Approximately Optimal Mechanisms:

Motivation, Examples, and Lessons Learned

Tim Roughgarden (Stanford)

Outline

1. *The optimal and approximately optimal mechanism design paradigms: Vickrey, Myerson, and beyond.*

2. Case study #1: do good single-item auctions require detailed distributional knowledge?

3. Case study #2: do good combinatorial auctions require complex bid spaces?

4. Conclusions

Example: The Vickrey Auction

Setup: Single-item auction, welfare-maximization.

- one seller with one item
- n bidders, bidder i has private valuation v_i
- goal: maximize welfare (i.e., winner = largest v_i)
- [Vickrey 61] solution = second-price auction
  - winner = highest bidder
  - price = 2^{nd}-highest bid
  - bidders follow dominant strategies => maximizes welfare

Looking ahead: what about selling multiple items?

Example: Myerson's Auction

Setup: Single-item auction, revenue-maximization.

- private valuations v_i drawn i.i.d. from known prior F
- goal: maximize seller's expected revenue
- [Myerson 81] solution = 2nd-price auction + reserve
  - reserve price r = monopoly price for F [i.e., $\text{argmax}_p \, p(1-F(p))$]
  - winner = highest bidder above r (if any)
  - price = maximum of r and 2^{nd}-highest bid

Looking ahead: what about heterogeneous bidders?

The Optimal Mechanism Design Paradigm

- define mechanism design space
  - e.g., sealed-bid auctions
- define desired properties
  - e.g., max welfare (ex post) or expected revenue (w.r.t. F)
- identify one or all mechanisms with properties

} how

- identify specific mechanisms or mechanism features that are potentially useful in practice
- optimal mechanism can serve as benchmark for comparing different "second-best" solutions

} why

The *Approximately* Optimal Mechanism Design Paradigm

- define design space, objective function
  - e.g., limited distributional knowledge; low-dimensional bid spaces
- define a *benchmark*
  - e.g., max welfare/revenue of an arbitrarily complex mechanism
- identify mechanisms that approximate benchmark

how

- identify specific mechanisms or mechanism features that are potentially useful in practice
- quantify cost of side constraints (e.g., "simplicity")
  - e.g., complex bids required⇔ best approx far from 100%

why

Two Case Studies

Case Study #1: revenue-maximization, non-i.i.d. bidders.

- Issue: Myerson's optimal auction requires detailed knowledge of valuation distributions.
- Question: is this essential for near-optimal revenue?

Case Study #2: welfare-maximization, multiple items.

- Issue: direct revelation (as in VCG) requires a complex bidding space (exponential in # of items).
- Question: is this essential for near-optimal welfare?

Many Applications of the Approximation Paradigm

- limited communication [Nisan/Segal 06], …
- limited computation [Lehmann/O'Callaghan/Shoham 99], [Nisan/Ronen 99], …
- unknown prior [Neeman 03], [Baliga/Vohra 03], [Segal 03], [Dhangwatnotai/Roughgarden/Yan 10], …
- worst-case revenue guarantees [Goldberg/Hartline/Karlin/Saks/Wright 06], …
- simple allocation rules [Chawla/Hartline/Kleinberg 07], [Hartline/Roughgarden 09], …
- simple pricing rules [Lucier/Borodin 10], [Paes Leme/Tardos 10], ….

Outline

1. The optimal and approximately optimal mechanism design paradigms: Vickrey, Myerson, and beyond.

2. *Case study #1: do good single-item auctions require detailed distributional knowledge?*

3. Case study #2: do good combinatorial auctions require complex bid spaces?

4. Conclusions

Optimal Single-Item Auctions

[Myerson 81]: characterized the optimal auction, as a function of the prior distributions $F_1,\ldots,F_n$.

- Step 1: transform bids to virtual bids: $b_i \rightarrow \varphi_i(b_i)$
  - formula depends on distribution: $\varphi_i(b_i) = b_i - [1 - F_i(b_i)] / f_i(b_i)$
- Step 2: winner = highest positive virtual bid (if any)
- Step 3: price = lowest bid that still would've won

I.i.d. case: 2^{nd}-price auction with monopoly reserve price.

General case: requires full knowledge of $F_1,\ldots,F_n$.

Motivating Question

Question: Does a near-optimal single-item auction require detailed distributional knowledge?

