Advanced Microeconomics III

Mechanism Design

Francisco Poggi

Introduction

- Before, we consider the problem of selling an object to a single agent (buyer).
  - We assumed that the agent was making *optimal* choices, given the mechanism.

- We are interested in applications where multiple agents have private information.
  - What outcomes can be *implemented* depends on what solution concept we consider.

Social Choice Problem

- An *environment* consists of:
  - (Finite) set of agents $i \in I$.
  - A set of outcomes Y.
  - Θ_i: set of possible types for agents i. Θ the Cartesian product.
  - $\hat{u}_i : Y \times \Theta \to \mathbb{R}$. utility of agent i given outcome and vector of types.

Social Choice Function

A *social choice function* (SCF) is a mapping $f : \Theta \to Y$.

- Examples:
  - Bilateral trade.
  - Auctions.
  - Public goods.
  - Elections.
  - Etc.

Mechanisms

- A *mechanism* $\Gamma = (S_1, S_2, ..., S_I, g)$ consists of a simultaneous game of incomplete information in which:
  - Each agent has a set of actions S_i. S the Cartesian product.
  - Agents are privately informed of their types.
  - $g : S \to Y$ is the outcome function.
- An environment and a mechanism define a game.
- A *strategy* for agent i in mechanism Γ is a map $\sigma_i : \Theta_i \to S_i$.
- **Question**: which SCF can be implemented given a solution concept? i.e. for which SCF f there exists a mechanism Γ and a strategy profile $\sigma = (\sigma_1, \sigma_2, ..., \sigma_I)$ such that
  - σ is a *solution/equilibrium* of the associated game.
  - $f(\theta) = g(\sigma(\theta))$ for every $\theta \in \Theta$.

Direct Revelation Mechanisms

A *Direct Revelation Mechanism* $g : \Theta \to Y$ is a mechanism in which agents are asked to report a type ($S_i = \Theta_i$) and the outcome function is given by g.

Overview

1 Dominant Strategies Implementation

2 Bayesian Implementation

3 Auctions

Dominant Strategies Implementation

Fix an environment and a mechanism Γ. We say that the strategy profile σ^* is a *dominant strategy solution* if for every agent i, type profile θ, and actions $s_i \in S_i$, and $s_{-i} \in S_{-i}$.

$$u_i(g(\sigma_i^*(\theta_i), s_{-i}), \theta) \geq u_i(g(s_i, s_{-i}), \theta)$$

- This is different than equilibrium in dominant strategies:
  - omitting the condition that the inequality must be sometimes strict is standard in mechanism design.
- The appeal of this solution concept is that is completely "belief-free".

Dominant Strategies Implementation

If there is a mechanism Γ with dominant strategy solution σ^* such that

$$f(\theta) = g(\sigma(\theta)) \quad \text{for all } \theta \in \Theta$$

Then we say that the social choice function f is *implemented in dominant strategies by* Γ.

- Γ is the *implementing mechanism*.
- f is *implementable in dominant strategies*.

Incentive Compatibility

- We say that a SCF f is *Dominant Strategy Incentive Compatible* (DSIC) if, for all $\theta \in \Theta$ $\theta'_i \in \Theta_i$ and $\theta'_{-i} \in \Theta_{-i}$,

$$\hat{u}_i(f(\theta_i, \theta'_{-i}), \theta) \geq \hat{u}_i(f(\theta'_i, \theta'_{-i}), \theta)$$

Revelation Principle

Revelation Principle

A social choice function f is implementable in dominant strategies if and only if f is DSIC.

- f is implementable in dominant strategies $\Rightarrow$ f is DSIC.
  - Otherwise, there is an agent i and type θ_i that would benefit from mimicking another type $\hat{\theta}_i$.
- f is DSIC $\Rightarrow$ f is implementable in dominant strategies.
  - Consider the DRM f.
- This Revelation Principle allows us to focus WLOG on DRM f such that f is DSIC.

Quasi-linear private-values setting

- In many applications, we assume the following structure:
  - $Y = X \times \mathbb{R}^N$ where
    - $x \in X$ is a non-monetary alternative.
    - $t = (t_1, ..., t_N)$ is a profile of monetary transfers.
    - t_i is the payment from agent i.
  - Quasi-linear utility and private values:

$$\hat{u}_i(y, \theta) = u_i(x, \theta_i) - t_i$$

- Examples include auctions and public goods provision.

