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

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."

Question: what if we want an input-by-input (distribution-free) guarantee?

- first study classical but ad hoc approach
- then systematic answer via novel connection to the Bayesian case

Auction Benchmarks (con'd)

Goal: prove results of the form:

"Theorem: for every valuation profile v:
auction A's revenue on v is at least $OPT(v)/\alpha$."
(for a hopefully small constant α)

Auction Benchmarks (con'd)

Goal: prove results of the form:

"Theorem: for every valuation profile v:
auction A's revenue on v is at least OPT(v)/α."
(for a hopefully small constant α)

Idea for OPT(v): sum of k largest vi's.

Problem: too strong, not useful.

- makes all auctions A look equally bad.
- every A has a bad v [no constant α possible]

The Obvious Idea Fails

Claim: no auction always has revenue at least a constant fraction of the sum of k largest vi's.

Proof sketch: by probabilistic method. Take k = n.

- pick each vi i.i.d. from distribution with CDF F(z) = 1-1/z on [1,∞) [density f(z) = 1/z2]
- expected revenue of every posted price pi ≥ 1 for bidder i = pi [1-(1-1/pi)] = 1.
- expected revenue of every auction: ≤ n
- expected sum of vi's: unbounded

The Fixed Price Benchmark

Solution: [Goldberg/Hartline/Karlin/Saks/Wright GEB 06]

- define OPT(v) := best *fixed-price* revenue:

$$F^{(2)}(v) := \max_{2 \leq i \leq k} i v_i \qquad \text{(assume sorted } v_i\text{'s)}$$

Justification?: for now, "seems to work".

- α-competitive auctions exist for small α
  - will prove today with $\alpha = 4$
- no auction has α smaller than 1 (or even 2.42)

Later: a fundamental explanation of why it works

Auction Benchmarks (con'd)

Goal: prove results of the form:

"Theorem: for every valuation profile v:
auction A's revenue on v is at least F(2)(v)/α."
(for a hopefully small constant α)

Example: the Vickrey auction achieves α=2 when n=2 (but no constant factor for large n).

A Competitive Auction

Theorem: [Fiat/Goldberg/Hartline/Karlin STOC 02]

There is a randomized auction for n-bidder n-item auctions that, for every input v, has expected revenue at least F(2)/4.

- works also for $k \geq 2$ items [exercise]

Recall:

$$F(2)(v) := \max_{2 \leq i \leq k} ivi$$ (assume sorted vi's)

Subroutine: Profit Extractor

Given: (truthful) bids v + revenue target R.:

- initialize S = all bidders
- while there is an i in S such that vi < R/ |S|:
  - remove such a bidder from S
- return final set S and charge all winners (if any) a price of p = R/ |S|

Note: allocation rule is monotone; prices are correct (p = min bid s.t. a winner still wins)

- => truthful by Myerson's Lemma

Profit Extractor (con'd)

Claim: ProfitExtract has revenue R if there is a common posted price that extracts R; and has revenue 0 otherwise.

Proof Sketch: 2nd statement is clear (ProfitExtract only uses common posted prices).

For 1st statement: suppose common posted price p works, define $T = \{ i \mid vi \geq p \}$. All such bidders can pay $R/|T|$. Inductively, no bidder of T ever gets deleted.

The RSPE Auction

- collect (truthful) bid v
- randomly partition v = (x,y)
  [each bidder placed 50%/50%, independently]
- let R1 = max revenue from x via common posted price;
  R2 = max revenue from y via common posted price
- ProfitExtract(x) with revenue target R2
- ProfitExtract(y) with revenue target R1

Example: n = 2, v = (1,1/2)

- F(2)(v) = 1
- expected revenue of RSPE = ½(½+0) = 1/4

The RSPE Auction

Claim #1: RSPE is a truthful auction.

Proof sketch: Say bidder i part of x. i can't change R2. For i, RSPE is same as ProfitExtract(x, R2), where truthful bidding is optimal.

Claim #2: RSPE's revenue is at least min{R1,R2}.

Proof : E.g., if R1 ≤ R2, ProfitExtract(y, R1) will successfully extract revenue R1.

- recall key property of ProfitExtract subroutine

13

The Final Lemma

Claim #3: For all inputs v, $E[\min\{R1,R2\}] \geq F^{(2)}/4$.

- expectation is over the random split $v=(x,y)$
- recall $F^{(2)}(v) := \max_{2 \leq i \leq k} i v_i$ (assume sorted v_i's)

Proof : Fix v. Let $i \geq 2$ satisfy $F^{(2)}(v) = i v_i$.

Let a,b = number of top i bidders in x,y.

