Chapter 10

Mechanism Design

Mechanism design is one of the important areas of economics. The 2007 Nobel prize in Economics went to three economists who laid its foundations. To quote from the article *Intelligent design*, published in *The Economist*, October 18th, 2007, mechanism design deals with the problem of 'how to arrange our economic interactions so that, when everyone behaves in a self-interested manner, the result is something we all like.' So these interactions are supposed to yield desired social decisions when each agent is interested in maximizing only his own utility.

In mechanism design one is interested in the ways of inducing the players to submit true information. To discuss it in more detail we need to introduce some basic concepts.

10.1 Decision problems

Assume a set of ***decisions*** D, a set $\{1, \ldots, n\}$ of players, and for each player

- a set of ***types*** Θ_i, and
- an ***initial utility function*** $v_i : D \times \Theta_i \rightarrow \mathbb{R}$.

In this context a type is some private information known only to the player, for example, in the case of an auction, player's valuation of the items for sale.

As in the case of strategy sets we use the following abbreviations:

- $\Theta := \Theta_1 \times \ldots \times \Theta_n$,

- $\Theta_{-i} := \Theta_1 \times \ldots \times \Theta_{i-1} \times \Theta_{i+1} \times \ldots \times \Theta_n$, and similarly with θ_{-i} where $\theta \in \Theta$,

- $(\theta'_i, \theta_{-i}) := \theta_1 \times \ldots \times \theta_{i-1} \times \theta'_i \times \theta_{i+1} \times \ldots \times \theta_n$.

In particular $(\theta_i, \theta_{-i}) = \theta$.

A ***decision rule*** is a function $f : \Theta \to D$. We call the tuple

$$(D, \Theta_1, \ldots, \Theta_n, v_1, \ldots, v_n, f)$$

a ***decision problem***.

Decision problems are considered in the presence of a ***central authority*** who takes decisions on the basis of the information provided by the players. Given a decision problem the desired decision is obtained through the following sequence of events, where f is a given, publicly known, decision rule:

- each player i receives (becomes aware of) his type $\theta_i \in \Theta_i$,

- each player i announces to the central authority a type $\theta'_i \in \Theta_i$; this yields a type vector $\theta' := (\theta'_1, \ldots, \theta'_n)$,

- the central authority then takes the decision $d := f(\theta')$ and communicates it to each player,

- the resulting initial utility for player i is then $v_i(d, \theta_i)$.

The difficulty in taking decisions through the above described sequence of events is that players are assumed to be ***rational***, that is they want to maximize their utility. As a result they may submit false information to manipulate the outcome (decision). We shall return to this problem in the next section. But first, to better understand the above notion let us consider some natural examples.

Given a sequence $a := (a_1, \ldots, a_j)$ of reals denote the least l such that $a_l = \max_{k \in \{1,\ldots,j\}} a_k$ by $\text{argsmax}\, a$.

Additionally, for a function $g : A \to \mathbb{R}$ we define

$$\text{argmax}_{x \in A} g(x) := \{y \in A \mid g(y) = \max_{x \in A} g(x)\}.$$

So $a \in \text{argmax}_{x \in A} g(x)$ means that a is a maximum of the function g on the set A.

Example 18 **[Sealed-bid auction]**

We consider a ***sealed-bid auction*** in which there is a single object for sale. Each player (bidder) simultaneously submits to the central authority his type (bid) in a sealed envelope and the object is allocated to the highest bidder.

We view each player's valutation as his type. More precisely, we model this type of auction as the following decision problem $(D, \Theta_1, \ldots, \Theta_n, v_1, \ldots, v_n, f)$:

- $D = \{1, \ldots, n\}$,
- for all $i \in \{1, \ldots, n\}$, $\Theta_i = \mathbb{R}_+$;

  $\theta_i \in \Theta_i$ is player's i valuation of the object,
- $v_i(d, \theta_i) := \begin{cases} \theta_i & \text{if } d = i \\ 0 & \text{otherwise} \end{cases}$
- $f(\theta) := \text{argsmax}\, \theta$.

