pricing mechanisms* are often constant approximations.

[Chawla, Hartline, Malec, Sivan ’10; Yan ’11]

Extensions

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

Thm: non-identical (possibly irregular) distributions, *posted pricing mechanisms* are often constant approximations.
[Chawla, Hartline, Malec, Sivan '10; Yan '11]

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: non-identical (possibly irregular) distributions, *posted pricing mechanisms* are often constant approximations.

[Chawla, Hartline, Malec, Sivan '10; Yan '11]

Proof technique:

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

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

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

Questions?

Part IIb: 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: 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)	**Problem:** Bayesian Single-item Auction (a.k.a., SD-AUCTION)
• 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.	• 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. item-pricing for $\mathbf{F}$.	**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,Hartline,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.
Proof: prophet inequality (tie-break by "$-p_i$"). [Chawla,Hartline,Malec,Sivan'10]

Multi-item Auctions

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

Multi-item Auctions

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

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

[Chawla, Hartline, Malec, Sivan ’10; Alaei ’11]

Multi-item Auctions

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

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

[Chawla, Hartline, Malec, Sivan ’10; Alaei ’11]

Approach:

1. *Analogy:* “single-dimensional analog”

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

Multi-item Auctions

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

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

[Chawla, Hartline, Malec, Sivan ’10; Alaei ’11]

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)

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

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

[Chawla, Hartline, Malec, Sivan '10; Alaei '11]

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 seq., offer posted prices.

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

[Chawla, Hartline, Malec, Sivan '10; Alaei '11]

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 seq., offer posted prices.

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

[Chawla, Hartline, Malec, Sivan '10; ; Alaei '11]

Sequential Posted Pricing Discussion

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

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

[Chawla, Hartline, Malec, Sivan '10; ; Alaei '11]

Discussion:

- *robust* to agent ordering, collusion, etc.
- *conclusive*:
  - competition not important for approximation.
  - unit-demand incentives similar to single-dimensional incentives.
- *practical*: posted pricings widely prevalent. (e.g., eBay)

Sequential Posted Pricing Discussion

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

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

[Chawla, Hartline, Malec, Sivan '10; ; Alaei '11]

Discussion:

- *robust* to agent ordering, collusion, etc.
- *conclusive*:
  - competition not important for approximation.
  - unit-demand incentives similar to single-dimensional incentives.
- *practical*: posted pricings widely prevalent. (e.g., eBay)

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

- analytically tractable and
- approximable.

Questions?

Part IIc: Approximation for prior-independent mechanism design.

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

The trouble with priors

The trouble with priors:

The trouble with priors

The trouble with priors:

- where does prior come from?

The trouble with priors

The trouble with priors:

- where does prior come from?
- is prior accurate?

The trouble with priors

The trouble with priors:

- where does prior come from?
- is prior accurate?
- prior-dependent mechanisms are non-robust.

The trouble with priors:

- where does prior come from?
- is prior accurate?
- prior-dependent mechanisms are non-robust.
- what if one mechanism must be used in many scenarios?

The trouble with priors

The trouble with priors:

- where does prior come from?
- is prior accurate?
- prior-dependent mechanisms are non-robust.
- what if one mechanism must be used in many scenarios?

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

Optimal Prior-independent Mechs

Optimal Prior-indep. Mech: (a.k.a., non-parametric implementation)

1. agents report value and prior,
2. shoot agents if disagree, otherwise
3. run optimal mechanism for reported prior.

Discussion:

- *complex*, agents must report high-dimensional object.
- *non-robust*, e.g., if agents make mistakes.
- *inconclusive*, begs the question.

Resource augmentation

First Approach: “resource” augmentation.

Resource augmentation

First Approach: “resource” augmentation.

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

[Bulow, Klemperer ’96]

Resource augmentation

First Approach: “resource” augmentation.

Thm: for iid, regular, single-item, the second-price 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-independent strategy.
- “bicriteria” approximation result.
- *conclusive:* competition more important than optimization.

Resource augmentation

First Approach: "resource" augmentation.

Thm: for iid, regular, single-item, the second-price 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-independent strategy.
- "bicriteria" approximation result.
- *conclusive:* competition more important than optimization.
- *non-general*: e.g., for k-unit auctions, need k additional bidders.

Special Case: $n = 1$

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

Special Case: $n = 1$

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

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

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

Special Case: $n = 1$

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

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

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

Special Case: $n = 1$

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

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

- each bidder in second-price views other bid as "random reserve".
- second-price 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 second-price revenue from two bidders is at least the optimal revenue from one bidder.

