6.254 : Game Theory with Engineering Applications
Lecture 19: Mechanism Design I

Asu Ozdaglar
MIT

April 29, 2010

Outline

- Mechanism design
- Revelation principle
    - Incentive compatibility
    - Individual rationality
- "Optimal" mechanisms

- Reading:
- Krishna, Chapter 5
- Myerson, "Optimal Auction Design," *Mathematics of Operations Research*, vol. 6, no. 1, pp. 58-73, 1981.

Introduction

- In the next 3 lectures, we will study Mechanism Design, which is an area in economics and game theory that has an engineering perspective.
- The goal is to design economic mechanisms or incentives to implement desired objectives (social or individual) in a strategic setting–assuming that the different members of the society each act rationally in a game theoretic sense.
- Mechanism design has important applications in economics (e.g., design of voting procedures, markets, auctions), and more recently finds applications in networked-systems (e.g., Internet interdomain routing, design of sponsored search auctions).

Auction Theory Viewpoint

- We first study the mechanism design problem in an auction theory context, i.e., we are interested in allocating a single indivisible object among agents.
- An auction is one of many ways that a seller can use to sell an object to potential buyers with unknown values.
- In an auction, the object is sold at a price determined by competition among buyers according to rules set by the seller (auction format), but the seller can use other methods.
- The question then is: what is the "best" way to allocate the object?
- Here, we consider the underlying allocation problem by abstracting away from the details of the selling format.

Model

- We assume a seller has a single indivisible object for sale and there are N potential buyers (or bidders) from the set $\mathcal{N} = \{1, \ldots, N\}$.
- Buyers have private values X_i drawn independently from the distribution F_i with associated density function f_i and support $\mathcal{X}_i = [0, w_i]$.
  - Notice that we allow for asymmetries among the buyers, i.e., the distributions of the values need not be the same for all buyers.
- We assume that the value of the object to the seller is 0.
- Let $\mathcal{X} = \prod_{j=1}^{N} \mathcal{X}_j$ denote the product set of buyers' values and let $\mathcal{X}_{-i} = \prod_{j \neq i} \mathcal{X}_j$.
- We define $f(x)$ to be the joint density of $x = (x_1, \ldots, x_N)$. Since values are independently distributed, we have $f(x) = f_1(x_1) \times \cdots \times f_N(x_N)$. Similarly, we define $f_{-i}(x_{-i})$ to be the joint density of $x_{-i} = (x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_N)$.

Mechanism

- A selling mechanism $(\mathcal{B}, \pi, \mu)$ has the following components:
  - A set of messages (or bids/strategies) $\mathcal{B}_i$ for each buyer i,
  - An allocation rule $\pi : \mathcal{B} \rightarrow \Delta$, where Δ is the set of probability distributions over the set of buyers $\mathcal{N}$,
  - A payment rule $\mu : \mathcal{B} \rightarrow \mathbb{R}^N$.
- An allocation rule specifies, as a function of messages $b = (b_1, \ldots, b_N)$, the probability $\pi_i(b)$ that i will get the object. Similarly, a payment rule specifies the payment $\mu_i(b)$ that i must make.
- Every mechanism defines a game of incomplete information among the buyers.
  - Strategies: $\beta_i : [0, w_i] \rightarrow \mathcal{B}_i$
  - Payoffs: Expected payoff for a given strategy profile and selling mechanism
- A strategy profile $\beta(\cdot)$ is a Bayesian Nash equilibrium of a mechanism if for all i and for all x_i, given the strategies β_{-i} of other buyers, $\beta_i(x_i)$ maximizes buyer i's expected payoff.

