G5212: Game Theory

Mark Dean

Spring 2017

Introduction to Adverse Selection

- In the today's lecture, we will do two things
  1. Give an example of why adverse selection problems are important
  2. Give an introduction to Mechanism Design -
     - A broader class of problems of which adverse selection is one example

An Insurance Example

- Consider the following scenario (and see if it reminds you of anything):
  - The population consists of people who have different probabilities of being sick
  - Probability of getting sick is uniformly drawn from 0 to 100%
  - Treatment costs $100
  - A company wants to offer insurance
  - Individuals know how likely they are to get sick, but insurance companies don't know
    - Or are not allowed to charge different prices based on this probability
- Note that this is an adverse selection problem
  - Patients know their type, insurance company does not
  - Insurance company moves first (offers insurance), then patients decide whether or not to take it

An Insurance Example

- Let's think of an insurance company which is
  - Kindly
  - Dumb
- So they just offer 'actuarilly fair' insurance contract
  - One single contract
  - Prices so firm breaks even
- Can this firm insure everyone?

An Insurance Example

- Can this firm insure everyone with this contract?
- No
  - If they insure everyone then the average probability of any one person getting sick is is 50%
    - Cost of the insurance contract is therefore $50
  - Who will buy such a contract?
    - Only people whose risk of illness is above 50%
  - But for this pool of people the risk of illness is 75%
    - Cost of acuarilly fair contract must rise to $75
  - Who will buy this contract?
    - Only people whose risk of illness is above 75%
  - And so on

An Insurance Example

- This is an example of **unravelling** in the insurance market
- For this type of contract there is no equilibrium in which anyone (apart form the worst type) gets insured
- If the insurees have private information then simple, fair contracts don't work
  - Need to do something else
- We can use brute force
  - e.g. individual mandates
- Or design smarter contracts

Mechanism Design

- Adverse selection is an example of a broader area of study: Mechanism design
  - Adverse selection the principal only deals with one agent at a time
  - Mechanism design problems may involve many agents
- General mechanism design problem
  - N agents
  - Each agent has a type $\theta_n \in \Theta_N$ ($\Theta = \Pi_{n \in N} \Theta_n$)
  - μ: probability distribution over Θ
  - There is a set of possible ouctomes Y_n for each agent ($Y = \Pi_{n \in N} Y_n$)
  - Principal has an objective function $y : \Theta \to Y$ which determines what oucome they would like given θ
  - Agents have preferences $u(y_i, \theta_i)$

Mechanism Design

- Principal's problem would be trivial if they observed θ
- But maybe they don't
- A mechanism is a set of possible messages M_i for each agent i, and a set of rules of the game described by g.
- The center commits to implement an allocation $g(m)$, where $m = (m_1, ..., m_n)$, and m_i is the message sent by i.
- Each agent i has information I_i, which contains θ_i.
- Using I_i, each agent i selects $m_i^* \in M_i$ according to some rule
- Implemented allocation is $g(m_1^*(I_1), ..., m_n^*(I_n))$.

Mechanism Design

- The center can be a government, a regulator, or a "principal" (seller, employer,...).
- The mechanism can be extremely complex (using bribes for revealing the truth, punishing caught liars,...).
- Questions we might want to ask:
  - Is y implementable? In other words, can we find a mechanism in which $y(\theta) \equiv g(m^*(\theta))$ for all θ?
    - Most cannot: They are not incentive-compatible.
    - Managing information generates distortions.
  - What is the best choice among different implementable solutions?
    - Maximize principal's objective under incentive-compatible constraints.
    - Find a second-best solution that minimizes economic inefficiencies.

Public Good

- Example: Public Good
- Suppose that the utility of agent i of consuming the public good G is a function of some private information θ_i: $U_i = U(G, \theta_i) - t_i$.
- If the social planner would know $(\theta_1, ..., \theta_n)$ and c, she would select the G^* that solves the FB problem:

$$G^* \in \arg\max_G \sum_{i=1}^{n} \{U(G, \theta_i)\} - cG.$$

Public Good

- But what if θ is not known?
- How could we implement this?
- We could ask everyone how much public good everyone wanted and implement the average
  - But then people who want a lot of the public good would have an incentive to exaggerate upwards
- We could ask them their willingness to pay and charge them that
  - But then everyone would have an incentive to exaggerate downwards
- Need a mechanism that *implements* y

