Adverse Selection

Mechanism Design

Mechanism design

- how can we aggregate individual preferences into a collective decision?
- especially if individuals' preferences are private information

Compared to the screening problem, we now consider multiple agents

- interests may conflict with each other
- there is increased competition that a seller may exploit
- will inefficiencies increase/decrease?

Mechanism Design – General Setup

The Environment

- n agents
- each agent i has private information (his type) $\theta_i \in \Theta_i$
- set of possible alternatives/outcomes $x \in X$
- each agent is expected-utility maximiser with vNM utility function

$$u_i(\theta, x) \in \mathbb{R}, \qquad \text{for } \theta \in \Theta = \Theta_1 \times \cdots \times \Theta_n \text{ and } x \in X.$$

- the type profile $\theta = (\theta_1, \ldots, \theta_n)$ is distributed according to F with density $f > 0$
- **notation:** we write

$$\theta_{-i} = (\theta_1, \ldots, \theta_{i-1}, \theta_{i+1}, \ldots, \theta_n) \quad \text{and} \quad (\theta_i, \theta_{-i}) = \theta$$

Mechanism Design – Setup – some terminology

Private Values

- i's preferences depend only on θ_i:

  $u_i(\theta, x) = u_i(\theta_i, x)$
- 'interdependent values' otherwise

Independent Types

- θ_i's distribution indep. of other types θ_{-i}:

  $f(\theta) = \prod_{i=1}^{n} f_i(\theta_i)$
- 'correlated' types otherwise

Quasi-linear Utilities

- outcomes $X = K \times \mathbb{R}^n$, where

  $k \in K$ some physical allocation,

  $t = (t_1, \ldots, t_n) \in \mathbb{R}^n$ transfers
- i's utility is linear in money (his transfer):

  $u_i(\theta, x) = v_i(\theta, k) + t_i$

Social Choice/Unrestricted Domain

- $X = \{a, b, \ldots\}$ finite set of alternatives
- θ_i gives ranking over alternatives:

  $a \theta_i b \;\Leftrightarrow\; a \succ_i b$
- Unrestricted domain if

  Θ_i contains all possible rankings over X

Mechanism Design – environment examples

Ex 1. Public good

- outcomes $(k, t) \in X = \{0, 1\} \times \mathbb{R}^n$

  $k \in \{0, 1\}$ with $k = 1$ if bridge is built

  $t_i \in \mathbb{R}$ transfer to agent i

- θ_i is i's willingness to pay for bridge

  $u_i(\theta, x) = \theta_i k + t_i$

Ex. 2 Allocation with externalities

- outcomes $(k, t) \in X = \{0, 1, \ldots, n\} \times \mathbb{R}^n$

  $$k = \begin{cases} 0 & \text{if nobody gets object} \\ i & \text{if agent } i \text{ gets object} \end{cases}$$

  $t_i \in \mathbb{R}$ transfer to agent i

- $\theta_i = (\theta_i^i, \theta_i^x)$ with utility

  $$u_i(\theta, x) = \begin{cases} t_i & \text{if } k = 0 \\ \theta_i^i + t_i & \text{if } k = i \\ -\theta_i^x + t_i & \text{if } k \notin \{0, i\} \end{cases}$$

Mechanism Design – Social Choice Functions

Our goal is generally to choose a good outcome $x \in X$ given the realised preferences $\theta \in \Theta$

Definition (Social Choice Function)

A **social choice function** (scf) $\xi\colon \Theta \to X$ assigns to each type profile $\theta \in \Theta$ an alternative $\xi(\theta) \in X$.

The problem of the mechanism designer is not 'lack of power'

- if the designer knew θ, she could always choose the 'optimal' outcome

The problem is 'just' the asymmetric information

Mechanism Design – Mechanisms (general/indirect)

Typically, social (collective) outcomes are determined through interaction in some institution

Definition (Mechanism)

A **mechanism** $\Gamma = (S_1, \ldots, S_n, g)$ consists of

- a strategy space S_i for each agent i
- an outcome function $g : S_1 \times \cdots \times S_n \to X$.

