Tutorial: Foundations of Non-truthful Mechanism Design

Part II: Non-truthful Sample Complexity

Part III: Simplicity, Robustness, the Revelation Gap

Tutor: Jason Hartline

Schedule:

Part II: 10-10:45am (http://ec20.sigecom.org/tech/tutorial)

Part III: 11-11:45am (http://ec20.sigecom.org/tech/tutorial)

Protocol:

During session, panelest will answer clarifying questions in chat.

In post-session Q/A, "raise hand" to ask question.

Tutorial Cochairs

Brendan Lucier

Sigal Oren

Panelists

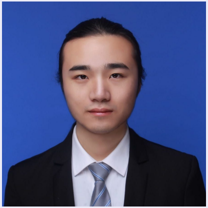

Yiding Feng

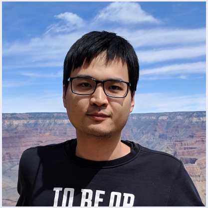

Yingkai Li

Foundations of Non-truthful Mechanism Design

http://jasonhartline.com/tutorial-non-truthful/

Jason Hartline

Northwestern University

hartline@northwestern.edu

EC Tutorial 2020

Context: The Revelation Principle

Mechanism Design: identify mechanism that has good equilibrium.

Revelation principle: if exists mechanism with good equilibrium, then exists mechanism with good truthtelling equilibrium. [Myerson '81]

Proof: truthful mechanism can simulate equilibrium strategies in non-truthful mechanism.

Consequence: literature focuses on truthful mechanisms.

Issues:

- practical mechanisms are not truthful.
- not without loss for simple or prior-independent mechanisms.
- non-trivial to undo the revelation principle.

Goal: theory for non-truthful mechanism design.

Example: Bad Welfare for Winner-pays-bid Mechanisms

Proposition (e.g., Lucier, Borodin '10)

winner-pays-bid highest-bids-win mechanisms can have very bad equilibria.

Example (Single-minded Combinatorial Auction)

Preferences:

- m items.
- $m + 2$ agents.
- agents $i \in \{1, \ldots, m\}$ values bundle $\mathsf{S}_i = \{i\}$ at $\mathsf{v}_i = 1$.
- agents $h \in \{m+1, m+2\}$ values bundle $\mathsf{S}_h = \{1, \ldots, m\}$ at $\mathsf{v}_h = 1$.

A Nash equilibrium:

- agents $h \in \{m+1, m+2\}$ bid $\mathsf{b}_h = 1$ (one wins, one loses)
- agents $i \in \{1, \ldots, m\}$ bid $\mathsf{b}_i = 0$ (all lose)
- all agent utilities $= 0$ for bids ≤ 1.

Nash welfare $= 1$; optimal welfare $= m$. **Goal for Part II:** $\mathrm{OPT} - \epsilon$

Sample Complexity in Mechanism Design

Story: Use past bid data to improve mechanism.

Definition (Truthful Sample Complexity)

Number of samples $N(\epsilon)$ from value distribution sufficient to identify truthful mechanism with expected performance at least $\text{OPT} - \epsilon$.

Observation: if designer ran truthful mechanism, can reoptimize truthful mechanism from truthful data.

Practical Issue: $> 99\%$ of mechanisms in real life are non-truthful.

- past bid data is non-truthful.
- need to design non-truthful auction.

Main Challenges:

- inference of values from bids requires strong assumptions on value distribution and mechanism.
- non-trivial to design Bayes-Nash equilibria in non-truthful mechanisms

Part II

Non-truthful Sample Complexity

Running Example

Running Example: three agents, highest-bids-win, winner-pays-bid

- Auction A: one unit.
- Auction B: two units.
- Auction C: mix 0.5A + 0.5B.

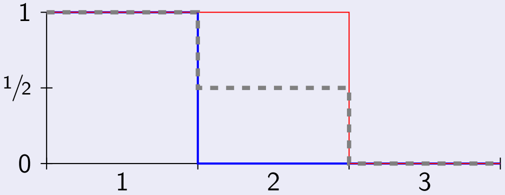

Qstn Given equilibrium bid data for C, estimate revenues of A and B?

Assumption: bids are in equilibrium, i.e, in best response to competing bid distribution.

Econometrics Observation

competing bid distribution is in observed data.

Approach:

1. given bid distribution, solve for bid strategy.
2. invert bid strategy to get agent's value from bid.