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

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

Recall: revenue curve

$$R(q) = q \cdot F^{-1}(1-q)$$

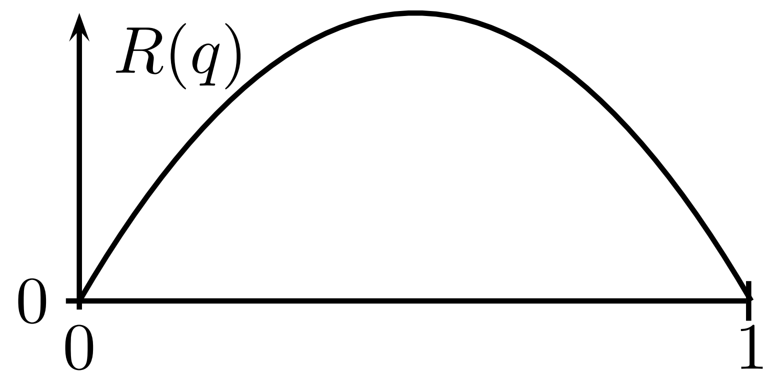

Special Case: $n = 1$

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

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

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

Recall: revenue curve

$$R(q) = q \cdot F^{-1}(1-q)$$

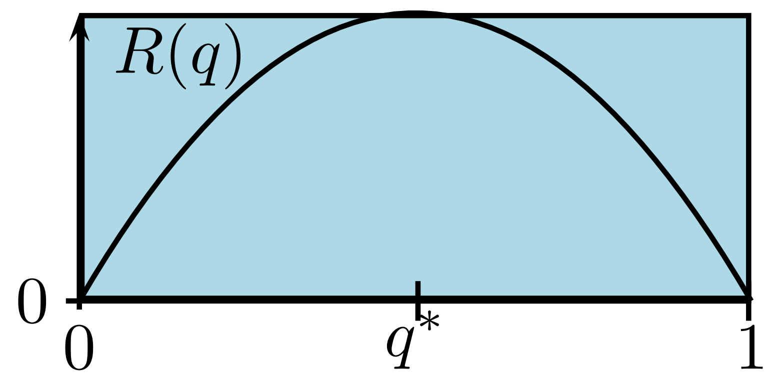

Special Case: $n = 1$

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

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

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

Recall: revenue curve

$$R(q) = q \cdot F^{-1}(1-q)$$

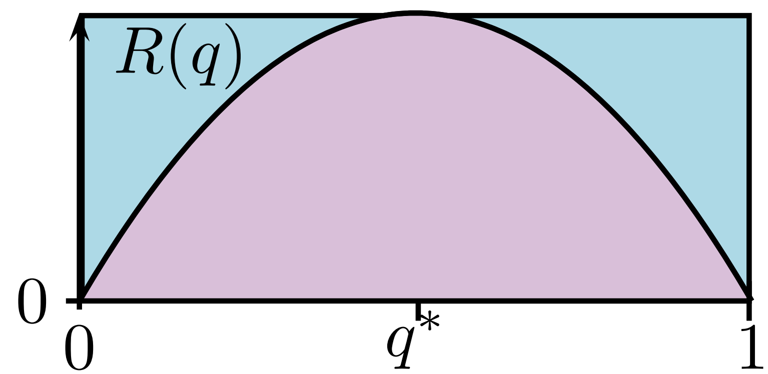

Special Case: $n = 1$

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

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

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

Recall: revenue curve

$$R(q) = q \cdot F^{-1}(1-q)$$

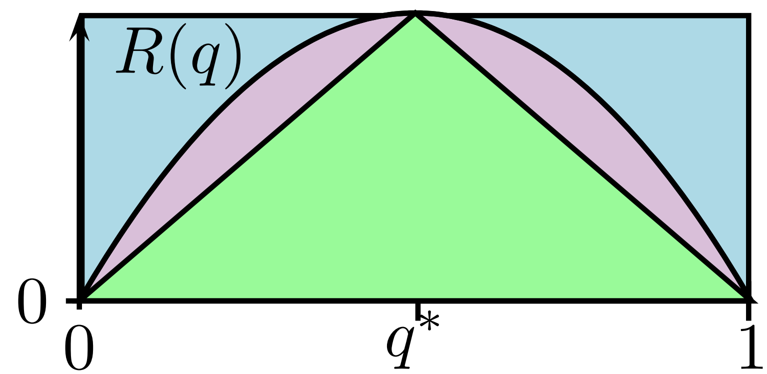

Special Case: $n = 1$

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

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

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

Recall: revenue curve

$$R(q) = q \cdot F^{-1}(1-q)$$

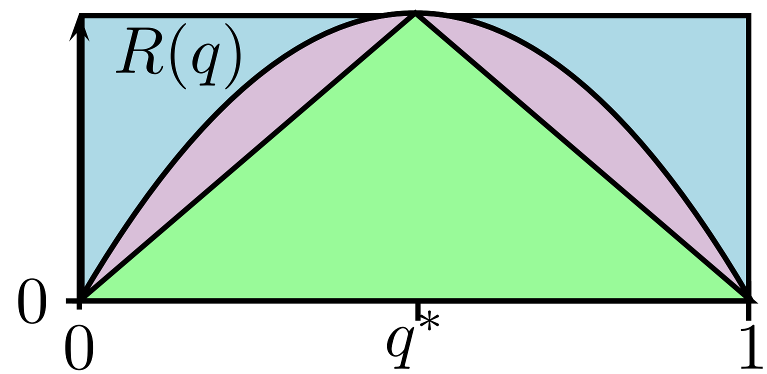

- So second-price on two bidders $\geq$ optimal revenue on one bidder.

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-independent

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-independent*.
- *conclusive*,
  - learn distribution from reports, not cross-reporting.
  - don't need precise distribution, only need single sample for approximation. (more samples can improve approximation/robustness.)
- *generic*, applies to general settings.

Extensions

Recent Extensions:

- non-identical distributions. [Dhangwatnotai, Roughgarden, Yan '10]
- position auctions, matroids, downward-closed environments. [Hartline, Yan '11; Ha, Hartline '11]
- multi-item auctions (multi-dimensional preferences). [Devanur, Hartline, Karlin, Nguyen '11; Roughgarden, Talgam-Cohen, Yan '12]

Extensions

Recent Extensions:

- non-identical distributions. [Dhangwatnotai, Roughgarden, Yan ’10]
- position auctions, matroids, downward-closed environments. [Hartline, Yan ’11; Ha, Hartline ’11]
- multi-item auctions (multi-dimensional preferences). [Devanur, Hartline, Karlin, Nguyen ’11; Roughgarden, Talgam-Cohen, Yan ’12]