  Here decision $d \in D$ indicates to which player the object is sold. Note that at this stage we only modeled the fact that the object is sold to the highest bidder (with the ties resolved in the favour of a bidder with the lowest index). We shall return to the problem of payments in the next section. □

Example 19 **[Public project problem]**

This problem deals with the task of taking a joint decision concerning construction of a ***public good***[1], for example a bridge.

It is explained as follows in the *Scientific Background* of the Royal Swedish Academy of Sciences Press Release that accompanied the Nobel prize in Economics in 2007:

> Each person is asked to report his or her willingness to pay for the project, and the project is undertaken if and only if the aggregate reported willingness to pay exceeds the cost of the project.

[1]In Economics public goods are so-called not excludable and nonrival goods. To quote from the book *N.G. Mankiw, Principles of Economics*, 2nd Editiona, Harcourt, 2001: "People cannot be prevented from using a public good, and one person's enjoyment of a public good does not reduce another person's enjoyment of it."

So there are two decisions: to carry out the project or not. In the terminology of the decision problems each player reports to the central authority his appreciation of the gain from the project when it takes place. If the sum of the appreciations exceeds the cost of the project, the project takes place. We assume that each player has to pay then the same fraction of the cost. Otherwise the project is cancelled.

This leads to the following decision problem:

- $D = \{0, 1\}$,
- each Θ_i is $\mathbb{R}_+$,
- $v_i(d, \theta_i) := d(\theta_i - \frac{c}{n})$,
- $f(\theta) := \begin{cases} 1 & \text{if } \sum_{i=1}^{n} \theta_i \geq c \\ 0 & \text{otherwise} \end{cases}$

Here c is the cost of the project. If the project takes place ($d = 1$), $\frac{c}{n}$ is the cost share of the project for each player. □

Example 20 [Taking an efficient decision] We assume a finite set of decisions. Each player submits to the central authority a function that describes his satisfaction level from each decision if it is taken. The central authority then chooses a decision that yields the maximal overall satisfaction.

This problem corresponds to the following decision problem:

- D is the given finite set of decisions,
- each Θ_i is $\{f \mid f : D \rightarrow \mathbb{R}\}$,
- $v_i(d, \theta_i) := \theta_i(d)$,
- the decision rule f is a function such that for all θ, $f(\theta) \in \text{argmax}_{d \in D} \sum_{i=1}^{n} \theta_i(d)$.

□

Example 21 [Reversed sealed-bid auction]

In the ***reversed sealed-bid auction*** each player offers the same service, for example to construct a bridge. The decision is taken by means of a sealed-bid auction. Each player simultaneously submits to the central authority his type (bid) in a sealed envelope and the service is purchased from the lowest bidder.

We model it in exactly the same way as the sealed-bid auction, with the only exception that for each player the types are now non-positive reals. So we consider the following decision problem:

- $D = \{1, \ldots, n\}$,
- for all $i \in \{1, \ldots, n\}$, $\Theta_i = \mathbb{R}_-$ (the set of non-positive reals);

  $-\theta_i$, where $\theta_i \in \Theta_i$, is player's i offer for the service,
- $v_i(d, \theta_i) := \begin{cases} \theta_i & \text{if } d = i \\ 0 & \text{otherwise} \end{cases}$
- $f(\theta) := \text{argsmax}\, \theta$.

Here decision $d \in D$ indicates from which player the service is bought. So for example $f(-8, -5, -4, -6) = 3$, that is, given the offers $8, 5, 4, 6$ (in that order), the service is bought from player 3, since he submitted the lowest bid, namely 4. As in the case of the sealed-bid auction, we shall return to the problem of payments in the next section. □

Example 22 [Buying a path in a network]

We consider a communication network, modelled as a directed graph $G := (V, E)$ (with no self-cycles or parallel edges). We assume that each edge $e \in E$ is owned by a player, also denoted by e. So different edges are owned by different players. We fix two distinguished vertices $s, t \in V$. Each player e submits the cost θ_e of using the edge e. The central authority selects on the basis of players' submissions the shortest $s - t$ path in G.

