Algorithmic Game Theory

Mechanisms for Revenue Maximization

Vangelis Markakis
markakis@gmail.com

Designing mechanisms for maximizing revenue for single parameter environments

Social Welfare vs Revenue

- Many reasons for focusing on social welfare:
  - In government auctions, revenue may not be the first priority
  - Also in competitive markets, greedily maximizing revenue may cause customers leave towards other sellers
- Strong positive results for social welfare maximization.
  - If we do not care about computational efficiency, we can **always** have a truthful mechanism that maximizes welfare (VCG)
  - We often have good approximations in polynomial time.
- Question: Similar results for **revenue** maximization?

Social Welfare vs Revenue

An illuminating example:

- Consider 1 item and only 1 bidder with private value v
- Only truthful mechanisms: set a price r **independent of the declared bid**
  - Any other pricing rule that depends on the bid is not **truthful**
  - These are called **posted price** mechanisms.
- If v ≥ r, the bidder will buy the item, sw = v and revenue = r
- If not, sw = revenue = 0

Social Welfare vs Revenue

- How do we maximize social welfare in this setting?
  - Easy, just set r = 0
  - All we care about for social welfare is that the bidder gets the item
  - We do not need to know the exact value of v
  - With more bidders, we also do not need to know the exact values to maximize welfare, only who is the highest bidder
  - Social welfare is quite special and relatively simple.

Social Welfare vs Revenue

- How do we maximize revenue?
  - Optimal revenue we can extract: equal to v
  - If we knew v, we would just set r:= v
  - But v is private information!
  - Optimal revenue really depends on the exact form of the valuation function
  - E.g., if we just set r = 100, then the mechanism does well only for bidders with v ≥ 100 (and not too large!).
    - For v < 100, it performs terribly

A Model for Revenue Maximization

Conclusions and modeling approach:

- Not easy to compare mechanisms
- We need to consider a **different model**
- Usual approach: Average case or Bayesian analysis

A Bayesian Model for Revenue Maximization

For single-parameter environments:

- Each bidder i has a value v_i which is private information
- For each bidder i, the value v_i is drawn from a **probability distribution F_i** on some interval $[0, v_{max}]$, with $v_{max} \neq +\infty$
  - $F_i(z) = Pr[v_i \leq z]$
- The distributions $F_1, F_2, ..., F_n$ are all independent
- Mechanism knows the distributions (but not the values)
  - Typically derived from historical data
- Objective: design an auction to maximize **expected** revenue

Goal: Characterize truthful mechanisms that maximize expected revenue.

A Bayesian Model for Revenue Maximization

Back to single item and single bidder

- Value v of the bidder drawn from distribution F
- Suppose we post a price r
- Expected revenue = $r \cdot \Pr[v \geq r] = r \cdot (1 - F(r))$
- It reduces to optimizing posted price r
    - Optimal price r is called **monopoly price** of F.
- E.g., if F is uniform in [0, 1], then $F(z) = z$
- Optimal mechanism: post $r = 1/2$ with expected revenue 1/4

A Bayesian Model for Revenue Maximization

Single-item auction with two bidders?

- This already gets more complex
- Can we start with something simple first?

2nd price auction with a **reserve** price

- Fix a reserve price r
- Allocation rule: If no bidder exceeds r, nobody gets the item. Otherwise, winner is the highest bidder
- Payment rule: max {reserve price, 2nd highest bid}

A Bayesian Model for Revenue Maximization

Single-item auction with two bidders:

- Reserve prices are used in practice to boost revenue
- Main **advantage**: much better revenue for the cases where 2nd highest bid is low
- Main **disadvantage**: in some cases nobody wins (no revenue)
  - Hopefully latter happens with small probability
- Is the optimal mechanism very far from such a format?

