Mechanism Design

Felix Munoz-Garcia

EconS 503 - Advanced Microeconomics II
Washington State University

Mechanism Design

- There are many situations in which some central authority wishes to implement a decision that depends on the private information of a set of players.
  - Government may wish to choose the design of a public-works project based on preferences of its citizens who have private information about their preferences.
  - Monopolistic firms may wish to determine a set of consumers' willingness to pay for different products it can produce with the goal of making as high a profit as possible.
- **Mechanism design** is a study of what kinds of mechanisms that the central authority can devise in order to reveal the private information of players.
- Central authority is mechanism designer.

Cake Cutting Problem

- Consider a mother of two children, who has to design a mechanism to make her kids share a cake equally. The mother is the central authority in this case.
- If the mother slices the cake into two equal pieces and distributes one piece to each of the kids, the solution is not necessarily acceptable to the kids because each kid will be left with the perception that they got the smaller of the two pieces.
- Can mom design a mechanism to make everyone happy?
  - Of course she can!

Cake Cutting Problem

- Consider the following mechanism:
  - Stage 1: One of the kids slices the cake into two pieces.
  - Stage 2: The other kid gets to choose which piece they want, with the leftover piece going to the first (slicing) child.

Cake Cutting Problem

- Child 1, who sliced the cake, will slice it exactly into two equal halves (in their eyes), as he knows that any other division will leave him with the smaller piece.
- Child 2 is happy because they got to choose the bigger piece (in their eyes).
- Thus, this mechanism implements the desirable outcome of the kids sharing the cake equally and furthermore, each kid has every reason to be happy about this mechanism.
  - At least until it is time to do the dishes.

Set up: Mechanism as Bayesian Games

- A set of players $N = \{1, 2, \ldots, n\}$.
- A set of public alternatives X that could represent many kinds of alternatives.
  - e.g., an alternative $x \in X$ could represent the attributes of a public good or service, like investment in education or in preserving the environment.
- The reason that X is called as public alternatives is the chosen alternative affects all the players in N,
  - e.g., in an auction, if one player gets a private good then the consequence is that everyone else does not.

Environment set up — Players

- Each player i privately observes his type $\theta_i \in \Theta_i$ which determines his preferences.
- Let $\theta = (\theta_1, \theta_2, \ldots, \theta_n)$ be the state of the world.
- State θ is drawn randomly from the state space $\Theta \equiv \Theta_1 \times \Theta_2 \times \cdots \times \Theta_n$.
- The draw of θ is according to some prior distribution $\phi(\cdot)$ over Θ.
- θ_i is player i's private information; $\phi(\cdot)$ is common knowledge.

Environment set up — Players

- Each player i has quasilinear preference:
  $v_i(x, m, \theta_i) = u_i(x, \theta_i) + m_i$
  - Alternatives have a "money-equivalent" value, and preferences are additive in money.
- m_i is the amount of money that is given to individual i. m_i can be negative meaning money is taken away from individual i.
- $u_i(x, \theta_i)$ is money-equivalent value of alternative $x \in X$ when i's type is θ_i.
- An outcome would be represented as $y = (x, m_1, \ldots, m_2)$.

Environment set up — Mechanism Designer

- The mechanism designer has the objective of achieving an outcome that depends on the types of players
- Assume that mechanism designer does not have a source of funds to pay the players.
  - Monetary payments have to be self-financed, which is $\sum_{i=1}^{n} m_i \leq 0$
- When $\sum_{i=1}^{n} m_i < 0$, it means that mechanism designer keeps some of the money that he raises from players.
- The set of outcomes is restricted as follows:

$$Y = \left\{ (x, m_1, \ldots, m_2) : x \in X, m_i \in \mathbb{R} \forall i \in N, \sum_{i=1}^{n} m_i \leq 0 \right\}$$

Environment set up — Mechanism Designer

- The mechanism designer's objective is given by a choice rule:

$$f(\theta) = (x(\theta), m_1(\theta), \ldots, m_n(\theta)),$$

  - where $x(\theta) \in X$ and $\sum_{i=1}^{n} m_i \leq 0$.
- $x(\theta)$ is the decision rule; $(m_1(\theta), \ldots, m_n(\theta))$ is the transfer rule.