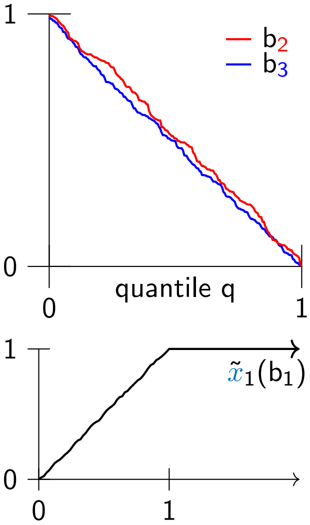

Bid Inversion

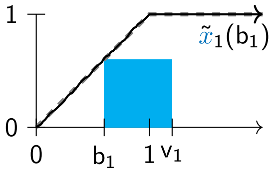

Example: How should agent 1 bid in Auction C?

- What's expected utility w. value v and bid b?

$$\mathbf{E}[\text{utility}(v, b)] = (v - b) \times \mathbf{Pr}[1 \text{ wins w. bid } b]$$
$$\approx (v - b) \times b = v\,b - b^2$$

- to maximize: take derivative $\frac{d}{db}$, set to zero, solve
- optimal to bid $b = {}^{v}/_{2}$ (bid half your value!)

Conclusion 1: Infer that agent with bid b has value $v = 2\,b$

Recall: Bids uniform on $[0, 1]$

Conclusion 2: Values are uniform on $[0, 2]$.

Revenue:

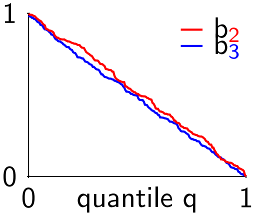

- value order statistics evenly divide interval:
  $\mathbf{E}\left[v_{(1)}\right] = {}^{3}/_{2}$; $\mathbf{E}\left[v_{(2)}\right] = 1$; $\mathbf{E}\left[v_{(3)}\right] = {}^{1}/_{2}$
- revenue equivalence:
  $\mathbf{Rev}[A] = \mathbf{E}\left[v_{(2)}\right] = 1$; $\mathbf{Rev}[B] = \mathbf{E}\left[2\,v_{(3)}\right] = 1$;
  $\mathbf{Rev}[C] = {}^{1}/_{2}\,\mathbf{Rev}[A] + {}^{1}/_{2}\,\mathbf{Rev}[B] = 1$

Section 2

Inference for I.i.d. Position Auctions

References:

1. Guerre, Perrigne, Vuong (2000) “Optimal nonparametric estimation of first-price auctions”
2. Chawla, Hartline, Nekipelov (2017) “Mechanism Redesign”

I.i.d. Position Auctions

Definition (I.i.d. Winner-pays-bid Position Auction)

m positions with weights $w_1 \geq \cdots \geq w_m$; m agents with iid values $v_j \sim F$

1. Agents submit bids.
2. Agents assigned to positions in decreasing order of bid.
3. Agent in position j wins with probability w_j.
4. Winners pay their bids.

Goal

From bids in position auction C, estimate revenue of position auction B.

Quantile Space, Revenue Curves, Expected Revenue

Definition (Inv. Demand Curve)

$v(\mathrm{q}) = F^{-1}(1 - \mathrm{q})$ is the value of an agent with quantile $\mathrm{q} \in [0, 1]$.

Definition (Revenue Curve)

$R(\hat{\mathrm{q}}) = \hat{\mathrm{q}}\, v(\hat{\mathrm{q}})$ is revenue from posting price with sale prob. $\hat{\mathrm{q}}$.

Example (Uniform Distribution)

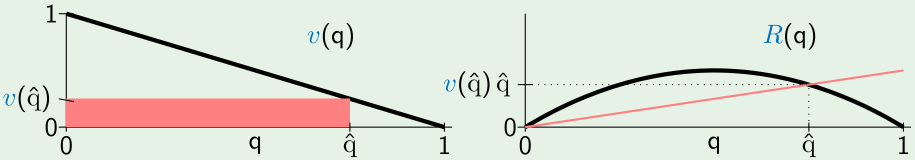

Def: Quantile allocation rule: $y(\mathrm{q}) = x(\mathrm{v}(\mathrm{q}))$

Thm: Expected revenue of alloc. y for agent w. R is: $-\int_0^1 R(\hat{\mathrm{q}})\, y'(\hat{\mathrm{q}})\, d\hat{\mathrm{q}}$

Pf: view y as cdf of critical quantile $\hat{\mathrm{q}}$ with density: $-y'$.

