Learning Near-Optimal Auctions: Statistical, Computational, and Strategic Challenges

Tim Roughgarden (Stanford University)

Revenue-Maximizing Auctions

Setup: 1 seller with 1 item, n bidders, bidder i has private valuation v_i.

Question: which auction maximizes (expected) revenue?

Issue: different auctions do better on different valuations. (Cf., VCG and welfare-maximization)

Bayesian assumption: bidders' valuations $v_1,\ldots,v_n$ drawn independently from distributions $F_1,\ldots,F_n$.

- F_i's known to seller, v_i's unknown

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: $b_i \to \varphi_i(b_i)$
  - 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 have won

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

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

Some Questions

Issue: where does this prior come from?

Modern answer: from data (e.g., past bids)

- [Ostrovsky/Schwarz 09] applied at Yahoo!

Question: How much data is necessary?

- "data" = samples from unknown distributions $F_1,\ldots,F_n$ (e.g., bids in previous auctions)
- goal = near-optimal revenue [e.g., $(1-\varepsilon)$-approx]
- formalism inspired by "PAC" learning theory [Vapnik/Chervonenkis 71, Valiant 84]

Part I: Statistical Aspects

(includes joint work with Jamie Morgenstern)

Some Related Work

Question: How much data is necessary for near-optimal revenue?

Asymptotic regime: [Neeman 03], [Segal 03], [Baliga/Vohra 03], [Goldberg/Hartline/Karlin/Saks/Wright 06]

Uniform bounds for finite-sample regime: [Elkind 07], [Dhangwatnotai/Roughgarden/Yan 10], [Cole/Roughgarden 14], [Chawla/Hartline/Nekipelov 14], [Medina/Mohri 14], [Cesa-Bianchi/Gentile/Mansour 15], [Dughmi/Han/Nisan 15], [Huang/Mansour/Roughgarden 15], [Morgenstern/Roughgarden 15,16], [Devanur/Huang/Psomas 16], [Roughgarden/Schrijvers 16], [Hartline/Taggart 17], [Gonczarowski/Nisan 17], [Bubeck/Devanur/Huang/Niazadeh 17], [Syrgkanis 17], [Cai/Daskalakis 17], ...

Other PAC-Type Models

- rationalizable choice functions [Kalai 03]
- rationalizable utility functions [Beigman/Vohra 06], [Zadimoghaddam/Roth 12]
- voting rules [Procaccia/Zohar/Peleg/Rosenschein 09], [Boutilier/Caragiannis/Haber/Lu/Procaccia/Sheffet 12]

The Model

Step 1: seller gets s samples $\mathbf{v}_1,\ldots,\mathbf{v}_s$ from

- 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 $\boldsymbol{F}$

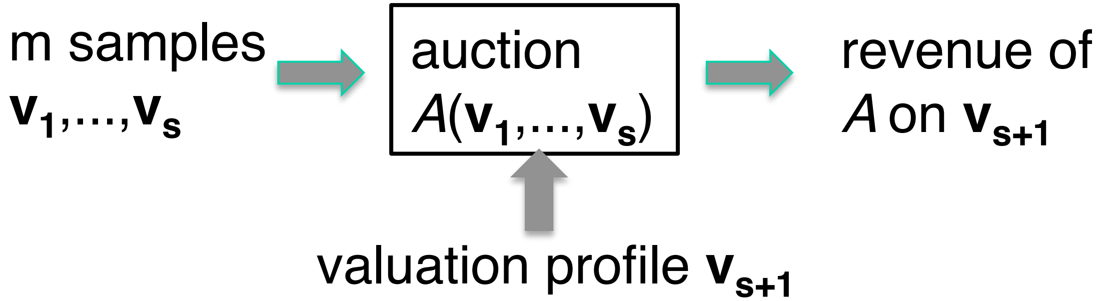

Goal: design A so $E_{v_1,\ldots,v_s}[E_{v_{s+1}}[\mathrm{Rev}(A(v_1,\ldots,v_s)(v_{s+1}))]]$ close to OPT

Desired Result

Agenda: [Morgenstern/Roughgarden 15,16] meta-theorem that for "simple" classes of mechanisms, can learn a near-optimal mechanism from few samples.

But what makes a mechanism "simple" or "complex"?

What Is...Simple?