Types of Implementation

- Types of implementation:
  - **Dominant strategy equilibrium:** m^* is the best strategy for every agent, regardless of what other agents do: for all I_i and m_{-i},

$$m_i^*(I_i) \in \arg\max_{m_i \in M_i} U_i(g(m_i, m_{-i})) \quad .$$

  - **Nash equilibrium** (only when agents know θ): if all other agents act according to m^* so does i: for all θ and i,

$$m_i^*(\theta) \in \arg\max_{m_i \in M_i} U_i(g(m_i, m_{-i}^*(\theta))).$$

  - **Bayesian equilibrium:** every agent i has a belief $\mu(I_{-i} \mid I_i)$ on others' information, conditional on observing his own; For all I_i,

$$m^*(I_i) \in \arg\max_{m_i} E\left(U_i\left(g(m_i, m^*(I_{-i}))\right) | I_i\right)$$

    where the expectation is over the belief μ.

Implementation in Dominant Strategies

- Gold standard is implementation in dominant strategies
- Very demanding, but robust ("detail-free").
  - Requires only rationality
  - Of course this is the only option for a problem when $n = 1$.
- If a mechanism (M, g) implements a social choice function y in dominant strategies, we can call y *strategy-proof*, or *non-manipulable.*
- Requires the mechanism to satisfy Incentive Compatibility (IC) constraint:

$$\forall i, \forall m_i' \in M_i, \forall m_{-i} : \; U(g(m_i^*, m_{-i}), \theta_i) \geq U(g(m_i', m_{-i}), \theta_i)$$

  - Where m_i^* is the equilibrium message

Implementation in Dominant Strategies

- The space of possible strategies is huge
- Here is an extremely handy theorem which is going to help us to narrow it down

Definition

A direct revelation mechanism is one in which the message space for each player is their type space. A truthful mechanism is a direct revelation mechanism in which everyone truthfully reports their type

Implementation in Dominant Strategies

Proposition (Revelation Principle)

If a social choice y can be implemented by some mechanism (M, g) in dominant strategies, then there exists a truthful direct revelation mechanism that implements y in dominant strategies.

- So we can focus on asking each agent to report what he knows, and on mechanisms in which each agent reports truthfully.

Implementation in Dominant Strategies

- Proof:
- Take a strategy-proof y. By definition, $\exists(g, M)$ such that each agent has a unique dominant strategy to play some $m^*(\theta_i)$, and for every profile we have

$$y(\theta_1, ..., \theta_n) = g\left(m^*(\theta_1), \ldots, m^*(\theta_n)\right).$$

- Since it is a dominant strategy for i to play $m^*(\theta_i)$, we have that $\forall(m_i, m_{-i})$

$$U_i\left(g(m^*(\theta_i), m_{-i})|\theta_i\right) \geq U_i\left(g(m_i, m_{-i})|\theta_i\right)$$

- In particular, $\forall\hat{\theta}_i, \hat{\theta}_{-i}$:

$$U_i \underbrace{\left(g(m^*(\theta_i), m^*(\hat{\theta}_{-i}))|\theta_i\right)}_{y(\theta_i, \hat{\theta}_{-i})} \geq U_i \underbrace{\left(g(m^*(\hat{\theta}_i), m^*(\hat{\theta}_{-i}))|\theta_i\right)}_{y(\hat{\theta}_i, \hat{\theta}_{-i})}$$

Mechanism Design

- So now construct a direct revelation mechanism such that

$$\hat{g}\left(\theta_1, \ldots, \theta_n\right) = y(\theta_1, ..., \theta_n)$$

- It follows directly that

$$\begin{aligned}
& U_i\left(\hat{g}(\theta_i, \theta_{-i})|\theta_i\right) \\
= \; & U_i\left(y(\theta_i, \theta_{-i})|\theta_i\right) \\
= \; & U_i\left(g(m^*(\theta_i), m^*(\theta_{-i}))|\theta_i\right) \\
\geq \; & U_i\left(g(m^*(\hat{\theta}), m^*(\theta_{-i}))|\theta_i\right) \\
= \; & U_i\left(y(\hat{\theta}_i, \theta_{-i})|\theta_i\right) \\
= \; & U_i\left(\hat{g}(\hat{\theta}_i, \theta_{-i})|\theta_i\right)
\end{aligned}$$