Classical Revenue Inference, Revisited

Inference Equation: for winner-pays-bid auction C:

$$\hat{v}(q) = \hat{b}_C(q) + \frac{y_C(q)\,\hat{b}'_C(q)}{y'_C(q)}$$

Notes:

- allocation rule y_C and derivative y'_C known. (from auction defn)
- estimated bid function $\hat{b}_C$ observed; derivative $\hat{b}'_C$ estimated.

Auction Theory: expected revenue of auction B:

$$\hat{R}_B = -\int_0^1 \hat{v}(q)\, q\, y'_B(q)\, dq$$

Estimators: for N samples from b

- empirical $\hat{b}_C$ has rate $\sqrt{N}$.
- standard $\hat{b}'_C$ estimator has rate worse than $\sqrt{N}$.
- $\Rightarrow$ revenue $\hat{R}_B$ estimator has rate worse than $\sqrt{N}$.

Direct Approach [Chawla, Hartline, Nekipelov '17]

Inference Equation: for winner-pays-bid auction C:

$$\hat{v}(q) = \hat{b}_C(q) + \frac{y_C(q)\,\hat{b}'_C(q)}{y'_C(q)}$$

Auction Theory: expected revenue of auction b:

$$\hat{R}_B = \int_0^1 \hat{v}(q)\,q\,y'_B(q)\,dq$$

Step 1: Combine:

$\hat{R}_B = \int_0^1 \left(\hat{b}_C(q) + \frac{y_C(q)\,\hat{b}'_C(q)}{y'_C(q)}\right) q\,y'_B(q)\,dq$

Step 2: Simplify with integration by parts (Define $W_{C,B}$):

$\hat{R}_B = \int_0^1 W_{C,B}(q)\,\hat{b}_C(q)\,dq$

Step 3: bound $\mathbf{E}\left[|R_B - \hat{R}_B|\right] = \mathbf{E}\left[|\int_0^1 W_{C,B}(q)\,(b_B(q) - \hat{b}_B(q))\,dq|\right]$

Step 4: estimator for N sorted bids is $\hat{R}_B = \sum_i W_{C,B}(\frac{i}{N+1})\,\hat{b}_{i,C}$

Section 3

General Reduction to I.i.d. Position Auctions

References:

1. Chawla, Hartline (2013) "Auctions with unique equilibria"
2. Chawla, Hartline, Nekipelov (2017) "Mechanism Redesign"
3. Hartline, Taggart (2019) "Sample Complexity for Non-truthful Mechanisms"

Definitions for Non-truthful Sample Complexity

Definition (Independent Single-Dimensional Environment)

- n agents, values $v_i \sim F_i$, $\boldsymbol{F} = F_1 \times \cdots \times F_n$.
- feasible allocations $\mathbf{x} = (x_1, \ldots, x_n) \in \mathscr{X} \subset [0,1]^n$

Definition (Batched Environment)

An batched environment for n populations and m stages is Cartesian product with nm agents. Cf. online environment.

Definition (I.i.d. Winner-pays-bid Position Auction)

m positions with weights $w_1 \geq \cdots \geq w_m$; m agents with iid values $v_j \sim F$

1. Agents submit bids.
2. Agents assigned to positions in decreasing order of bid.
3. Agent in position j wins with probability w_j.
4. Winners pay their bids.

Theorems for Non-truthful Sample Complexity

Theorem (Chawla, Hartline '13)

i.i.d. winner-pays-bid position auction: BNE is unique, symmetric, efficient.

Theorem (Chawla, Hartline, Nekipelov '17)

For i.i.d. position auctions B and C and values in $[0, 1]$*:*

> ϵ *error in welfare/revenue estimate of auction B with* $N(\epsilon) = \tilde{O}(m^4/\epsilon^2)$ *samples from BNE bids from auction C.*

Theorem (Hartline, Taggart '19)

Batched non-iid single-dimensional mechanism design $(1 - \epsilon)$*-approx. reduces to i.i.d. position auction with batch size* $M(\epsilon) = n/\epsilon^3$*.*

Corollary (Batch, Sample Complexity)

ϵ *revenue/welfare loss w. batch, sample size* $M(\epsilon) = n^4/\epsilon^3$, $N(\epsilon) = \tilde{O}(n^{16}/\epsilon^{14})$

Batch ⇒ IID Position Auction

Theorem (Hartline, Taggart '16,'19)

Batched non-iid single-dimensional mechanism design $(1-\epsilon)$-approx. reduces to i.i.d. position auction with batch size $M(\epsilon) = n/\epsilon^3$.