A mechanism $\Gamma = (S, g)$ together with the environment induces a Bayesian game:

$G_\Gamma = (n,\ \{S_i\}_{i \leq n},\ \{\tilde{u}_i\}_{i \leq n},\ \Theta,\ F)\,,$ with payoffs $\tilde{u}_i(\theta, s_1, \ldots, s_n) = u_i(\theta, g(s_1, \ldots, s_n))$.

Mechanism Design – Incentive Compatibility

We have several solution concepts: Let $(s_i^*)_{i=1}^n$ be a strategy profile, where $\forall i\colon\ s_i : \Theta_i \to S_i$

- Dominant strategy equilibrium: for all i, θ_i, s_i:

$$u_i((\theta_i, \theta_{-i}), g(s_i^*(\theta_i), s_{-i})) \geq u_i((\theta_i, \theta_{-i}), g(s_i, s_{-i})) \quad \forall \theta_{-i}, s_{-i}$$

- Ex-post equilibrium: for all i, θ_i, s_i:

$$u_i((\theta_i, \theta_{-i}), g(s_i^*(\theta_i), s_{-i}^*(\theta_{-i}))) \geq u_i((\theta_i, \theta_{-i}), g(s_i, s_{-i}^*(\theta_{-i}))) \quad \forall \theta_{-i}$$

- Bayes-Nash equilibrium: for all i, θ_i, s_i:

$$\int_{\Theta_{-i}} u_i((\theta_i, \theta_{-i}), g(s_i^*(\theta_i), s_{-i}^*(\theta_{-i}))) \ \mathrm{d}F_{-i}(\theta_{-i}|\theta_i)$$
$$\geq \int_{\Theta_{-i}} u_i((\theta_i, \theta_{-i}), g(s_i, s_{-i}^*(\theta_{-i}))) \ \mathrm{d}F_{-i}(\theta_{-i}|\theta_i)$$

Mechanism Design – Participation Constraints

Mechanism Design – Implementation

Definition

We say that mechanism $\Gamma = (S, g)$ [...]-**implements** scf ξ
if there exists a [...]-equilibrium strategy profile $(s_i^*)_{i=1}^n$ such that

$$g(s_1^*(\theta_1), \ldots, s_n^*(\theta_n)) = \xi(\theta) \quad \text{for all } \theta \in \Theta.$$

where [...]$\in$ {dominant strategy, ex-post, Bayes}

- Full implementation: every equilibrium results in $\xi(\theta)$
- Partial implementation: there is an equilibrium that results in $\xi(\theta)$

We focus on partial implementation

Mechanism Design – Revelation Principle (direct mechanisms)

Theorem (Revelation Principle)

For any mechanism $\Gamma = (S, g)$ and [...]-equilibrium strategy profile $(s_i^*)_{i=1}^n$ that implements scf ξ, there exists a ***direct*** *mechanism $\hat{\Gamma} = (\Theta, \xi)$ such that* ***truthtelling*** *is a [...]equilibrium.*

- Only ensures that there is AN equilibrium
- In different (indirect) mechanisms sharing the same direct mechanism other equilibria may arise

Mechanism Design – Revelation Principle – proof

proof of revelation principle for dominant strategy case

Mechanism Design – The Gibbard-Satterthwaite Theorem

Recall from micro 3:

Definition (Dictatorial)

An scf $\xi : \Theta_1 \times \cdots \times \Theta_n \to X$ is **dictatorial** if there is an agent $d \in \{1, \ldots, n\}$ such that $\xi(\theta_d, \theta_{-d})$ is always the favourite outcome of type θ_d.

Theorem (Gibbard-Satterthwaite)

Suppose $|X| \geq 3$ and for all i, Θ_i contains all possible preference rankings over X. If scf ξ with $\xi(\Theta) = X$ is strategy proof, then it is dictatorial.

With unrestricted preferences, there is not a lot we can do...