Simple vs. Optimal Theorem [Hartline/Roughgarden 09] (extending [Chawla/Hartline/Kleinberg 07]): in single-parameter settings, independent but not identical private valuations:

expected revenue of VCG with monopoly reserves ≥ ½ •(OPT expected revenue)

What Is...Simple?

[Babaioff/Immorlica/Lucier/Weinberg 14] for a single buyer, k items, additive and independent valuations:

better of selling the grand bundle or selling items separately $\geq$ constant •(OPT expected revenue)

- [Yao 15], [Rubinstein/Weinberg 15], [Cai/Zhao 17] extend to subadditive valuations, multiple buyers.

Pseudodimension: Examples

Proposed simplicity measure of a class C of mechanisms: *pseudodimension* of the real valued functions (from valuation profiles to revenue) induced by C.

Examples:

- Vickrey auction, anonymous reserve — O(1)
- Vickrey auction, bidder-specific reserves — O(n log n)
- grand bundling/selling items separately — O(k log k)
- virtual welfare maximizers — unbounded

Pseudodimension: Implications

Theorem: [Haussler 92], [Anthony/Bartlett 99] if C has low pseudodimension, then it is easy to learn from data the best mechanism in C.

- obtain $s = \tilde{\Omega}(H^2 \varepsilon^{-2} d)$ samples $\mathbf{v}_1, \ldots, \mathbf{v}_s$ from F, where d = pseudodimension of C, valuations in [0,H]
- let M^* = mechanism of C with maximum total revenue on the samples ("ERM" learning algorithm)

Guarantee: with high probability, expected revenue of M^* (w.r.t. F) within ε of optimal mechanism in C.

Part II: Computational Aspects
(includes joint work with Josh Wang)

Offline Problem: Definition

ERM algorithm: pick the auction that is best on the samples.

Question: how to do this efficiently?

Input: valuation profiles $\mathbf{v}_1,\ldots,\mathbf{v}_T$.

Output: auction of C with maximum total revenue (i.e., $\text{argmax}_C \, \Sigma_t \, \text{Rev}(A, \mathbf{v}_i)$)).

Computing Multiple Reserves

Harder: Vickrey with *bidder-specific* reserves.

- 1 item, n bidders, reserve prices $r_1,\ldots,r_n$
- winner i: highest bidder who clears reserve r_i
- price: max{2^{nd}-highest bid clearing its reserve, r_i}

Note: complexity of offline problem not obvious.

  – exponential (in n) reserve price vectors to check

Idea: optimize each reserve r_i separately.

