VCG Mechanism

Yonsei CS Theory Student Group

Presenter: Changyeol Lee (Yonsei Univ.)

Mechanism Design?

The *reverse engineering of games*

The *art of designing the rules of a game to achieve a specific desired outcome*

Game

- $N = \{1, \cdots, n\}$ is a set of *players*
- Θ_i is the set of all possible *types* (or *private information*) of player i.
  - Each player i has a type (private information) $\theta_i \in \Theta_i$.
  - $\boldsymbol{\theta} = (\theta_1, \cdots, \theta_n)$ is the *type profile*.
  - The set of all type profiles is given by $\boldsymbol{\Theta} = \Theta_1 \times \cdots \times \Theta_n$.
- S_i is the set of all possible *strategies* of player i.
  - Each player i decides his strategy $s_i: \Theta_i \to S_i$ according to his private information.
  - $\boldsymbol{s} = \left(s_1(\cdot), \cdots, s_n(\cdot)\right)$ is the *strategy profile*.
  - The set of all strategy profiles is given by $\boldsymbol{S} = S_1 \times \cdots \times S_n$.

Game

- Let Y be the set of all possible *outcomes*.
- $g: \boldsymbol{S} \to Y$ is an *outcome function*.
- $u_i: Y \times \Theta_i \to \mathbb{R}$ is the *utility function*.
  - We assume that u_i does not depend on the private information of the other players
  - WLOG, every player wants to maximize the utility

- Common knowledge: $N, \boldsymbol{\Theta}, \boldsymbol{S}, Y, g, \{u_i\}_{i=1,\cdots,n}$

There can be a prior knowledge on the distribution of $\boldsymbol{\Theta}$.

This is a (incomplete information) *strategic form game*.

Dominant Strategy

- Let $\boldsymbol{S}_{-i} = S_1 \times \cdots S_{i-1} \times S_{i+1} \times \cdots \times S_n$ and $\boldsymbol{s}_{-i} = (s_1, \cdots, s_{i-1}, s_{i+1}, \cdots, s_n)$. Similarly, let $\boldsymbol{\theta}_{-i} = (\theta_1, \cdots, \theta_{i-1}, \theta_{i+1}, \cdots, \theta_n)$
- A strategy $s_i^*(\cdot) \in S_i$ is a *dominant strategy for agent i* if, for every other strategy $s_i(\cdot) \in S_i$, for any type profile $\boldsymbol{\theta} = (\theta_i, \boldsymbol{\theta}_{-i}) \in \boldsymbol{\Theta}$, for all $\boldsymbol{s}_{-i} \in \boldsymbol{S}_{-i}$,

$$u_i\left(\underbrace{g\left(s_i^*(\theta_i), \boldsymbol{s}_{-i}(\boldsymbol{\theta}_{-i})\right)}_{\text{outcome of playing a dominant strategy}}, \theta_i\right) \geq u_i\left(\underbrace{g\left(s_i(\theta_i), \boldsymbol{s}_{-i}(\boldsymbol{\theta}_{-i})\right)}_{\text{outcome of playing another strategy}}, \theta_i\right).$$

sticking to a dominant strategy makes his utility greater than (or equal to)
that of doing another strategy no matter what other players do.

Dominant Strategy Equilibrium

- A strategy profile $\boldsymbol{s}^*(\cdot) = \left(s_1^*(\cdot), \cdots, s_n^*(\cdot)\right)$ is a *dominant strategy equilibrium* if, the strategy $s_i^*(\cdot)$ is a dominant strategy for agent i.
  - sometimes, we just say this strategy profile is a dominant strategy.

Mechanism Design Problem

- Setting: : $N, \mathbf{\Theta}, Y, \{u_i\}_{i=1,\cdots,n}$
- *Mechanism* $\mathcal{M} = (\boldsymbol{S} = (S_1, \cdots, S_n), g)$
- The *principal* (or designer) designs $\boldsymbol{S}$ and g.
- Each *agent* (or player) plays a game induced by the mechanism.
  - Recall. Each player i decides what to do, given $(N, \mathbf{\Theta}, \boldsymbol{S}, Y, g, \{u_i\})$ and his type θ_i.
- What properties do we want $\mathcal{M}$ to satisfy?

  Does it lead to the *desired* result?