Reformulation: How much data is necessary and sufficient to justify revenue-optimal auction theory?

- "data" = samples from unknown $F_1,\ldots,F_n$
  - formalism inspired by learning theory [Valiant 84]
  - Yahoo! example: [Ostrovsky/Schwarz 09]
- benchmark: Myerson's optimal auction for $F_1,\ldots,F_n$
  - want expected revenue at least $(1-\varepsilon)$ times benchmark

Answer: governed by degree of heterogeneity of bidders.

Formalism: Single Buyer

Step 1: seller gets s samples $v_1,\ldots,v_s$ from unknown F

Step 2: seller picks a price $p = p(v_1,\ldots,v_s)$

Step 3: price p applied to a fresh sample v_{s+1} from F

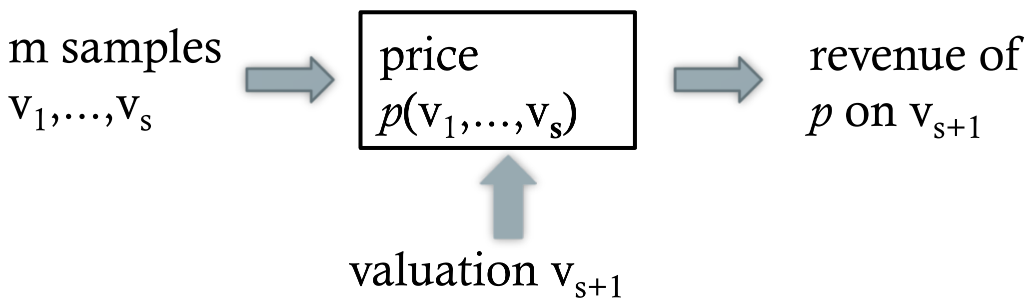

Goal: design p so that $E_{v_1,\ldots,v_s}[p(v_1,\ldots,v_s)\bullet(1-F(p(v_1,\ldots,v_s)))]$
is close to $\max_p[p\bullet(1-F(p))]$ (no matter F is)

Results for a Single Buyer

1. no assumption on *F*: no finite number of samples yields non-trivial revenue guarantee (for every F)
2. if *F* is "regular": with s=1, setting $p(v_1) = v_1$ yields a ½-approximation (consequence of [Bulow/Klemperer 96])
3. for regular *F*, arbitrary ε : $\approx (1/\varepsilon)^3$ samples necessary and sufficient for $(1-\varepsilon)$-approximation [Dhangwatnotai/Roughgarden/Yan 10], [Huang/Mansour/Roughgarden 14]
4. for *F* with a montone hazard rate, arbitrary ε : $\approx (1/\varepsilon)^{3/2}$ samples necessary and sufficient for $(1-\varepsilon)$-approximation [Huang/Mansour/Roughgarden 14]

Formalism: Multiple Buyers

Step 1: seller gets s samples $\mathbf{v_1},\ldots,\mathbf{v_s}$ from $\boldsymbol{F} = F_1 \times \cdots \times F_n$

- each $\mathbf{v_i}$ an n-vector (one valuation per bidder)

Step 2: seller picks single-item auction $A = A(\mathbf{v_1},\ldots,\mathbf{v_s})$

Step 3: auction A is run on a fresh sample $\mathbf{v_{s+1}}$ from $\boldsymbol{F}$

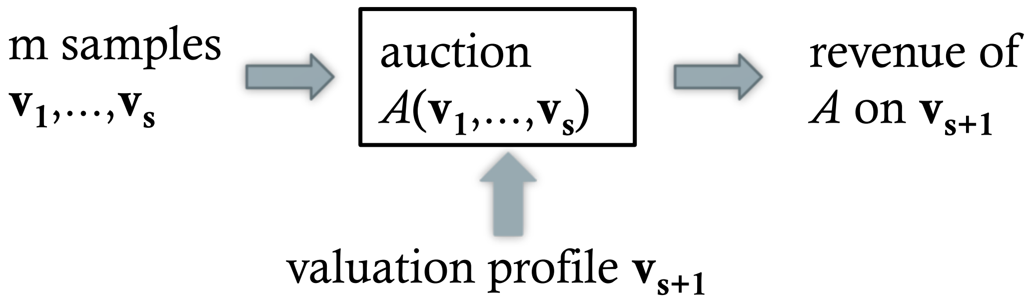

Goal: design A so $E_{v_1,\ldots,v_s}[E_{v_{s+1}}[\mathrm{Rev}(A(v_1,\ldots,v_s)(v_{s+1}))]]$ close to OPT

Positive Results

One sample (s=1) still suffices for ¼-approximation

- 2nd-price auction with reserves = samples
- consequence of [Hartline/Roughgarden 09]

Polynomial (in ε^{-1} only) samples still suffice for $(1-\varepsilon)$-approximation if bidders are i.i.d.

