Bayesian and Worst-Case Revenue-Maximizing Auctions [Part I]

Click to edit Master subtitle style

Tim Roughgarden (Stanford)

Revenue-Maximizing Auctions

Goal: prove results of the form (e.g., for revenue):

"Theorem: auction A is (approximately) optimal."

Revenue-Maximizing Auctions

Goal: prove results of the form (e.g., for revenue):

"Theorem: auction A is (approximately) optimal."

Auction model: first focus on multi-unit auctions

- n bidders, k identical goods, unit demand
- allocation rule: $\vec{bi's} \longrightarrow \vec{xi's}$
- payment rule: bi's pi's [all i: pi ≤ bixi]
- truthful (i.e., truthful bidding [bi=vi] dominant)
  - each i faces bid-independent posted price

On Truthful Auctions

Easy Lemma: a multi-unit auction is truthful if and only if it is equivalent to the following:

- separately for each $i=1,2,\ldots,n$:
  - choose (possibly randomized) "take-it-or-leave-it" offer at price $t_i(b_{-i})$ ["posted price"]

Example: the Vickrey auction (with $k=1$) corresponds to $t_i(b_{-i}) = \max_{j\neq i} b_j$ for each i

Fact: [Myerson 81] these assumptions are "WLOG".

- "Revelation Principle" + "Revenue Equivalence Thms"

Bayesian Profit Maximization

Example: 1 bidder, 1 item, v ~ known distribution F

- truthful auctions = posted prices p
- expected revenue of p: $p(1-F(p))$
  - given F, can solve for optimal p^*
  - e.g., $p^* = ½$ for v ~ uniform[0,1]
- but: what about k,n >1 (with i.i.d. vi's)?

Bayesian Profit Maximization

Example: 1 bidder, 1 item, v ~ known distribution F

- truthful auctions = posted prices p
- expected revenue of p: $p(1-F(p))$
  - given F, can solve for optimal p^*
  - e.g., $p^* = ½$ for v ~ uniform[0,1]
- but: what about k,n >1 (with i.i.d. vi's)?

Theorem: [Myerson 81] auction with max expected revenue is Vickrey with above reserve p^*.

- note p^* is *independent of k and n*

Two Definitions

Implementable Allocation Rule: is a function x (from bids to winners/losers) that admits a payment rule p such that (x,p) is truthful.

- i.e., truthful bidding [bi:=vi] always maximizes a bidder's (expected) utility

Two Definitions

Implementable Allocation Rule: is a function x (from bids to winners/losers) that admits a payment rule p such that (x,p) is truthful.

- i.e., truthful bidding [bi:=vi] always maximizes a bidder's (expected) utility

Monotone Allocation Rule: for every fixed bidder i, fixed other bids b-i, probability of winning only increases in the bid bi.

- example: highest bidder wins
- non-example: 2nd-highest bidder wins

Myerson's Lemma

Myerson's Lemma: [1981; also Archer-Tardos 01]
an allocation rule x is implementable if and only if it is monotone.

Myerson's Lemma

Myerson's Lemma: [1981; also Archer-Tardos 01]
an allocation rule x is implementable if and only if it is monotone.

Moreover: for every monotone allocation rule x, there is a unique payment rule p such that (x,p) is truthful and losers always pay 0.

Myerson's Lemma

Myerson's Lemma: [1981; also Archer-Tardos 01]
an allocation rule x is implementable if and only if it is monotone.

Moreover: for every monotone allocation rule x, there is a unique payment rule p such that (x,p) is truthful and losers always pay 0.

Explicit formula for pi(b):

- keep b-i fixed, increase z from 0 to bi
- consider breakpoints y1,...,yq at which xi jumps
- set pi(b) := Σj yj • [jump in xi at yj]

Digression: Sponsored Search

Explicit formula for pi(b):

- keep b-i fixed, increase z from 0 to bi
- consider breakpoints y1,...,yq at which xi jumps
- set pi(b) := Σj yj • [jump in xi at yj]

Example - Sponsored search:

- n bidders with valuations vi per click
- k slots, worth α1 ≥ α2 ≥ ... ≥ αk clicks
- note: "rank by bid" is monotone allocation rule
- total payment from ith ranked bidder:

  pi(b) := Σj>i vj • [αj-αj+1]

Myerson's Lemma (Proof)

Proof: let x be monotone, fix i and b-i. Write x(z), p(z) for xi(z, b-i), pi(z, b-i).

