From Bayesian Auctions to Approximation Guarantees

Tim Roughgarden (Stanford)

based on joint work with:

Jason Hartline (Northwestern)
Shaddin Dughmi, Mukund Sundararajan (Stanford)

Auction Benchmarks

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

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

Auction model: focus on multi-item auctions

- n bidders, k identical goods, unit demand
- allocation rule: b_i's $\longrightarrow$ x_i's
- payment rule: b_i's $\longrightarrow$ p_i's [all i: $p_i \leq b_i x_i$]
- truthful (i.e., truthful bidding [b_i=v_i] dominant)
    - each i faces bid-independent posted price

Auction Benchmarks (con'd)

Goal: prove results of the form:

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

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

Problem: too strong, not useful.

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

The Fixed Price Benchmark

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

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

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

Usual justification: "seems to work".

- α-competitive auctions exist for small α
  - assuming no "dominant bidder"
- no auction has α smaller than 1 (or even 2.42)

Question: is there a fundamental explanation?

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 n >1 bidders (with i.i.d. v_i's)?

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

- note p^* is independent of k and n

Myerson's Theorem (Step 1)

Step 1: each allocation rule [v_i's $\longrightarrow$ x_i's] has unique candidate truthful payment rule [v_i's $\longrightarrow$ p_i's].

Reason: for every i, fixed v_{-i}, truthfulness implies:

marginal benefit of lying $\longrightarrow$ $$\left\{ v_i \cdot \frac{d}{dv_i} x_i(v_i, v_{-i}) = \frac{d}{dv_i} p_i(v_i, v_{-i}) \right\}$$ $\longleftarrow$ marginal cost of lying

Then: integrate to get (the only possible) payments.

Check: these are indeed truthful iff x_i always nondecreasing.

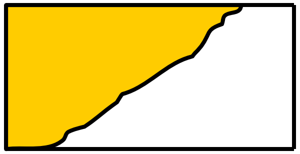

Myerson's Theorem (Step 2)

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 $p_i(v_i, v_{-i})$ w.r.t. v_i (according to F)
- write p_i in terms of x_i, simplify:

$$= E_{v_i}[\varphi(v_i) \cdot x_i(v_i, v_{-i})]$$

where

"virtual valuation" → $$\varphi(v_i) = v_i - \frac{1-F(v_i)}{f(v_i)}$$

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

Myerson's Theorem (Step 3)

Step 3: optimize pointwise (over v_i's).

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

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

E.g., $\varphi(v_i) = 2v_i - 1$ when $F = \mathrm{Unif}[0,1]$

- to maximize: *for each vector* $\boldsymbol{v}$*, set* x_i*'s to maximize sum of virtual valuations.*
  - fine print: need F to be "regular" for this to be truthful
- multi-item auctions: award k items to the top k $\varphi(v_i)$'s that are also positive
  - i.e., Vickrey with reserve price $\varphi^{-1}(0)$ [for all k,n]

Bulow-Klemperer ('96)

Observation: for every F, $E[\varphi(v_i)] = 0$.

- proof #1: consider Vickrey with $k = n = 1$
- proof #2: integrate $\varphi(v_i) = v_i - (1-F(v_i)/f(v_i))$

Corollary [BK96]: for $k = 1$, every $n \geq 1$, every F:

Vickrey's revenue	$\geq$	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.

Bulow-Klemperer (Proof)

Proof idea:

- OPT's expected revenue [n bidders]:

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

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

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

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

Application: Search Auctions

Theorem 1: [Dughmi/Roughgarden/Sundararajan 07,08]

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

Theorem 2 [DRS]: for every F and every $k,n \geq 1$:

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

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

Idea: bound Vickrey's revenue from final k bidders.

Opt Fixed-Price via Myerson

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

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

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

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

Corollary 1: *for all F and all **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 A_F:

A's expected revenue $\geq$ (A_F's expected revenue)/α

Proof: inequality holds for every v:

A's revenue on v $\geq$ RB(v)/α $\geq$ A_F's revenue on v

Interpretation: ignorance of F costs only α factor.

Money-Burning Mechanisms

New Objective: *residual surplus:*

$$\max \Sigma_i v_i x_i - p_i$$

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

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

Example: $k = 1$, $n = 2$, $v_1 > v_2$

- Vickrey residual surplus = $v_1 - v_2$
- random allocation (0 payments) = $(v_1 + v_2)/2$

Maximizing Residual Surplus

Goal: "optimal" prior-free mechanism to max residual surplus in multi-item auctions.

Question: what is a well-motivated benchmark?

Ad Hoc Approach: Guess. E.g., residual surplus earned by optimal *p-lottery* for **v**:

- given v_i's, pick a fixed price p
- random subset of bidders with $v_i > p$ win at price p
  - $p = 0$: random allocation;
  - $p = v_{k+1}$: k-Vickrey

Maximizing Residual Surplus

Systematic Approach: Characterize ex post behaviors of Bayesian optimal mechanisms.

Theorem [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:
  - all bidders with $v_i > p$ get an item at price p
  - random subset of bidders with $q < v_i \leq p$ awarded remaining items at price q

A Money-Burning Benchmark

So: for every valuation profile, *define*

RSB(v) :=$\max_{p,q}$ [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 [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 (+ non-trivial work) for this

Take-Home Points

Moral: Bayesian auction design can usefully inform worst-case approximation guarantees.

- revenue guarantees for the VCG mechanism
  - Bayesian, but guarantees independent of distribution
- generic framework for prior-free benchmarks
  - characterize ex post Bayesian optimal behavior
  - simultaneously compete with all such mechanisms
  - max revenue (optimal fixed price); money burning (optimal (p,q)-lottery); more to come...?