Below we denote by $G(\theta)$ the graph G augmented with the costs of edges as specified by θ. That is, the cost of each edge i in $G(\theta)$ is θ_i.

This problem can be modelled as the following decision problem:

- $D = \{p \mid p \text{ is a } s - t \text{ path in } G\}$,
- each Θ_i is $\mathbb{R}_+$;

  θ_i is the cost incurred by player i if the edge i is used in the selected path,

- $v_i(p, \theta_i) := \begin{cases} -\theta_i & \text{if } i \in p \\ 0 & \text{otherwise} \end{cases}$

- $f(\theta) := p$, where p is the shortest $s - t$ path in $G(\theta)$.

  In the case of multiple shortest paths we select, say, the one that is alphabetically first.

Note that in the case an edge is selected, the utility of its owner becomes negative. This reflects the fact we focus on incurring costs and not on benefits. In the next section we shall introduce taxes and discuss a scheme according to which each owner of a selected path is *paid* by the central authority an amount exceeding the incurred costs. □

Let us return now to the decision rules. We call a decision rule f ***efficient*** if for all $\theta \in \Theta$ and $d' \in D$

$$\sum_{i=1}^{n} v_i(f(\theta), \theta_i) \geq \sum_{i=1}^{n} v_i(d', \theta_i),$$

or alternatively

$$f(\theta) \in \text{argmax}_{d \in D} \sum_{i=1}^{n} v_i(d, \theta_i).$$

This means that for all $\theta \in \Theta$, $f(\theta)$ is a decision that maximizes the ***initial social welfare***, defined by $\sum_{i=1}^{n} v_i(d, \theta_i)$.

It is easy to check that the decision rules used in Examples 18–22 are efficient. Take for instance Example 22. For each $s - t$ path p we have $\sum_{i=1}^{n} v_i(p, \theta_i) = -\sum_{j \in p} \theta_j$, so $\sum_{i=1}^{n} v_i(p, \theta_i)$ reaches maximum when p is a shortest $s - t$ path in $G(\theta)$, which is the choice made by the decision rule f used there.

10.2 Direct mechanisms

Let us return now to the subject of manipulations. A problem with our description of the sealed-bid auction is that we intentionally neglected the fact that the winner should pay for the object for sale. Still, we can imagine in this limited setting that player i with a strictly positive valuation of the object somehow became aware of the types (that is, bids) of the other players.

Then he should just submit a type strictly larger than the other types. This way the object will be allocated to him and his utility will increase from 0 to θ_i.

The manipulations are more natural to envisage in the case of the public project problem. A player whose type (that is, appreciation of the gain from the project) exceeds $\frac{c}{n}$, the cost share he is to pay, should manipulate the outcome and announce the type c. This will guarantee that the project will take place, irrespectively of the types announced by the other players. Analogously, player whose type is lower than $\frac{c}{n}$ should submit the type 0 to minimize the chance that the project will take place.

To prevent such manipulations we use taxes. This leads to mechanisms that are constructed by combining decision rules with taxes (transfer payments). Each such mechanism is obtained by modifying the initial decision problem $(D, \Theta_1, \ldots, \Theta_n, v_1, \ldots, v_n, f)$ to the following one:

- the set of decisions is $D \times \mathbb{R}^n$,
- the decision rule is a function $(f, t) : \Theta \to D \times \mathbb{R}^n$, where $t : \Theta \to \mathbb{R}^n$ and $(f, t)(\theta) := (f(\theta), t(\theta))$,
- the ***final utility function*** for player i is a function $u_i : D \times \mathbb{R}^n \times \Theta_i \to \mathbb{R}$ defined by

  $$u_i(d, t_1, \ldots, t_n, \theta_i) := v_i(d, \theta_i) + t_i.$$

  (So defined utilities are called ***quasilinear***.)

