LECTURE 8

JD R-M

FOR TODAY

i. Revelation Principle
ii. Envelope Theorem
iii. Optimal Auction (Myerson 1981)
iv. Auctions/Monopolies Duality (Bulow and Roberts 1989)
v. Auctions versus Negotiations (Bulow and Klemperer 1996)

PRIMITIVES

- There is 1 seller (called agent 0) and $n \geq 1$ buyers
- Set of agents is $I = \{0,1,2,\cdots,n\}$
- Each buyer i has a valuation over the goods that is $\theta_i \in \Theta$ with a distribution $\mu \in \Delta(\Theta)$
- Assume that the seller values the good at $\theta_0 = 0$ and for each type $\theta \in \Theta, \theta \geq 0$
- Outside option for everyone is 0
- Set of possible outcomes is $O = I \times \Re^n$: e.g. an outcomes $o = (j, (p_i)_{i=1}^n)$ is a winner j and the payments made by each buyer i, p_i

PAYOFFS

- Fix an outcome $o = (j, (p_i)_{i=1}^{n})$
- The payoff to a buyer i to the outcome above when he has a valuation θ_i is

$$(1)\ u_i(o, \theta_i) = \chi(i = j)\theta_i - p_i$$

- The payoff to the seller would be

$$(2)\ \pi(o) = \sum_{i=1}^{n} p_i$$

MECHANISM DESIGN OBJECTIVE

- The seller picks a game for the agents to play and he might as well pick one where all buyers participate
- Given the primitives discussed, the seller picks a tuple $\Gamma = (g, (A_i)_{i=1}^{n})$ for

I. For each buyer i, A_i is a non-empty action set

II. $\mathrm{g: A \equiv \prod_{i=1}^{n} A_i \rightarrow O}$

- The seller also picks an equilibrium in the game. I assume that he picks a Bayes Nash Equilibrium (BNE).

BAYES NASH EQUILIBRIUM

- A buyer i strategy for a game Γ picked by the seller is a function $\sigma_i: \Theta \to A_i$
- A BNE is a tuple $\sigma = (\sigma_i)_{i=1}^n$ such that for each buyer i and type θ_i,

$$(3)\ E\left[u_i\left(\theta_i, g\left(\sigma_i(\theta_i), \sigma_{-i}(\theta_{-i})\right)\right) | \theta_{-i}\right] = \max_{b \in A_i} E\left[u_i\left(\theta_i, g\left(b, \sigma_{-i}(\theta_{-i})\right)\right) | \theta_{-i}\right]$$

BACK TO MECHANISM OBJECTIVE

- The seller also picks an equilibrium σ
- Note that a mechanism, i.e. (Γ, σ), implies a mapping $f_{(\Gamma,\sigma)}: \Theta^n \rightarrow O$
- But there’s a HUGE problem: there are too many games to consider.
- Q: How can one narrow the focus?
- A: The Revelation Principle!

REVELATION PRINCIPLE

- Proposition: For every mechanism (Γ, σ), there exists an alternative mechanism $\left(\hat{\Gamma}, \hat{\sigma}\right)$ s.t.

i. For each buyer $i, \hat{A}_i = \Theta$,

ii. $\forall i, \theta_i, \sigma_i(\theta_i) = \theta_i$ is a BNE,

iii. $f_{(\Gamma,\sigma)} = f_{(\hat{\Gamma},\hat{\sigma})}$

PROOF OF THE REVELATION PRINCIPLE

- Fix some (Γ, σ).
- Let $\hat{g} = f_{(\Gamma,\sigma)}$, then $f_{(\Gamma,\sigma)} = f_{(\hat{\Gamma},\hat{\sigma})}$.
- Now, for every buyer i $and\ type\ \theta_i$

$$\max_{\theta\in\Theta} E[u_i(\theta_i, \hat{g}(\theta, \theta_{-i}))|\theta_{-i}] \geq E[u_i(\theta_i, \hat{g}(\theta_i, \theta_{-i}))|\theta_{-i}]$$