Environment set up — Mechanism Designer

Example 1:

- Let $X = [0, \bar{x}]$ be the size of a water treatment plant.
- The plant will benefit some citizens and may displease others.
- The citizens are the group of players, N.
- Player i's willingness to pay from $x \in X$ of the plant is $u_i(x, \theta_i)$
- The utilitarian mechanism designer maximizes the sum of the players' valuations by choosing the value of x.
- So his decision rule $x(\theta)$ would maximize $\sum_{i=1}^{n} u_i(x, \theta_i)$.

Environment set up

Example 2:

- A good: a license to use a certain portion of the electromagnetic spectrum for cell coverage.
- The license can be allocated to one of a group of cellular carriers $i \in N$.
- $x_i \in \{0, 1\}$ indicates whether player i receives the license $(x_i = 1)$ or not $(x_i = 0)$.
- The possible set of alternatives is

$$X = \{(x_1, \ldots, x_n)\}$$

  - such that $x_i \in \{0, 1\}$ and $\sum_{i=1}^{n} x_i = 1$
- Player i's willingness to pay for the license is $u_i(x, \theta_i) = \theta_i x_i$.
- The mechanism designer maximizes the sum of the players' valuations by choosing x.
- So his decision rule $x(\theta)$ would maximize $\sum_{i=1}^{n} u_i(x, \theta_i)$.

The Mechanism Game

- The mechanism designer desires to implement a choice rule $f : \Theta \rightarrow Y$.
- The problem is that the mechanism designer's choice rule depends on the unobserved state Θ.
- Two ways that the mechanism designer have to solve the problem.
- The first method is to ask each player directly. But are they willing to share their true preference?

The Mechanism Game

- Example 3 (MWG 23.B.1): Consider an abstract case, where we are given a set of alternatives $X = \{x, y, z\}$ and two players.
- Suppose that agent 1 has one possible type, so that $\Theta_1 = \{\bar{\theta}_1\}$ and that agent 2 has two possible types, so that $\Theta_2 = \{\theta_2', \theta_2''\}$.
- The agents' possible preference orderings are given as

	$\bar{\theta}_1$	θ_2'	θ_2''
Best	x	z	y
Middle	y	y	x
Worst	z	x	z

The Mechanism Game

- Suppose that the agents wish to implement the ex post efficient social choice function $f(\cdot)$ with

$$f(\bar{\theta}_1, \theta_2') = y \qquad \text{and} \qquad f(\bar{\theta}_1, \theta_2'') = x$$

- If so, then agent 2 must be relied upon to truthfully reveal his preferences.
- It is apparent, however, that he will not find it in his interest to do so. When $\theta_2 = \theta_2''$, agent 2 will wish to lie and claim that his type is θ_2'.
  - In the words of a famous American TV Doctor, "Everyone lies."
  - Let's look at the alternative way.

The Mechanism Game

- In the second way, the mechanism designer designs some clever game that ends up revealing the players' private information..
- The rules of the game endows each player i with an action set A_i.
- Following the choice $a_i \in A_i$ by each player, there is some outcome function $g\left(a_1, \ldots, a_n\right)$.
- That makes a choice of an outcome $y \in Y$.
- The payoffs of player i over outcomes is $v_i\left(g\left(s\right), \theta_i\right)$.

The Mechanism Game

Definition

A **mechanism**, $\Gamma = \{A_1, A_2, \ldots, A_n, g(\cdot)\}$ is a collection of n action sets $A_1, A_2, \ldots, A_n$ and an outcome function $g : A_1 \times A_2 \times \cdots \times A_n \longrightarrow Y$. A pure strategy for player i in the mechanism Γ is a function that maps types into actions, $s_i : \Theta_i \longrightarrow A_i$. The payoffs of the players are given by $v_i(g(s), \theta_i)$.

