A biased sample of

Recent Advances and Techniques in Algorithmic Mechanism Design

Part 2: Bayesian Mechanism Design

SHADDIN DUGHMI – MSR REDMOND

BRENDAN LUCIER – MSR NEW ENGLAND

Prologue:

An Introduction to Bayesian Mechanism Design

Bayesian Mechanism Design

Algorithmic Mechanism Design: a central authority wants to achieve a global objective in a computationally feasible way, but participant values/preferences are private.

Bayesian Algorithmic Mechanism Design: If the authority/participants have information about the distribution of private values, does this lead to better mechanisms?

For Example:

Historical market data

Domain-specific knowledge

Presumption of natural inputs

Example: selling a single item

Problem: Single-item auction

1 object to sell

n potential buyers, with values $\boldsymbol{v} = v_1, v_2, \dots, v_n$ for the object.

Buyer objective: maximize utility = value - price

Design Goals:

a) Maximize social welfare (value of winner)

b) Maximize revenue (payment of winner)

Example: selling a single item

Vickrey auction:

Each player makes a bid for the object.

Sell to player with highest bid.

Charge winner an amount equal to the next-highest bid.

Properties:

- Vickrey auction is *dominant strategy truthful*.
- Optimizes social welfare (highest-valued player wins).
- Revenue is equal to the 2nd-highest value.

Example: selling a single item

First-price auction:

Each player makes a bid for the object.

Sell to player with highest bid.

Charge winner an amount equal to *his own* bid.

First-price auction is *not truthful*.

How should players bid? What is "rational"?

How much social welfare is generated?

How much revenue is generated?

Bayes-Nash Equilibrium

Bayesian Setting: buyer values are drawn independently from a known product distribution $\boldsymbol{F} = F_1 \times F_2 \times \cdots \times F_n$.

Players bid to maximize expected utility, given distribution $\boldsymbol{F}$.

Definition: a *strategy* s maps values to bids: $b = s(v)$.

A strategy profile $\boldsymbol{s} = (s_1, s_2, \dots, s_n)$ is a Bayes-Nash equilibrium for distribution $\boldsymbol{F}$ if, for each i and v_i, $s_i(v_i)$ maximizes the expected utility of player i, given that others play $\boldsymbol{s}$ and $\boldsymbol{v} \sim \boldsymbol{F}$.

$$E_{v\sim F}[u_i(s_i(v_i), s_{-i}(v_{-i})) \mid v_i]$$

First-Price Auction: Equilibria

Example: First-price auction, two bidders, values iid from U[0,1].

Claim: strategy $s(v) = \frac{v}{2}$ is a symmetric Bayes-Nash equilibrium.

Proof: Suppose player 1 plays $s_1(v_1) = \frac{v_1}{2}$.

How should player 2 bid, given his value v_2?

E[2's utility] = $(v_2 - b_2) \times \Pr[b_2 > b_1]$

$= (v_2 - b_2) \times \Pr\left[b_2 > \frac{v_1}{2}\right]$

$= (v_2 - b_2) \times 2\mathrm{b}_2$

$= 2(v_2 b_2 - b_2^2)$

Take derivative with respect to b_2 and set to 0. Solution is $b_2 = \frac{v_2}{2}$, so $s(v_2) = \frac{v_2}{2}$ is utility-maximizing.

First-Price Auction: Equilibria

Example: First-price auction, two bidders, values iid from U[0,1].

Claim: strategy $s(v) = \frac{v}{2}$ is a symmetric Bayes-Nash equilibrium.

Corollary 1: Player with highest value always wins, so the first-price auction maximizes social welfare.

Corollary 2:

$$\text{Expected revenue} = \frac{1}{2} \times E[\max\{v_1, v_2\}] = \frac{1}{2} \times \frac{2}{3} = \frac{1}{3}$$

Note: same social welfare and revenue as the Vickrey auction!

Characterization of BNE

Notation: Suppose that players are playing strategy profile $\boldsymbol{s}$.

$x_i(v_i)$ - probability of allocating to bidder i when he declares v_i

$p_i(v_i)$ - expected payment of bidder i when he declares v_i

where expectations are with respect to the distribution of others' values.

Theorem [Myerson'81]: For single-parameter agents, a mechanism and strategy profile are in BNE iff:

a) $x_i(v_i)$ is monotone non-decreasing,

b) $p_i(v_i) = v_i x_i(v_i) - \int_0^{v_i} x_i(z)dz + p_i(0)$ (normally $p_i(0) = 0$)

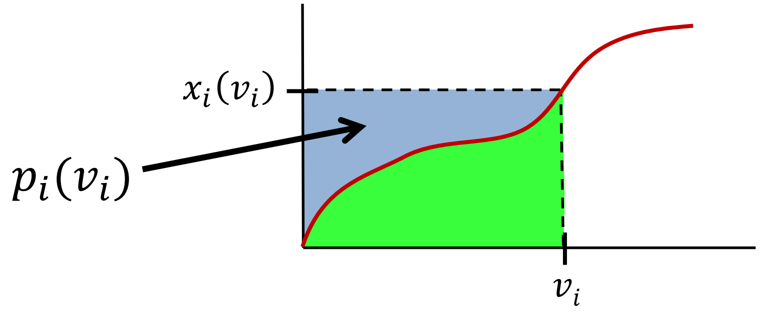

Implication (Revenue Equivalence): Two mechanisms that implement the same allocation rule at equilibrium will generate the same revenue.

