TRUTH JUSTICE ALGOS

Mechanism Design: Multi-Dimensional Mechanism Design

Teachers: Ariel Procaccia and Alex Psomas (this time)

SO FAR

- Revelation Principle
- Single parameter environments
  - Second price auctions
  - Myerson’s lemma
  - Myerson’s optimal auction
  - Cremer-McLean auction for correlated buyers
  - Prophet inequalities
  - Bulow-Klemperer

TODAY

- Multidimensional environments
  - The VCG mechanism
  - Challenges
  - Revenue optimal auctions are strange

COMBINATORIAL AUCTIONS

- n buyers
- m items
  - Result holds for arbitrary feasibility sets
- Agent i has an (arbitrary) value function v_i from subsets of items to non-negative reals
  - Complementarities: $v_i(\{1\}) = 1, v_i(\{2\}) = 2, v_i(\{1,2\}) = 100$
  - Substitutes: $v_i(\{1\}) = 10, v_i(\{2\}) = 10, v_i(\{1,2\}) = 12$
  - Or anything else really: $v_i(\{1\}) = 10, v_i(\{2\}) = 10, v_i(\{1,2\}) = 0$

VCG

- Goal: find a truthful (DSIC) and feasible mechanism that maximizes social welfare $argmax_{feasible\ S=(S_1,\ldots,S_n)} \sum_i v_i(S_i)$

VCG

- Let's try our single dimensional approach
- Commit to the allocation rule x that maximizes social welfare, and then find appropriate payments
- $x(\vec{b}) = argmax_S \sum_i b_i(S_i)$
  - b_i here is a reported *valuation function,* not a number
- Let S^* be the social welfare maximizing allocation
- What about payments?
- Perhaps try to prove another "Myerson's Lemma"?
- Not clear what monotonicity is...
- Not clear what the "critical bid" is (bids are functions)…

VICKREY-CLARKE-GROVES

- Key idea: "**externality**"
- How much pain does your existence cause?
- Aka, how much does your presence hurt everyone else's value?
- That's how much you should pay
- In retrospect, same idea as second price auction
    - I pay exactly the maximum social welfare when I'm not there

VICKREY-CLARKE-GROVES

$$p_i(\vec{b}) = \max_S \sum_{j \neq i} b_j(S_j) - \sum_{j \neq \mathrm{i}} b_j(S_j^*)$$

$$\underbrace{\max_S \sum_{j \neq i} b_j(S_j)}_{\text{Without } i} \qquad \underbrace{\sum_{j \neq \mathrm{i}} b_j(S_j^*)}_{\text{With } i}$$

- $p_i(\vec{b}) \geq 0$
  - Why?
- Claim: VCG is DSIC

VICKREY-CLARKE-GROVES

Proof:

- Fix i and b_{-i}
- Let S^* be the social welfare maximizing allocation when i submits b_i

$$u_i = \underbrace{v_i(S_i^*) + \sum_{j \neq \mathrm{i}} b_j(S_j^*)}_{\text{(A)}} - \underbrace{\max_S \sum_{j \neq i} b_j(S_j)}_{\text{(B)}}$$

- Observation: (B) is out of i's control
- Thought experiment: i can directly pick any S^*
  - As opposed to picking b_i (that affects S^*)
- What would she pick?
  - Maximize (A)
- What is our mechanism doing?

CHALLENGES

- Preference elicitation
  - Bidders report their valuation function
  - For m items each bidder has to report more or less 2^m numbers
- Ok, so no direct revelation mechanisms
  - Keep everything on a "need to know" basis
- Canonical approach is the English/ascending auction
  - Have an increasing price p_i for every item.
  - When there is a single bidder left give it to them for that price

FUN STORY

[CRAMTON, SCHWARTZ 2000]

- Auction for spectrum rights in the US
- Iterative auction where companies bid for licenses (items) to broadcast over specific bands of the electromagnetic in certain areas
- Bids are public
- FCC raised ~23 billion USD from 1994 to 1998 via 16 such auctions

FUN STORY

- Mercury PCS and High Plains Wireless are having a bidding war for license #264 for Lubbock, Texas.
- In the meantime, High Plains Wireless is currently the highest bidder for license #013 for Amarillo, Texas (not much competition for this license)
- Mercury PCS outbids High Plains Wireless in #013 in round 68 (first time Mercury PCS made a bid for that area)
  - The last three digits of the bid: 264
- Next bid of Mercury PCS on #264 had last three digits “013”
- In other words “Stay away from 264, otherwise 013 will cost you a lot more”