The Mechanism Game

Definition

The strategy profile $s^*(\cdot) = (s_1^*(\cdot), \ldots, s_n^*(\cdot))$ is a **Bayesian Nash equilibrium** of the mechanism $\Gamma = \{A_1, \ldots, A_n, g(\cdot)\}$ if for every $i \in N$ and for every $\theta_i \in \Theta_i$

$$\begin{aligned} & E_{\theta_{-i}}\left[v_i\left(g\left(s_i^*(\theta_i), s_{-i}^*(\theta_{-i})\right), \theta_i\right) | \theta_i\right] \\ \geq \; & E_{\theta_{-i}}\left[v_i\left(g\left(a_i', s_{-i}^*(\theta_{-i})\right), \theta_i\right) | \theta_i\right] \quad \text{for all } a_i' \in A_i \end{aligned}$$

- That is, if player i believes that other players are playing according to $s_{-i}^*(\theta)$ then he maximizes his expected payoff by following the behavior prescribed by $s_i^*(\theta_i)$ regardless of which type player i is

The Mechanism Game

- The mechanism designer designs a mechanism in which $s_i^* \longrightarrow A_i$ such that the outcome is exactly what the mechanism designer desires given each θ_i.
- for all $\theta \in \Theta, g\left(s_1^*\left(\theta_1\right), s_2^*\left(\theta_2\right), \ldots, s_n^*\left(\theta_n\right)\right) = f\left(\theta\right)$

Definition

A mechanism Γ **implements** the choice rule $f\left(\cdot\right)$ if there exists a Bayesian Nash equilibrium of the mechanism Γ, $\left(s_1^*\left(\theta_1\right), s_2^*\left(\theta_2\right), \ldots, s_n^*\left(\theta_n\right)\right)$, such that $g\left(s_1^*\left(\theta_1\right), s_2^*\left(\theta_2\right), \ldots, s_n^*\left(\theta_n\right)\right) = f\left(\theta\right)$ for all $\theta \in \Theta$.

The Mechanism Game

- That is, it instead implements $f(\cdot)$ after knowing the true θ, the mechanism does what the mechanism designer wants to do: $g(a(\theta)) = f(\theta)$.
- It is a partial implementation because it requires that the desired outcome be an equilibrium, but allows for other, undersirable equilibrium outcomes as well
- The implementation without "bad equilibria" is called full implementation.

The Mechanism Game

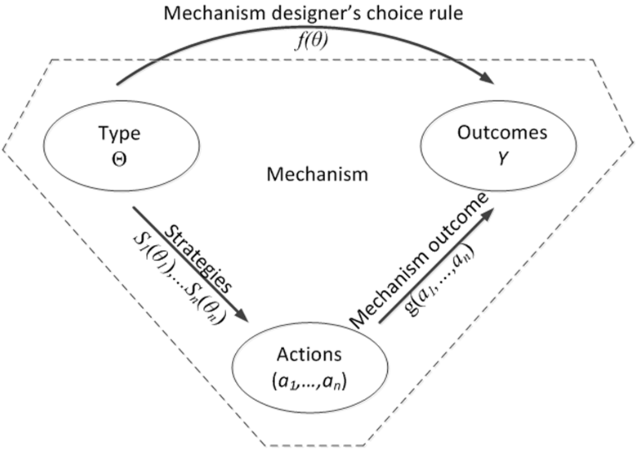

The Revelation Principle

- The mechanism game is a Bayesian game.
- It is useful when the mechanism designer cannot get players to reveal their types.
- There is a particular mechanism which is also a Bayesian game in which the mechanism designer asks players directly to reveal their types in order to implement $f(\cdot)$.
- The mechanism designer implements $f(\hat{\theta})$, with $\hat{\theta}$ is announced by the players.

Definition

$\Gamma = \{\Theta_1, \ldots, \Theta_n, f(\cdot)\}$ is a **direct revelation mechanism** for choice rule $f(\cdot)$ if $A_i = \Theta_i$ for all $i \in N$ and $g(\theta) = f(\theta)$ for all $\theta \in \Theta$.