Bayesian Truthfulness

How should we define truthfulness in a Bayesian setting?

Bayesian incentive compatibility (BIC): every agent maximizes his expected utility by declaring his value truthfully.

- Expectation is over the distribution of other agents' values, as well as any randomization in the mechanism.

That is, a mechanism is BIC for distribution $\boldsymbol{F}$ if the truth-telling strategy $s(v) = v$ is a Bayes-Nash equilibrium.

Prior-Independent Mechanisms

In general, a mechanism can explicitly depend on distribution $\boldsymbol{F}$.

However, the mechanisms is then tied to this distribution.

- What if we want to reuse the mechanism in another setting?
- What if F is unavailable / incorrect / changing over time?

Prior-Independent Mechanism: does not explicitly use F to determine allocation or payments.

Desirable in practice: robust, can be deployed in multiple settings, possible when prior distribution is not known.

Big Research Questions

For a given interesting/complex/realistic mechanism design setting, can we:

1. Construct computationally feasible BIC mechanisms that (approximately) maximize social welfare?
2. Describe/compute/approximate the revenue-optimal auction?
3. Show that simple/natural mechanisms generate high social welfare and/or revenue at equilibrium?
4. Design prior-independent mechanisms that approximately optimize revenue for every distribution?
5. Extend the above to handle budgets, online arrivals, correlations, …?

Outline

Intro to Bayesian Mechanism Design

Social Welfare and Bayesian Mechanisms

- Truthful Reductions and Social Welfare
- Designing mechanisms for equilibrium performance

Revenue and Bayesian Mechanisms

- Introduction to Revenue Optimization
- Prophet inequality and simple mechanisms
- Prior-independent mechanism design

Part 1:

Truthful Reductions and Social Welfare

Bayesian Truthfulness

One lesson from the first part of the tutorial:

- Many approximation algorithms are not dominant strategy truthful.
- Designing a dominant strategy truthful mechanism is complicated!

Question: Is the problem of designing truthful algorithms easier in the Bayesian setting?

The dream: a general method for converting an arbitrary *approximation* algorithm for social welfare into a BIC mechanism.

This section: such transformations are possible in the Bayesian setting! (And are not possible for IC in the prior-free setting.)

Example

Problem: Single-Parameter Combinatorial Auction

Set of m objects for sale
n buyers
Buyer i wants bundle $S_i \subseteq \{1,2,\dots,m\}$, known in advance
Buyer i's value for S_i is v_i, drawn from distribution F_i

Goal: maximize social welfare.

Possible Solution: VCG Mechanism

- Allocate optimal solution, charge agents their externalities.
- Problem: NP-hard to find optimal solution (set packing).
- Can't plug in an approximate solution – no longer truthful!

What about Bayesian truthfulness?

Bayesian Incentive Compatibility

Recall: $x_i(v_i)$ - probability of allocating to bidder i when he declares v_i.
$p_i(v_i)$ - expected payment of bidder i when he declares v_i.

Theorem [Myerson'81]: A single-parameter mechanism is BIC iff:
a) $x_i(v_i)$ is monotone non-decreasing, and
b) $p_i(v_i) = v_i x_i(v_i) - \int_0^{v_i} x_i(z)dz$