Property 1

Dominant Strategy Incentive Compatible

Implementation in dominant strategies

- A <u>*social choice function*</u> (<u>SCF</u>) $f: \boldsymbol{\Theta} \to Y$ maps a type profile to an outcome.
- Mechanism $(\boldsymbol{S}, g)$ <u>*implements f in dominant strategies*</u> if there is a strategy profile $\boldsymbol{s}^*$ that is a dominant strategy equilibrium of the induced game s.t.

$$g\left(s_1^*(\theta_1), \cdots, s_n^*(\theta_n)\right) = f(\theta_1, \cdots, \theta_n)$$

for all type profile $(\theta_1, \cdots, \theta_n) \in \boldsymbol{\Theta}$.

Q. Consider a specific SCF. What are the characteristics of mechanism that can implement it in dominant strategies?

Direct Revelation Principle

- Difficult to consider infinite number of mechanisms
  - infinite ways to define strategies and the outcome function

- The *direct revelation principle* allows us to restrict our attention to *direct mechanisms*!

Direct Revelation Mechanism

- We say a mechanism $(\boldsymbol{S}, g)$ is a *direct (revelation) mechanism* if $\boldsymbol{S} = \boldsymbol{\Theta}$
  - The strategy is to directly report a candidate type in θ_i.
  - This implies that g is a SCF ($g\colon \boldsymbol{\Theta} \rightarrow Y$).

- **Direct revelation principle**. Let $\mathcal{M} = (\boldsymbol{S}, g)$ implements some f in dominant strategies. Then there exists a direct mechanism $\mathcal{M}'$ that implements f in dominant strategies.

Proof (by construction)

- Let $\boldsymbol{s}^*$ be a strategy profile that is a dominant strategy equilibrium in $\mathcal{M}$.
- We define $\mathcal{M}' = (\boldsymbol{\Theta}, g')$ as follow:
  - the strategy is to directly report one of θ_i; (Thus, strategy is no longer a function;)
  - $g' : \boldsymbol{\Theta} \rightarrow Y$ be the outcome function where $g'(\theta_1, \cdots, \theta_n) = g\big(s_1^*(\theta_1), \cdots, s_n^*(\theta_n)\big)$.
- $\mathcal{M}'$ implements f in dominant strategies since $\boldsymbol{\theta}$ is a dominant strategy equilibrium that satisfies $g'(\boldsymbol{\theta}) = g\big(\boldsymbol{s}^*(\boldsymbol{\theta})\big) = f(\boldsymbol{\theta})$.
- Note. $\mathcal{M}' = (\boldsymbol{\Theta}, f)$.

Dominant Strategy Incentive Compatible

- **Direct revelation principle**. Let $\mathcal{M} = (\boldsymbol{S}, g)$ implements some f in dominant strategies. Then there exists a direct mechanism $\mathcal{M}'$ that **truthfully implements** f in dominant strategies.

- We say a mechanism is ***dominant-strategy incentive-compatible*** (**DSIC**) if truth telling is a dominant strategy for every agent.

Is DSIC enough?

Does DSIC mechanism lead to a desired situation?

Gibbard-Satterthwaite (Impossibility) Theorem

Gibbard-Satterthwaite Theorem

- The SCF f is *dictatorial* if there is an agent i s.t. for all type profile $\boldsymbol{\theta} \in \boldsymbol{\Theta}$,

$$f(\boldsymbol{\theta}) \in \{y' \in Y : u_i(y', \theta_i) \geq u_i(y, \theta_i), \forall y \in Y\}$$

agent i prefers the outcome y'

the SCF always chooses the outcome that agent i prefers

- Simply, i is a dictator if $u_i(f(\boldsymbol{\theta}), \theta_i) \geq u_i(y, \theta_i)$ for all $\boldsymbol{\theta}$ and y.
- **Gibbard-Satterthwaite Theorem**. Suppose $|Y| \geq 3$, agents can have any preference, and f is an onto mapping. Then $(\boldsymbol{\Theta}, f)$ is DSIC if and only if f is dictatorial.

Restriction on Environment

Uni-dimensional type. Quasilinear utility.

