29-3-22 Lecture 06

yesterday

- equilibria & revenue in different auction formats

⇒ all three auction formats give the same expected revenue to the auctioneer!

today

- Generalise this result
- Intro to Moral Hazard problems.

Mechanism Design – Revenue Maximisation – Optimal Auctions

Theorem

In the auction problem, any Bayesian incentive compatible mechanism that implements $q(\theta) = (q_1(\theta), \dots, q_n(\theta))$ gives each agent i payoff

$$V_i(\theta_i) = V_i(\underline{\theta}_i) + \int_{\underline{\theta}_i}^{\theta_i} \underbrace{\int_{\Theta_{-i}} q_i(s, \theta_{-i})\, \overbrace{\mathrm{d}F_{-i}(\theta_{-i})}^{\prod_{j=1, j\neq i}^{n} f_j(\theta_j)\, d\theta_j}}_{\equiv\, Q_i(s)} \,\mathrm{d}s,$$

and expected transfer

$$-\underbrace{T_i(\theta_i)}_{\int_{\Theta_{-i}} t_i(\theta_i, \theta_{-i})\, dF_{-i}(\theta_{-i})} = \theta_i Q_i(\theta_i) - V_i(\theta_i).$$

- In any BIC mechanism, the allocation rule almost pins down the transfers (up the constants $V_i(\underline{\theta}_i)$)

$T_1(\theta_1)$ in our example with $n=2$, $\theta_1, \theta_2 \sim \mathcal{U}[0,1]$, $c=0$

① 2^{nd} price auction ($\hat{=}$ English auction)

bidders bid own valuation and pay other's bid if they win

$$\Rightarrow \quad t_1(\theta_1, \theta_2) = \begin{cases} -\theta_2 & \text{if } \theta_1 \geq \theta_2 \\ 0 & \text{otherwise} \end{cases}$$

$$\Rightarrow \quad T_1(\theta_1) = \int_{\underline{\theta}_2}^{\bar{\theta}_2} t_1(\theta_1, \tilde{\theta}_2)\, f_2(\tilde{\theta}_2)\, d\tilde{\theta}_2 = \int_0^1 -\tilde{\theta}_2 \cdot \mathbb{1}_{\tilde{\theta}_2 \leq \theta_1}\, d\tilde{\theta}_2 = \int_0^{\theta_1} -\tilde{\theta}_2\, d\tilde{\theta}_2 = \underline{-\frac{1}{2}\theta_1^2}$$

② All-pay auction: bidders pay their bid always. equilibrium bid $s^{APA}(\theta_1) = -\frac{1}{2}\theta_1^2$

$$\Rightarrow t_1(\theta_1, \theta_2) = -\frac{1}{2}\theta_1^2 \quad \forall \theta_2 \quad \Longrightarrow \quad T_1(\theta_1) = \int_{\underline{\theta}_2}^{\bar{\theta}_2} -\frac{1}{2}\theta_1^2 f_2(\tilde{\theta}_2)\, d\tilde{\theta}_2 = \underline{-\frac{1}{2}\theta_1^2}$$

Mechanism Design – Revenue Maximisation – Optimal Auctions

The seller's expected revenue from mechanism (q, t) is

$$\int_{\Theta} \left[\sum_{i=1}^{n} (-t_i(\theta) - cq_i(\theta)) \right] \mathrm{d}F(\theta) \qquad = \qquad \sum_{i=1}^{n} \int_{\Theta_i} (-T_i(\theta_i) - cQ_i(\theta_i))\, \mathrm{d}F_i(\theta_i).$$

It follows from previous result that:

Theorem (Revenue Equivalence)

Any two equilibria of any two auctions that yield (i) identical allocation probabilities $q_i(\cdot)$ and (ii) identical interim utility for type $\underline{\theta}_i$ of each bidder i give the seller the same expected revenue.

