TRUTH JUSTICE ALGOS

Mechanism Design III: Simple single item auctions

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

CORRECTION IN THE DEFINITION OF MHR

- $\phi(v) = v - \frac{1-F(v)}{f(v)}$
- D is MHR if $\frac{1-F(v)}{f(v)}$ is monotone non increasing.

TODAY

- Cremer-McLean for correlated buyers
- Prophet Inequalities
- Bulow-Klemperer

BEYOND INDEPENDENCE

- Myerson: Optimal auction for independent bidders.
- What if the bidders' values are correlated?
  - Very realistic!
- We'll see a 2 agent instance of a result of Cremer and McLean [1998]
  - They show how to extract the full social welfare under very mild conditions on the correlation

CREMER-MCLEAN

v_1/v_2	1	2	3
1	1/6	1/12	1/12
2	1/12	1/6	1/12
3	1/12	1/12	1/6

Poll 1

How much revenue does a second price auction make (in expectation)?

1. 8/6
2. 10/6
3. 12/6
4. 14/6

CREMER-MCLEAN

v_1/v_2	1	2	3
1	1/6	1/12	1/12
2	1/12	1/6	1/12
3	1/12	1/12	1/6

Poll 2

What's the maximum possible revenue an auction can make?

1. 8/6
2. 10/6
3. 12/6
4. 14/6

CREMER-MCLEAN

v_1/v_2	1	2	3
1	1/6	1/12	1/12
2	1/12	1/6	1/12
3	1/12	1/12	1/6

- $P_{i,j} = \Pr[v_2 = j \mid v_1 = i]$

P =

1/2	1/4	1/4
1/4	1/2	1/4
1/4	1/4	1/2

- $E[utility\ of\ v_1 = 1\ from\ SP] = 0$
- $E[utility\ of\ v_1 = 2\ From\ SP] = 1/4 \cdot 1 = 1/4$
- $E[utility\ of\ v_1 = 3\ from\ SP] = 1/4 \cdot 2 + 1/4 \cdot 1 = 3/4$

CREMER-MCLEAN

- Observation: P has full rank
- Therefore, $P \cdot (x_1, x_2, x_3)^T = (0, {}^1/_4, {}^3/_4)^T$ has a solution:
  - $x_1 = -1, x_2 = 0, x_3 = 2$

The magic part

- Consider the following bet B_1 for player 1:
  - I pay you 1 if $v_2 = 1$
  - Nothing happens if $v_2 = 2$
  - You pay me 2 if $v_2 = 3$

CREMER-MCLEAN

- Consider the following bet B_1 for player 1: (a) I pay you 1 if $v_2 = 1$, (b) Nothing happens if $v_2 = 2$, (c) You pay me 2 if $v_2 = 3$
- What's the expected value for taking this bet if $v_1 = 1$?
  - $1/2 \cdot 1 + 1/4 \cdot 0 + 1/4 \cdot (-2) = 0$
- What if $v_1 = 2$? $-1/4$
- What if $v_1 = 3$? $-3/4$
- Similar bet B_2 for player 2
- Auction: Player i is offered bet B_i. After the bet we'll run a second price auction
  - $E[utility\ of\ v_1 = 1] = E[utility\ of\ B_1] + E[utility\ from\ SP] = 0$
  - $E[ut. of\ v_1 = 2] = -1/4 + 1/4 = 0$
  - $E[ut. of\ v_1 = 3] = -3/4 + 3/4 = 0$

CREMER-MCLEAN

- Since buyers always have zero utility, and the item is always sold, the seller must be extracting all of the social welfare
- Expected revenue = 14/6
- Wth just happened???
- That’s a pretty weird auction!
- This “prediction” is very unlikely to be observed in practice.

MYERSON IS WEIRD

- $n = 2. D_1 = U[0,1], \mathrm{D}_2 = \mathrm{U}[0,100]$
- $\phi_1(v_1) = 2v_1 - 1, \phi_2(v_2) = 2v_2 - 100$
- Optimal auction
  - When $v_1 \leq 1/2$ and $v_2 \geq 50$: Sell to 2 for 50
  - When $v_1 > 1/2$ and $v_2 < 50$: Sell to 1 for ½
  - When $0 < 2v_1 - 1 < 2v_2 - 100$: Sell to 2 for (99+$2v_1$)/2 (slightly over 50)
  - When $0 < 2v_2 - 100 < 2v_1 - 1$: Sell to 1 for $(2v_2 - 99)/2$ (slightly over ½)
- Wth is this???
- Impossible to explain, unless you go through all of Myerson's calculations!