Quasilinear Environment w/ WBB

- Consider we want to allocate some goods to agents.
- Let X be the set of possible allocations.
  - Assume $x \in X$ is a vector with n elements, describing the allocation to each agent.
- Each agent i's type θ_i is a one dimensional.
  - i.e., each agent i has a (private) valuation $\theta_i(x)$ for each outcome $x \in X$.
  - $\theta_i: X \to \mathbb{R}$
- Each agent transfers $p_i(\boldsymbol{\theta})$ amount of money to the mechanism.
  - Given an allocation and his type, his utility is defined as $u_i(x, \theta_i) = \theta_i(x) - p_i(\boldsymbol{\theta})$.
  - $\sum_i p_i(\boldsymbol{\theta}') \geq 0$, i.e., *weakly budget balanced* (WBB)

quasilinear utility

Mechanism Perspective

- Setting: : $N, \{\Theta_i\}_{i=1,\cdots,n}, \left(X, \{P_i\}_{i=1,\cdots,n}\right), \{u_i\}_{i=1,\cdots,n}$
- Mechanism $\mathcal{M} = \left(\Theta_1, \cdots, \Theta_n, f = \left(x, \{p_i\}_{i=1,\cdots,n}\right)\right)$
- Mechanism announces f, i.e., how to allocate and how much to transfer.
- Each agents i reports θ_i'. Let $\boldsymbol{\theta}' = (\theta_1', \cdots, \theta_n')$ be the reported type profile.
- Mechanism do the allocation $x(\boldsymbol{\theta}') \in X$
- Each agent i pays $p_i(\boldsymbol{\theta}')$ to the mechanism.
  - Again, we need to satisfy WBB, i.e., $\sum_i p_i(\boldsymbol{\theta}') \geq 0$.

No dictator in quasilinear setting w/ WBB

- Fix any mechanism and let x^* be the allocation rule and the payment rule p^*.
- Suppose d is a dictator. Then for all $\boldsymbol{\theta} \in \boldsymbol{\Theta}$ and for all possible x and p,

$$\theta_d\big(x^*(\boldsymbol{\theta})\big) - p_d^*(\boldsymbol{\theta}) \geq \theta_d\big(x(\boldsymbol{\theta})\big) - p_d(\boldsymbol{\theta}).$$

Consider the following payment rule p'.

- Case $\sum_i p_i^*(\boldsymbol{\theta}) > 0$.

$$p_i'(\boldsymbol{\theta}) \coloneqq p_i^*(\boldsymbol{\theta})\ if\ i \neq d, \qquad p_d'(\boldsymbol{\theta}) \coloneqq p_d^*(\boldsymbol{\theta}) - \textstyle\sum_i p_i^*(\boldsymbol{\theta})$$

With (x^*, p'), we have $\theta_d\big(x^*(\boldsymbol{\theta})\big) - p_d'(\boldsymbol{\theta}) = \theta_d\big(x^*(\boldsymbol{\theta})\big) - p_d^*(\boldsymbol{\theta}) + \sum_i p_i^*(\boldsymbol{\theta})$

No dictator in quasilinear setting w/ WBB

- Fix any mechanism and let x^* be the allocation rule and the payment rule p^*.
- Suppose d is a dictator. Then for all $\boldsymbol{\theta} \in \boldsymbol{\Theta}$ and for all possible x and p,

$$\underline{\theta_d\big(x^*(\boldsymbol{\theta})\big) - p_d^*(\boldsymbol{\theta})} \geq \theta_d\big(x(\boldsymbol{\theta})\big) - p_d(\boldsymbol{\theta}).$$

Consider the following payment rule p'.

- Case $\sum_i p_i^*(\boldsymbol{\theta}) = 0$. For any $j \neq d$,

$$p_i'(\boldsymbol{\theta}) \coloneqq p_i^*(\boldsymbol{\theta}) \; if \; i \neq d, j, \quad p_d'(\boldsymbol{\theta}) \coloneqq p_d^*(\boldsymbol{\theta}) + \epsilon, \;\; p_j'(\boldsymbol{\theta}) \coloneqq p_j^*(\boldsymbol{\theta}) - \epsilon$$

With (x^*, p'), we have $\theta_d\big(x^*(\boldsymbol{\theta})\big) - p_d'(\boldsymbol{\theta}) = \underline{\theta_d\big(x^*(\boldsymbol{\theta})\big) - p_d^*(\boldsymbol{\theta})} + \epsilon$

Desired Situation

Social Welfare Maximization

Efficient (welfare-maximizing)

- We say a mechanism is ***efficient*** if it maximizes the welfare of the society.
  - In some literature, they say "social optimal"
  - We say a mechanism is *optimal* if it maximizes the revenue of the principal.

- Under this "quasilinear" environment,
  is there an efficient DSIC mechanism?
- **YES!**

Efficient, DSIC Mechanism?