Tricky Example

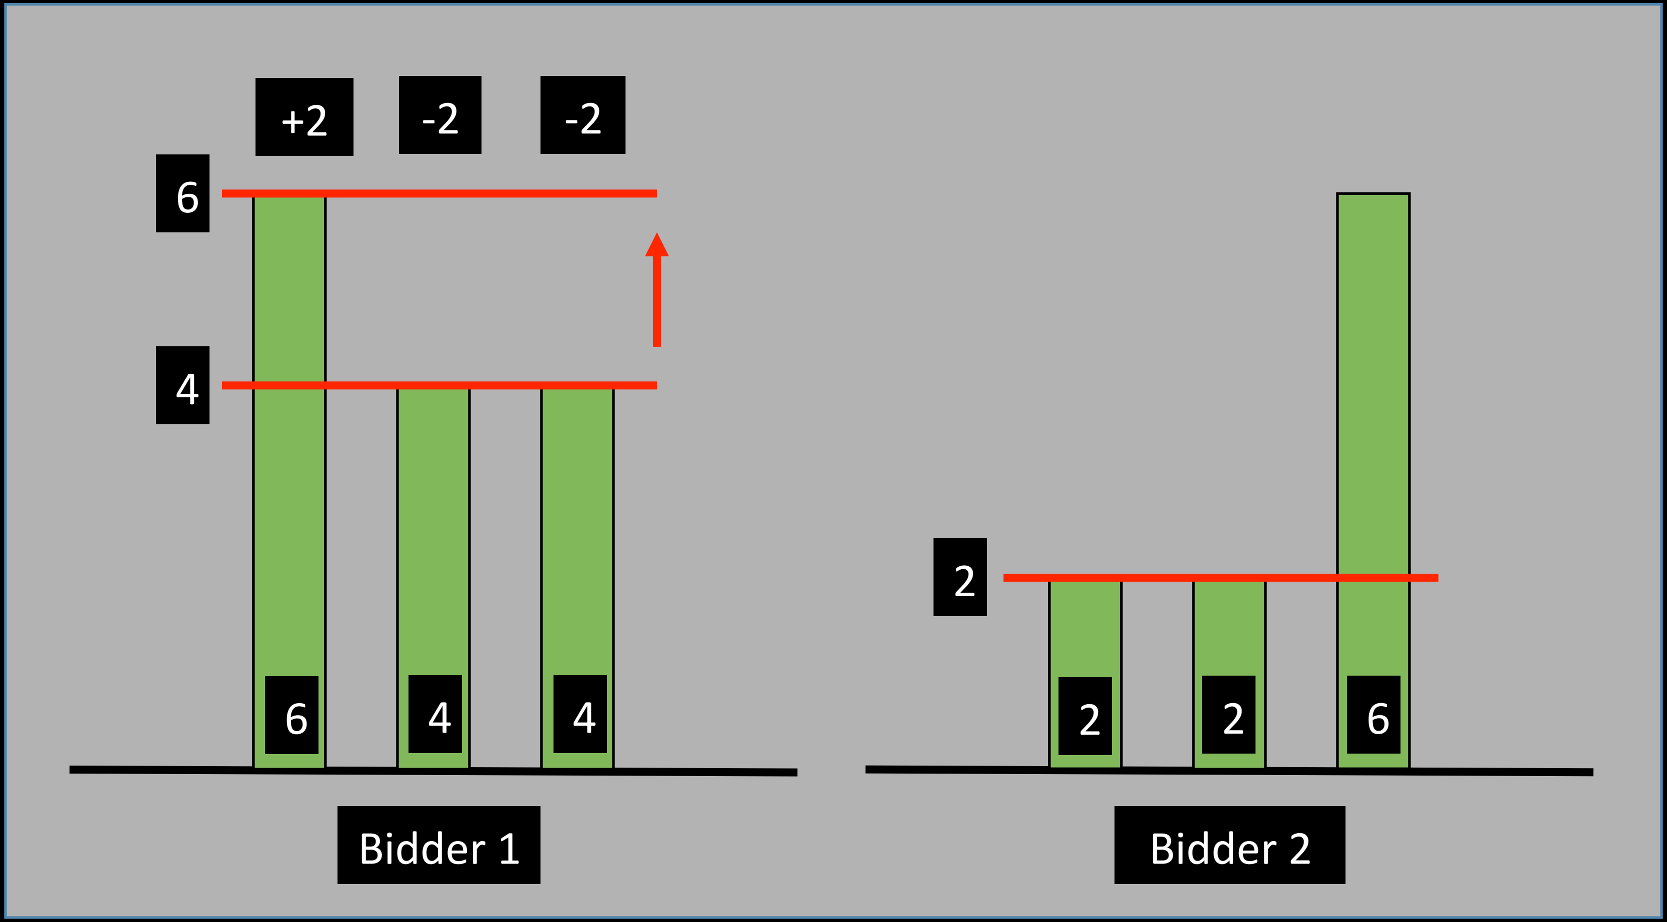

Computing Multiple Reserves

Harder: Vickrey with *bidder-specific* reserves.

- 1 item, n bidders, reserve prices $r_1,\ldots,r_n$
- winner i: highest bidder who clears reserve r_i
- price: max{2^{nd}-highest bid clearing its reserve, r_i}

Theorem: [Roughgarden/Wang 16, Paes Leme/Pal/Vassilvitskii 16] NP- and APX-hard.

An Approximation Algorithm

Option 1: use all-zero reserves.

– revenue = sum of 2^{nd}-highest bids

Option 2: optimize each bidder separately.

- for bidder i, day t, possible reserve z, define $q_{it}(z)$ = any extra revenue (beyond second-highest bid) extracted from i on day t if $r_i = z$
  – non-zero if and only if i is the highest bidder on day t, z is between first- and second-highest bids
- set r_i to maximize $\Sigma_t\, q_{it}(z)$ [try all T possibilities]

Approximation Guarantee

Theorem: [Roughgarden/Wang 16] The better of Options 1 and 2 earns revenue at least 50% of the optimal reserves.