ANOTHER FUN STORY

[MCMILLAN. SELLING SPECTRUM RIGHTS]

- What a rookie mistake…
- Just run a sealed-bid auction…
- New Zealand 1990: Television broadcasting licenses auctioned off via simultaneous sealed-bid Vickrey auctions
- Government projected it would make 250 million
- Made 34
- Even funnier:
  - A company bid 7 million NZ$ for a license. Ended up paying 5000
  - Another license had highest bid 100,000. Second highest: 6.
  - A university student bid NZ$1 for a television license for a small city; no one else bid anything so he won and paid nothing

CHALLENGES

- Computational tractability
  - Even when valuation functions are very simple, maximizing welfare can be NP-hard
  - Example: Single Minded buyers.
    - Bidder i has value v_i for some subset S_i, and zero for everything else
      - Maximizing welfare is essentially the same as finding the largest Independent Set
        - » Vertices are bidders
        - » Edges are items
        - » S_i = edges adjacent to vertex/bidder i

CHALLENGES

- Computational tractability
  - What about approximations?
- Say you have an algorithm that can find in polynomial time an α approximation to the optimal (social welfare maximizing) allocation
  - In fact we do have such algorithms for single minded bidders ($\sqrt{m}$ approximation)
- Almost-VCG payments are not incentive compatible
  - We need to be able to maximize social welfare in order to compute the externality of each agent
  - If we can only do this approximately, the payments are not the VCG payments, and truthfulness cannot be guaranteed

CHALLENGES

- Horrendous for revenue
  - As opposed to second price which was pretty good (by Bulow-Klemperer)
- 2 bidders, A and B, and 2 items
- A wants both or nothing
  - $v_A(\{1,2\}) = 1, v_A(S) = 0$ for all $S \neq \{1,2\}$
- B only wants item 1
  - $v_B(\{1,2\}) = v_B(\{1\}) = 1, v_B = 0$ otherwise

POLL

- 2 bidders, A and B, and 2 items
- A wants both or nothing
  - $v_A(\{1,2\}) = 1, v_A(S) = 0$ for all $S \neq \{1,2\}$
- B only wants item 1
  - $v_B(\{1,2\}) = v_B(\{1\}) = 1, v_B = 0$ otherwise

Question

What's the allocation? What's the payment?

1. A gets both, pays 1
2. A gets both, pays 0
3. B gets {1}, pays 1
4. B gets both, pays 1

VCG IS HORRIBLE FOR REVENUE

- Suppose a third bidder C shows up
  - $v_C(\{1,2\}) = v_C(\{2\}) = 1$, and zero otherwise
- What will VCG do?
- Social welfare maximizing solution is to give item 1 to B and item 2 to C
- What are the payments?
- Zero!

THE PLOT THICKENS: REVENUE

- For now one buyer
- m items
- Focus on additive buyers
  - $v(S) = \sum_{j \in S} v_j$
  - Can we efficiently compute VCG in this setting?
- Again, Bayesian approach
  - D_j = distribution for item j

COMPUTATIONAL PROBLEM

- Can we at least *compute* the optimal auction?
- Yes! Linear program!
- $x(v)$ = variable for the allocation of the agent when she reports v
- $p(v)$ = variable for the payment of the agent when she reports v

COMPUTATIONAL PROBLEM

- Objective:

$$\max \sum_{v} p(v) \cdot \Pr[value = v]$$

- Constraints:
  - IC: $\forall v, v': vx(v) - p(v) \geq vx(v') - p(v')$
  - IR: $\forall v: vx(v) - p(v) \geq 0$
  - Feasibility: $\forall v: 1 \geq x(v) \geq 0$
- Great!
- What does the optimal auction look like?
- That's a good question…

SELLING SEPARATELY IS NOT OPTIMAL

- 2 items
  - v_i = value for item i
- Additive buyer
  - Value for both items is $v_1 + v_2$
- $D_1 = D_2 = U\{1,2\}$
- Idea #1: Run optimal auction for each item
  - Posting a price for item i of 1 makes 1
  - Posting a price for item i of 2 makes $2 \cdot 1/2 = 1$
- Total (expected) revenue $= 2$

SELLING SEPARATELY IS NOT OPTIMAL

- Selling separately makes 2
- Idea #2: Bundle the items!
  - What if we post a price of 3$ for both items (as a bundle)?
  - You pay $3 and get both, or you get nothing.
- $\Pr[v_1 + v_2 \geq 3] = 3/4$
  - $\Pr[(1,2)] = \Pr[(2,1)] = \Pr[(2,2)] = 1/4$
- Expected revenue $= 3/4 \cdot 3 = 2.25 > 2$
- Selling separately is not optimal!