$$= E\left[u_i\left(\theta_i, g(\sigma_i(\theta_i), \sigma_{-i}(\theta_{-i}))\right)|\theta_{-i}\right] = \max_{b\in A_i} E\left[u_i\left(\theta_i, g(b, \sigma_{-i}(\theta_{-i}))\right)|\theta_{-i}\right]$$

$$\geq \max_{b\in\sigma_i(\Theta)} E\left[u_i\left(\theta_i, g(b, \sigma_{-i}(\theta_{-i}))\right)|\theta_{-i}\right] = \max_{\theta\in\Theta} E[u_i(\theta_i, \hat{g}(\theta, \theta_{-i}))|\theta_{-i}]\ \square$$

MYERSON'S PROBLEM

- This last result implies that we can focus on a mapping $f: \Theta^n \rightarrow O$, but in order to implement the mechanism, the seller must ensure that buyers would be willing to participate in the game of the type $(\hat{\Gamma}, f)$ and that reporting one's type IS an equilibrium in such game.
- Assume that $\Theta = [0,1]$ and types are distributed given a CDF F admitting a PDF $f \gg 0$
- The seller might as well pick $(x, t): \Theta^n \rightarrow \Delta(I) \times \Re^n$ (odds for each agent wins the good and payments)

RESTRICTIONS

- Define for each buyer $i, \tilde{\theta}_i, q_i(\tilde{\theta}_i) = E_{\theta_{-i}}[x_i(\tilde{\theta}_i, \theta_{-i})], p_i(\tilde{\theta}_i) = E_{\theta_{-i}}[t_i(\tilde{\theta}_i, \theta_{-i})]$
- The first restriction faced by the seller is that buyers must be willing to participate in the game, so

$$(Participation\ Constraint)\ \forall i \geq 1, \theta_i, \theta_i q_i(\theta_i) - p_i(\theta_i) \geq 0$$

- The second constraint is that telling the truth IS optimal, so

$$(Individual\ Rationality) \forall i \geq 1, \theta_i, \theta_i q_i(\theta_i) - p_i(\theta_i) = V_i(\theta_i) = \max_{x \in [0,1]} \theta_i q_i(x) - p_i(x)$$

- The seller's problem is then to solve $\max_{x,t} \sum_{i=1}^{n} E[p_i(\theta_i)]\ s.t. (4)\ and\ (5)\ hold.$

ENVELOPE THEOREM

- The problem seems very difficult to solve, but this is not so.
- Envelope Theorem: Assume that one maximizes a function $f: [0,1] \times X \rightarrow \Re$ such that for each $x \in X$, $f(.,x)$ is absolutely continuous, continuously differentiable, and $\partial_t f(.,x)$ is a bounded function. Then, if $\forall t \in [0,1], V(t) = \max_{x \in X} f(t,x)$ then

$$(6)\ \forall t \in [0,1], V(t) = V(0) + \int_0^t \partial_t f\big(s, x(s)\big) ds\,, \forall t, f\big(t, x(t)\big) = V(t)$$

USING THE ENVELOPE THEOREM

- Notice that in Myerson's problem $f(t, x) = t\, q(x) - p(x)$, so individual rationality states

$$(7)\ \forall i \geq 1, \theta_i, \theta_i q_i(\theta_i) - p_i(\theta_i) = V_i(0) + \int_0^{\theta_i} q(x)dx$$

- Thus, $p_i(\theta_i) = \theta_i q_i(\theta_i) - \int_0^{\theta_i} q(x)dx - \mathrm{V_i}(0)$
- Note that the seller might as well pick $p_i(0) = 0 = -V_i(0)$ since any positive surplus given to buyers with this low type just comes as a cost to the seller and nothing more.
- Further notice that now $p_i(\theta_i) = \theta_i q_i(\theta_i) - \int_0^{\theta_i} q(x)dx$: thus, payments are simply derived from the allocation rule (i.e. what matters is WHO GETS THE GOOD.)