Main idea:

- batched env. is m i.i.d. single-dimensional auctions with n agents.
- convert to n position auctions on m i.i.d. agents.

Definition (Surrogate Ranking Mechanism)

population $i \in [n]$; stage $j \in [m]$; surrogate values $\{\Phi_i^1 \geq \cdots \geq \Phi_i^m\}_{i\in[n]}$.

1. solicit bids: $\{\mathrm{b}_i^j\}_{i\in[n]}^{j\in[m]}$;
2. compute ranks of each agent ij among population i bids $\{\mathrm{b}_i^j\}^{j\in[m]}$: r_i^j.
3. maximize surrogate welfare in each stage j: $\mathbf{x}^j = \operatorname{argmax}_{\mathbf{x}\in\mathscr{X}} \sum_i \Phi_i^{\mathrm{r}_i^j} \mathrm{x}_i$
4. charge winners their bids.

Note Optimal surrogate values are expected order statistics.

Part III

Simplicity, Robustness, & the Revelation Gap

Prior-independent Mechanism Design

Motivation: understand mechanisms that are robust to variation in distribution of preferences.

Cf. [Wilson '87] [Bergemann, Morris '05] [Carroll '15]

Prior-independent Mechanism Design

$$\min_{\mathscr{M} \in \text{MECH}} \max_{F \in \text{DIST}} \frac{\mathbf{E}_{\mathbf{v} \sim F}[\text{OPT}_F(\mathbf{v})]}{\mathbf{E}_{\mathbf{v} \sim F}[\mathscr{M}(\mathbf{v})]}$$

Notation

- MECH: family of mechanisms.
- DIST: family of type distributions.
- $\mathbf{v} = (v_1, \ldots, v_n)$: profile of private types.
- OPT_F: optimal mechanism for type distribution F.
- $\mathscr{M}(\mathbf{v})$: welfare/revenue of mechanism on private types $\mathbf{v}$.

Revelation Principle vs. Prior-independence

Mechanism Design: identify mechanism that has good equilibrium.

Revelation principle: if exists mechanism with good equilibrium, then exists mechanism with good truthtelling equilibrium. [Myerson '81]

Observation: the construction of the revelation principle breaks prior-independence.

Question: are non-truthful mechanisms better than truthful mechanisms for prior-independent mechanism design?

Section 4

Revelation Gap

References:

1. Feng, Hartline (2018) "An End-to-End Argument in Mechanism Design (Prior-Independent Auctions for Budgeted Agents)"
2. Feng, Hartline, Li (202?) "A Revelation Gap for Pricing from Samples"
3. Hartline (202?) "Mechanism Design and Approximation" Chapter 5

Pricing from Samples

Model: Pricing from Samples

- single item, single buyer.
- buyer has private valuation $\mathsf{v} \sim F$.
- F is monotone hazard rate (MHR), i.e., $\frac{f(z)}{1-F(z)}$ is non-decreasing.
- seller has access to a single sample $s \sim F$.

Goal: approximate the optimal revenue (when F is known)

Revelation Gap for MHR distribution

Theorem (Allouah, Besbes '19)

For monotone hazard rate distributions, the prior-independent approx. of truthful pricing from a sample is between 1.543 *and* 1.575.

Theorem

For monotone hazard rate distributions, the prior-independent approx. of (non-truthful) pricing from a sample is between 1.073 *and* 1.296.

Corollary

For monotone hazard rate distributions, the revelation gap for pricing from a sample is between 1.19 *and* 1.47.

Revelation Gap for MHR distribution

Theorem (Lower Bound)

For uniform distributions (including pointmasses), the prior-independent approximation of pricing from a sample is at least 1.07.

Theorem (Upper Bound)

For monotone hazard rate distributions, exists non-truthful mechanism with prior-independent approximation ratio at most 1.296.

Definition (Sample Pricing Mechanism)

1. Let the agent decide to participate or not.
2. A participating agent receives the item and pays $\alpha \cdot s$.

Buyer behavior

Participates if $v \geq \alpha \cdot w$, where $w = \mathbf{E}_{s \sim F}[s]$.