Expected Revenue for Single-Parameter Bidders

- We focus on single-parameter bidders, monotone allocations and Myerson's truthful payments.
  - Due to truthfulness, bids = true values.
- Maximize $\mathbb{E}_{v_1 \sim F_1, \ldots, v_n \sim F_n} \left[\sum_{i=1}^{n} \boldsymbol{p}_i(\boldsymbol{v}) \right] = \sum_{i=1}^{n} \mathbb{E}_{\boldsymbol{v}_{-i}} \left[\mathbb{E}_{v_i} [\boldsymbol{p}_i(v_i, \boldsymbol{v}_{-i})] \right]$
  - Due to independence, we focus on single bidder i.
- We recall (dfn of expectation and Myerson's payments):

$$\mathbb{E}_{v_i}[\boldsymbol{p}_i(v_i, \boldsymbol{v}_{-i})] = \int_0^{v_{\max}} \boldsymbol{p}_i(v_i, \boldsymbol{v}_{-i}) f(v_i) dv_i \qquad p_i(v_i, \boldsymbol{v}_{-i}) = \int_0^{v_i} z \cdot \boldsymbol{x}_i'(z, \boldsymbol{v}_{-i}) dz$$

- Therefore:

$$\mathbb{E}_{v_i}[\boldsymbol{p}_i(v_i, \boldsymbol{v}_{-i})] = \int_0^{v_{\max}} \boldsymbol{p}_i(v_i, \boldsymbol{v}_{-i}) f(v_i) dv_i = \int_0^{v_{\max}} \left[\int_0^{v_i} z \cdot \boldsymbol{x}_i'(z, \boldsymbol{v}_{-i}) dz \right] f(v_i) dv_i$$

Expected Revenue for Single-Parameter Bidders

- Reversing the order of integration:

$$\int_0^{v_{\max}} \left[\int_0^{v_i} z \cdot x_i'(z, \mathbf{v}_{-i}) dz\right] f_i(v_i) dv_i = \int_0^{v_{\max}} \left[\int_z^{v_{\max}} f_i(v_i) dv_i\right] z \cdot x_i'(z, \mathbf{v}_{-i}) dz$$

$$= \int_0^{v_{\max}} (1 - F_i(z)) \cdot z \cdot x_i'(z, \mathbf{v}_{-i}) dz.$$

- Integration by parts and simplification:

$$\int_0^{v_{\max}} \underbrace{(1 - F_i(z)) \cdot z}_{f} \cdot \underbrace{x_i'(z, \mathbf{v}_{-i})}_{g'} dz$$

$$= \underbrace{(1 - F_i(z)) \cdot z \cdot x_i(z, \mathbf{v}_{-i})|_0^{v_{\max}}}_{=0-0} - \int_0^{v_{\max}} x_i(z, \mathbf{v}_{-i}) \cdot (1 - F_i(z) - z f_i(z)) dz$$

$$= \int_0^{v_{\max}} \underbrace{\left(z - \frac{1 - F_i(z)}{f_i(z)}\right)}_{:=\varphi_i(z)} x_i(z, \mathbf{v}_{-i}) f_i(z) dz = \mathbb{E}_{v_i}\left[\varphi(v_i) \cdot \boldsymbol{x}_i(\boldsymbol{v})\right]$$

Virtual Valuations

We transform valuations to virtual valuations, that include information about valuation distribution.

Definition: For an agent i, with

- actual value v_i,
- distribution F_i,
- probability density function f_i,

the virtual valuation at v_i is:

$$\phi_i(v_i) = v_i - \frac{1 - F_i(v_i)}{f_i(v_i)}$$

"information rent" for agent i

Optimal revenue extracted from i

Monopoly price of F where virtual valuation is 0:

$$\Big(r(1 - F(r))\Big)' = 0 \Leftrightarrow r - \frac{1 - F(r)}{f(r)} = 0 \Leftrightarrow \varphi(r) = 0$$

Virtual Valuations

Example: uniform distribution on [0, 1] for player i:

- distribution function: $F_i(z) = z$
- density function: $f_i(z) = 1$
- virtual valuation: $\phi_i(v_i) = v_i – (1-v_i)/1 = 2v_i - 1$

Observations:

- Virtual valuations can also take negative values, even though $v_i \geq 0$
- For any distribution, $\phi_i(v_i) \leq v_i$

Summary: $\mathbb{E}_{\boldsymbol{v}\sim F}\left[\boldsymbol{p}_i(\boldsymbol{v})\right] = \mathbb{E}_{\boldsymbol{v}\sim F}\left[\varphi(v_i) \cdot \boldsymbol{x}_i(\boldsymbol{v})\right]$

Expected Revenue Equals Expected Virtual Welfare

Main result for revenue maximization:

Consider a single-parameter domain with valuation distributions $F_1, F_2, \ldots, F_n$ and let $F = F_1 \times F_2 \times \ldots \times F_n$ be the product distribution.