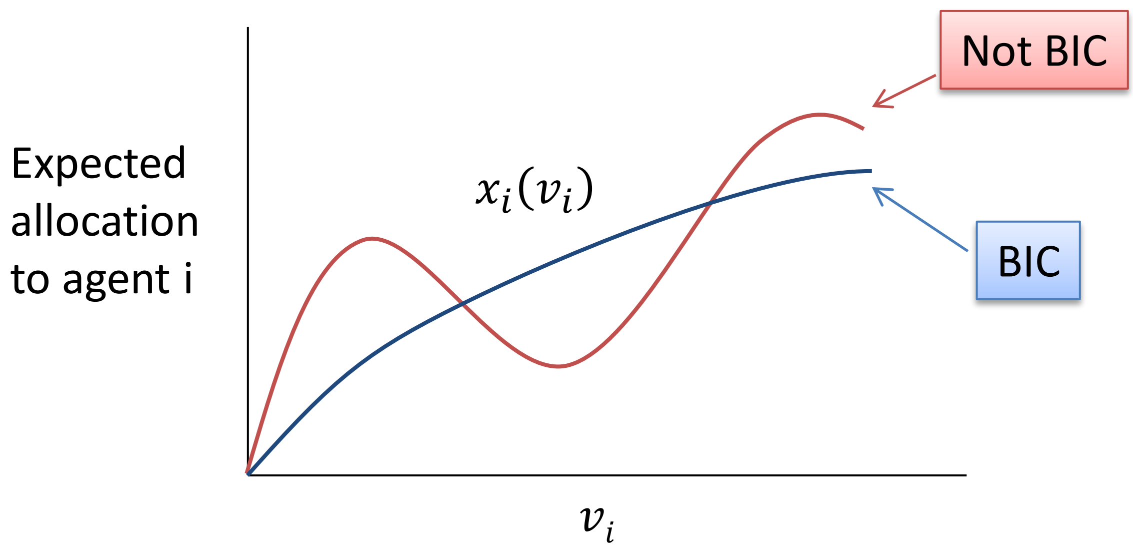

Conclusion: To convert an algorithm into a BIC mechanism, we must monotonize its allocation curves. (Given monotone curves, the prices are determined).

Monotonizing Allocation Rules

Example:

Focus on a single agent i. v_i is either 1 or 2, with equal probability.
Some algorithm A has the following allocation rule for agent i:

v_i	$\Pr[v_i]$	$x_i(v_i)$	$\sigma(v_i)$	$x_i(\sigma(v_i))$
1	0.5	0.7	2	0.3
2	0.5	0.3	1	0.7

Note: $x_i(\cdot)$ is non-monotone, so our algorithm is not BIC.

Idea: we would like to swap the expected outcomes for $v_i = 1$ and $v_i = 2$, without completely rewriting the algorithm.

How to do it: whenever player i declares $v_i = 1$, "pretend" that he reported $v_i = 2$, and vice-versa. Pass the permuted value (say $\sigma(v_i)$) to the original algorithm.

Possible problem: maybe this alters the algorithm for the other players?

No! Other agents only care about the distribution of v_i, which hasn't changed!

Monotonizing Allocation Curves

More Generally:

Focus on each agent i separately.
Suppose there is a finite set V of possible values for i, all equally likely.

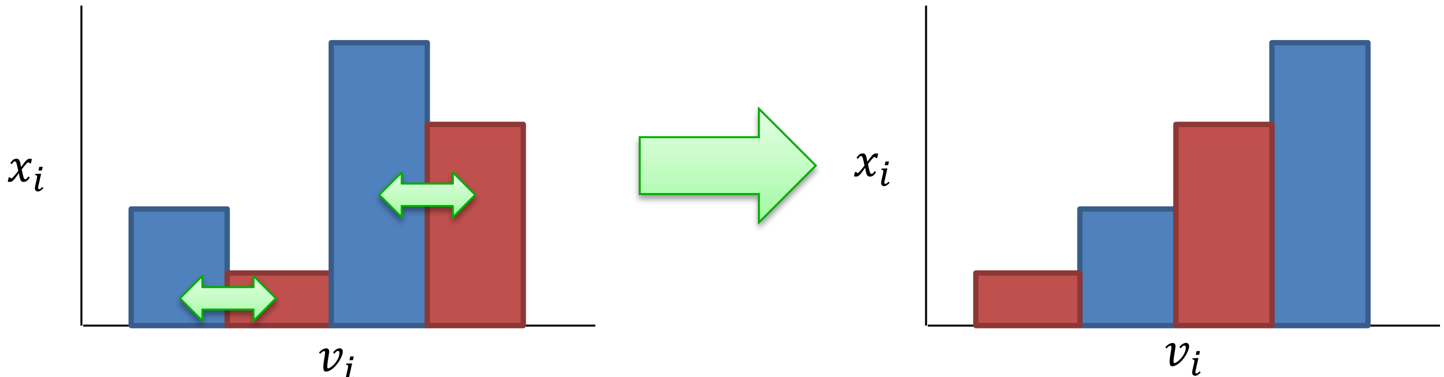

Idea: permute the values of V so that $x_i(\cdot)$ is non-decreasing.
Let this permutation be σ_i.
On input $(v_1, v_2, \dots, v_n)$, return $A(\sigma_1(v_1), \sigma_2(v_2), \dots, \sigma_n(v_n))$.

Claim: This transformation can only increase the social welfare.
Also, since all v_i are equally likely, F_i is stationary under σ_i. So other agents are unaffected, and we can apply this operation to each agent independently!

Monotonizing Allocation Curves

Theorem: Any algorithm can be converted into a BIC mechanism with no loss in expected welfare. Runtime is polynomial in size of each agent's type space.

[Hartline, L. '10, Hartline, Kleinberg, Malekian '11, Bei, Huang'11]

- Applies to general (multi-dimensional) type spaces as well!
- Works for algorithms tailored to the distribution, not just worst-case approximations.
- If agent values aren't all equally likely, or if the allocation rules aren't fully specified (algorithm is black-box), can approximate by sampling.
- For continuous types, number of samples needed (and hence runtime) depends on dimension of type space.