Not hopeless if preferences are more restricted:

- voting/social-choice literature typically focuses on single-peaked preferences
- we will consider quasi-linear utilities and (mostly) private values

Mechanism Design – quasi-linear utility and private value

- Outcomes: $X = K \times \mathbb{R}^n$: $k \in K$ allocation and $(t_1, \ldots, t_n) \in \mathbb{R}^n$ transfers
- Utilities: $u_i(\theta, x) = v_i(\theta_i, k) + t_i$

Note:

- $v_i(\theta_i, k)$ measures the value of allocation k in terms of money
- Utility is transferable across agents through money
- Agents are risk-neutral with respect to money

Mechanism Design – quasi-linear utilities and efficiency

Definition (Pareto efficiency)

An outcome $x = (k, t_1, \ldots, t_n) \in X$ is Pareto efficient if there is no other $x' = (k', t'_1, \ldots, t'_n) \in X$ such that:

$$\sum_{i=1}^{n} t'_i = \sum_{i=1}^{n} t_i \qquad \text{and} \qquad v_i(\theta_i, k') + t'_i \geq v_i(\theta_i, k) + t_i$$

for all i, with strict inequality for at lease one i.

Proposition

An scf $\xi = (k, t)$ is Pareto efficient if and only if for all $\theta \in \Theta$:

$$\sum_{i=1}^{n} v_i(\theta_i, k(\theta)) \geq \sum_{i=1}^{n} v_i(\theta_i, k') \ \ \forall k'.$$

Mechanism Design – VCG Mechanisms

Definition

A Vickrey-Clarke-Groves (VCG) mechanism is given by (k^*, t) where k^* is efficient and

$$t_i(\theta) = \sum_{j \neq i} v_j(\theta_j, k^*(\theta)) + h_i(\theta_{-i}),$$

for some collection of functions $(h_i)_i$ where each h_i is independent of θ_i

Theorem

Truthtelling is a dominant-strategy equilibrium of any VCG mechanism.

Mechanism Design – VCG Mechanism – proof

Mechanism Design – VCG Mechanisms

A special case of VCG mechanisms is the pivot mechanism (or Clarke mechanism):

Definition (pivot mechanism)

A pivot mechanism is a VCG mechanism with

$$h_i(\theta_{-i}) = -\sum_{j \neq i} v_j(\theta_j, k^*_{-i}(\theta_{-i})),$$

where $k^*_{-i}(\theta_{-i})$ is an efficient alternative for the $n-1$ agents different from i

- Each agent pays the externality imposed on other agents:

$$t_i(\theta) = \sum_{j \neq i} v_j(\theta_j, k^*(\theta)) - \sum_{j \neq i} v_j(\theta_j, k^*_{-i}(\theta_{-i})).$$

- If adding agent i with type θ_i does not change allocation, then $t_i = 0$
- The second-price auction is a pivot mechanism

Mechanism Design – VCG Mechanisms

- Is there an ex-post efficient mechanism that is DIC but not a VCG mechanism?
- If the environment is 'rich' enough, the answer is no:

  Let $\mathcal{V}$ denote the set of all possible functions from K to $\mathbb{R}$

Theorem

If for all agents i, the set of preferences is such that $\{v_i(\theta_i, \cdot)\}_{\theta_i \in \Theta_i} = \mathcal{V}$, then every direct mechanism in which truthtelling is a dominant strategy is a VCG-mechanism.

Mechanism Design – DIC and efficiency

Ex post efficiency and DIC is 'almost equivalent' to VCG mechanism

That is great because...

- these are simple to characterise
- we can simply check for the best VCG mechanism in each situation

However,...

- they potentially require large transfers
- we have ignored participation constraints
- they are generally not budget balanced: $\sum_i t_i(\theta) \neq 0$

Mechanism Design – Bayesian incentive compatibility

What if we weaken our solution concept and look at Bayesian Mechanism Design?