Note: $R1 \geq a v_i$ and $R2 \geq b v_i$

- So: just need $E[\min\{a,b\}] \geq i/4$.

The Final Lemma (con'd)

Need to show : For every i ≥ 2, a random split into a,b (with a+b=i) satisfies E[min{a,b}] ≥ i/4.

Case i = 2: min{a,b} is either 0 (50% probability) or 1 (50% probability) => E[min{a,b}] = 1/2 = i/4.

Case i = 3: min{a,b} is either 0 (25% probability) or 1 (75% probability) => E[min{a,b}] = 3/4 =

Opt Fixed-Price via Myerson

Recall question: meaning of the optimal fixed-price revenue for (non-Bayesian) auctions?

$$RB(v) := \max_{i \leq k} \; i v_i \qquad \text{(assume sorted vi's)}$$

Recall: "seems to work" (even with apples vs. oranges).

Myerson: *for all F,* Vickrey + a reserve is optimal.

Corollary 1: *for all F and all* $\boldsymbol{v}$, ex post behavior of optimal auction for F is to charge a fixed price.

- namely: max{reserve price, (k+1)th highest bid of **v**}

Opt Fixed-Price via Myerson

Corollary 2: If auction A is α-competitive w.r.t benchmark RB, then it is *simultaneously competitive with all Bayesian optimal auctions!*

I.e.: For every F, corresponding opt auction AF:

A's expected revenue ≥ (AF's expected revenue) / α

Proof: inequality holds for every v:

From Old to New Results

So far: re-interpretation of old results for worst-case profit-maximization in multi-item auctions.

Next: new applications [Hartline/Roughgarden 08, 09]

- beyond multi-item auctions
- beyond identical bidders
- novel objectives ("money burning")

Analysis Template

Moral: Bayesian auction design yields strong worst-case performance benchmarks.

- characterize ex post Bayesian optimal behavior
  - for all distributions of interest, all valuation profiles
- want to simultaneously compete with all such behaviors (on each valuation profile)
- automatic corollary: competitive with expected performance of every Bayesian-optimal auction

Digression

Alternative: can directly aspire to the "automatic corollary": competitive with expected performance of every Bayesian-optimal auction

Theorem: [Dhangwotnotai/Roughgarden/Yan 09] There is a single auction with expected revenue is at least 25% of OPT for:

- all DC environments, if the Fi's satisfy the monotone hazard rate condition
- all matroid environments, if Fi's are IID, regular

Money-Burning Mechanisms

New Objective: *residual surplus:*

$$\max \Sigma i \; vixi - pi$$

Motivation: welfare maximization with private values, non-transferable payments (e.g., time).

- queueing; computational payments (e.g. for spam)

Example: $k = 1$, $n = 2$, $v1 > v2$

- Vickrey residual surplus = $v1 - v2$
- random allocation (0 payments) = $(v1 + v2)/2$

Maximizing Residual Surplus

To derive benchmark: Characterize ex post behaviors of Bayesian optimal mechanisms.

Theorem 1 [Hartline/Roughgarden STOC 08]: *for all F and all v,* ex post behavior of optimal auction for F is to use a *(p,q)-lottery.*

- "optimal" = max expected residual surplus (for F)

Maximizing Residual Surplus

To derive benchmark: Characterize ex post behaviors of Bayesian optimal mechanisms.

Theorem 1 [Hartline/Roughgarden STOC 08]: *for all F and all v,* ex post behavior of optimal auction for F is to use a *(p,q)-lottery.*

- "optimal" = max expected residual surplus (for F)
- (p,q)-lottery: [assume ≤ k bidders have vi > p]
  - all bidders with vi > p get an item at price p
  - random subset of bidders with q < vi ≤ p awarded remaining items at price q

A Money-Burning Benchmark

So: for every valuation profile, *define*

RSB(v) :=maxp,q [resid. surplus of (p,q)-lottery on v]

Reason: If auction A is α-competitive w.r.t benchmark RSB, then it is *simultaneously competitive with all Bayesian optimal auctions!*

- same trivial proof as before

A Money-Burning Mechanism

Theorem 2 [HR08]: there is a (prior--free) auction A that is O(1)-competitive with RSB(v) (for all v).

A Money-Burning Mechanism

Theorem 2 [HR08]: there is a (prior--free) auction A that is O(1)-competitive with RSB(v) (for all v).

Key Lemma: for every v + (p,q)-lottery L, there is a p'-lottery with ≥ ½ of L's residual surplus on v.

Key Lemma #2: there is an auction A that is O(1)-competitive with optimal p-lottery (for each v).

- only one parameter ⊳ can use random sampling techniques for this

A Prior-Free Auction

Idea: (given v) randomly sample some bids, use statistics to formulate prices for the rest.