We call $(D \times \mathbb{R}^n, \Theta_1, \ldots, \Theta_n, u_1, \ldots, u_n, (f, t))$ a ***direct mechanism*** and refer to t as the ***tax function***.

So when the received (true) type of player i is θ_i and his announced type is θ_i', his final utility is

$$u_i((f, t)(\theta_i', \theta_{-i}), \theta_i) = v_i(f(\theta_i', \theta_{-i}), \theta_i) + t_i(\theta_i', \theta_{-i}),$$

where θ_{-i} are the types announced by the other players.

In each direct mechanism, given the vector θ of announced types, $t(\theta) := (t_1(\theta), \ldots, t_n(\theta))$ is the vector of the resulting payments that the players have to make. If $t_i(\theta) \geq 0$, player i ***receives*** from the central authority $t_i(\theta)$, and if $t_i(\theta) < 0$, he ***pays*** to the central authority $|t_i(\theta)|$.

The following definition then captures the idea that taxes prevent manipulations. We say that a direct mechanism with tax function t is ***incentive compatible*** if for all $\theta \in \Theta$, $i \in \{1, \ldots, n\}$ and $\theta'_i \in \Theta_i$

$$u_i((f, t)(\theta_i, \theta_{-i}), \theta_i) \geq u_i((f, t)(\theta'_i, \theta_{-i}), \theta_i).$$

Intuitively, this means that announcing one's true type (θ_i) is better than announcing another type (θ'_i). That is, false announcements, i.e., manipulations do not pay off.

From now on we focus on specific incentive compatible direct mechanisms. Each ***Groves mechanism*** is a direct mechanism obtained by using a tax function $t := (t_1, \ldots, t_n)$, where for all $i \in \{1, \ldots, n\}$

- $t_i : \Theta \rightarrow \mathbb{R}$ is defined by $t_i(\theta) := g_i(\theta) + h_i(\theta_{-i})$, where
- $g_i(\theta) := \sum_{j \neq i} v_j(f(\theta), \theta_j)$,
- $h_i : \Theta_{-i} \rightarrow \mathbb{R}$ is an arbitrary function.

Note that $v_i(f(\theta), \theta_i) + g_i(\theta) = \sum_{j=1}^{n} v_j(f(\theta), \theta_j)$ is simply the initial social welfare from the decision $f(\theta)$. In this context the ***final social welfare*** is defined as $\sum_{i=1}^{n} u_i((f, t)(\theta), \theta_i)$, so it equals the sum of the initial social welfare and all the taxes.

The importance of Groves mechanisms is then revealed by the following crucial result due to T. Groves.

Theorem 45 *Consider a decision problem* $(D, \Theta_1, \ldots, \Theta_n, v_1, \ldots, v_n, f)$ *with an efficient decision rule* f*. Then each Groves mechanism is incentive compatible.*

Proof. The proof is remarkably straightforward. Since f is efficient, for all $\theta \in \Theta$, $i \in \{1, \ldots, n\}$ and $\theta'_i \in \Theta_i$ we have

$$\begin{aligned} u_i((f, t)(\theta_i, \theta_{-i}), \theta_i) &= \sum_{j=1}^{n} v_j(f(\theta_i, \theta_{-i}), \theta_j) + h_i(\theta_{-i}) \\ &\geq \sum_{j=1}^{n} v_j(f(\theta'_i, \theta_{-i}), \theta_j) + h_i(\theta_{-i}) \\ &= u_i((f, t)(\theta'_i, \theta_{-i}), \theta_i). \end{aligned}$$

□

When for a given direct mechanism for all θ' we have $\sum_{i=1}^{n} t_i(\theta') \leq 0$, the mechanism is called ***feasible*** (which means that it can be realized without external financing) and when for all θ' we have $\sum_{i=1}^{n} t_i(\theta') = 0$, the mechanism is called ***budget balanced*** (which means that it can be realized without a deficit).