BUNDLING IS NOT OPTIMAL

- $D_1 = D_2 = U\{0,1,2\}$
- Selling separately gives revenue $4/3$
  - Price of 0 gives 0
  - Price of 1 gives $1 \cdot 2/3 = 2/3$
  - Price of 2 gives $2 \cdot 1/3 = 2/3$
- Bundling gives same revenue
  - Optimal price ends up being price of $2
  - It suffices to check that 0$, 1$, 3$ are not better

BUNDLING IS NOT OPTIMAL

- $D_1 = D_2 = U\{0,1,2\}$
- Selling separately and Bundling give revenue $4/3$
  - Henceforth, $\mathrm{SRev} = \mathrm{BRev} = 4/3$
- How about this:
  - You can pay \$2 and get one item (whichever you want)
  - You can pay \$3 for both

v_1/v_2	0	1	2
0	0\$	0\$	2\$
1	0\$	0\$	3\$
2	2\$	3\$	3\$

- Expected revenue $= 13/9 > 4/3$

OPTIMALITY REQUIRES RANDOMIZATION

- $D_1 = U\{1,2\}, D_2 = U\{1,3\}$
- Every deterministic auction (i.e. allocation is either 0 or 1 for each item) sets a price for every subset of items
  - One can check that in this instance the optimal deterministic auction makes revenue 2.5
- Randomized auction:
  - Pay \$4 and get both items for sure
  - Pay \$2.5 and: (1) get the first item, (2) flip a (fair) coin for the second item
- You'll pay \$4 every time $v_2 = 3$
- If $v_1 = 2$ and $v_2 = 1$ you'll buy the gamble
- Expected revenue $= 2.65 > 2.5$
- Optimal auction can be randomized!

REVENUE NON MONOTONICITY

- $Rev(X \times X) =$ Optimal expected revenue from single additive buyer with two items whose values are distributed i.i.d. according to the random variable X
- Y stochastically dominates X
  - $\forall x, \Pr[Y \geq x] \geq \Pr[X \geq x]$
  - For example, $U[0,2]$ stoch. dominates $U[0,1]$
- Theorem: [Hart, Reny 2012]
  - There exist X, Y such that Y stochastically dominates X, and

$$Rev(X \times X) > Rev(Y \times Y)$$

REVENUE NON MONOTONICITY

- Theorem: [Hart, Reny 2012]
  - There exist X, Y such that Y stochastically dominates X, and
$$Rev(X \times X) > Rev(Y \times Y)$$
- Intuitively:
  - You compute the optimal revenue for $D_1 = D_2 = U[0,1]$. It's some number R.
  - Then I tell you "oops, I messed up! $D_1 = D_2 = U[1,2]$! Find the new optimal revenue R'".
  - It could be that $R' < R$!
    - That's not the case in this example, but still!

REVENUE NON MONOTONICITY

- Theorem: [Hart, Reny 2012]
  - There exist X, Y such that Y stochastically dominates X, and

$$Rev(X \times X) > Rev(Y \times Y)$$

$$X = \begin{cases} 10 \text{ with probability } 4/15 \\ 46 \text{ with probability } 1/90 \\ 47 \text{ with probability } 1/3 \\ 80 \text{ with probability } 7/30 \\ 100 \text{ with probability } 7/45 \end{cases}$$

$$Y = \begin{cases} 10 \text{ with probability } \frac{2399}{4000} \\ 13 \text{ with probability } 1/9000 \\ 46 \text{ with probability } 1/90 \\ 47 \text{ with probability } 1/3 \\ 80 \text{ with probability } 7/30 \\ 100 \text{ with probability } 7/45 \end{cases}$$

SO FAR

- Revelation Principle
- Single parameter environments
  - Second price auctions
  - Myerson’s lemma
  - Myerson’s optimal auction
  - Cremer-McLean auction for correlated buyers
  - Prophet inequalities
  - Bulow-Klemperer
- Multiparameter environments
  - The VCG mechanism
  - Challenges
  - Revenue optimal auctions are strange