For every truthful mechanism (**x**, **p**)

$$E_{\mathbf{v}\sim F}\left[\Sigma_i\, p_i(\mathbf{v})\right] = E_{\mathbf{v}\sim F}\left[\Sigma_i\, \phi_i(v_i) \cdot x_i(\mathbf{v})\right]$$

Expected revenue

Expected virtual welfare

Surprisingly, finding the revenue-optimal mechanism reduces to maximizing the expected **virtual welfare**

Maximizing Virtual Welfare

- Although we care about payments, we reduced the problem to designing an appropriate allocation rule!
- How do we maximize expected virtual welfare?
  - Forget about the expectation and maximize pointwise.
  - For each profile **v** = $(v_1, v_2,\ldots, v_n)$, maximize $\Sigma_i\ \phi_i(v_i) \cdot x_i(\mathbf{v})$
  - This is simply a welfare maximization problem
    - With $\phi_i(v_i)$ playing the role of v_i
    - We apply Myerson's Lemma, but for the virtual values.
    - Allocation rule must be **monotone** (wrt bids / valuations v_i), as required for truthfulness.
- Whenever we can solve welfare maximization efficiently, we can also do it for the virtual welfare.

Maximizing Virtual Welfare

For the single-item auction:

- Give the item to bidder with the highest virtual value.
- Actually, not always...
- Recall: a virtual value can take negative values
- Give it to bidder with the highest **positive** virtual value
- Sometimes, the item is not allocated to anyone.
- Example: Let F_i be the uniform distribution on [0, 1]
  - $\phi_i(v_i) = 2v_i - 1$
  - Allocation rule: give it to the highest bidder whose bid exceeds 1/2 (reserve price), if such bidder exists

Monotonicity of Virtual Welfare Maximization

- Is the allocation rule that maximizes the virtual welfare **monotone** (wrt. bids)?
  - If yes, then we are done by Myerson's lemma
  - Unfortunately this depends on the distributions

Definition: A distribution is called **regular** if the corresponding virtual valuation function is non-decreasing

- Examples: the uniform distribution and many other common distributions satisfy this
- Non-regular distributions: multi-modal distributions or with heavy tails

Monotonicity of Virtual Welfare Maximization

Observation: If we have regular distributions for all bidders, then the virtual welfare maximizing rule is monotone

Optimal mechanism for revenue maximization

Assumptions: **Independent** and **regular** distributions

- Collect the bids and transform each b_i into its corresponding virtual bid $\phi_i(b_i)$
- Choose an allocation $(x_1, x_2, \ldots, x_n)$ that maximizes the virtual welfare $\Sigma_i \phi_i(b_i) \cdot x_i$
- Charge each bidder according to Myerson's payment formula

Expected Revenue Maximization

Let’s apply this to single-item auctions with i.i.d bidders

Implementing the revenue-optimal mechanism

- Collect the bids and transform each b_i into its corresponding virtual bid $\phi_i(b_i)$
- Allocation: since the virtual valuation function is non-decreasing, for i.i.d. bidders, the highest virtual value corresponds to the highest bidder
  - Thus: we allocate the item to the highest bidder i, as long as $\phi_i(b_i) \geq 0$, otherwise, there is no winner
- Payment: need to find the threshold bid, where does the jump in the allocation occur?

Expected Revenue Maximization

- Consider i.i.d. bidders with the uniform distribution on [0, 1]
- $\phi_i(z) = 2z - 1$ for every bidder i
- Let i be the winner, and fix a profile $\mathbf{b}_{-i}$ for the other bidders
- The jump in the allocation can happen either at the 2nd highest bid or at 1/2