Each Groves mechanism is uniquely determined by the functions $h_1, \ldots,$ h_n. A special case, called ***pivotal mechanism*** is obtained by using

$$h_i(\theta_{-i}) := -\max_{d \in D} \sum_{j \neq i} v_j(d, \theta_j).$$

So then

$$t_i(\theta) := \sum_{j \neq i} v_j(f(\theta), \theta_j) - \max_{d \in D} \sum_{j \neq i} v_j(d, \theta_j).$$

Hence for all θ and $i \in \{1, \ldots, n\}$ we have $t_i(\theta) \leq 0$, which means that each player needs to make the payment $|t_i(\theta)|$ to the central authority. In particular, the pivotal mechanism is feasible.

10.3 Back to our examples

When applying Theorem 45 to a specific decision problem we need first to check that the used decision rule is efficient. We noted already that this is the case in Examples 18–22. So in each example Theorem 45 applies and in particular the pivotal mechanism can be used. Let us see now the details of this and other Groves mechanisms for these examples.

Sealed-bid auction

To compute the taxes we use the following observation.

Note 46 *In the sealed-bid auction we have for the pivotal mechanism*

$$t_i(\theta) = \begin{cases} -\max_{j \neq i} \theta_j & \text{if } i = \mathit{argsmax}\, \theta. \\ 0 & \text{otherwise} \end{cases}$$

□

So the highest bidder wins the object and pays for it the amount $\max_{j \neq i} \theta_j$, i.e., the second highest bid. This shows that the pivotal mechanism for the sealed-bid auction is simply the second-price auction proposed by W. Vickrey. By the above considerations this auction is incentive compatible.

In contrast, the first-price sealed-bid auction, in which the winner pays the price he offered, is not incentive compatible. Indeed, suppose that the true types are (4,5,7) and that players 1 and 2 bid truthfully. If player 3 bids truthfully, he wins the object and his payoff is 0. But if he bids 6, he increases his payoff to 1.

Bailey-Cavallo mechanism

Second-price auction is a natural approach in the set up when the central authority is a seller, as the tax corresponds then to payment for the object for sale. But we can also use the initial decision problem simply to determine which of the player values the object most. In such a set up the central authority is merely an arbiter and it is meaningful then to reach the decision with limited taxes.

Below, given a sequence $\theta \in \mathbb{R}^n$ of reals we denote by θ^* its reordering from the largest to the smallest element. So for example, for $\theta = (1, 4, 2, 3, 0)$ we have $(\theta_{-2})^*_2 = 2$ since $\theta_{-2} = (1, 2, 3, 0)$ and $(\theta_{-2})^* = (3, 2, 1, 0)$.

In the case of the second-price auction the final social welfare, i.e., $\sum_{j=1}^{n} u_j((f,t)(\theta), \theta_j)$, equals $\theta_i - \max_{j \neq i} \theta_j$, where $i = \text{argsmax}\theta$, so it equals the difference between the highest bid and the second highest bid.

We now discuss a modification of the second-price auction which yields a larger final social welfare. To ensure that it is well-defined we need to assume that $n \geq 3$. This modification, called ***Bailey-Cavallo mechanism***, is achieved by combining each tax $t'_i(\theta)$ to be paid in the second-price auction with

$$h'_i(\theta_{-i}) := \frac{(\theta_{-i})^*_2}{n},$$

that is, by using

$$t_i(\theta) := t'_i(\theta) + h'_i(\theta_{-i}).$$

Note that this yields a Groves mechanism since by the definition of the pivotal mechanism for specific functions $h_1, \ldots, h_n$

$$t'_i(\theta) = \sum_{j \neq i} v_j(f(\theta), \theta_j) + h_i(\theta_{-i}),$$

and consequently

$$t_i(\theta) = \sum_{j \neq i} v_j(f(\theta), \theta_j) + (h_i + h'_i)(\theta_{-i}).$$