Open Question: non-downward-closed environments?

Extensions

Recent Extensions:

- non-identical distributions. [Dhangwatnotai, Roughgarden, Yan '10]
- position auctions, matroids, downward-closed environments. [Hartline, Yan '11; Ha, Hartline '11]
- multi-item auctions (multi-dimensional preferences). [Devanur, Hartline, Karlin, Nguyen '11; Roughgarden, Talgam-Cohen, Yan '12]

Open Question: non-downward-closed environments?

Questions?

Part IId: Computational Tractability in Bayesian mechanism design

(where the optimal mechanism may be computationally intractable)

Problem: Single-minded combinatorial auction

- n agents,
- m items for sale.
- Agent i wants only bundle $S_i \subset \{1, \ldots, m\}$.
- Agent i's value v_i drawn from F_i.

Goal: auction to maximize *social surplus* (a.k.a., welfare).

Example 5: single-minded combinatorial auction

Problem: Single-minded combinatorial auction

- n agents,
- m items for sale.
- Agent i wants only bundle $S_i \subset \{1, \ldots, m\}$.
- Agent i's value v_i drawn from F_i.

Goal: auction to maximize *social surplus* (a.k.a., welfare).

Question: What is optimal mechanism?

Optimal Combinatorial Auction

Optimal Combinatorial Auction: Vickrey-Clarke-Groves (VCG):

1. allocate to maximize reported surplus,
2. charge each agent their “critical value”.

Optimal Combinatorial Auction

Optimal Combinatorial Auction: Vickrey-Clarke-Groves (VCG):

1. allocate to maximize reported surplus,
2. charge each agent their “critical value”.

Discussion:

- distribution is irrelevant (for welfare maximization).
- Step 1 is NP-hard *weighted set packing* problem.
- Cannot replace Step 1 with approximation algorithm.

BNE reduction

Question: Can we convert any algorithm into a mechanism without reducing its social welfare?

BNE reduction

Question: Can we convert any algorithm into a mechanism without reducing its social welfare?

Recall: BNE $\Leftrightarrow$ allocation rule $x_i(v_i)$ is monotone in v_i.

Question: Can we convert any algorithm into a mechanism without reducing its social welfare?

Recall: BNE $\Leftrightarrow$ allocation rule $x_i(v_i)$ is monotone in v_i.

Challenge: $x_i(v_i)$ for alg $\mathcal{A}$ with $\mathbf{v}_{\text{-}i} \sim \mathbf{F}_{\text{-}i}$ may not be monotone.

Question: Can we convert any algorithm into a mechanism without reducing its social welfare?