Vickrey-Clarke-Groves (VCG) Mechanism

VCG Mechanism

- Part 1. Define an allocation rule
- Given reported type profile $\boldsymbol{\theta}'$, let x^* be the allocation rule defined by

$$x^*(\boldsymbol{\theta}') \in \underset{x \in X}{\operatorname{argmax}} \sum_{i=1}^{n} \theta_i'\big(x(\boldsymbol{\theta}')\big).$$

- If all agents report truthfully, this allocation makes the mechanism efficient.
  - We say this allocation is efficient.
  - More specifically, ex-post efficient.

VCG Mechanism

- Part 2. Define a payment rule
- Given reported type profile $\boldsymbol{\theta}'$, let p_i be the payment rule defined by

$$p_i(\boldsymbol{\theta}') = \max_{x \in X} \sum_{j \neq i} \theta'_j\big(x(\boldsymbol{\theta}'_{-i})\big) - \sum_{j \neq i} \theta'_j\big(x^*(\boldsymbol{\theta}')\big).$$

Total welfare in efficient allocation of θ'_{-i} (assuming agent i was not included at the first place)

Total welfare except for the welfare that agent i makes in the efficient allocation

- This payment is the "externality" caused by an agent i.

VCG Mechanism

- Part 3. Truth telling is a dominant strategy for every agent.
- Suppose not. Fix i. This implies that θ_i is not a dominant strategy.
- There exists a strategy profile $\boldsymbol{\theta}'$ such that, given $\boldsymbol{\theta}'_{-i}$, it is better to report θ_i' instead of θ_i, i.e.,

$$u_i(x^*(\theta_i, \boldsymbol{\theta}'_{-i}), \theta_i) < u_i(x^*(\theta_i', \boldsymbol{\theta}'_{-i}), \theta_i)$$

$$\theta_i\big(x^*(\theta_i, \boldsymbol{\theta}'_{-i})\big) - p_i(\theta_i, \boldsymbol{\theta}'_{-i}) < \theta_i\big(x^*(\theta_i', \boldsymbol{\theta}'_{-i})\big) - p_i(\theta_i', \boldsymbol{\theta}'_{-i})$$

VCG Mechanism

$$\theta_i\left(x^*(\theta_i, \boldsymbol{\theta}'_{-i})\right) - p_i(\theta_i, \boldsymbol{\theta}'_{-i}) < \theta_i\left(x^*(\theta'_i, \boldsymbol{\theta}'_{-i})\right) - p_i(\theta'_i, \boldsymbol{\theta}'_{-i})$$

- Recall $p_i(\boldsymbol{\theta}') = \underline{\max_{x \in X} \sum_{j \neq i} \theta'_j\left(x(\boldsymbol{\theta}'_{-i})\right)} - \sum_{j \neq i} \theta'_j\left(x^*(\boldsymbol{\theta}')\right)$.
- We substitute $p_i(\theta_i, \boldsymbol{\theta}'_{-i})$ and $p_i(\theta'_i, \boldsymbol{\theta}'_{-i})$.
- Note that $\underline{\max_{x \in X} \sum_{j \neq i} \theta'_j\left(x(\boldsymbol{\theta}'_{-i})\right)}$ does not depend on agent i's report.

VCG Mechanism

$$\theta_i\left(x^*(\theta_i, \boldsymbol{\theta}'_{-i})\right) + \sum_{j \neq i} \theta'_j\left(x^*(\theta_i, \boldsymbol{\theta}'_{-i})\right)$$

Total welfare in the **efficient** allocation of (θ_i, θ'_{-i})

$$< \theta_i\left(x^*(\theta'_i, \boldsymbol{\theta}'_{-i})\right) + \sum_{j \neq i} \theta'_j\left(x^*(\theta'_i, \boldsymbol{\theta}'_{-i})\right)$$

Total welfare in **some** allocation of (θ_i, θ'_{-i}) since

- Recall $x^*(\boldsymbol{\theta}') \in argmax_{x \in X} \sum_{i=1}^{n} \theta'_i\left(x(\boldsymbol{\theta}')\right)$.

$$x^*(\theta_i, \boldsymbol{\theta}'_{-i}) \in argmax_{x \in X} \left[\theta_i\left(x(\theta_i, \boldsymbol{\theta}'_{-i})\right) + \sum_{j \neq i} \theta'_j\left(x(\theta_i, \boldsymbol{\theta}'_{-i})\right)\right]$$