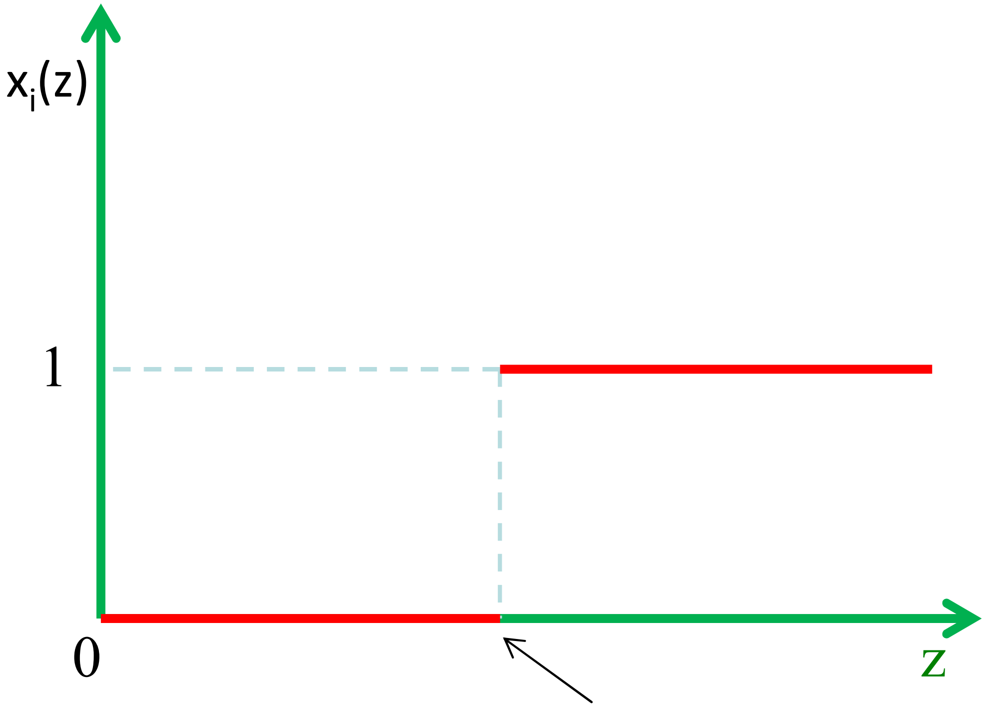

Either 2nd highest bid or 1/2

2 cases to consider:

Case 1: at least one other bidder has a positive virtual bid

Case 2: no other bidder has a positive virtual bid

Payment = max{2nd highest bid, ½}

This is a 2nd price auction with reserve price = 1/2

Expected Revenue Maximization

More generally:

- Consider a single-item auction
- Suppose we have i.i.d. bidders with a regular distribution
- Let ϕ be the common virtual valuation function

Optimal mechanism: 2nd price auction with reserve = $\phi^{-1}(0)$

- i.e., the eBay format is optimal (with appropriate opening bid)
- Surprising that the optimal mechanism has such a simple format

Single-Item Auctions with Non I.I.D. Bidders

- Things become complicated when bidders are not i.i.d.
- For example, suppose bidders' valuations are drawn independently but from from different regular distributions
- The revenue-optimal auction does not resemble any format used in practice
- It is also not easy to interpret as a natural rule to follow and does not have a practical appeal
- Current research: Identify simple auction rules for which we can prove they are near-optimal in terms of expected revenue
  - Based again on virtual valuations and on using prophet inequalities for estimating the derived revenue

Prophet Inequality and Simple Single-Item Auctions

- Let $F_1, \ldots, F_n$ be independent distributions, let $X_1, \ldots, X_n$ be realizations from $F_1, \ldots, F_n$, and let $X^* = \max_i \{ X_i \}$.
  - Let $t : \text{Prob}[X^* \geq t] = 1/2$ (or simply $t = E[X^*]/2$)
  - Then, accepting an arbitrary $X_i \geq t$ (if any) guarantees an expected reward of $\geq E[X^*]/2$.
- Choose t s.t $\mathbb{P}_{\boldsymbol{v}}\left[\max_i \varphi(v_i)^+ \geq t\right] = 1/2$ (or $t = \mathbb{E}_{\boldsymbol{v}}\left[\max_i \varphi(v_i)^+\right]/2$)
  - Threshold t can be computed (or estimated), given $F_1, \ldots, F_n$
- Give the item to arbitrary bidder i with $\phi_i(v_i) \geq t$, if any, at (i's reserve) price r_i defined as $\phi_i(r_i) = t$.
  - If many candidate winners, any monotone selection works. E.g., highest bidder.
  - Also applies if bidders arrive online and offers are take-it-or-leave-it.