Proof Sketch

Theorem (Upper Bound)

For monotone hazard rate distributions, exists non-truthful mechanism with prior-independent approximation ratio at most 1.296.

Proof sketch:

1. Lower bound probability of participating in two cases:

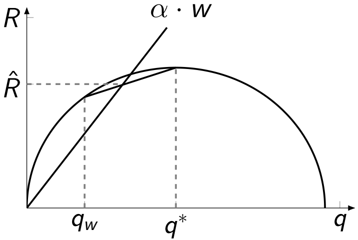

Case 1: $\alpha\, w \geq v(\hat{q}^\star)$

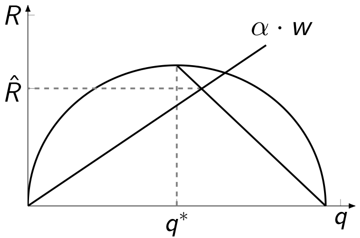

Case 2: $\alpha\, w \leq v(\hat{q}^\star)$

2. Best $\alpha = 0.78$ gives approximation at most 1.39.
   (1.296 approx follows from better analysis using curvature) □

Section 5

Implementation Theory

References:

1. Jackson (2001) "A crash course in implementation theory"
2. Caillaud, Robert (2005) "Implementation of the revenue-maximizing auction by an ignorant seller"

Mechanism Design for an Ignorant Seller [cf. Jackson '01]

Proposition (Informal)

Anything commonly known by the agents, the mechanism can be assumed to know.

Definition (Report-the-prior Mechainsm)

1. Solicit prior.
2. "shoot agents if they disagree".
3. Run optimal mechanism for reported prior.

Discussion:

1. possesses an optimal equilibrium.
2. possesses other equilibria (but there are tricks for removing them).
3. begs the question.

Revenue Maximization with a Prior [Myerson '81]

Consider selling a single-item to agents with values $\mathbf{v} \sim \boldsymbol{F}$.

Definition (Ascending Virtual Price Mechanism)

Given monotone virtual value function $\boldsymbol{\phi} = (\phi_1, \ldots, \phi_n)$

1. raise a virtual price ϕ from 0
   (where agent i's price is $\hat{v}_i = \phi_i^{-1}(\phi)$)
2. when one bidder remains, sell at her price.

Theorem

For any distribution $\boldsymbol{F}$, there are ϕ for which the ascending virtual price mechanism is revenue optimal.

Mechanism Design for an Ignorant Seller [Caillaud, Robert '05]

Definition (Belief Free Ascending Mechanism, BFA)

1. run ascending mechanism w. uniform price ϕ until one agent remains.
2. remaining agent i can offer to increase the price to $p \geq \phi$.
3. a random agent j is allowed to challenge at price $q > p$.
4. if no challenge: i pays p; if challenge: i pays Δ
5. if i accepts challenge: i pays p to seller and $q - p$ to challenger
6. if i rejects challenge: challenger j pays $p - \phi$ to seller.

Thm: BFA admits a revenue-optimal equilibrium.

Proof.

The following is an equilibrium:

- Agents remain in ascending auction until, for i: $\phi_i^{-1}(\phi) > \mathrm{v}_i$.
- Remaining agent i offers $p = \phi_i^{-1}(\phi)$.
- If $p < \phi_i^{-1}(\phi)$ then challenger j challenges $q = \phi_i^{-1}(\phi)$
- Agent i accepts challenges $q < \mathrm{v}_i$.

Conclusion

Conclusion

- Strange non-truthful mechanisms for ignorant sellers.
- Need to consider prior-independent non-truthful carefully.

Directions

- single-agent sample-based pricing [e.g., Feng, Hartline, Li]
- e.g., restrict to single-round, winner-pays-bid mechanisms.

Tutorial: Foundations of Non-truthful Mechanism Design

Tutor: Jason Hartline

Part I: Equilibrium Analysis

1. Single-dimensional Environments
2. Revenue Equivalence and Applications
3. Robust Analysis of Equilibria

Part II: Non-truthful Sample Complexity

1. Counterfactual Estimation
2. Inference for I.i.d. Position Auctions
3. General Reduction to I.i.d. Position Auctions

Part III: Simplicity, Robustness, & the Revelation Gap

1. Revelation Gap
2. Implementation Theory