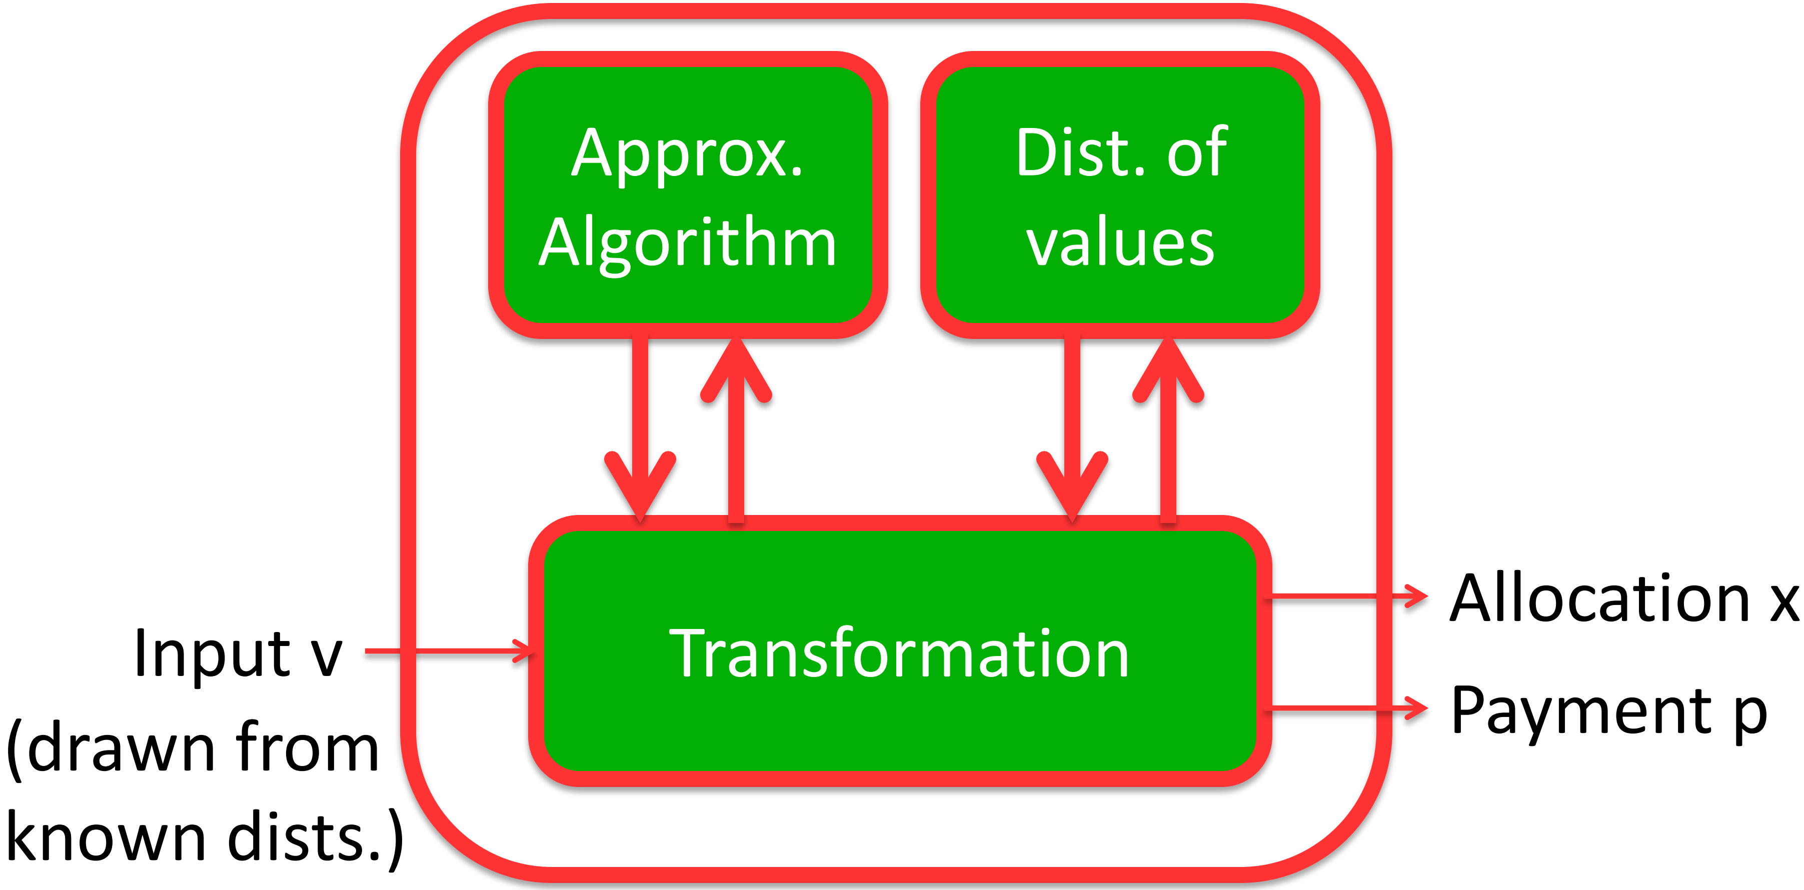

We can view this mechanism construction as a *black-box transformation* that converts arbitrary algorithms into mechanisms.