In fact, this modification is a Groves mechanism if we start with an arbitrary Groves mechanism. In the case of the second-price auction the resulting mechanism is feasible since for all $i \in \{1, \ldots, n\}$ and θ we have $(\theta_{-i})^*_2 \leq \theta^*_2$ and as a result, since $\max_{j \neq i} \theta_j = \theta^*_2$,

$$\sum_{i=1}^{n} t_i(\theta) = \sum_{i=1}^{n} t'_i(\theta) + \sum_{i=1}^{n} h'_i(\theta_{-i}) = \sum_{i=1}^{n} \frac{-\theta^*_2 + (\theta_{-i})^*_2}{n} \leq 0.$$

Let, given the sequence θ of submitted bids (types), π be the permutation of $1, \ldots, n$ such that $\theta_{\pi(i)} = \theta^*_i$ for $i \in \{1, \ldots, n\}$ (where we break the ties by selecting players with the lower index first). So the ith highest bid is by player $\pi(i)$ and the object is sold to player $\pi(1)$. Note that then

- $(\theta_{-i})^*_2 = \theta^*_3$ for $i \in \{\pi(1), \pi(2)\}$,
- $(\theta_{-i})^*_2 = \theta^*_2$ for $i \in \{\pi(3), \ldots, \pi(n)\}$,

so the above mechanism boils down to the following payments by player $\pi(1)$:

- $\frac{\theta^*_3}{n}$ to player $\pi(2)$,
- $\frac{\theta^*_2}{n}$ to players $\pi(3), \ldots, \pi(n)$,
- $\theta^*_2 - \frac{2}{n}\theta^*_3 - \frac{n-2}{n}\theta^*_2 = \frac{2}{n}(\theta^*_2 - \theta^*_3)$ to the central authority.

To illustrate these payments assume that there are three players, A, B, and C whose true types (valuations) are 18, 21, and 24, respectively. When they bid truthfully the object is allocated to player C. In the second-price auction player's C tax is 21 and the final social welfare is $24 - 21 = 3$.

In constrast, in the case of the Bailey-Cavallo mechanism we have for the vector $\theta = (18, 21, 24)$ of submitted types $\theta^*_2 = 21$ and $\theta^*_3 = 18$, so player C pays

- 6 to player B,
- 7 to player A,
- 2 to the tax authority.

So the final social welfare is now $24 - 2 = 22$. Table 10.1 summarizes the situation.

player	type	tax	u_i
A	18	7	7
B	21	6	6
C	24	-15	9

Table 10.1: The Bailey-Cavallo mechanism

Public project problem

Let us return now to Example 19. To compute the taxes in the case of the pivotal mechanism we use the following observation.

Note 47 *In the public project problem we have for the pivotal mechanism*

$$t_i(\theta) = \begin{cases} 0 & \text{if } \sum_{j \neq i} \theta_j \geq \frac{n-1}{n}c \text{ and } \sum_{j=1}^{n} \theta_j \geq c \\ \sum_{j \neq i} \theta_j - \frac{n-1}{n}c & \text{if } \sum_{j \neq i} \theta_j < \frac{n-1}{n}c \text{ and } \sum_{j=1}^{n} \theta_j \geq c \\ 0 & \text{if } \sum_{j \neq i} \theta_j \leq \frac{n-1}{n}c \text{ and } \sum_{j=1}^{n} \theta_j < c \\ \frac{n-1}{n}c - \sum_{j \neq i} \theta_j & \text{if } \sum_{j \neq i} \theta_j > \frac{n-1}{n}c \text{ and } \sum_{j=1}^{n} \theta_j < c \end{cases}$$

□

To illustrate the pivotal mechanism suppose that there are three players, A, B, and C whose true types are 6, 7, and 25, and $c = 30$, respectively. When these types are announced the project takes place and Table 10.2 summarizes the taxes that players need to pay and their final utilities. The taxes were computed using Note 47.