- tight for our algorithm
- extends to matroid environments

Key Lemma: revenue earned by $r^*_1,\dots,r^*_n$ at most

$$\underbrace{\Sigma_t \,(2^{nd}\text{-highest bid on day } t)}_{\text{achieved by all-zero reserves (Option 1)}} + \underbrace{\Sigma_i \,\Sigma_t \, q_{it}(r^*_i)}_{\text{maximized by Option 2}}$$

Open Questions

1. Achieve a better-than-½ approximation (with a different algorithm).
2. Non-trivial results for non-matroid settings.
3. Non-trivial results for t-level auctions, $t \geq 2$. [Morgenstern/Roughgarden 15]
   - winner = whoever clears the most reserves (break ties by bid)
   - Vickrey w/bidder-specific reserves ⇔ 1-level auctions

Learning Reserves Online

Online version of problem:

For t=1,2,...,T:

- online algorithm chooses auction A_t in C
  - randomization allowed (and essential)
- adversary chooses valuation profile $\mathbf{v_t}$

Goal: small expected α-*regret*: ([Kakade/Kalai/Ligett 09])

$$\alpha \bullet [\underbrace{(\text{argmax}_C \Sigma_t \text{Rev}(A,\mathbf{v_t})}_{\text{best in hindsight (offline problem)}}/T] - [(\Sigma_t \text{Rev}(A_t,\mathbf{v_t}))/T]$$

Extension to Online Algorithm

Option 1: use all-zero reserves.

Option 2: optimize each bidder separately.

- set r_i to maximize $\Sigma_{s<t}\, q_{is}(z) + \text{noise}_{iz}$
  - only need to consider t different values for z
  - inspired by FTPL [Kalai/Vempala 05]

Output: randomize 50/50 between options.

Theorem: [RW16] has vanishing ½-regret.

- moral reason: based on a "maximal-in-range" offline algorithm (exactly optimizes upper bound on OPT)

Black-Box Reductions

Question: general “offline-to-online” reduction?:

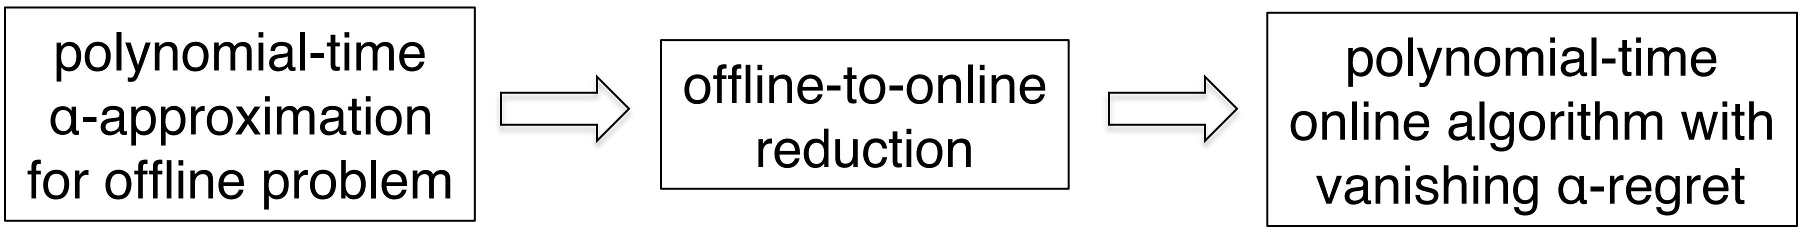

Comment: reverse reduction holds (see [Roughgarden/Wang 16] and [Daskalakis/Syrgkanis 16]).