Swapping trick: if (x,p) is truthful, p satisfies:

- [take true value = z, false bid = z + ε]:
  $z \circ x(z) - p(z) \geq z \circ x(z + \varepsilon) - p(z + \varepsilon)$

- [take true value = z + ε, false bid = z]:
  $(z + \varepsilon) \circ x(z + \varepsilon) - p(z + \varepsilon) \geq (z + \varepsilon) \circ x(z) - p(z)$

Myerson's Lemma (Proof)

Proof: let x be monotone, fix i and b-i. Write x(z), p(z) for xi(z, b-i), pi(z, b-i).

Swapping trick: if (x,p) is truthful, p satisfies:

- [take true value = z, false bid = z + ε]:
  $$z \circ x(z) - p(z) \geq z \circ x(z + \varepsilon) - p(z + \varepsilon)$$

- [take true value = z + ε, false bid = z]:
  $$(z + \varepsilon) \circ x(z + \varepsilon) - p(z + \varepsilon) \geq (z + \varepsilon) \circ x(z) - p(z)$$

Thus: p(z + ε) - p(z) lies between

$z \circ [x(z + \varepsilon) - x(z)]$ and $(z + \varepsilon) \circ [x(z + \varepsilon) - x(z)]$

Myerson's Lemma (Proof con'd)

The story so far: $p(z + \varepsilon) - p(z)$ lies between

$$z \circ [x(z + \varepsilon) - x(z)] \text{ and } (z + \varepsilon) \circ [x(z + \varepsilon) - x(z)]$$

So: dividing by ε and taking ε to zero, get

- $p'(z) = z \circ x'(z)$ [if x differentiable at z] or
- jump in p at z = z ∘ [jump in x at z]

Integrating from 0 to bi, get:

$$pi(b) := \Sigma j\ yj \bullet [\text{jump in } xi \text{ at } yj]$$

15

Myerson's Revenue Formula

Step 2: formula for revenue via "virtual valuations".

- for *any* truthful auction (x,p)

Write: for fixed i,v-i, expected revenue from i:

- integrate over pi(vi,v-i) w.r.t. vi (according to F)
- write pi in terms of xi, simplify:

Myerson's Revenue Formula

Step 2: formula for revenue via "virtual valuations".

- for *any* truthful auction (x,p)

Write: for fixed i,v-i, expected revenue from i:

- integrate over pi(vi,v-i) w.r.t. vi (according to F)
- write pi in terms of xi, simplify:

$$= E_{vi}[\phi(vi) \cdot xi(vi,v-i)]$$

where

$$\phi(vi) = vi - \frac{1-F(vi)}{f(vi)}$$

"virtual valuation" → ϕ

E.g., $\phi(vi) = 2vi-1$ when F = Unif[0,1]

Proof of Myerson's Theorem

So far: expected revenue of any (x,p):

E.g., $\phi(v_i) = 2v_i - 1$ when F = Unif[0,1]

$$= E_{\mathbf{v}}[\Sigma_i \ \phi(v_i) \cdot x_i(\mathbf{v})]$$

- to maximize: *for each vector* ***v****, set xi's to maximize sum of virtual valuations.*
- multi-item auctions: award k items to the top k $\phi(v_i)$'s that are also positive

The Fine Print

So far: to maximize expected revenue $E_{\mathbf{v}}[\Sigma_i \phi(v_i) \cdot x_i(\mathbf{v})]$: for each vector $\mathbf{v}$, set x_i's to maximize sum of virtual valuations.

Question: is this allocation rule monotone?

The Fine Print

So far: to maximize expected revenue $E_{\mathbf{v}}[\Sigma_i \phi(v_i) \cdot x_i(\mathbf{v})]$: for each vector **v**, set x_i's to maximize sum of virtual valuations.

Question: is this allocation rule monotone?

- yes if ϕ is nondecreasing; otherwise not
- *regular* distribution F := ϕ is nondecreasing
  - most prominent non-example: bimodal distributions
- Myerson's theory can be extended to irregular distributions - important but more complicated

Bulow-Klemperer ('96)

Observation: for every F, E[ϕ (vi)] = 0.

- proof #1: consider Vickrey with k = n = 1
- proof #2: integrate ϕ (vi) = vi - (1-F(vi)/ f(vi))

Bulow-Klemperer ('96)

Observation: for every F, E[ϕ (vi)] = 0.