The Revelation Principle

- The straightforward direct revelation mechanism will actually have an equilibrium that implements the mechanism designer's intended outcome.

Definition

The choice rule $f(\cdot)$ is **truthfully implementable in Bayesian Nash equilibrium** if for all θ the direct revelation mechanism $\Gamma = \{\Theta_1, \ldots, \Theta_I, f(\cdot)\}$ has a Bayesian Nash equilibrium $s_i^*(\theta_i) = \theta_i$ for all i. Equivalently, for all i,

$$E_{\theta_{-i}}\left[v_i\left(f\left(\theta_i, \theta_{-i}\right), \theta_i\right) | \theta_i\right] \geq E_{\theta_{-i}}\left[v_i\left(f\left(\hat{\theta}_i, \theta_{-i}\right), \theta_i\right) | \theta_i\right]$$
$$\text{for all } \hat{\theta}_i \in \Theta_i.$$

The Revelation Principle

- That is, $f(\cdot)$ is truthfully implementable in Bayesian Nash equilibrium if truthtelling is a Bayesian Nash equilibrium strategy in the direct revelation mechanism.
- If every player i believes that all other players are reporting their types truthfully, then player i is also willing to report truthfully.

The Revelation Principle

- Example 4 (MWG 23.B.7): Consider a first-price sealed-bid auction where two potential buyers have valuations θ_i that are drawn from a uniform distribution on $[0, 1]$.
  - Recall that the equilibrium bidding function for each player i is $b_i(\theta_i) = \frac{1}{2}\theta_i$.
- When facing the direct revelation mechanism $\Gamma = \{\Theta_1, ..., \Theta_I, f(\cdot)\}$, buyer 1's optimal announcement $\hat{\theta}_1$ when he has type θ_1 solves

$$\max_{\hat{\theta}_1} \left(\theta_1 - \frac{1}{2}\hat{\theta}_1\right) \Pr(\theta_2 \leq \hat{\theta}_1)$$

$$= \max_{\hat{\theta}_1} \left(\theta_1 - \frac{1}{2}\hat{\theta}_1\right) \hat{\theta}_1$$

The Revelation Principle

- The first-order condition is

$$\theta_1 - \hat{\theta}_1 = 0 \implies \hat{\theta}_1 = \theta_1$$

- Hence, truth telling is buyer 1's optimal strategy given that buyer 2 always tells the truth.
  - A similar conclusion follows for buyer 2.
- Thus, the social choice function implemented by the first-price sealed-bid auction (in a Bayesian Nash equilibrium) can also be truthfully implemented (in a Bayesian Nash equilibrium) through a direct revelation mechanism!

The Revelation Principle

Proposition: (The Revelation Principle for Bayesian Nash Implementation) A choice rule $f(\cdot)$ is implementable in Bayesian Nash equilibrium if and only if it is truthfully implementable in Bayesian Nash equilibrium.

The Revelation Principle

Proof:

- **IF part:** By definition, if $f(\cdot)$ is truthfully implementable in Bayesian Nash equilibrium then it is implementable in Bayesian Nash equilibrium using the direct revelation mechanism.

- **ONLY IF part:** Suppose that there exists some mechanism $\Gamma = (A_1, \ldots, A_n, g(\cdot))$ that implements $f(\cdot)$ using the equilibrium strategy profile $s^*(\cdot) = (s_1^*(\cdot), \ldots, s_n^*(\cdot))$ and $g(s^*(\cdot)) = f(\cdot)$, so that for every $i \in N$ and $\theta_i \in \Theta_i$,

$$\begin{aligned} & E_{\theta_{-i}}\left[v_i\left(g\left(s_i^*(\theta_i), s_{-i}^*(\theta_{-i})\right), \theta_i\right) | \theta_i\right] \\ \geq \quad & E_{\theta_{-i}}\left[v_i\left(g\left(a_i', s_{-i}^*(\theta_{-i})\right), \theta_i\right) | \theta_i\right] \quad \text{for all } a_i' \in A_i \end{aligned}$$

which means that no player i wishes to deviate from $s_i^*(\cdot)$.