Black-Box Reductions

Question: general "offline-to-online" reduction?:

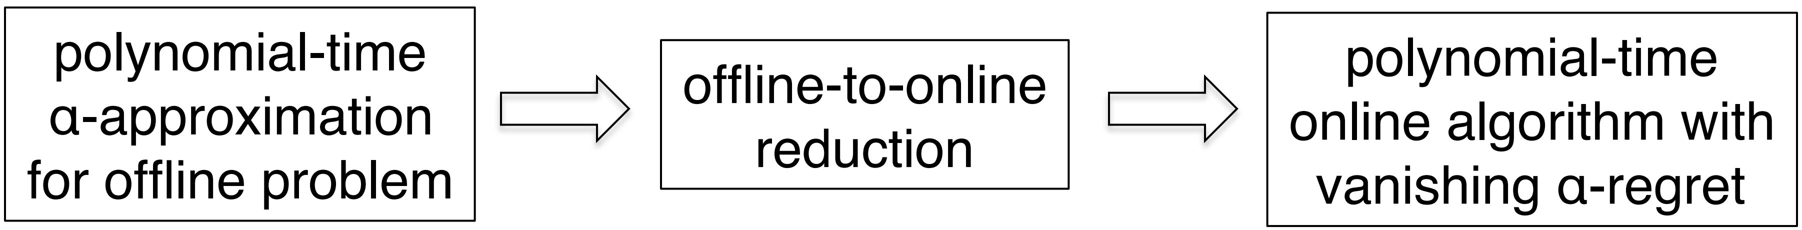

Answer: no. [Hazan/Koren 16]

Black-Box Reductions

Question: general "offline-to-online" reduction?:

polynomial-time α-approximation for offline problem → offline-to-online reduction → polynomial-time online algorithm with vanishing α-regret

Answer: yes, for:

- linear problems [Kakade/Kalai/Ligett 09]
- for certain problems, "maximal-in-range" approximation algorithms [Dudik/Haghtalab/Luo/Schapira/Syrgkanis/Wortman Vaughan 17] (generalizes [Roughgarden/Wang 16])
- submodular max/min: [Golovin/Krause/streeter 09], [Hazan/Kale 12], [Fujita/Hatano/Takimoto 13], [Roughgarden/Wang 17]

Part III: Strategic Aspects
(includes joint work with Okke Schrijvers)

Persistent, Strategic Bidders

Assumption to far: bidders bid truthfully (myopic, or only participate in one auction).

Worry: if bidders participate in multiple auctions and are strategic, might not be truthful.

Recent developments: [Amin/Rostamizadeh/Syed 13, 14], [Devanur/Peres/Sivan 15], [Ashlagi/Daskalakis/Haghtalab 16], [Immorlica/Lucier/Pountourakis/Taggart 17], [Kanoria/Nazerzadeh 17], [Braverman/Mao/Schneider/Weinberg 17]

Classical Online Prediction

Classical model:

- want to predict a sequence of binary events
  - rain vs. sunny, stock goes up vs. goes down
- at each t=1,2,...,T:
  - each of n "experts" reports belief b_i in [0,1]
  - learning algorithm makes prediction p in [0,1]
  - binary event x revealed
  - algorithm incurs loss L(x,p), expert i incurs loss L(x, b_i)

Classical Guarantee

Weight-Based Algorithm: (e.g., "Multplicative Weights")

- maintain a weight for each expert
  - weight increasing function of past accuracy
- take weighted average of predictions

Theorem: [Littlestone/Warmuth 94, Freund/Schapire 97] for a suitable weight update function (e.g. "Multiplicative Weights"), time-averaged loss of weight-based prediction at most that of best expert in hindsight (+o(1) as T→infinity).

Strategic Online Prediction

Our model:

- at each t=1,2,...,T:
  - each of n "experts" has a true belief b_i in [0,1]
  - each expert strategically reports belief r_i in [0,1]
  - learning algorithm makes prediction p in [0,1]
  - binary event x revealed
  - algorithm incurs loss L(x,p), expert i incurs loss L(x, b_i)

Assumption: each expert wants to maximize her own weight (reputation/credibility/rating/etc.).

- assume weight-update-based algorithm

Summary of Results

Results: [Roughgarden/Schrijvers 17]

- if L = squared loss, then standard multiplicative weights algorithm is incentive-compatible
- if L = absolute loss, standard MW is not IC
- if L = absolute loss, *no* online weight-based-algorithm matches regret guarantee of MW
- but can use weight update based on spherical scoring rule to get within small constant factor

Open Questions

1. Tight sample complexity bounds for multi-parameter problems (beyond [Cai/Daskalakis 17]).
2. Better/more general approximation algorithms for computing good reserve prices offline.
3. More general offline-to-online reductions (or impossibility results for natural problems).
4. With persistent and strategic bidders, when can (approximately) optimal revenue be obtained?