OPTIMAL AUCTIONS ARE WEIRD

- Weirdness inevitable if you want optimality
- Weirdness inevitable if you're 100% confident in the model
- Take away: Optimality requires complexity
- In the remainder: ask for simplicity and settle for approximately optimal auctions.

CRITIQUE #1: TOO COMPLEX

A (cool) detour: Prophet inequalities!

PROPHET INEQUALITY

- n treasure boxes.
- Treasure in box i is distributed according to known distribution D_i
- In stage i you open box i and see the treasure (realization of the random variable) x_i
- After seeing x_i you either take it, or discard it forever and move on to stage $i + 1$
- What should you do?
- Our goal will be to compete against a prophet who knows the realizations of the D_is

PROPHET INEQUALITY

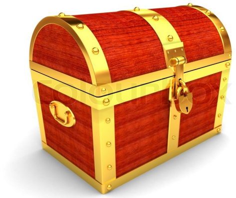 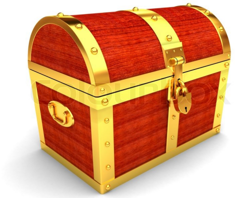 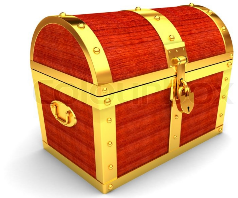 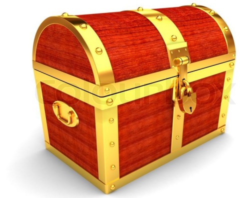

$D_1 = U[0,60]$ $D_1 = Exp[1/60]$ $D_1 = N[1,1]$ $D_1 = U[0,100]$

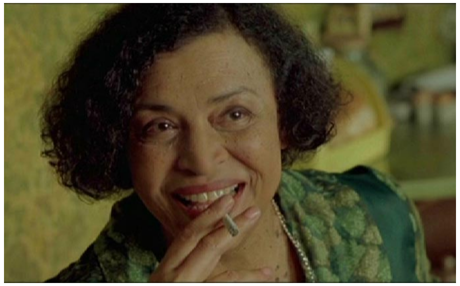

~~$x_1 = 54$~~ $x_2 = 52$ ✓

Our value is 52, Prophet gets 61

PROPHET INEQUALITY

- Optimal policy: Solve it backwards!
  - If we get to the last box, we should clearly take x_n
  - For the second to last, we should take x_{n-1} if it's larger than $E[x_n]$
  - We should take x_{n-2} only if it's larger than the expected value of the optimal policy starting at $n-1$, i.e. $\Pr\left[x_{n-1} > E[x_n]\right] \cdot E[x_{n-1} | x_{n-1} > E[x_n]] + \Pr\left[x_{n-1} \leq E[x_n]\right] \cdot E[x_n]$
  - And so on…
- Ok, that's pretty complicated…
- Any simpler policies?
  - Focus on policies that set a single threshold t and accept x_i if it's above t, otherwise reject
  - How good are those?

PROPHET INEQUALITY

- **Theorem**: There exists a single threshold t^* such that the policy that accepts x_i when $x_i \geq t^*$ gives expected reward at least $\frac{1}{2}E[\max_i x_i]$, i.e. at least half of what the prophet makes (in expectation).

PROPHET INEQUALITY

Proof

- $Z^+ = \max\{z, 0\}$
- Given a "threshold policy" with threshold t, let $q(t) = \Pr[policy\ accepts\ no\ prize]$
- Large t: large $q(t)$, but big rewards
- Small t: small $q(t)$, but small rewards
- $E[reward] \geq q(t) \cdot 0 + \left(1 - q(t)\right) \cdot t$
- A little too pessimistic…
- When $x_i \geq t$ we'll count x_i, not t