Extensions

- Impossibility of general lossless black-box reductions when the social objective is to minimize makespan.

  [Chawla, Immorlica, L. '12]

- Impossibility of general lossless black-box truthful-in-expectation reductions for social welfare in prior-free setting.

  [Chawla, Immorlica, L. '12]

Open:

More efficient methods when type space is very large, or continuous with high dimension?

Part 2:

Simple Mechanisms and the Price of Anarchy

Example

Problem: k-Size Combinatorial Auction

Set of m objects for sale
n buyers
Buyer i has a value for each bundle $S \subseteq \{1, \dots, m\}$ of *size at most k*
Specified by a valuation function: $v_i(S)$
Valuation function v_i drawn from distribution F_i

Goal: maximize social welfare.

Possible Solution 1: VCG Mechanism

- Problem: NP-hard to find optimal solution (set packing).

Possible Solution 2: BIC Reduction

- Type space has high dimension. Exponential runtime in general.
- Construction is specific to the prior distribution $\boldsymbol{F}$

Question: is there a simple, prior-independent mechanism that approximates social welfare, if we don't insist on Bayesian truthfulness?

A Simple Approximation

Greedy algorithm:

- Allocate sets greedily from highest bid value to lowest.
  - Assumes either succinct representation of valuation functions or appropriate query access.

Notes:

Recall: sets of size at most k

- Worst-case (k+1)-approximation to the social welfare
- *Not truthful* (with any payment scheme)

Question: how well does the greedy algorithm perform as a mechanism?

A Greedy Mechanism

Greedy first-price mechanism:

- Elicit bid functions $b_1, \dots, b_n$ from the players
- Allocate sets greedily from highest bid value to lowest.
- Each winning bidder pays his bid for the set received.
  - If player i wins set A_i, he pays $p_i = b_i(A_i)$.

Notes:

- Greedy mechanism is *prior-independent*.
- Since the mechanism is not truthful, we would like to maximize the social welfare at *every* BNE, for *every* prior distribution F.
  - In other words: we want to bound the *Bayesian Price of Anarchy*
- **Important caveat**: unlike truthfulness, the burden of finding/computing an equilibrium is shifted to the agents.

Analysis

Claim: For any $\boldsymbol{F}$, the social welfare of any BNE of the greedy first-price mechanism is a (k+2) approximation to the optimal expected social welfare.

Main idea: (shared by many similar proofs)

- Choose some $\boldsymbol{F}$ and a Bayes Nash equilibrium of the mechanism.

- Consider a deviation by one player aimed at winning a valuable set.
  1. Either this deviation "succeeds" and a high-valued set was won, resulting in high utility…
  2. …or it fails, because it was "blocked" by another player's bid.

- But the player can't increase utility by deviating (equilibrium)!
- So either (2) occurs often (blocking player has high value) or the player's utility was already high (deviating player has high value).

- Summing up over players, and taking expectation over types, we conclude that the total welfare must be large.

Notes

Conclusion: the “natural” greedy algorithm performs almost as well at BNE as it does when agents simply report their true values.

Theorem: For any combinatorial auction problem that allows single-minded bids, a β-approximate greedy algorithm with first-price payments obtains a $(\beta + o(1))$ approximation to the social welfare at every BNE.

[L.,Borodin’10]

Another natural payment method: critical prices

- If a bidder wins set S, he pays the smallest amount he could have declared for set S and still won it.
- A similar analysis holds for critical prices (with a slightly different bound, and some additional assumptions).

Related Work

Combinatorial auctions via independent item bidding.

[Christodoulou, Kovács, Schapira ’08, Bhawalkar, Roughgarden ’11, Hassidim,Kaplan,Mansour,Nisan’11]

Analysis of Generalized Second-Price auction for Sponsored Search.

[Paes Leme,Tardos’10, L.,Paes Leme’11, Caragiannis,Kaklamanis,Kanellopoulos,Kyropoulou‘11]

Price of anarchy of sequential auctions.

[Paes Leme, Syrgkanis, Tardos’12, Syrgkanis’12]

A general “smoothness” argument for analyzing Bayesian Price of Anarchy.

[Roughgarden ’12, Syrgkanis’12]

Interlude:
Intro to Revenue Maximization

Selling a single item, Revisited

Problem: Single-item auction

1 object to sell

n buyers

Value for buyer *i* is v_i drawn from distribution F_i.

Goal: Maximize revenue

What is the optimal mechanism?

Characterization of BNE

Recall:

Theorem [Myerson'81]: A single-parameter mechanism and strategy profile are in BNE if and only if:

a) $x_i(v_i)$ is monotone non-decreasing,

b) $p_i(v_i) = v_i x_i(v_i) - \int_0^{v_i} x_i(z)dz$

Solution 1: Write out the incentive compatibility constraints, apply Myerson's characterization, express as an LP, and solve.

But: not very informative; may not be able to solve efficiently in general.