- proof #1: consider Vickrey with k = n = 1
- proof #2: integrate ϕ (vi) = vi - (1-F(vi)/ f(vi))

Corollary [BK96]: for k = 1, every n, every regular F:

Vickrey's revenue OPT's revenue

Bulow-Klemperer ('96)

Observation: for every F, E[ϕ (vi)] = 0.

- proof #1: consider Vickrey with k = n = 1
- proof #2: integrate ϕ (vi) = vi - (1-F(vi)/f(vi))

Corollary [BK96]: for k = 1, every n, every regular F:

Vickrey's revenue $\geq$ OPT's revenue

Bulow-Klemperer ('96)

Observation: for every F, E[ϕ (vi)] = 0.

- proof #1: consider Vickrey with k = n = 1
- proof #2: integrate ϕ (vi) = vi - (1-F(vi)/ f(vi))

Corollary [BK96]: for k = 1, every n, every regular F:

Vickrey's revenue	≥	OPT's revenue
[with (n+1) i.i.d. bidders]		[with n i.i.d. bidders]

Interpretation: small increase in market size more important than running optimal auction.

24

Bulow-Klemperer (Proof)

Proof idea:

- OPT's expected revenue [n bidders]:

  $$E_{\mathbf{v}}[\max\{\max_{i\leq n} \phi(v_i), 0\}]$$

- Vickrey's expected revenue [(n+1) bidders]:

  $$E_{\mathbf{v}}[\max\{\max_{i\leq n} \phi(v_i), \phi(v_{n+1})\}]$$

- condition on $\phi(v_1),\ldots,\phi(v_n)$, use observation that $E[\phi(v_i)] = 0$

Application: Search Auctions

Theorem 1: [Dughmi/Roughgarden/Sundararajan EC 09]

The BK theorem extends to multi-unit auctions (add k new bidders); search auctions (ditto); matroid domains (add a new matroid basis).

Theorem 2 [DRS09]: for every regular F and k,n:

Vickrey's revenue $\geq$ (1-k/n) · OPT's revenue

[with n i.i.d. bidders] [with n i.i.d. bidders]

Interpretation: efficiency and revenue well-aligned with i.i.d. bidders and modest competition.

More General Problems

Downward-closed environment:

- n bidders, independent valuations from F1,…,Fn
- collection C of feasible sets of bidders
- subsets of feasible sets must be again feasible
  - k-unit auction: C = all subsets of at most k bidders

Example: single-minded combinatorial auctions

- set of m different goods
- each bidder i wants a known bundle Si of goods
- feasible subsets = bidders with disjoint bundles

Myerson still works

In a general downward-closed environment:

- allocation rule still implementable iff monotone
- revenue formula remains valid:
  expected revenue under x = $E_{\mathbf{v}}[\Sigma_i \phi_i(v_i) \cdot x_i(\mathbf{v})]$
- if each F_i is regular, OPT just maximizes total virtual value pointwise (for each **v**)

Myerson still works

In a general downward-closed environment:

- allocation rule still implementable iff monotone
- revenue formula remains valid:
  expected revenue under x = $E_{\mathbf{v}}[\Sigma_i \phi_i(v_i) \cdot x_i(\mathbf{v})]$
- if each F_i is regular, OPT just maximizes total virtual value pointwise (for each **v**)

Issue: OPT much more complicated/unintuitive

- single-item auction: might sell to non-highest bidder!
- need detailed information about ϕ_i's (not just $\phi^{-1}(0)$)

Simple vs. Optimal Auctions

Question: is there a simple (+ practically implementable) auction that's almost as good?

Theorem: [Hartline/Roughgarden EC 09] Yes!

Expected revenue of VCG + "monopoly reserve prices" is at least 50% of OPT for:

- all DC environments, if the Fi's are such that $\frac{1-F_i(v_i)}{f_i(v_i)}$ nonincreasing in vi (*monotone hazard rate)*
- all matroid environments, if Fi's are regular
  - i.e., only assume vi - [(1-Fi(vi))/fi(vi)] is

VCG + Monopoly Reserves

Input: general DC environment, unknown vi's drawn from known regular Fi's

- set reserve ri = argmaxp p(1-Fi(p)) [each i]
- delete all bidders with vi < ri
- run "VCG": maximize sum of vi's (i.e., welfare) of remaining bidders, subject to feasibility
- this allocation rule is monotone; charge suitable prices to make truthful (=max{VCG prices, ri's})