Direct Mechanisms and Revelation Principle

- A mechanism could be very complicated since we made no assumptions on the message sets $\mathcal{B}_i$.
- A special class of mechanisms, referred to as direct mechanisms, are those for which the set of messages is the same as the set of types (or values), i.e., $\mathcal{B}_i = \mathcal{X}_i$ for all i.
- These mechanisms are called "direct" since every buyer is asked directly to report a value.
- Formally a direct mechanism (Q, M) consists of the following components:
  - A function $Q : \mathcal{X} \to \Delta$, where $Q_i(x)$ is the probability that i will get the object,
  - A function $M : \mathcal{X} \to \mathbb{R}^N$, where $M_i(x)$ is the payment by buyer i.
- If it is a Bayesian Nash equilibrium for each buyer to report (or reveal) their type x_i correctly, we say that the direct mechanism has a truthful equilibrium.
- We refer to the pair $(Q(x), M(x))$ as the outcome of the mechanism.

Revelation Principle

- The following key result, referred to as the revelation principle, allows us to restrict our attention to direct mechanisms.
- More specifically, it shows that the outcomes resulting from any equilibrium of any mechanism can be replicated by a truthful equilibrium of some direct mechanism.

Proposition (Revelation Principle)

Given a mechanism $(\mathcal{B}, \pi, \mu)$ and an equilibrium β of that mechanism, there exists a direct mechanism (Q, M), in which

(i) *it is a Bayesian Nash equilibrium for each buyer to report his value truthfully,*

(ii) *the outcomes are the same as in equilibrium β of the original mechanism.*

Proof: This follows simply by defining the functions $Q : \mathcal{X} \to \Delta$ and $M : \mathcal{X} \to \mathbb{R}^N$ as $Q(x) = \pi(\beta(x))$,and $M(x) = \mu(\beta(x))$. Instead of buyers submitting message $b_i = \beta(x_i)$, the mechanism asks the buyer to report their value and makes sure the outcome is the same as if they had submitted $\beta_i(x_i)$.

Revelation Principle

The basic idea behind revelation principle is as follows:

- Suppose that in mechanism $(\mathcal{B}, \pi, \mu)$, each agent finds that, when his type is x_i, choosing $\beta_i(x_i)$ is his best response to others' strategies.
- Then, if we have a mediator who says "Tell me your type x_i and I will play $\beta_i(x_i)$ for you," each agent will find truth telling to be an optimal strategy given that all other agents tell the truth.
- In other words, a direct mechanism does the "equilibrium calculations" for the buyers automatically.

Incentive Compatibility

- For a given direct mechanism (Q, M), we define

$$q_i(z_i) = \int_{\mathcal{X}_{-i}} Q_i(z_i, x_{-i}) f_{-i}(x_{-i}) dx_{-i},$$

to be the probability that i will get the object when he reports his value to be z_i and all other buyers report their values truthfully.

- Similarly, we define

$$m_i(z_i) = \int_{\mathcal{X}_{-i}} M_i(z_i, x_{-i}) f_{-i}(x_{-i}) dx_{-i}$$

to be the expected payment of i when his report is z_i and all other buyers tell the truth.

- The expected payoff of buyer i when his true value is x_i and he reports z_i, assuming all others tell the truth, can be written as

$$q_i(z_i) x_i - m_i(z_i).$$

Incentive Compatibility

Definition

We say that the direct revelation mechanism (Q, M) is incentive compatible (IC) if

$$q_i(x_i)x_i - m_i(x_i) \geq q_i(z_i)x_i - m_i(z_i) \qquad \text{for all } i, x_i, z_i.$$

We refer to the left-hand side of this relation as the equilibrium payoff function denoted by $U_i(x_i)$, i.e.,

$$U_i(x_i) = \max_{z_i \in \mathcal{X}_i} \{q_i(z_i)x_i - m_i(z_i)\}.$$

Properties under IC:

- Since U_i is a maximum of a family of affine functions, it follows that U_i is a convex function.
- Moreover, it can be seen that incentive compatibility is equivalent to having for all z_i and x_i

$$U_i(z_i) \geq U_i(x_i) + q_i(x_i)(z_i - x_i). \qquad (1)$$

- This follows by writing for all z_i and x_i

$$\begin{aligned} q_i(x_i)z_i - m_i(x_i) &= q_i(x_i)x_i - m_i(x_i) + q_i(x_i)(z_i - x_i) \\ &= U_i(x_i) + q_i(x_i)(z_i - x_i). \end{aligned}$$