In our 3 examples: (i) identical allocation rule? Yes, $q_1(\theta_1, \theta_2) = \mathbb{1}_{\theta_1 \geq \theta_2} = 1 - q_2(\theta)$

(ii) same $V_i(\underline{\theta}_i)$? Yes, $V_i(0) = 0$ in all 3 auction

Mechanism Design – Revenue Maximisation – Optimal Auctions

Were the three auctions optimal for the seller?

example with $\theta_1, \theta_2 \sim U[0,1]$, $c = 0$

$J_i(\theta_i) = \theta_i - \frac{1-\theta_i}{1} = 2\theta_i - 1$

What constitutes an optimal auction?

$J_i(\theta_i) = \theta_i - \frac{1 - F_i(\theta_i)}{f_i(\theta_i)}$

Theorem (Optimal auction)

Suppose $n \geq 2$ and each bidder's virtual valuation J_i is increasing in θ_i. Any alloaction rule $q^* : \Theta \to [0,1]^n$ satisfying

$$q_i^*(\theta) > 0 \quad \textit{only if} \quad J_i(\theta_i) = \max_j J_j(\theta_j) > c,$$

agent with highest virtual value if virtual value exceeds cost.

$$\sum_{i=1}^{n} q_i^*(\theta) < 1 \quad \textit{only if} \quad \max_j J_j(\theta_j) \leq c$$

and the implied transfers with $T_i(\underline{\theta}_i) = \underline{\theta}_i Q_i^*(\underline{\theta}_i)$ (i.e. $V_i(\underline{\theta}_i) = 0$) is an incentive-compatible individually-rational mechanism that maximises revenue.

proof: exercise | Myerson 1981 | Mailath p. 334, MWG p. 889

Mechanism Design – Further Topics

Some issues that we have not covered:

- Collusion → multiple bidders coordinate on joint deviation
- Interdependent valuations (for example, common-value auctions)
- Correlated types
- Evidence / Verification → different types can only report some other types
- Dynamic problems (multiple stages) principal can pay to learn true typ
- Limited commitment for principal

Moral Hazard

Moral Hazard – Intro

We now consider models with 'hidden action'

- Principal commits to payment schedule
- Agent takes an action
- Principal observes (imperfect) signal about action and pays according to schedule

Examples:

- Insurance contract
- Employment contract
- Rental contract

Moral Hazard – Intro

General setup:

- Agent chooses effort level $e \in E \subset \mathbb{R}_+$
- The profit is $\pi \in [\underline{\pi}, \bar{\pi}] = \Pi \subset \mathbb{R}$.
- Distribution of profit depends on effort: $\pi \sim F(\cdot|e)$ with density $f(\pi|e) > 0$

  $F(\cdot|e)$ is ordered by first-order stochastic dominance:

  If $e'' > e'$, then $F(\pi|e'') \leq F(\pi|e')$ $\forall \pi \in \Pi$.

  $\Rightarrow$ more effort $\Rightarrow$ higher output

- Principal observes only π and commits to pay the agent a wage $w(\pi)$
- **Payoffs:** agent: $v(w) - c(e)$ principal: $\pi - w$

  $v(\cdot)$ is increasing and concave, effort cost $c(\cdot)$ is increasing and convex

  ↳ agent is risk-averse

  principal is risk neutral

The principal's problem is to choose an effort level e and wage scheme $w(\cdot)$ to solve

$$\max_{e,w(\cdot)} \int_{\underline{\pi}}^{\bar{\pi}} (\pi - w(\pi))\, f(\pi|e)\, \mathrm{d}\pi \quad \text{such that}$$

· obedience constraint

$$e = \operatorname*{argmax}_{e' \in E} \int_{\underline{\pi}}^{\bar{\pi}} v\,(w(\pi))\, f(\pi|e')\, \mathrm{d}\pi - c(e'), \tag{IC}$$

· participation constraint

$$\int_{\underline{\pi}}^{\bar{\pi}} v\,(w(\pi))\, f(\pi|e)\, \mathrm{d}\pi - c(e) \geq 0 \tag{IR}$$