. ■

Mechanism Design

- Implementation in dominant strategies works really well sometimes. VCG mechanism in the provision of a public good.
- M is the set of possible utility functions, and $g = (G, t_1, ..., t_n)$.
  - G is the level of public good
  - t_i is the transfer to person i
- Government wants to choose G to maximize

$$\sum_i u_i(G) - cG$$

- Preferences are quasi linear

$$u_i(G) - t_i$$

Mechanism Design

- VCG mechanism works as follows:
  1. Each person reports their type $\hat{u}_i$
  2. Government chooses G to maximize $\sum_i \hat{u}_i(G) - cG$
  3. Charges each person

$$t_i(\hat{u}) = cG(\hat{u}) - \sum_{j \neq i} \hat{u}_j(G(\hat{u})).$$

Mechanism Design

- Claim: it is optimal to report $\hat{u}_i = u_i$
- Note that payoff for i is

$$\begin{aligned} & u_i(G(\hat{u}_i, \hat{u}_{-i})) - t_i \\ = \ & u_i(G(\hat{u}_i, \hat{u}_{-i})) - \left(cG(\hat{u}_i, \hat{u}_{-i}) - \sum_{j \neq i} \hat{u}_j(G(\hat{u}_i, \hat{u}_{-i})) \right) \\ = \ & u_i(G(\hat{u}_i, \hat{u}_{-i})) + \sum_{j \neq i} \hat{u}_j(G(\hat{u}_i, \hat{u}_{-i})) - cG(\hat{u}_i, \hat{u}_{-i}) \end{aligned}$$

Mechanism Design

- Say that for some $\hat{u}_i \neq u_i$ we had

$$\begin{aligned} & u_i(G(\hat{u}_i, \hat{u}_{-i})) + \sum_{j \neq i} \hat{u}_j(G(\hat{u}_i, \hat{u}_{-i})) - cG(\hat{u}_i, \hat{u}_{-i}) \\ > \quad & u_i(G(u_i, \hat{u}_{-i})) + \sum_{j \neq i} \hat{u}_j(G(u_i, \hat{u}_{-i})) - cG(u_i, \hat{u}_{-i}) \end{aligned}$$

- This violates the fact that $G(u_i, \hat{u}_{-i})$ was chosen in order to maximize the latter expression

Mechanism Design

- A simple example
  - $n = 2$
  - $G \in \{0, 1\}$ (build a bridge or not)
  - $c = 1$,
- Type of each player is the utility from the bridge:
  - $u_i(1)$
  - Assumue this is drawn from some distribution F
  - Normalize $u_i(0) = 0$
- Utility of a player who is charged t_i is

$$u_i(G) - t_i$$

Mechanism Design

- VGC mechanism in this case
  - Each player announces $u_i'(1)$
  - If $u_1'(1) + u_2'(1) \geq 1$ then bridge is built
  - Taxes are

$$
\begin{aligned}
t_1 &= c - u_2'(1) \\
t_2 &= c - u_1'(1)
\end{aligned}
$$

  - If the bridge gets build,
  - If $u_1'(1) + u_2'(1) < 1$ bridge is not built, zero taxes

Mechanism Design

- Claim: truthtelling is the optimal strategy
- Focus on player 1, treat $u'_2(1)$ as fixed
- Assume player 1 announces $u_1(1)$
- Two cases
    - $u'_2(1) + u_1(1) < 1$
    - $u'_2(1) + u_1(1) > 1$

Mechanism Design

- Case 1:
  - $u_2'(1) + u_1(1) < 1$
  - Bridge does not get built under truthtelling
  - Player 1 gets a utility of 0
- Can they do better by lying?
  - The only way to get the bridge built is by announcing $u_1'(1) > 1 - u_2'(1)$
  - This would provide payoff

$$
\begin{aligned}
& u_1(1) - t_1 \\
= \ & u_1(1) - (1 - u_2'(1)) \\
= \ & u_1(1) + u_2'(1) - 1 < 0
\end{aligned}
$$

Mechanism Design

- Case 2:
  - $u_2'(1) + u_1(1) > 1$
  - Bridge does gets built under truthtelling
  - Player 1 gets a utility of

$$\begin{aligned} & u_1(1) - t_1 \\ = \; & u_1(1) - (1 - u_2'(1)) \\ = \; & u_1(1) + u_2'(1) - 1 > 0 \end{aligned}$$

- Can they do better by lying?
  - Notice that changing their announcement does not change their tax rate assuming bridge gets built
  - So the only thing player 1 can do to change their payoff is to announce a utility

$$u_1(1) < 1 - u_2'(1)$$

  - Bridge won't get built
  - Utility of 0