- only need to learn monopoly price

Take-away: for these cases,

- modest amount of data (independent of n) suffices
- modest distributional dependence suffices

Negative Results

Theorem: [Cole/Roughgarden 14] at least $\approx n / \sqrt{\varepsilon}$ samples are necessary for (1- ε)-approximation.

- for every sufficiently small constant ε
- even when distributions guaranteed to be truncated exponential distributions (monotone hazard rate)

Corollary (of proof): near-optimal auctions require detailed knowledge of the valuation distributions.

Motivating Question Revisited

Question: does a near-optimal single-item auction require detailed distributional knowledge?

Answer: if and only if bidders are heterogeneous.

Reformulation: How much data is necessary and sufficient to justify revenue-optimal auction theory?

Answer : polynomial in ε^{-1} but also linear in the "amount of heterogeneity" (i.e., # of distinct F_i's)

Outline

1. The optimal and approximately optimal mechanism design paradigms: Vickrey, Myerson, and beyond.

2. Case study #1: do good single-item auctions require detailed distributional knowledge?

3. *Case study #2: do good combinatorial auctions require complex bid spaces?*

4. Conclusions

The VCG Mechanism

Setup: n bidders, m non-identical goods

- bidder i has private valuation $v_i(S)$ for each subset S of goods [$\approx 2^m$ parameters]
- welfare of allocation $S_1,S_2,\ldots,S_n$: $\sum_i v_i(S_i)$
  - goal is to allocate goods to maximize this quantity

VCG mechanism: [Vickrey 61, Clarke 71, Groves 73]

- each player reports full valuation (yikes!)
- compute welfare-maximizing allocation
- charge payments to incentive truthful revelation

Motivating Question

Question: When can simple auctions perform well in complex settings?

- "simple" = low-dimensional bid space
  - polynomial in m, rather than exponential in m
- benchmark = VCG welfare (i.e., maximum-possible)
- want equilibrium welfare close to benchmark
- example interpretation: is package bidding essential?

Answer: governed by structure of bidders' valuations (e.g., extent to which items are complements).

A Simple Auction: Selling Items Separately

Simultaneous First-Price Auction (S1A): [Bikhchandani 99]

- each bidder submits one bid per item
  - m parameters instead of 2^m
- each item sold separately in a first-price auction

Question: when do S1A's have near-optimal equilibrium welfare?

- seems unlikely if items are complements (exposure problem)
- expect inefficiency even with known gross substitutes valuations (demand reduction)
- expect inefficiency even with m=1 (i.e., first-price single-item auction) with non-iid valuations