Quasi-linear private-values setting

- In quasi-linear private-values, the outcome and the SCF have two components:
  - $\alpha : \Theta \to X$ allocation rule.
  - $\tau : \Theta \to \mathbb{R}^N$ transfers rule.

- **Note**: in private-values settings θ_{-i} should be interpreted as the report by i's opponents.

Quasi-linear private-values setting

- In dominant strategy implementation, it does not matter for i whether reports of others coincide with the truth or not.
- Moreover, the solution concept is robust to any distributions of true types, so this does not need to be specified.
- A natural question is whether there are other SCF that can be implemented when we relax the solution concept.

Overview

1. Dominant Strategies Implementation
2. Bayesian Implementation
3. Auctions

Bayesian Implementation

- A Bayesian environment consists of and environment plus a distribution over types $\Phi \in \Delta(\Theta)$, with density ϕ when applicable.

- We assume that
  1. Agents are EU maximizers.
  2. Types are independently distributed.
  3. Quasi-linear utility with private values.

- First, we prove a revelation principle for Bayesian implementation without assuming 2 or 3.

Bayesian Nash equilibrium

- Consider a Bayesian environment and a mechanism Γ.
- A strategy profile σ^* is a *Bayesian Nash equilibrium* if for every agent i and type θ_i,

$$\sigma_i^*(\theta_i) \in \arg \max_{s_i \in S_i} \quad E_{\theta_{-i}} \left[\hat{u}_i(g(s_i, \sigma_{-i}^*(\theta_{-i}), (\theta_i, \theta_{-i})) \mid \theta_i \right]$$

Implementation

Given a Bayesian environment, a mechanism Γ implements a social choice function f if there exists a BNE σ^* of the associated game such that $f(\theta) = g(\sigma^*(\theta))$ for all θ.

- By the revelation principle, we can restrict attention WLOG to DRM.

Revelation Principle for Bayesian Implementation

A social choice function f is called *incentive compatible* (IC) if for all i and θ_i,

$$\theta_i \in \arg\max_{\hat{\theta}_i \in \Theta_i} E_{\theta_{-i}}\left[\hat{u}_i(f(\hat{\theta}_i, \theta_{-i}), (\theta_i, \theta_{-i})) \,\middle|\, \theta_i\right]$$

Revelation Principle

A mechanism that implements f exists if and only if f is IC.

- In other words, implementability is a property of the SCF, no need to check any equilibria of any games.

Revelation Principle for Bayesian Implementation

Proof.

- f is IC $\Rightarrow$ Exists mechanism that implements f
  - Consider the direct mechanism associated with f.
  - Define the *truth-telling strategy* σ_i^* of i:

$$\sigma_i^*(\theta_i) = \theta_i \qquad \text{for all } \theta_i \in \Theta_i$$

  - From IC of f it is immediate that $(\sigma_1^*, ..., \sigma_N^*)$ is an equilibrium of Γ.
  - Hence, f is implemented by Γ.

Revelation Principle for Bayesian Implementation

Proof. (cont).

- Exists mechanism that implements $f \Rightarrow f$ is IC
  - Proof by contrapositive.
  - Suppose that f is not IC.
  - Then there exists i, θ_i and $\hat{\theta}_i$ such that:

$$E_{\theta_{-i}}\left[\hat{u}_i(f(\hat{\theta}_i, \theta_{-i}), (\theta_i, \theta_{-i})) \;\middle|\; \theta_i\right] > E_{\theta_{-i}}\left[\hat{u}_i(f(\theta_i, \theta_{-i}), (\theta_i, \theta_{-i})) \;\middle|\; \theta_i\right]$$

  - Suppose that there exists a game Γ and an equilibrium σ such that $g(\sigma(\theta)) = f(\theta)$
  - Then type θ_i if agent i has an incentive to *mimic* $\hat{\theta}_i$, i.e. deviate to action $\sigma_i(\hat{\theta}_i)$.
  - This contradicts the fact that σ was an equilibrium.