Virtual Value

Notation: when value v drawn from distribution F, we write

$F(z) = \Pr[v \leq z]$, the *cumulative distribution function*

$f(z) = dF(z)/dz$, the *probability density function*

Myerson's Lemma: In BNE, $E[\sum_i p_i(v_i)] = E[\sum_i \phi_i(v_i) x_i(v_i)]$

Where $\phi_i(v_i)$ is the *virtual value function:*

Hazard Rate

$$\phi_i(v_i) = v_i - \boxed{\frac{1 - F_i(v_i)}{f_i(v_i)}}$$

Proof: Write expectation as an integration over payment densities, apply Myerson characterization of payments, and simplify.

Virtual Value

Myerson's Lemma: In BNE, $E[\sum_i p_i(v_i)] = E[\sum_i \phi_i(v_i) x_i(v_i)]$

Expected revenue is equal to *expected virtual welfare*.

Idea: to maximize revenue, allocate to the player with highest *virtual* value.

Problem: if function ϕ_i is not monotone, then allocating to the player maximizing $\phi_i(v_i)$ may not be a monotone allocation rule.

Solution: restrict attention to cases where ϕ_i is monotone.

Definition: distribution F is *regular* if its virtual valuation function ϕ is monotone.

Myerson's Auction

Theorem: If each F_i is regular, the revenue-optimal auction allocates to the bidder with the highest positive virtual value.

Example: Agents are i.i.d. regular, distribution F.

- All players have the same virtual value function ϕ.
- If all virtual values are negative, no winner.
- Otherwise, winner is player with maximum $\phi(v_i)$.
- Since F is regular, this is the player with maximum v_i.

Conclusion: For iid regular bidders, Myerson optimal auction is the Vickrey auction with *reserve price* $r = \phi^{-1}(0)$.

Natural and straightforward to implement!

Multi-parameter Settings

The Myerson optimal auction (i.e. maximize virtual surplus) extends to all single-parameter mechanism design problems.

Our understanding of the revenue-optimal auction for multi-parameter settings is far less complete.

Recent developments: computability of the revenue-optimal auction (for a given F) for certain multi-parameter auction problems.

[Cai,Daskalakis,Weinberg'12,Daskalakis,Weinberg'12, Alaei,Fu,Haghpanah,Hartline,Malekian'12]

Part 3:

Revenue, Prophet Inequalities, and Simple Mechanisms

Example

Myerson's Auction: A non-identical example:

Two bidders, not identical: v_1~U[0,2], v_2~U[0,3].

$$\phi_1(v_1) = v_1 - \frac{1 - F_1(v_1)}{f_1(v_1)} = v_1 - \frac{1 - (v_1/2)}{1/2} = 2v_1 - 2$$

$$\phi_2(v_2) = v_2 - \frac{1 - F_2(v_2)}{f_2(v_2)} = v_2 - \frac{1 - (v_2/3)}{1/3} = 2v_2 - 3$$

Myerson Optimal Auction:

Player 1 wins if $\phi_1(v_1) > \max\{\phi_2(v_2), 0\}$, i.e. $v_1 > 1$ and $v_1 > v_2 - \frac{1}{2}$

Player 2 wins if $\phi_2(v_2) > \max\{\phi_1(v_1), 0\}$, i.e. $v_2 > \frac{3}{2}$ and $v_2 > v_1 + \frac{1}{2}$

Seems overly complex. How well could we do with a simpler auction?

A Simpler Auction

Vickrey Auction with Reserves:

Offer each bidder a reserve price r_i

Sell to highest bidder who meets his reserve.

Question: How much revenue do we lose by using a Vickrey auction rather than the optimal (Myerson) auction?

Informal Theorem: In many settings, revenue is within a constant factor of the optimal.

[Hartline, Roughgarden'09, Chawla, Hartline, Malec, Sivan'10]

Monotone Hazard Rate

Recall: $\phi_i(v_i) = v_i - \frac{1-F_i(v_i)}{f_i(v_i)}$

F_i is regular if $\phi_i(v_i)$ is non-decreasing.

F_i satisfies the Monotone Hazard Rate assumption (MHR) if $\frac{1-F_i(v_i)}{f_i(v_i)}$ is non-increasing.

Lemma: if F_i is MHR, and
$r = \phi^{-1}(0)$ is the Myerson reserve,
then $v \leq \phi(v) + r$ for all $v \geq r$.

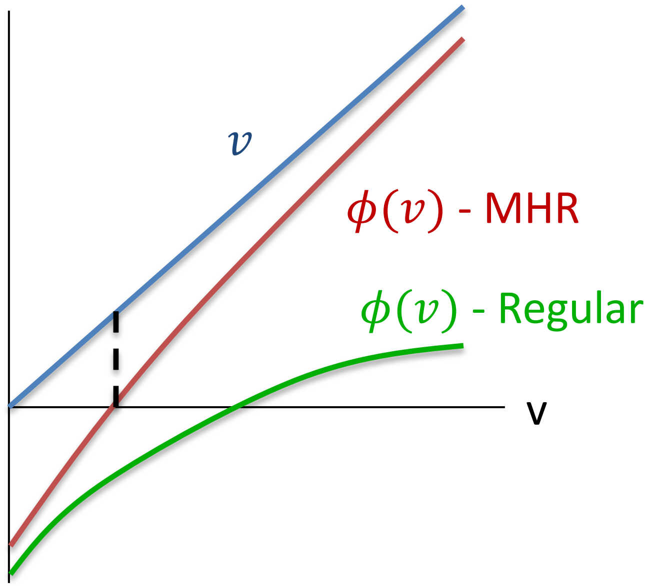

Monotone Hazard Rate

Theorem: If all F_i satisfy MHR, then the revenue of the Vickrey auction with reserves $r_i = \phi_i^{-1}(0)$ is a 2-approximation to the optimal revenue.