The Revelation Principle

- However, when player i is asked his type, if he pretends that his type is $\hat{\theta}_i$ rather than θ_i, then $a_i' = s_i^*\left(\hat{\theta}_i\right)$.

Thus

$$
\begin{aligned}
& E_{\theta_{-i}}\left[v_i\left(g\left(s_i^*\left(\theta_i\right), s_{-i}^*\left(\theta_{-i}\right)\right), \theta_i\right)|\theta_i\right] \\
\geq \quad & E_{\theta_{-i}}\left[v_i\left(g\left(s_i^*\left(\hat{\theta}_i\right)\right), s_{-i}^*\left(\theta_{-i}\right), \theta_i\right)|\theta_i\right] \quad \text{for every } \hat{\theta}_i \in \Theta_i
\end{aligned}
$$

Because $g\left(s^*\left(\theta\right)\right) = f\left(\theta\right)$ for all $\theta \in \Theta$,

$$
\begin{aligned}
E_{\theta_{-i}}\left[v_i\left(f\left(\theta_i, \theta_{-i}\right), \theta_i\right)|\theta_i\right] \quad & \geq \quad E_{\theta_{-i}}\left[v_i\left(f\left(\hat{\theta}_i, \theta_{-i}\right), \theta_i\right)|\theta_i\right] \\
\text{for every } \hat{\theta}_i \quad & \in \quad \Theta_i
\end{aligned}
$$

- This is just the condition for $f\left(\cdot\right)$ to be truthfully implementable in Bayesian Nash equilibrium.

The Revelation Principle

- If the mechanism designer cannot implement $f(\cdot)$ directly then there is no mechanism in the world that can.
- The designed mechanism and direct revelation mechanism are equivalent.
  - In equilibrium the players know that the mechanism implements $f(\cdot)$, and they choose to stick to it.
  - So they announce their types truthfully and have the mechanism designer implement $f(\cdot)$ directly.

Dominant Strategies Implementation

Definition

The strategy profile $s^*(\cdot) = (s_1^*(\cdot), \ldots, s_n^*(\cdot))$ is a **dominant strategy equilibrium** of the mechanism $\Gamma = \{A_1, A_2, \ldots, A_n, g(\cdot)\}$ if for every $i \in N$ and for every $\theta_i \in \Theta_i$

$$\begin{aligned} v_i\left(g\left(s_i^*(\theta), a_{-i}\right), \theta\right) &\geq v_i\left(g\left(a_i', a_{-i}\right), \theta\right) \\ \text{for all } a_i' &\in A_i \text{ and for all } a_{-i} \in A_i \end{aligned}$$

- Is there a mechanism Γ that implements $f(\cdot)$ in dominant strategies?

Dominant Strategies Implementation

- Since a dominant strategy equilibrium is a special case of a Bayesian equilibrium, the revelation principle applies.
- So we only check that $f(\cdot)$ is implementable in dominant strategies directly to see if $f(\cdot)$ is implementable in dominant strategies. That is

$$\begin{aligned} v_i\left(f\left(\theta_i, \theta_{-i}\right), \theta_i\right) &\geq v_i\left(f\left(\hat{\theta}_i, \theta_{-i}\right), \theta_i\right) \\ \text{for all } \hat{\theta}_i &\in \Theta_i, \text{ and for all } \theta_{-i} \in \Theta_{-i} \end{aligned}$$

Vickrey-Clarke-Groves Mechanism

- Recall that our quasilinear preferences are additive in money, $v_i(x, m_i, \theta_i) = u_i(x, \theta_i) + m_i$.
- There is a nice feature of this quasilinear environment: Monetary transformation can benefit the whole group.
- Imagine player i with θ_i and player j with θ_j such that $u_i(x', \theta_i) > u_i(x, \theta_i)$, $u_j(x, \theta_i) > u_j(x', \theta_j)$
  - Intuitively, player i would prefer alternative x' and player j would prefer alternative x.
- Furthermore, let

$$u_i\left(x', \theta_i\right) - u_i\left(x, \theta_i\right) > u_j\left(x, \theta_i\right) - u_j\left(x', \theta_j\right)$$

which implies that the gains received by player i when implementing alternative x' are greater than the gains received by player j when implementing alternative x.