SOLVING THE PROBLEM, PART 1

- Notice that the envelope condition and $V_i(0) = 0$ ensures that the participation constraint also holds
- What's left to see is what is the revenue maximizing allocation.
- Note that

$$(8) r(x,t) = \sum_{i=1}^{n} E[p_i(\theta_i)] = \sum_{i=1}^{n} \int_0^1 \left[\theta_i q_i(\theta_i) - \int_0^{\theta_i} q(x)dx\right] df(\theta_i)\theta_i$$

$$= \sum_{i=1}^{n} \int_0^1 f(\theta_i)\theta_i q_i(\theta_i) - [1 - F(\theta_i)]q(\theta_i)d\theta_i$$

$$\sum_{i=1}^{n} \int_0^1 f(\theta_i)q_i(\theta_i)\left[\theta_i - \frac{1 - F(\theta_i)}{f(\theta_i)}\right] d\theta_i = \sum_{i=1}^{n} \int_0^1 f(\theta_i)q_i(\theta_i)v(\theta_i)d\theta_i$$

SOLVING THE PROBLEM, PART 2

- Notice that the seller cannot gain from treating different buyers differently, then it is without loss of generality that $\forall i \geq 1, q_i(.) = q(.)$ and revenues equal to

$$(9)\ r(x,t) = \int_0^1 nf(\theta)q(\theta)v(\theta)d\theta$$

- Now assume that $v(.)$ is a strictly increasing function and observe that $v(0) = -\frac{1}{f(0)} < 0, v(1) = 1$, and it is a continuous function, then there exists a type $p^M \in (0,1)$ such that $v(p^M) = 0$ the seller is better off not selling the good to anyone if all consumers have a valuation below p^M
- Assume that $v(.)$ is increasing, then the seller might as well give the good to whomever values the good the most. It turns out that a typical auction that implements this mechanism, in expectation, is either a first or second price auction with a reserve price of p^M

BACK TO MONOPOLIES

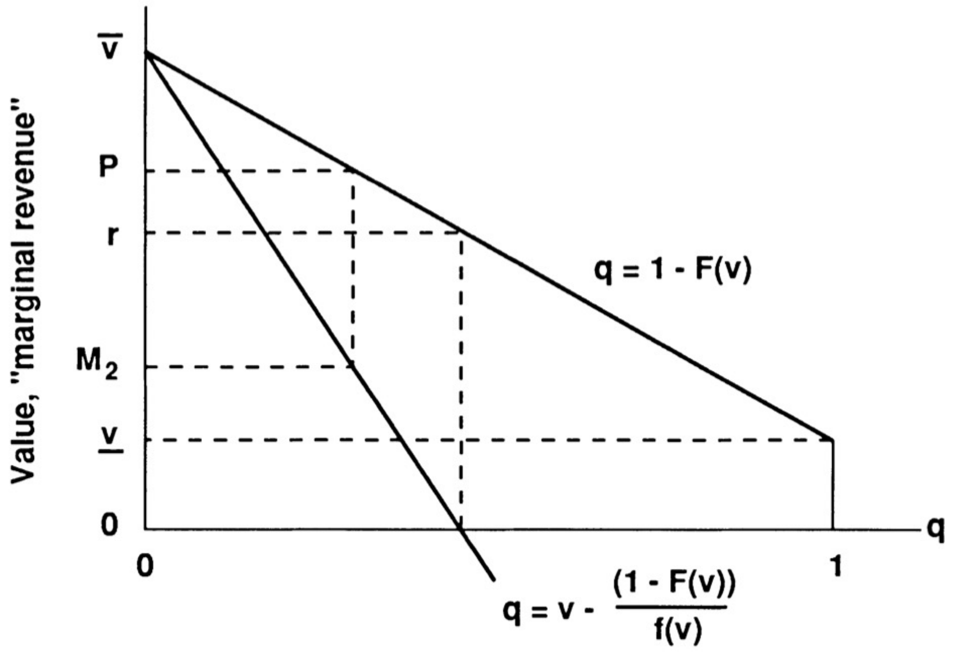

FIG. 1.—Construction of an optimal auction

- Looking back at the optimal monopoly pricing section of the course, we derived that the optimal monopoly price solves