[Hartline, Roughgarden'09]

Proof: $\boldsymbol{x}(\boldsymbol{v})$, $R(\boldsymbol{v})$ – allocation rule / revenue of Vickrey auction.
$\boldsymbol{x}^*(\boldsymbol{v})$, $R^*(\boldsymbol{v})$ – allocation rule / revenue of Myerson auction.

By Myerson's Lemma: $E[R(\boldsymbol{v})] = E[\sum_i \phi_i(v_i) x_i(v_i)]$
Winners in Vickrey pay at least their reserve: $E[R(\boldsymbol{v})] \geq E[\sum_i r_i x_i(v_i)]$
So $2E[R(\boldsymbol{v})] \geq E[\sum_i (r_i + \phi_i(v_i)) x_i(v_i)]$

$\geq E[\sum_i v_i x_i(v_i)]$ (MHR)

$\geq E[\sum_i v_i x_i^*(v_i)]$ (Vickrey SW > Myerson SW)

$\geq E[R^*(\boldsymbol{v})]$ (Myerson SW > Myseron Rev)

Aside: Prophet inequality

A Gambling Game:
n prizes $z_1, \dots, z_n$, each prize chosen from distribution F_i
Prizes revealed to the gambler one at a time.
After prize i is revealed, the gambler must either

accept prize z_i and leave the game, or
abandon prize z_i permanently and continue.

Goal: maximize value of prize accepted

Optimal strategy: backward induction.
Simple strategy: pick threshold t, accept first prize with value at least t.

Theorem [Prophet Inequality]: Choosing t such that Pr[accept any prize] = ½ yields expected winnings at least $\frac{1}{2} \max_i z_i$.

[Samuel,Cahn'84]

Prophet inequality

Vickrey Auction with Prophet Reserves:

For n bidders and regular distributions, choose a value R and set all reserves equal to $r_i = \phi_i^{-1}(R)$.

Theorem: If R is chosen so that Pr[no sale] = 1/2, then the Vickrey auction with reserve prices $r_1, r_2, \dots, r_n$ obtains a 2-approximation to the optimal revenue.

[Chawla, Hartline, Malec, Sivan '10]

Proof: Direct application of Prophet inequality.

Our problem: choose threshold R, so that arbitrary virtual value $\geq R$ is a good approximation to the maximum virtual value.

Prophet inequality: choose threshold t, so that first prize $\geq t$ is a good approximation to the maximum prize.

Other applications

Theorem: Single-item auction with anonymous reserve and selling to max-valued bidder yields a 4-approximation to the optimal revenue.

[Hartline,Roughgarden'09]

Theorem: GSP auction with bidder values drawn i.i.d. from a regular distribution, with appropriate reserve, is a 6-approximation of optimal revenue at any BNE.

[L.,Paes Leme,Tardos'12]

Selling Multiple Items

Problem: Unit-Demand Pricing

n objects to sell.
1 buyer, wants at most one item.
Value for item i is $v_i \sim F_i$

Problem: Single-Item Auction

1 object to sell.
n buyers.
Value of bidder i is $v_i \sim F_i$

Goal: Set Prices to Maximize revenue

- For single-item auction, Vickrey with “prophet inequality” reserves gives a ½ approximation to optimal revenue.
- Structurally the problems are very similar. Can we apply similar techniques to the unit-demand auction?

Prophet Inequality Again

Theorem: Setting prophet reserve prices in the unit-demand pricing problem gives a 2-approximation to optimal revenue.

[Chawla, Hartline, Malec, Sivan'10]

Proof Sketch: Compare with single-item auction.

- Imagine splitting the single multi-demand bidder into multiple single-parameter agents, one per item, but can only serve one.
- **Claim**: Optimal revenue in single-item auction $\geq$ Optimal revenue in unit-demand pricing. (Why? Increased competition!)
- **Claim**: Revenue for unit-demand pricing with prophet reserves is at least half of optimal revenue for single-item auction.
  - Analysis same as for single-item auction!

Extending to Multiple Bidders

Unit-demand Auction Problem:

n agents, m items. Each agent wants at most one item.

Agent i has value $v_{ij} \sim F_{ij}$ for item j

Goal: maximize revenue.

Sequential Posted Price Mechanism:

- Agents arrive in (possibly arbitrary) sequence
- Offer each agent a list of prices for the items
- Each agent chooses his utility-maximizing item

Extending to Multiple Bidders

Theorem (Informal): In the unit-demand setting with values drawn independently for bidders and items, for various settings, a sequential posted price mechanism obtains a constant approximation to the optimal revenue.