□

Bayesian Incentive Compatibility

- From now on, we consider independent types with quasi-linear utilities and private values (assumptions 2 and 3).
- A DRM is a pair (Q, t) where $Q : \Theta \to \Delta(X)$ and $t : \Theta \to \mathbb{R}^N$.
- Let

$$\bar{Q}_i(\hat{\theta}_i)(x) := \int_{\Theta_{-i}} Q(\hat{\theta}_i, \theta_{-i})(x) \; dF_{-i}(\theta_{-i})$$

  - This denotes the interim expected lottery over X when agent i reports $\hat{\theta}_i$ and all other agents report truthfully.
  - Notice that the distribution does not depend on the true type θ_i. This is because of the independence assumption.
- Similarly, let

$$\bar{t}(\hat{\theta}_i) := \int_{\Theta_{-i}} t_i(\hat{\theta}_i, \theta_{-i}) \; dF_{-i}(\theta_{-i})$$

  - This denotes the expected transfer from i that reports $\hat{\theta}_i$.

Bayesian Incentive Compatibility

- A DRM (Q, t) is *Bayesian Incentive Compatible (BIC)* if for all i and θ_i

$$u_i(\bar{Q}_i(\theta_i), \theta_i) - \bar{t}_i(\theta_i) \geq u_i(\bar{Q}_i(\hat{\theta}_i), \theta_i) - \bar{t}_i(\hat{\theta}_i) \qquad \forall \hat{\theta}_i \in \Theta_i$$

- By virtue of the Revelation Principle, we restrict attention to BIC DRMs.

Interim Individual Rationality

- A DRM (Q, t) is *interim individually rational* if, for all i, all θ_i,

$$U_i(\theta_i) := u_i(\bar{Q}(\theta_i), \theta_i) - \bar{t}_i(\theta_i) \geq 0$$

- $U_i(\theta_i)$ is the *interim* utility of type θ_i of agent i.

Payoff Equivalence

- Incentive compatibility implies that

$$U_i(\theta) = \max_{\hat{\theta}_i \in \Theta_i} \quad u_i(\bar{Q}_i(\hat{\theta}_i), \theta_i) - \bar{t}_i(\hat{\theta}_i)$$

- Applying the Envelope Theorem:

$$U_i(\theta_i) = U_i(0) + \int_0^{\theta_i} u_{i2}(\bar{Q}_i(\tilde{\theta}), \tilde{\theta})\ d\tilde{\theta}$$

Revenue Equivalence

Theorem

Let (Q, t) and (Q', t') be two BIC mechanisms such that $\bar{Q}(\theta_i) = \bar{Q}'(\theta_i)$ for all i and θ_i. Then there exist C_i such that $\bar{t}(\theta_i) = \bar{t}'(\theta_i) + C_i$ for all θ and all i.

- **Note**: First price auction, second price auction, English auction, and Dutch auction generate the same allocation and give zero to each of the lowest bidder types.
- By revenue equivalence, they must all generate the same revenue for the seller.

Overview

1. Dominant Strategies Implementation

2. Bayesian Implementation

3. Auctions

Auctions

- Buyers: $i = 1, ...N$
- Single indivisible object.
- Buyer i values the object θ_i.
- Independent valuations: θ_i distributed with cdf F_i and pdf f_i.
- Seller knows F_i.

Auctions

- Auction setting:

$$X = \left\{ (x_1, ..., x_N) \in [0,1]^N : \sum_{j=1}^{N} x_j \leq 1 \right\}$$

$$u_i(x, \theta_i) = \theta_i \cdot x_i$$

Revenue Maximizing Auctions

- For any auction (and in any linear-utility environment with voluntary participation) we can pose the question:
  - Among all scf f that can be implemented with voluntary participation, what is the one that maximizes expected revenue $R(f)$?

$$\max_f \quad R(f) \qquad s.t. \qquad f \text{ is IC and } U_i(\theta_i) \geq \bar{u}_i(\theta_i) = 0$$

- (We normalize outside value of each type to zero.)