Recall: BNE $\Leftrightarrow$ allocation rule $x_i(v_i)$ is monotone in v_i.

Challenge: $x_i(v_i)$ for alg $\mathcal{A}$ with $\mathbf{v}_{\text{-}i} \sim \mathbf{F}_{\text{-}i}$ may not be monotone.

Approach:

- Run $\mathcal{A}(\sigma_1(v_1), \ldots, \sigma_n(v_n))$.

BNE reduction

Question: Can we convert any algorithm into a mechanism without reducing its social welfare?

Recall: BNE $\Leftrightarrow$ allocation rule $x_i(v_i)$ is monotone in v_i.

Challenge: $x_i(v_i)$ for alg $\mathcal{A}$ with $\mathbf{v}_{\text{-}i} \sim \mathbf{F}_{\text{-}i}$ may not be monotone.

Approach:

- Run $\mathcal{A}(\sigma_1(v_1), \ldots, \sigma_n(v_n))$.
- σ_i calculated from *max weight matching* on i's type space.

BNE reduction

Question: Can we convert any algorithm into a mechanism without reducing its social welfare?

Recall: BNE $\Leftrightarrow$ allocation rule $x_i(v_i)$ is monotone in v_i.

Challenge: $x_i(v_i)$ for alg $\mathcal{A}$ with $\mathbf{v}_{\text{-}i} \sim \mathbf{F}_{\text{-}i}$ may not be monotone.

Approach:

- Run $\mathcal{A}(\sigma_1(v_1), \ldots, \sigma_n(v_n))$.
- σ_i calculated from *max weight matching* on i's type space.
  - stationary with respect to F_i.
  - $x_i(\sigma_i(v_i))$ monotone.
  - welfare preserved.

Example: σ_i

Example:

$F_i(v_i)$	v_i	$x_i(v_i)$
.25	1	0.1
.25	4	0.5
.25	5	0.4
.25	10	1.0

Example: σ_i

Example:

$F_i(v_i)$	v_i	$x_i(v_i)$	$\sigma_i(v_i)$
.25	1	0.1	1
.25	4	0.5	5
.25	5	0.4	4
.25	10	1.0	10

Example: σ_i

Example:

$F_i(v_i)$	v_i	$x_i(v_i)$	$\sigma_i(v_i)$	$x_i(\sigma_i(v_i))$
.25	1	0.1	1	0.1
.25	4	0.5	5	0.4
.25	5	0.4	4	0.5
.25	10	1.0	10	1.0

Example:

$F_i(v_i)$	v_i	$x_i(v_i)$	$\sigma_i(v_i)$	$x_i(\sigma_i(v_i))$
.25	1	0.1	1	0.1
.25	4	0.5	5	0.4
.25	5	0.4	4	0.5
.25	10	1.0	10	1.0

Note:

- σ_i is from max weight matching between v_i and $x_i(v_i)$.
- σ_i is stationary.
- σ_i (weakly) improves welfare.

BNE reduction discussion

Thm: Any algorithm can be converted into a mechanism with no loss in expected welfare. Runtime is polynomial in size of agent's type space.
[Hartline, Lucier '10; Hartline, Kleinberg, Malekian '11; Bei, Huang '11]

Discussion:

- applies to all algorithms not just worst-case approximations.
- BNE incentive constraints are solved independently.
- works with multi-dimensional preferences too.

Extensions

Extension:

- impossibility for dominant strategy reduction.

  [Chawla, Immorlica, Lucier '12]

Extensions

Extension:

- impossibility for dominant strategy reduction.

[Chawla, Immorlica, Lucier ’12]

Open Questions:

- non-brute-force in type-space? e.g., for product distributions?
- other objectives, e.g., makespan? [Chawla, Immorlica, Lucier ’12]

Extension:

- impossibility for dominant strategy reduction.

[Chawla, Immorlica, Lucier ’12]

Open Questions:

- non-brute-force in type-space? e.g., for product distributions?
- other objectives, e.g., makespan? [Chawla, Immorlica, Lucier ’12]

Questions?

Part II Conclusions

Conclusions:

- approximation pinpoints salient characteristics of good mechanisms.
- reserve-price-based auctions are approximately optimal.
- posted-pricings are approximately optimal.
- good mechanisms can be designed without prior information.
- good algorithms can be converted into good mechanisms.

Part II Conclusions

Conclusions:

- approximation pinpoints salient characteristics of good mechanisms.
- reserve-price-based auctions are approximately optimal.
- posted-pricings are approximately optimal.
- good mechanisms can be designed without prior information.
- good algorithms can be converted into good mechanisms.

Questions?