- Eq. (1) implies that for all x_i, $q_i(x_i)$ is a subgradient of the function U_i at x_i.
- Thus at every point that U_i is differentiable,

$$U_i'(x_i) = q_i(x_i).$$

- Since U_i is convex, this implies that q_i is a nondecreasing function.
- Moreover, we have

$$U_i(x_i) = U_i(0) + \int_0^{x_i} q_i(t_i)dt_i. \tag{2}$$

- This shows that, up to an additive constant, the expected payoff to a buyer in an IC direct mechanism (Q, M) depends only on the allocation rule Q.
- From the preceding relations, one can also infer that incentive compatibility is *equivalent to* the function q_i being nondecreasing.

Revenue Equivalence

The payoff equivalence derived in the previous slide leads to the following general revenue equivalence principle.

Proposition (Revenue Equivalence)

If the direct mechanism (Q, M) is incentive compatible, then for all i and x_i, the expected payment is given by

$$m_i(x_i) = m_i(0) + q_i(x_i)x_i - \int_0^{x_i} q_i(t_i)dt_i.$$

Thus the expected payments in any two IC mechanisms with the same allocation rule are equivalent up to a constant.

Revenue Equivalence

Remarks:

- Given two BNE of two different auctions such that for each i:
    - For all $(x_1, \ldots, x_N)$, probability of i getting the object is the same,
    - They have the same expected payment at 0 value.

  These equilibria generate the same expected revenue for the seller.
- This generalizes the result from last time:
    - Revenue equivalence at the symmetric equilibrium of standard auctions (object allocated to buyer with the highest bid).

Individual Rationality (participation constraints)

- A seller cannot force a bidder to participate in an auction which offers him less expected utility than he could get on his own.
- If he did not participate in the auction, the bidder could not get the object, but also would not pay any money, so his payoff would be zero.
- We say that a direct mechanism (Q, M) is individually rational (IR) if for all i and x_i, the equilibrium expected payoff satisfies $U_i(x_i) \geq 0$.
- If the mechanism is IC, then from Eq. (2), individual rationality is equivalent to $U_i(0) \geq 0$.
- Since $U_i(0) = -m_i(0)$, individual rationality is equivalent to

$$m_i(0) \leq 0.$$

Optimal Mechanisms

- Our goal is to design the optimal mechanism that maximizes the expected revenue among all mechanisms that are IC and IR.
- Without loss of generality we can focus on direct revelation mechanisms.
- Consider the direct mechanism (Q, M).
- We can write the expected revenue to the seller as:

$$E[R] = \sum_{i \in N} E[m_i(X_i)], \qquad \text{where}$$

$$\begin{aligned} E[m_i(X_i)] &= \int_0^{w_i} m_i(x_i) f_i(x_i) dx_i \\ &= m_i(0) + \int_0^{w_i} q_i(x_i) x_i f_i(x_i) dx_i - \int_0^{w_i} \int_0^{x_i} q_i(t_i) dt_i f_i(x_i) dx_i \end{aligned}$$

- Changing the order of integration in the third term, we obtain

$$\begin{aligned} E[m_i(X_i)] &= m_i(0) + \int_0^{w_i} (x_i - \frac{1 - F_i(x_i)}{f_i(x_i)}) q_i(x_i) f_i(x_i) dx_i \\ &= m_i(0) + \int_{\mathcal{X}} (x_i - \frac{1 - F_i(x_i)}{f_i(x_i)}) Q_i(x) f(x) dx. \end{aligned}$$

Optimal Mechanism Design Problem

- The optimal mechanism design problem can be written as

$$\begin{aligned} \text{maximize} \quad & E[R] \\ \text{subject to} \quad & IC(\Leftrightarrow q_i \text{nondecreasing}) + IR(\Leftrightarrow m_i(0) \leq 0) \end{aligned}$$