- We will focus in the independent case: $f(\theta) = \prod_i f_i(\theta_i)$

Recall: truthtelling is a Bayes-Nash equilibrium if for all i and all θ_i:

$$\mathbb{E}_{\theta_{-i}}\left[v_i(\theta_i, k(\theta_i, \theta_{-i})) + t_i(\theta_i, \theta_{-i})\right] \geq \mathbb{E}_{\theta_{-i}}\left[v_i(\theta_i, k(\hat{\theta}_i, \theta_{-i})) + t_i(\hat{\theta}_i, \theta_{-i})\right] \quad \forall \hat{\theta}_i \qquad \text{(BIC)}$$

We hope that we can exploit weakened IC requirement (now only in expectation over θ_{-i}) to eliminate some undesirable features of VCG mechanisms.

- ...and indeed we can
- ...at first sight

Mechanism Design – Expected Externality Mechanism

- Let k^* be an ex-post efficient allocation rule
- Consider the following transfers:

$$t_i(\theta_i, \theta_{-i}) = \mathbb{E}_{\tilde{\theta}_{-i}}\left[\sum_{j\neq i} v_j(\tilde{\theta}_j, k^*(\theta_i, \tilde{\theta}_{-i}))\right] + h_i(\theta_{-i}),$$

with

$$h_i(\theta_{-i}) = -\frac{1}{n-1}\sum_{j\neq i}\mathbb{E}_{\tilde{\theta}_{-j}}\left[\sum_{\ell\neq j} v_\ell(\tilde{\theta}_\ell, k^*(\theta_j, \tilde{\theta}_{-j}))\middle|\,\theta_j\right].$$

Definition (Expected Externality Mechanism)

The mechanism (k^*, t) defined above is called Expected Externality Mechanism.

Mechanism Design – Expected Externality Mechanism

Proposition

The Expected Externality Mechanism is budget balanced and truthtelling is BIC.

That is

$$\begin{aligned}
&\sum_{i=1}^{n} t_i(\theta) \\
&= \sum_{i=1}^{n} \mathbb{E}_{\tilde{\theta}_{-i}} \left[\sum_{j \neq i} v_j(\tilde{\theta}_j, k^*(\theta_i, \tilde{\theta}_{-i})) \right] - \frac{1}{n-1} \sum_{i=1}^{n} \sum_{j \neq i} \mathbb{E}_{\tilde{\theta}_{-j}} \left[\sum_{\ell \neq j} v_\ell(\tilde{\theta}_\ell, k^*(\theta_j, \tilde{\theta}_{-j})) \right] \\
&= 0.
\end{aligned}$$

Mechanism Design – Expected Externality Mechanism

- Expected Externality mechanism achieves budget balance
- but did we really gain that much?

  the following result suggests no:

Theorem

Fix an ex-post efficient allocation rule k^* and a BIC mechanism that implements k^*. Then there exist constants h_i such that the VCG mechanism with transfer rule

$$t_i(\theta) = \sum_{j \neq i} v_j(\theta_j, k^*(\theta)) + h_i$$

gives each player the same interim payoff.

Mechanism Design – BIC and Efficiency

To sum up:

- VCG mechanisms give us a pretty complete picture of the expected utilities that can be achieved in incentive compatible and efficient mechanisms
- With expected externality mechanism we can achieve budget balance ex post

But...

- We still completely ignored participation constraints
- ...and that is generally problematic as we see now

Mechanism Design – Bilateral Trade

Question: Is efficient trade possible when both sides have private information?

- single indivisible good
- one buyer with $\theta \in [\underline{\theta}, \bar{\theta}]$ drawn from F
- one seller with production cost $c \in [\underline{c}, \bar{c}]$ drawn from G
- trade is efficient sometimes: $\underline{c} < \bar{\theta}$ but not always: $\underline{\theta} < \bar{c}$

Theorem (Myerson-Satterthwaite)

There is no ex-post efficient, budget balanced, BIC mechanism that satisfies interim IR for buyer and seller.