Vickrey-Clarke-Groves Mechanism

- There is any amount of money $k > 0$, satisfying

$$u_i\left(x', \theta_i\right) - u_i\left(x, \theta_i\right) > k > u_j\left(x, \theta_i\right) - u_j\left(x', \theta_j\right).$$

- So both players will better off if we replace x with x' and transfer k from player i to player j.

Vickrey-Clarke-Groves Mechanism

Proposition: In the quasilinear environment, given a state of the world $\theta \in \Theta$, an alternative $x^* \in X$ is Pareto optimal if and only if it is a solution to

$$\max_{x \in X} \sum_{i=1}^{I} u_i \left(x, \theta_i \right).$$

Proof: If an alternative *a* did not maximize this sum, then there was another x' that did. Then money transfers among players that would ensure the gains of some players more than compensate for the losses of others.

Vickrey-Clarke-Groves Mechanism

Definition

We call a decision rule $x^*(\cdot)$ the **first-best decision rule** if for all $\theta \in \Theta$, $x^*(\theta)$ is Pareto optimal.

$$x^*(\theta) \in \arg\max_{x \in X} \sum_{i=1}^{I} u_i(x, \theta_i) \quad \forall \theta \in \Theta.$$

- When faced with the Pareto optimal choice rule $(x^*(\cdot), m_1(\cdot), \ldots, m_n(\cdot))$, will truth-telling be a dominant strategy for each player in the direct revelation mechanism?
- No, when $m_i(\hat{\theta}_i, \hat{\theta}_{-i}) \equiv 0$ $(\hat{\theta}_i$ is announced by player $i)$. The reason is that each player i only maximizes his own payoff, not the total surplus.
- This problem could be solved by having a clever transfer rule $m_i(\hat{\theta}_i, \hat{\theta}_{-i})$ to let player internalize the externality.

Vickrey-Clarke-Groves Mechanism

Definition

Given announcements $\hat{\theta}$, the choice rule $f\left(\hat{\theta}\right) = \left(x^*\left(\hat{\theta}\right), m_1\left(\hat{\theta}\right), \ldots, m_n\left(\hat{\theta}\right)\right)$ is a **Vickrey-Clarke-Groves (VCG) mechanism** if $x^*(\cdot)$ is the first-best decision rule and if for all $i \in N$

$$m_i\left(\hat{\theta}\right) = \sum_{j \neq i} u_j\left(x^*\left(\hat{\theta}_i, \hat{\theta}_{-i}\right), \hat{\theta}_j\right) + h_i\left(\hat{\theta}_{-i}\right)$$

where $h_i\left(\hat{\theta}_{-i}\right)$ is an arbitrary function of $\hat{\theta}_{-i}$.

Vickrey-Clarke-Groves Mechanism

Proposition: Any VCG mechanism is truthfully implementable in dominant strategies.

- In the VCG mechanism every player i solves

$$
\begin{aligned}
& \max_{\hat{\theta}_i \in \Theta_i} u_i\left(x^*\left(\hat{\theta}_i, \hat{\theta}_{-i}\right), \theta_i\right) + m_i\left(\hat{\theta}_i, \hat{\theta}_{-i}\right) \\
= \; & \max_{\hat{\theta}_i \in \Theta_i} \underbrace{u_i\left(x^*\left(\hat{\theta}_i, \hat{\theta}_{-i}\right), \theta_i\right) + \sum_{j \neq i} u_j\left(x^*\left(\hat{\theta}_i, \hat{\theta}_{-i}\right), \hat{\theta}_j\right)}_{\text{total surplus}} + h_i\left(\hat{\theta}_{-i}\right)
\end{aligned}
$$

- $h_i\left(\hat{\theta}_{-i}\right)$ does not affect $i's$ choice.
- player i indeed maximizes total surplus according to his type and others' announced types.
- So player i would tell the truth $\hat{\theta}_i = \theta_i$.