Moral Hazard – Observable effort

Suppose effort is observable to the principal (and contractible) so that wage is $w(e, \pi)$

- The principal can enforce effort e by setting $w(e', \pi) = \begin{cases} 0 & \text{for all } e' \neq e \\ w_e(\pi) & \text{for } e' = e \end{cases}$

  ↳ 'forcing contract'

Principal's problem is to choose effort level e and wage function $w_e(\cdot)$ to

$$\max_{e, w_e(\cdot)} \int_{\underline{\pi}}^{\bar{\pi}} (\pi - w_e(\pi)) f(\pi|e) \, d\pi$$

$$\text{such that} \quad \int_{\underline{\pi}}^{\bar{\pi}} v(w_e(\pi)) f(\pi|e) \, d\pi - c(e) \geq 0 \qquad \text{(IR)}$$

$$\max_{e,w_e(\cdot)} \int_{\underline{\pi}}^{\bar{\pi}} (\pi - w_e(\pi))\, f(\pi|e)\, \mathrm{d}\pi \longrightarrow \text{linear objective in } w_e(\cdot)$$

$$\text{such that } \lambda \left(\int_{\underline{\pi}}^{\bar{\pi}} v\,(w_e(\pi))\, f(\pi|e)\, \mathrm{d}\pi - c(e) \right) \geq 0 \longrightarrow \text{concave constraint in } w$$

We can separat this optimisation into 2 steps:
1. what is the optimal $w_e(\cdot)$ for any given e?
2. what is the optimal e?

} for this example we will only consider step 1.

FOC: $\frac{\partial}{\partial w} = -1\, f(\pi|e) + \lambda\, v'(w_e(\pi))\, f(\pi|e) = 0$

$\iff v'(w_e(\pi)) = 1/\lambda$ for all π $\implies$ agent is fully insured in optimal contract

$\implies v'(\cdot)$ is the same at all output levels π

$\to$ if $v(\cdot)$ is strictly concave, then $w_e(\pi)$ must be same for all π

Moral Hazard – Risk-neutral agent

Suppose effort is unobservable but the agent is risk neutral: $v(w) = w$

- The principal could simply 'sell the firm' to the agent:

first best effort level

$$w(\pi) = \pi - \underbrace{\max_e \left\{ \int_{\underline{\pi}}^{\bar{\pi}} \tilde{\pi} f(\tilde{\pi}|e)\, \mathrm{d}\tilde{\pi} - c(e) \right\}}_{\text{expect maximal surplus of the firm}} = \pi - \left(\int_{\underline{\pi}}^{\bar{\pi}} \pi f(\pi|e^{FB})\, \mathrm{d}\pi - c(e^{FB}) \right)$$

- The agent's expected payoff when choosing some e' is

$$\int_{\underline{\pi}}^{\bar{\pi}} \pi f(\pi|e')\, \mathrm{d}\pi - c(e') - \left(\int_{\underline{\pi}}^{\bar{\pi}} \pi f(\pi|e^{FB})\, \mathrm{d}\pi - c(e^{FB}) \right)$$

→ what is the optimal choice of e' for the agent? → agent optimally chooses e^{FB}

– what does agent get? → agent gets 0 in expectation because principal "sells firm" at price equal to maximal ex-ante surplus.

Moral Hazard – Risk-neutral agent

We saw: if the agent is risk-neutral and wages are unrestricted, we get efficient outcome

- Agent chooses first-best effort level e^{FB}
- Principal extracts all surplus by 'selling' the firm at expected value

Next, we consider frictions that may induce inefficiencies:

1. Limited Liability:

   if agent has limited funds, requiring $w(\theta) \geq \underline{w}$, selling the firm infeasible

2. Risk-averse agent:

   if $v(w)$ is concave, tradeoff between incentive-provision and risk-sharing arises