Mechanism Design – Bilateral Trade – proof of Myerson Satterthwaite

Direct Mechanism:

- $q(\theta, c) \in [0, 1]$ = prob. of trade
- $t_B(\theta, c)$ transfer to buyer $\quad t_S(\theta, c)$ transfer to seller

The buyer's expected utility from report $\hat{\theta}$ is

$$\int_{\underline{c}}^{\bar{c}} \Big(\theta q(\hat{\theta}, c) + t_B(\hat{\theta}, c)\Big) \, \mathrm{d}G(c)$$

Define:

- $Q_B(\hat{\theta}) = \int_{\underline{c}}^{\bar{c}} q(\hat{\theta}, c) \, \mathrm{d}G(c)$ and $T_B(\hat{\theta}) = \int_{\underline{c}}^{\bar{c}} t_B(\hat{\theta}, c) \, \mathrm{d}G(c)$ for buyer
- $Q_S(\hat{c}) = \int_{\underline{\theta}}^{\bar{\theta}} q(\theta, \hat{c}) \, \mathrm{d}F(\theta)$ and $T_S(\hat{c}) = \int_{\underline{\theta}}^{\bar{\theta}} t_S(\theta, \hat{c}) \, \mathrm{d}F(\theta)$ for seller

Mechanism Design – Bilateral Trade – proof of Myerson Satterthwaite

Incentive compatibility (Bayesian):

$$\theta Q_B(\theta) + T_B(\theta) \geq \theta Q_B(\hat{\theta}) + T_B(\hat{\theta}) \qquad (BIC_{buyer})$$

$$T_S(c) - cQ_S(c) \geq T_S(\hat{c}) - cQ_S(\hat{c}) \qquad (BIC_{seller})$$

Individual rationality (interim):

$$\theta Q_B(\theta) + T_B(\theta) \geq 0 \qquad (IR_{buyer})$$

$$T_S(c) - cQ_S(c) \geq 0 \qquad (IR_{seller})$$

Budget Balance holds if $t_B(\theta, c) + t_S(\theta, c) \leq 0$, we will require a weaker condition:

$$\int_{\underline{\theta}}^{\bar{\theta}} \int_{\underline{c}}^{\bar{c}} (t_B(\theta, c) + t_S(\theta, c)) \, \mathrm{d}F(\theta) \, \mathrm{d}G(c) \leq 0 \qquad (BB)$$

no mechanism with ex-post efficient trade ($q(\theta, c) = \mathbb{1}_{\{\theta > c\}}$) satisfies these conditions

Mechanism Design – Bilateral Trade – proof of Myerson Satterthwaite

We can apply screening results to expected terms Q and T to conclude

Lemma

Suppose (q, t_B, t_S) *satisfies* BIC_{buyer} *and* BIC_{seller}, *then*

1. $Q_B(\theta)$ *is non-decreasing*
2. $Q_S(c)$ *is non-increasing*
3. $V_B(\theta) = V_B(\underline{\theta}) + \int_{\underline{\theta}}^{\theta} Q_B(s)\,\mathrm{d}s$
4. $V_S(c) = V_s(\bar{c}) + \int_{c}^{\bar{c}} Q_S(s)\,\mathrm{d}s$

Mechanism Design – Bilateral Trade – proof of Myerson Satterthwaite

Since we are interested in ex-post efficient allocations: recall the following theorem:

With constants h_B and h_S, VCG implies the following transfer rules:

$$t_B(\theta, c) = \begin{cases} -c + h_B & \text{if } \theta > c \\ h_B & \text{otherwise} \end{cases} \quad \text{and} \quad t_S(\theta, c) = \begin{cases} \theta + h_S & \text{if } \theta > c \\ h_S & \text{otherwise.} \end{cases}$$

The (interim) expected utility of the buyer is then

$$V_B(\theta) = \int_{\underline{c}}^{\bar{c}} \Big((\theta - c)\, \mathbb{1}_{\{\theta > c\}} + h_B \Big)\, \mathrm{d}G(c) = \int_{\underline{c}}^{\bar{c}} (\theta - c)\, \mathbb{1}_{\{\theta > c\}}\, \mathrm{d}G(c) + h_B$$