[Chawla, Hartline, Malec, Sivan'10]

Proof: similar to the single-bidder pricing problem.

Take-away: setting high prices in accordance with the prophet inequality reduces competition, thereby simplifying analysis.

Extensions

Multi-unit auctions with budget-constrained agents.

[Chawla, Malec, Malekian'11]

General reductions from multi-parameter auctions to single-agent pricing problems.

[Alaei'11]

Future Work:

Extend the class of multi-parameter auctions for which we can obtain constant-factor approximations to revenue.

Part 4:
Prior-Independent Revenue Maximization

Priors vs. Additional Bidders

Question: How useful is knowing the prior distribution?

Theorem: for iid, regular, single-item auctions, the Vickrey auction on $n+1$ bidders (and no reserve) generates higher expected revenue than the optimal auction on n bidders.

[Bulow, Klemperer'96]

If the mechanism designer doesn't have access to prior distribution, he can do just as well by recruiting one more bidder.

Special Case: 1 Bidder

Theorem: The Vickrey auction with 2 bidders generates at least as the optimal revenue from a single bidder, for regular distributions.

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

For single bidder, consider Revenue as a function of probability of sale.

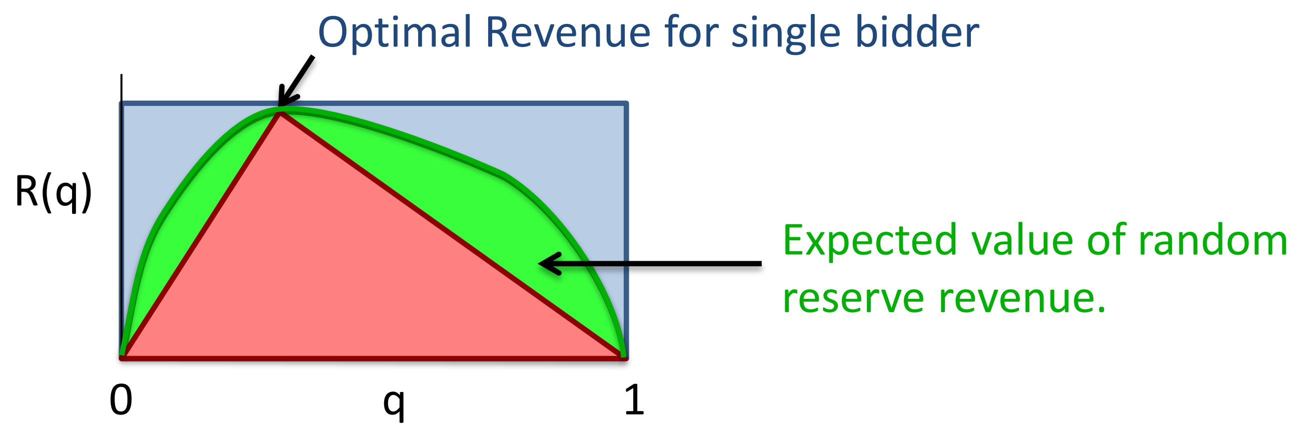

- Vickrey auction: each bidder views the other as a randomized reserve.
- Vickrey revenue = 2 x E[random reserve revenue]
- E[random reserve revenue] ≥ ½ optimal reserve revenue

Example: Digital Goods

Problem: Digital Goods

n identical objects to sell, n buyers.

Each buyer wants at most one object.

Each buyer has value $v_i \sim F$.

Goal: Maximize revenue

Optimal auction: Offer each agent Myerson reserve $\phi^{-1}(0)$.

How well can we do with a prior-independent mechanism?

Example: Digital Goods

Single-Sample Mechanism:

1. Pick an agent i at random
2. Offer every other agent price v_i
3. Do not sell to agent i

Theorem: For iid, regular distributions, the single sample auction with $n + 1$ bidders is a 2-approximation to the optimal revenue with n bidders.

[Dhangwatnotai, Roughgarden, Yan'10]

Proof: Follows from the geometric argument for n=1.

Further Work

- Non-identical distributions [Dhangwatnotai, Roughgarden, Yan'10]
- Online Auctions [Babaioff, Dughmi, Kleinberg, Slivkins'12]
- Matroids, other complex feasibility constraints [Hartline, Yan'11]
- Alternative approach: Limited-Supply Mechanisms [Roughgarden, Talgam-Cohen, Yan'12]

Summary

- We surveyed recent results in Bayesian mechanism design.
- Social Welfare:
  - General transformations from approximation algorithms to BIC mechanisms.
  - Mechanisms with simple greedy allocation rules tend to have good social welfare at Bayes-Nash equilibria.
- Revenue:
  - Optimal auctions tend to be complex; simple auctions can often obtain constant approximation factors (even in multi-parameter settings).
  - It is sometimes possible to approximate the optimal revenue with a prior-independent mechanism, e.g. via sampling techniques.