Payment in VCG Mechanism

Clarke Mechanism

or Pivot Mechanism

$$p_i(\boldsymbol{\theta}') = \max_{x \in X} \sum_{j \neq i} \theta_j'\big(x(\boldsymbol{\theta}_{-i}')\big) - \sum_{j \neq i} \theta_j'\big(x^*(\boldsymbol{\theta}')\big).$$

- Assuming (natural assumptions), $p_i \geq 0$, which implies no subsidy needed.
- Note. The first term can be replaced by some function $h_i(\boldsymbol{\theta}_{-i}')$ that does not depend on i's type.

Groves Mechanism

$$p_i(\boldsymbol{\theta}') = h_i(\boldsymbol{\theta}_{-i}') - \sum_{j \neq i} \theta_j'\big(x^*(\boldsymbol{\theta}')\big)$$

Example

Auctions, Bilateral Trade

Auction Design

- Setting
  - Single item
  - Utility of each bidder: $\theta_i - \theta_i'$ if wins, and 0 otherwise.
  - Bidders can bid any amount of money.
    - but no reason to bid higher than his true valuation (type).
- How to design an auction so that it is efficient DSIC?
  - Recall, we only need to consider a direct mechanism.
  - Given bids, how should we allocate the item and how much should bidders pay?

Auction Design

- Allocation?
  - To the highest bidder. (Efficient)
- Payment?
- Let i be the winner, i.e., a bidder with highest bid.
  - The total welfare in an efficient allocation of $\boldsymbol{\theta}'_{-i}$ = the second highest bid.
  - The total welfare except for i in an efficient allocation of θ' = 0.
- Let i be a loser. He pays 0.
- Payment: The winner pays the second highest bid; other pays 0.

Auction Design

- This is a special case of Clarke mechanism and thus efficient and DSIC.
  - and WBB also.
- This is called Vickrey auction or second-price (sealed-bid auction).

First-Price Auction

- To whom do I allocate? To the highest bidder.
- Payment? Winner pays the **highest** bid.

- This is not VCG mechanism.
  - The payment here does not fit VCG payment function.
  - This does not mean that there is no efficient DSIC mechanism.
- Indeed, there is no DSIC mechanism for this auction.

Combinatorial Auction

- setting:
  - m non-identical items
  - each agent has 2^m private valuation
- Define efficient allocation and Clarke payment
- This auction is efficient and DSIC and WBB.

Bilateral Trade

- Setting:
  - one item, buyer and seller (two agents), utility same as before
  - private valuation $\theta_b, \theta_s \in [0,1]$, respectively.
- Allocation:
  - if $\theta_b' > \theta_s'$, give the item to the buyer; otherwise give the item to the seller
- Payment:
  - $p_b(\theta_b', \theta_s') = \theta_s' \cdot \mathbb{I}[\theta_b' > \theta_s'], p_s(\theta_b', \theta_s') = \theta_b' \cdot \mathbb{I}[\theta_s' \geq \theta_b']$

Seller pays to keep the good.

No reason to join the mechanism!

More about VCG

Individual Rationality

Individual Rationality (IR)

- Let $u_i(\theta_i)$ be the utility of withdrawing from the mechanism.
  - In Vickrey auction, $u_i(\theta_i) = 0$
  - In bilateral trade, for the seller, $u_s(\theta_s) = \theta_s$

- A mechanism is ex-post individual rational if $u_i(f(\boldsymbol{\theta}), \theta_i) \geq u_i(\theta_i)$ for all i and for all $\boldsymbol{\theta} = (\theta_i, \boldsymbol{\theta}_{-i})$.

Clarke Mechanism

- Under some natural assumptions and $u_i(\cdot) = 0$,
  the Clarke Mechanism is ex-post individual rational.
- Combinatorial Auction corresponds to this mechanism.

- Clarke Mechanism is ex-post efficient, DSIC, WBB and ex-post IR.