Optimal Auctions

- By the Revelation Principle we can focus on DRM.
  - $q : \Theta \to [0,1]^N$,
  - $\sum_i q_i(\theta) \leq 1$
  - $t : \Theta \to \mathbb{R}^N$

$$U_i(\theta_i) = E_{\theta_{-i}}[\theta_i q_i(\theta) - t_i(\theta)] = \theta_i \bar{q}_i(\theta_i) - \bar{t}_i(\theta_i)$$

- Where

$$\bar{q}_i(\theta_i) = E_{\theta_{-i}}[q_i(\theta)]$$
$$\bar{t}_i(\theta_i) = E_{\theta_{-i}}[t_i(\theta)]$$

Maximization Problem

- Choose the DRM (q, t) that maximizes expected revenue subject to
  - Bayesian Incentive Compatibility
  - Interim Individual Rationality

- (Seller's value for the object is normalized to zero.)

Expected Total Revenue

$$
\begin{aligned}
E[R] &= E_\theta \sum_{i=1}^{N} t_i(\theta) \\
&= \sum_{i=1}^{N} E_\theta[t_i(\theta)] \\
&= \sum_{i=1}^{N} E_{\theta_i}[\bar{q}_i(\theta_i)\theta_i - U_i(\theta_i)]
\end{aligned}
$$

Expected Revenue from single bidder

- By payoff-equivalence:

$$U_i(\theta_i) = U_i(0) + \int_0^{\theta_i} \bar{q}_i(s)ds$$

- So, (recall from the single buyer case)

$$\begin{aligned} E[R_i] :&= E_{\theta_i}[\bar{q}_i(\theta_i)\theta_i - U_i(\theta_i)] \\ &= \int_0^1 \left[\bar{q}_i(r)r - U_i(0) - \int_0^r \bar{q}_i(s)\ ds\right] f_i(r)\ dr \\ &= E_{\theta_i}\left[\bar{q}_i(\theta_i) \cdot VS_i(\theta_i)\right] - U_i(0) \end{aligned}$$

Total Expected Revenue

$$E[R] := E_\theta\left[\sum_{i=1}^{N} q_i(\theta)\left[\theta_i - \frac{1-F_i(\theta_i)}{f_i(\theta_i)}\right]\right] - \sum_{i=1}^{N} U_i(0)$$

- Seller chooses the functions q_i and the constants $U_i(0)$ to maximize the expression subject to:
  - Monotonicity.
  - IIR.
- At the optimum, $U_i(0) = 0$ for all $i \in I$.
- All IIR constraints are satisfied by the envelope condition.

Ignoring Monotonicity

$$\max_{q\nearrow} \quad E_\theta \left[\sum_{i=1}^{N} q_i(\theta) \left[\theta_i - \frac{1 - F_i(\theta_i)}{f_i(\theta_i)} \right] \right]$$

- As before, we
    - ignore monotonicity,
    - maximize separately for all $\theta \in \Theta$
    - check if the allocation rule satisfies monotonicity.

Ignoring Monotonicity

$$\max_{q} \quad \sum_{i=1}^{N} q_i(\theta) \left[\theta_i - \frac{1 - F_i(\theta_i)}{f_i(\theta_i)} \right]$$

- The optimal q is:

$$q_i(\theta) = \begin{cases} 1 & \text{if } VS_i(\theta_i) > VS_j(\theta_j) \ \forall j \neq i \text{ and } VS_i(\theta_i) \geq 0 \\ 0 & \text{otherwise.} \end{cases}$$

- (Ties are not important.)
- This allocation rule is monotone if VS_i is nondecreasing.
- A sufficient condition (often assumed) is that hazard rate is increasing.

Properties of optimal auctions

- Downward distortions: the seller might inefficiently retain the object.
  - This happens when VS are all negative but θ_i is positive for some i.
- For symmetric bidders with nondecreasing hazard rate, the allocation rule is efficient conditional on sale.
- For asymmetric bidders, the object might be allocated to a bidder different than the one that values the good the most.
- In the symmetric case, the optimal auction can be implemented by any of the standard auction formats (FPSB, SPSB, English, Dutch) with a reserve price.

Dominant Strategy Implementation

- The first price auction with an optimal reserve price maximizes, in equilibrium, the revenue of the seller.
- The same allocation and revenue can be obtained with a second price auction. However, the equilibrium in the second price auction is in dominant strategies!
- Manelli and Vincent (2010) provide conditions under which SCF that are BIC can also be implemented in Dominant Strategies.