Valuation Classes

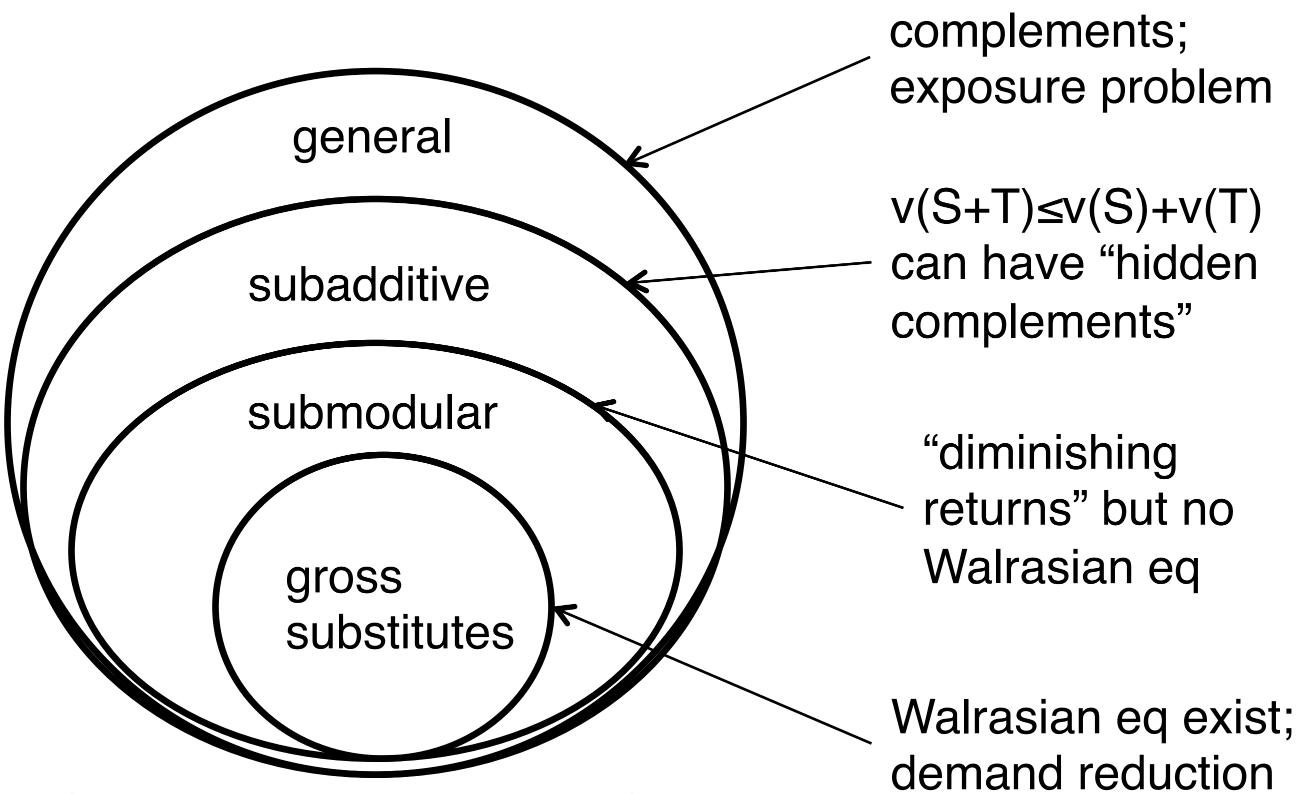

When Do S1A's Work Well?

General valuations: S1A's can have equilibria with welfare arbitrarily smaller than the maximum possible. [Hassidim/Kaplan/Mansour/Nisan 11]

Subadditive valuations: [$v_i(S+T) \leq v_i(S)+v_i(T)$] every equilibrium of a S1A is at least 50% of the maximum possible. [Feldman/Fu/Gravin/Lucier 13]

- full-info Nash equilibria or Bayes-Nash equilibria
- 63% when valuations are submodular [Syrgkanis/Tardos 13]

Take-away: S1A's work reasonably well if and only if there are no complements.

Digression on Approximation Ratios

Recall: main motivations for approximation approach:

- identify specific mechanisms or mechanism features that are potentially useful in practice
- quantify cost of side constraints (e.g., "simplicity")
  - e.g., complex bids required ⇔ best approx far from 100%

Also: (if you insist on a literal interpretation)

- by construction, can't achieve 100% of benchmark
- non-asymptotic => bounded below 100%
  - possible escapes: large markets, parameterized approximation

Negative Results

Theorem: [Roughgarden 14]

- With subadditive bidder valuations, no simple auction guarantees equilibrium welfare better than 50% OPT.
  - “simple”: bid space dimension ≤ polynomial in # of goods
- With general valuations, no simple auction guarantees non-trivial equilibrium welfare.

Take-aways:

1. In these cases, S1A’s optimal among simple auctions.
2. With complements, complex bid spaces (e.g., package bidding) necessary for welfare guarantees.

Outline

1. The optimal and approximately optimal mechanism design paradigms: Vickrey, Myerson, and beyond.

2. Case study #1: do good single-item auctions require detailed distributional knowledge?

3. Case study #2: do good combinatorial auctions require complex bid spaces?

4. Conclusions

Conclusions

Thesis: approximately optimal mechanisms significantly extend reach of optimal mechanism design theory.

Example #1: identify when distributional knowledge is essential for near-optimal revenue-maximization.

- open: beyond single-item auctions

Example #2: identify when high-dimensional bid spaces are essential for near-optimal welfare-maximization.

- open: better understanding of specific formats