Vickrey-Clarke-Groves Mechanism

- **Pivotal mechanism** suggested by Clarke (1971) is a particular VCG mechanism.
- It is obtained by setting

$$h_i\left(\hat{\theta}_{-i}\right) = -\sum_{j \neq i} u_j\left(x^*_{-i}\left(\hat{\theta}_{-i}\right), \hat{\theta}_j\right),$$

where

$$x^*_{-i}\left(\hat{\theta}_{-i}\right) \in \arg\max_{x \in X} \sum_{j \neq i} u_j\left(x, \hat{\theta}_j\right)$$

is the optimal choice of x for a society from which player i was absent. Thus

$$m_i\left(\hat{\theta}\right) = \sum_{j \neq i} u_j\left(x^*\left(\hat{\theta}_i, \hat{\theta}_{-i}\right), \hat{\theta}_j\right) - \sum_{j \neq i} u_j\left(x^*_{-i}\left(\hat{\theta}_{-i}\right), \hat{\theta}_j\right)$$

Vickrey-Clarke-Groves Mechanism

- Pivotal mechanism lets player i make his announcement that affects the outcome had he not been part of society.
- There are relevant cases:
  - Case 1: $x^*\left(\hat{\theta}_i, \hat{\theta}_{-i}\right) = x^*_{-i}\left(\hat{\theta}_{-i}\right)$ where player i's announcement does not change what would have happened if he were not part of society. Then the mechanism specifies a transfer of zero to i.
  - Case 2: $x^*\left(\hat{\theta}_i, \hat{\theta}_{-i}\right) \neq x^*_{-i}\left(\hat{\theta}_{-i}\right)$ where player i is pivotal that his announcement changes what would have happened without him. His transfer ends up taxing him for the externality that his announcement imposes on the other players.

Example: allocation of an indivisible private good

- Returning to Example 2, where the mechanism designer is trying to determine who to give a license to use a certain portion of the electromagnetic spectrum for cell coverage.
- An object can be allocated to one of N players.
- The value of owning the private good for player i is given by $u_i(x, \theta_i) = \theta_i x_i$.
- The first-best allocation solves

$$\max_{(x_1, \ldots x_n) \in \{0,1\}^n} \sum_{i=1} \theta_i x_i \quad \text{subject to} \quad \sum_i x_i = 1,$$

Example: allocation of an indivisible private good

- Which results in allocating the good to the player i^* with the highest valuation: $i^* \in \arg\max x_i \theta_i$, and

$$x_i^*(\theta) = \begin{cases} 1 & \text{if } i = i^* \\ 0 & \text{otherwise} \end{cases}$$

- The pivotal mechanism then has transfers

$$\begin{aligned} m_i\left(\hat{\theta}\right) &= \sum_{j \neq i} u_j\left(x^*\left(\hat{\theta}\right), \hat{\theta}_{-i}\right) - \sum_{j \neq i} u_j\left(x^*_{-i}\left(\hat{\theta}_{-i}\right), \hat{\theta}\right) \\ &= \begin{cases} -\left\{\max_{j \neq i^*} \hat{\theta}_j\right\} 1 & \text{if } i = i^* \\ 0 & \text{otherwise} \end{cases} \end{aligned}$$

Example: allocation of an indivisible private good

- That is, every player $i \neq i^*$ is not pivotal and his presence does not affect the allocation.
- Therefore $m_i\left(\hat{\theta}\right) = 0$.
- Player i^* is pivotal: without him, the object would go to the player with the second-highest valuation.
- The total surplus would be $\max_{j \neq i^*} \theta_j$.
- This is the externality player i^* imposes on the others by being present, and how much he has to pay in the pivotal mechanism.
- Notice that this mechanism is identical to the second-price sealed-bid auction.