- We define the virtual valuation of a buyer with value x_i as

$$\Psi_i(x_i) = x_i - \frac{1 - F_i(x_i)}{f_i(x_i)}.$$

- We say that the design problem is regular when the virtual valuation $\Psi_i(x_i)$ is strictly increasing in x_i.
- We next show that under this regularity assumption, we can without loss of generality neglect the IC and the IR constraints.
- The seller should choose Q and M to maximize

$$\sum_{i \in \mathcal{N}} m_i(0) + \int_{\mathcal{X}} \Big(\sum_{i \in \mathcal{N}} \Psi_i(x_i) Q_i(x) \Big) f(x) dx.$$

Optimal Mechanism

The following is an optimal mechanism:

- Allocation Rule:

$$Q_i(x) > 0 \Longleftrightarrow \Psi_i(x_i) = \max_{j \in \mathcal{N}} \Psi_j(x_j) \geq 0.$$

- Payment Rule:

$$M_i(x) = Q_i(x)x_i - \int_0^{x_i} Q_i(z_i, x_{-i})dz_i.$$

We finally show that this mechanism satisfies IC and IR.

- We have $M_i(0, x_{-i}) = 0$ for all x_{-i} implying that $m_i(0) = 0$, and therefore satisfying IR.
- By the regularity assumption, for any $z_i < x_i$, we have $\Psi_i(z_i) < \Psi_i(x_i)$. This implies that $Q_i(z_i, x_{-i}) \leq Q_i(x_i, x_{-i})$ for all x_{-i}, and therefore $q_i(z_i) \leq q_i(x_i)$, i.e., q_i is nondecreasing. Hence, IC is also satisfied.

Optimal Mechanism

- The optimal expected revenue is given by

$$E[\max\{\Psi_1(x_1), \ldots, \Psi_N(x_N), 0\}],$$

i.e., it is the expectation of the highest virtual valuation provided it is nonnegative.

- We define

$$y_i(x_{-i}) = \inf\{z_i | \Psi_i(z_i) \geq 0, \Psi_i(z_i) \geq \Psi_j(x_j) \quad \text{for all } j \neq i\},$$

i.e., it is the smallest value for i that wins against x_{-i}.

- Using this, we can write

$$Q_i(z_i, x_{-i}) = \begin{cases} 1 & \textit{if} \quad z_i > y_i(x_{-i}) \\ 0 & \textit{if} \quad z_i < y_i(x_{-i}) \end{cases}$$

Optimal Mechanism

- We have

$$\int_0^{x_i} Q_i(z_i, x_{-i}) = \begin{cases} x_i - y_i(x_{-i}) & \textit{if} \quad x_i > y_i(x_{-i}) \\ 0 & \textit{if} \quad x_i < y_i(x_{-i}) \end{cases}$$

  implying that

$$M_i(x) = \begin{cases} y_i(x_{-i}) & \textit{if} \quad Q_i(x) = 1 \\ 0 & \textit{if} \quad Q_i(x) = 0 \end{cases}$$

- This implies that:
  - Only the winning buyer pays,
  - He pays the smallest value that would result in his winning.

Optimal Mechanism – Symmetric Case

- Suppose that distributions of values are identical across buyers, i.e., for all i, we have $f_i = f$. This implies that for all i, we have $\Psi_i = \Psi$.
- Note that in this case, we have

$$y_i(x_{-i}) = \max\{\Psi^{-1}(0), \max_{j \neq i} x_j\}.$$

Proposition

Assume that the design problem is regular and symmetric. Then a second price auction (Vickrey) with reservation price $r^* = \Psi^{-1}(0)$ is an optimal mechanism.

- Note that, unlike first and second price auctions, the optimal mechanism is not efficient, i.e., object does not necessarily end up with the person who values it most.

6.254 Game Theory with Engineering Applications
Spring 2010