PROPHET INEQUALITY

$$E[reward] = t\big(1 - q(t)\big) +$$

$$\sum_i E[x_i - t | x_i \geq t \;\&\; x_j < t, \forall j \neq i] \cdot \Pr[x_i \geq t \& \; x_j < t, \forall j \neq i]$$

$$= t\big(1 - q(t)\big) +$$

$$\sum_i E[x_i - t | x_i \geq t] \cdot \Pr[x_i \geq t] \cdot \Pr\big[\, x_j < t, \forall j \neq i\big]$$

$$= t\big(1 - q(t)\big) + \sum_i \; E[(x_i - t)^+] \cdot \Pr[x_j < t, \forall j \neq i]$$

$$\geq t\big(1 - q(t)\big) + q(t) \sum_i E[(x_i - t)^+]$$

(we used that $q(t) = \Pr\big[x_j < t \,, \forall j\big] \leq \Pr[x_j < t, \forall j \neq \mathrm{i}]$)

PROPHET INEQUALITY

$$E[reward] \geq t\big(1 - q(t)\big) + q(t) \sum_i E[(x_i - t)^+]$$

$$E[\max_i x_i] = E[t + \max_i (x_i - t)]$$
$$= t + E[\max_i (x_i - t)]$$
$$\leq t + E[\max_i (x_i - t)^+]$$
$$\leq t + \textstyle\sum_i E[(x_i - t)^+]$$

$$t^*: q(t^*) = {}^1\!/_2$$

$$E[reward] \geq \frac{t^*}{2} + \frac{1}{2} \sum_i E[(x_i - t^*)^+] \geq \frac{1}{2} E[\max_i x_i]$$

BACK TO AUCTIONS

- $Rev = E[\sum_i \phi_i(v_i)x_i(v_i)] = E[\max_i \phi_i(v_i)^+]$
- Pick t^* such that $\Pr[\max_i \phi_i(v_i)^+ \geq t^*] = 1/2$
- Give item to bidder i if $\phi_i(v_i) \geq t^*$
- Prophet inequality gives

$$E[reward] = E[\sum_i \phi_i(v_i)x_i(v_i) \geq \frac{1}{2} E[\max_i \phi_i(v_i)^+]$$

- More concretely:
  - $r_i = \phi_i^{-1}(t^*)$
  - Remove all bidders with $b_i < r_i$
  - Run a second price with the remaining bidders

CRITIQUE #2: TOO MUCH DEPENDENCE ON THE DISTRIBUTION

- Optimal auction depends on the distribution
- Wasn't the whole point of the Bayesian approach that this is unavoidable?
- We'll assume that $v_i \sim D_i$ (in the analysis), but our auctions will **not** depend on the $D_i s$
  - "Prior independent" mechanism design

PRIOR INDEPENDENT MECHANISMS

- Sounds pretty optimistic…
- Existence of a good prior independent auction A for (say) regular distributions implies that a **single** auction can compete with all the (uncountably many) optimal auctions, *tailored* to each distribution, **simultaneously**!
- Pretty wild!
- Any candidates?
  - Second price auction!

BULOW-KLEMPERER THEOREM

- $OPT(n, D)$= Expected revenue of optimal auction with n i.i.d. buyers from D.
- $V(n, D)$ = Expected revenue of Vickrey with n i.i.d. buyers from D.
- Theorem (1996): For all regular D we have

$$V(n + 1, D) \geq OPT(n, D)$$

- In more modern language: "The competition complexity of single-item auctions with regular distributions is 1"
  - The competition complexity of n bidders with additive valuations over m independent, regular items is at least $log m$ and at most $\mathrm{n} + 2m - 2$ [EFFTW 17]

BULOW-KLEMPERER THEOREM

- Theorem (1996): For all regular D we have
$$V(n+1, D) \geq OPT(n, D)$$
- Intuitively: It is better to increase competition by a single buyer than invest in learning the underlying distribution!

BULOW-KLEMPERER THEOREM

Proof:

- Let A be the following auction for $n+1$ buyers from D:
  - Run $OPT(n, D)$ on buyers $1, \ldots, n$
  - If the item is not sold, give it for free to buyer $n+1$
- Obvious observation 1: $Rev(A) = OPT(n, D)$
- Obvious observation 2: A always allocates the item.

BULOW-KLEMPERER THEOREM

- Non obvious:
- The second price auction is the revenue maximizing auction over all auctions that always allocate the item.
  - Why?
- Therefore

$$V(n+1, D) \geq Rev(A) = OPT(n, D)$$

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