Mechanism Design – Bilateral Trade – proof of Myerson Satterthwaite

Considering the ex-ante expected utility of the buyer

$$\int_{\underline{\theta}}^{\bar{\theta}} V_B(\theta)\,\mathrm{d}F(\theta) = \underbrace{\int_{\underline{\theta}}^{\bar{\theta}} \int_{\underline{c}}^{\bar{c}} (\theta - c)\,\mathbb{1}_{\{\theta > c\}}\,\mathrm{d}G(c)\,\mathrm{d}F(\theta)}_{=\text{ex-ante surplus from efficient trade } \equiv \mathcal{S}} + h_B = \mathcal{S} + h_B.$$

Same steps for the seller

$$\int_{\underline{c}}^{\bar{c}} V_S(c)\,\mathrm{d}G(c) = \mathcal{S} + h_S.$$

However, by Budget Balance (we don't inject money from outside) it must be that

$$\int_{\underline{\theta}}^{\bar{\theta}} V_B(\theta)\,\mathrm{d}F(\theta) + \int_{\underline{c}}^{\bar{c}} V_S(c)\,\mathrm{d}G(c) \leq \mathcal{S}.$$

Mechanism Design – Bilateral Trade – proof of Myerson Satterthwaite

Mechanism Design – Bilateral Trade – Recap

- Ex-post efficient trade is not feasible
- Note: what we showed implies that trade is ex-post inefficient in every equilibrium of any bargaining game with voluntary particiation

Mechanism Design – Revenue Maximisation

The auction problem:

- single indivisible object
- seller cost c
- n potential buyers with type θ_i
- utility $\theta_i q_i + t_i$
- types are independently distributed on $[\underline{\theta}_i, \bar{\theta}_i]$ according to F_i with density $f_i > 0$
- feasible allocation probabilities: $q_i(\theta) \in [0, 1]$ with $\sum_{i=1}^n q_i(\theta) \leq 1$

Seller commits to mechanism $(q, t)\colon \Theta \to [0, 1]^n \times \mathbb{R}^n$ to maximise revenue

Mechanism Design – Revenue Maximisation – Optimal Auctions

Let's compare different auction formats for the example $n = 2, \quad \theta_i \overset{iid}{\sim} U([0,1]), \quad c = 0$

1. **First-price auction**
   - Each bidder makes a bid $b = \beta(\theta_i)$. The highest bid wins. The winner pays his bid.
   - Find the symmetric equilibrium bid function β^* (hint: linear function)
2. **English auction (=ascending-clock auction)**
   - A price is publicly displayed.
     It increases continuously from $p_0 = 0$. Bidder i drops out when the price reaches $\rho(\theta_i)$.
     When i drops out (first), $j \neq i$ wins and pays ρ_i.
   - What is the weakly dominant stopping strategy $\rho(\theta_i)$
3. **All-pay auction (=contest)**
   - Each bidder makes a bid $b = \alpha(\theta_i)$. The highest bid wins. Each bidder pays his bid.
   - Find the symmetric equilibrium bid function α^*

All: **What is the expected revenue of the seller?**

Theorem

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

$$V_i(\theta_i) = V_i(\underline{\theta}_i) + \int_{\underline{\theta}_i}^{\theta_i} \int_{\Theta_{-i}} q_i(s, \theta_{-i}) \, \mathrm{d}F_{-i}(\theta_{-i}) \, \mathrm{d}s,$$

and expected transfer

$$-T_i(\theta_i) = \theta_i Q_i(\theta_i) - V_i(\theta_i).$$

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

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.

Were the three auctions optimal for the seller?

What constitutes an optimal auction?

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

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

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

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

Mechanism Design – Further Topics

Some issues that we have not covered:

- Collusion
- Interdependent valuations (for example, common-value auctions)
- Correlated types
- Evidence / Verification
- Dynamic problems (multiple stages)
- Limited commitment for principal