Prophet Inequality and Simple Single-Item Auctions

- Choose t s.t $\mathbb{P}_{\boldsymbol{v}}\left[\max_i \varphi(v_i)^+ \geq t\right] = 1/2$ (or $t = \mathbb{E}_{\boldsymbol{v}}\left[\max_i \varphi(v_i)^+\right]/2$)
  - Threshold t can be computed (or estimated), given $F_1, \ldots, F_n$
- Give the item to arbitrary bidder i with $\phi_i(v_i) \geq t$, if any, at (i's reserve) price r_i defined as $\phi_i(r_i) = t$.
  - If many candidate winners, choose the highest bidder.
- Prophet inequality implies ≥ **50%** of optimal revenue!
  - Simple, virtual valuations determine reserves, not the winner.
  - However, reserves are still player-dependent.
- Open Problem: how much of optimal revenue we can recover with anonymous reserve prices, if bidders are independent but not identically distributed.

Prior-Independent Auctions

- Design auctions that extract significant fraction of optimal revenue without resorting to knowledge of valuation distributions $F_1, \ldots, F_n$
  - Distributions are used in the analysis of the auction, not in its design.
- Expected revenue of Vickrey auction with n+1 i.i.d. bidders from any regular distribution F ≥ expected revenue of optimal auction (Vickrey auction with optimal reserve price derived with knowledge of F) with n i.i.d. bidders from F.

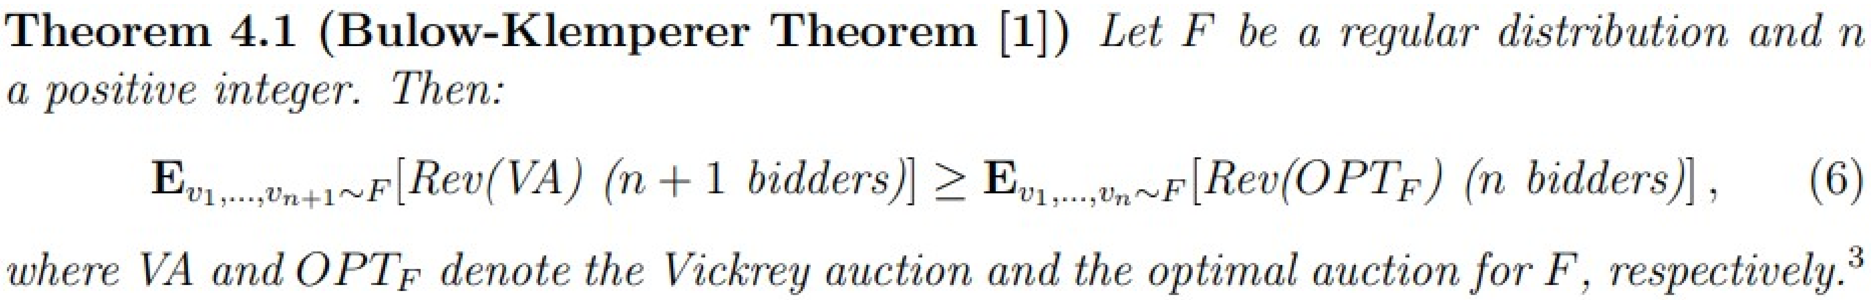

Theorem 4.1 (Bulow-Klemperer Theorem [1]) *Let F be a regular distribution and n a positive integer. Then:*

$$\mathbf{E}_{v_1,\ldots,v_{n+1}\sim F}[Rev(VA)\ (n+1\ bidders)] \geq \mathbf{E}_{v_1,\ldots,v_n\sim F}[Rev(OPT_F)\ (n\ bidders)]\,, \qquad (6)$$

where VA and OPT_F denote the Vickrey auction and the optimal auction for F, respectively.[3]

Multi-Parameter Revenue Maximization

- A much harder problem!
- Recall Myerson’s lemma does not hold any more for more complex valuations
- Not easy to characterize truthful mechanisms when the valuation functions depend on multiple private parameters of the bidders
- Very active research field even for auctions with a small number of items