- let S = uniform random sample (50/50 each bidder)
- let p* = optimal p-lottery for S
- run p*–lottery on rest of bidders

A Prior-Free Auction

Idea: (given v) randomly sample some bids, use statistics to formulate prices for the rest.

- let S = uniform random sample (50/50 each bidder)
- let p* = optimal p-lottery for S
- run p*–lottery on rest of bidders

Extra detail: with constant probability, run a standard (k=1) Vickrey auction.

  - just in case vmax dwarfs all other bids

28

Analysis Sketch

Need: constant fraction of residual surplus of optimal p-lottery, namely (for some m):

$$\frac{k}{m}[v1 + v2... + ... + vm] - kvm+1$$

(assume sorted vi's)

a.k.a.:

$$\frac{k}{m}[(v1 - v2) + 2(v2 - v3) + ... + m(vm - vm+1)]$$

Analysis Sketch

Need: constant fraction of residual surplus of optimal p-lottery, namely (for some m):

$$\frac{k}{m}[(v1 - v2) + 2(v2 – v3) + \ldots + m(vm – vm+1)]$$

Analysis Sketch

Need: constant fraction of residual surplus of optimal p-lottery, namely (for some m):

$$\frac{k}{m}[(v1 - v2) + 2(v2 - v3) + \ldots + m(vm - vm+1)]$$

Balanced Sampling Lemma: with constant probability: for *every* i ≥ 2, # of top i bidders inside, outside sample differ by factor of ≤ 5.

- crude Chernoff/Union Bounds don’t work
- leverages [Feige et al 05] biased random walk analysis

Analysis Template (reprise)

Moral: Bayesian auction design immediately suggests worst-case performance benchmarks.

- characterize ex post Bayesian optimal behavior
  - for all distributions of interest, all valuation profiles
- want to simultaneously compete with all such behaviors (on each valuation profile)
- if possible, are automatically competitive with performance of every Bayesian-optimal auction

Beyond Multi-Item Auctions

Example: n bidders (valuations vi), feasible subsets of winners = independent sets of some matroid.

- e.g., spanning trees; multi-item = uniform matroid

Beyond Multi-Item Auctions

Example: n bidders (valuations v_i), feasible subsets of winners = independent sets of some matroid.

- e.g., spanning trees; multi-item = uniform matroid

Myerson's Revenue Formula: given an IID regular distribution F (with virtual value ϕ), expected revenue of an auction: $E_{\mathbf{v}}[\Sigma_i \phi(v_i) \cdot x_i(\mathbf{v})]$

To maximize: given v, choose independent set maximizing $\Sigma_i \phi(v_i)$ [e.g., via greedy algorithm]

A Prior-Free Benchmark

So: ex post behavior of a Bayesian optimal auction:

- choose independent set with max $\Sigma_i \phi(v_i)$

Equivalent: the “VCG mechanism” with a common reserve price $\phi^{-1}(0)$.

- regularity => v_i's and $\phi(v_i)$'s are ordered the same

Upshot: prior-free benchmark RB(v) := max revenue achievable via VCG with a common reserve.

- [HR]: can be 8-competitive with this benchmark
- randomize between VCG and ProfitExtract

More Prior-Free Benchmarks

Beyond matroids?: e.g., each bidder wants a bundle of goods, can only allocate each good once.

- the optimal mechanism is complicated

But: [Hartline/Roughgarden EC 09] VCG mechanism with a common reserve is a 2-approximation.

  - needs somewhat stronger distributional assumption
  - offers simple and provably good alternative to the (complex) optimal auction
  - justifies VCG + optimal reserve benchmark in general

Beyond Symmetric Bidders

Asymmetric bidders (Bayesian): different prior F_i (and corresponding ϕ_i) for each bidder i.

- Myerson's formula: $E_{\mathbf{v}}[\Sigma_i \phi_i(v_i) \cdot x_i(\mathbf{v})]$
- arbitrary F_i's => all prices can arise ex post
- *ordered* F_i's => optimal prices *monotone*
  - "ordered" = ϕ_i^{-1}'s can be consistently ordered

Prior-free version: only know bidder ordering .

- RB(v) := max revenue via monotone prices
- Open: can you O(1)-compete with this?

Conclusions

Take-home point: new template for generating meaningful worst-case auction benchmarks.

- automatic: simultaneous competitive guarantee with all Bayesian-optimal auctions
- enables new theorems for money-burning mechanisms, asymmetric allocations and/or bidders

Open Questions:

- thoroughly understand “single-parameter” problems
- need more progress for the non-IID case
- multi-parameter? (e.g., combinatorial auctions)