Limitations

Limitation

- Some games have no dominant strategy equilibrium.
  We need "Bayesian" equilibrium.
  - Rock & Scissor & Paper, First-price auction
- (Very) Restricted Environment
- Computational Limitation.
  - Unless $P = NP$, no polynomial time algorithm finds an allocation/payment rule.
  - Also, for agents, even hard to list all the valuations to the mechanism.
- Different objectives
  - optimal (revenue-maximizing) mechanism

Thank you for your attention

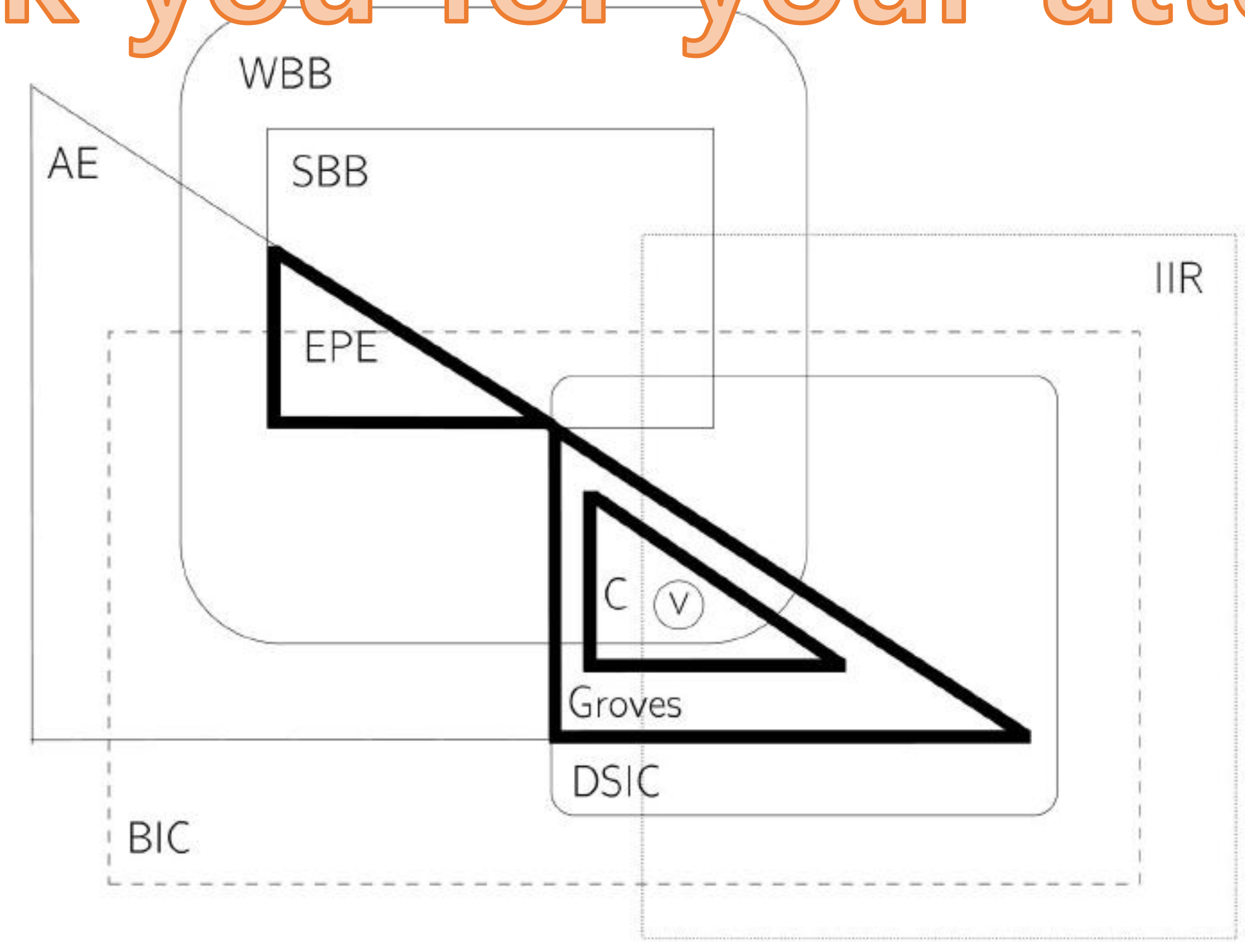

AE : Allocative Efficient
SBB : Strict Budget Balanced
WBB : Weak Budget Balanced
BIC : Bayesian Incentive Compatible
DSIC : Dominant Strategy Incentive Compatible
EPE : Ex-Post Efficient
IIR : Interim Individually Rational
C : Clarke Mechanism
V : Vickrey Auction

Fig. 19.2: Mechanism design space in quasilinear environment

Textbook of Y. Narahari.
"Game Theory and Mechanism Design"

사진찍기..?