$$(10)\ p^M - \frac{1 - F(p^M)}{f(p^M)} = c$$

- So if the marginal costs is $c = 0$, then the optimal reserve price EQUALS the static monopoly price
- Indeed, the standard monopoly problem can be seen as an auction problem where $n = 1$
- Bulow and Roberts (1989) further illustrate how the analogy discussed above can be fully generalized. They even find that optimal auctions can be re-written in terms of marginal revenues and costs.

GENERAL ANSWER

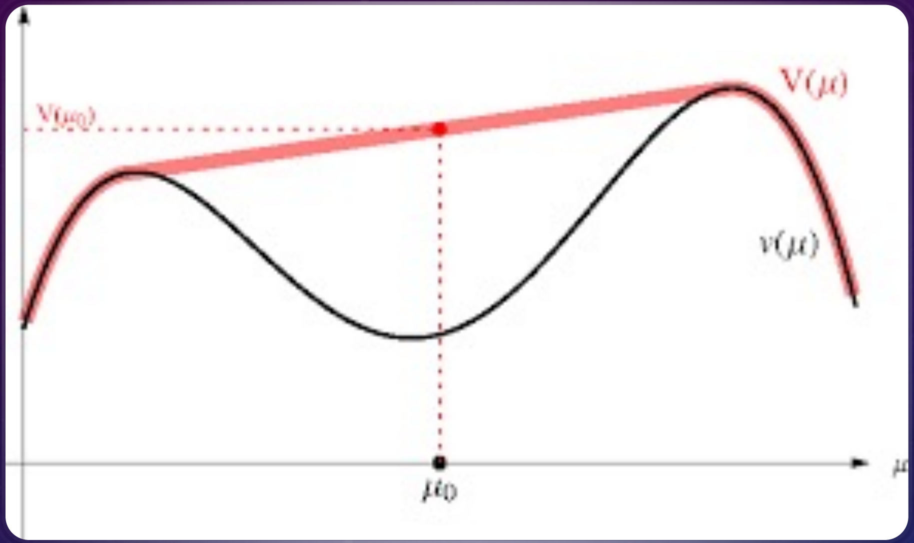

- I made two unsubstantiated claims and I must now fess up.

1. There is no reason to believe that $F_i = F$ for each buyer
2. The virtual value function $\nu(.)$ DOES NOT have to be increasing

- Myerson found that the second issue is rather straightforward. Rather than dealing with the real virtual value, he proposed that the seller uses "ironed out virtual values" (technique is described in the left).
- Next, the right prognosis is the seller should give the good to the buyer with the highest "ironed out virtual value".
- But why?
- Intuitively, virtual values embed the costs of persuading a buyer to report his type truthfully: i.e. it can be the case that high types can demand more compensation in order to reveal their private information. Thus, it they need not be the seller's most lucrative trading partners.

MORE DETAILS TO KEEP IN MIND…

- Before you leave, there is one last nuance to answer
- Q: If you have the option between running the optimal auction with the right reserve price with n buyers or run an English Auction with NO reserve price, which one should you choose?
- A: When the virtual value function is increasing and all buyers have valuations above the seller's , it is immediate that the second auction is better
- Intuition: the returns to the optimal auction is the expected maximum virtual value of buyers times the odds of a sale; whereas the second option's return is the maximum virtual value of all previous buyers (the same as before when the is a trade) AND another buyer's virtual value.