player	type	tax	u_i
A	6	0	-4
B	7	0	-3
C	25	-7	8

Table 10.2: The pivotal mechanism for the public project problem

Suppose now that the true types of players are 4, 3 and 22, respectively and, as before, $c = 30$. When these types are also the announced types, the project does not take place. Still, some players need to pay a tax, as Table 10.3 illustrates.

player	type	tax	u_i
A	4	-5	-5
B	3	-6	-6
C	22	0	0

Table 10.3: The pivotal mechanism for the public project problem

Reversed sealed-bid auction

Note that the pivotal mechanism is not appropriate here. Indeed, we noted already that in the pivotal mechanism all players need to *make* a payment to the central authority, while in the context of the reversed sealed-bid auction we want to ensure that the lowest bidder *receives* a payment from the authority and other bidders neither pay nor receive any payment.

This can be realized by using the Groves mechanism with the following tax definition:

$$t_i(\theta) := \sum_{j \neq i} v_j(f(\theta), \theta_j) - \max_{d \in D \setminus \{i\}} \sum_{j \neq i} v_j(d, \theta_j).$$

The crucial difference between this mechanism and the pivotal mechanism is that in the second expression we take a maximum over all decisions in the set $D \setminus \{i\}$ and *not* D.

To compute the taxes in the reversed sealed-bid auction with the above mechanism we use the following observation.

Note 48

$$t_i(\theta) = \begin{cases} -\max_{j \neq i} \theta_j & \text{if } i = \textit{argsmax}\, \theta. \\ 0 & \text{otherwise} \end{cases}$$

□

This is identical to Note 46 in which the taxes for the pivotal mechanism for the sealed bid auction were computed. However, because we use here negative reals as bids the interpretation is different. Namely, the taxes are now positive, i.e., the players now *receive* the payments. More precisely, the winner, i.e., player i such that $i = \text{argsmax}\, \theta$, receives the payment equal to the second lowest offer, while the other players pay no taxes.

For example, when $\theta = (-8, -5, -4, -6)$, the service is bought from player 3 who submitted the lowest bid, namely 4. He receives for it the amount 5. Indeed, $3 = \text{argsmax}\, \theta$ and $-\max_{j \neq 3} \theta_j = -(-5) = 5$.

Buying a path in a network

As in the case of the reversed sealed-bid auction the pivotal mechanism is not appropriate here since we want to ensure that the players whose edge was selected *receive* a payment. Again, we achieve this by a simple modification of the pivotal mechanism. We modify it to a Groves mechanism in which

- the central authority is viewed as an agent who procures an $s - t$ path and pays the players whose edges are used,
- the players have an incentive to participate: if an edge is used, then the final utility of its owner is ≥ 0.

Recall that in the case of the pivotal mechanism we have

$$t'_i(\theta) = \sum_{j \neq i} v_j(f(\theta), \theta_j) - \max_{p \in D(G)} \sum_{j \neq i} v_j(p, \theta_j),$$

where we now explicitly indicate the dependence of the decision set on the underlying graph, i.e., $D(G) := \{p \mid p \text{ is a } s - t \text{ path in } G\}$.

We now put instead

$$t_i(\theta) := \sum_{j \neq i} v_j(f(\theta), \theta_j) - \max_{p \in D(G \setminus \{i\})} \sum_{j \neq i} v_j(p, \theta_j).$$

The following note provides the intuition for the above tax. We abbreviate here $\sum_{j \in p} \theta_j$ to $cost(p)$.

Note 49

$$t_i(\theta) = \begin{cases} cost(p_2) - cost(p_1 - \{i\}) & \text{if } i \in p_1 \\ 0 & \text{otherwise} \end{cases}$$

where p_1 is the shortest $s - t$ path in $G(\theta)$ and p_2 is the shortest $s - t$ path in $(G \setminus \{i\})(\theta_{-i})$.

Proof. Note that for each $s - t$ path p we have

$$-\sum_{j \neq i} v_j(p, \theta_j) = \sum_{j \in p - \{i\}} \theta_j.$$

Recall now that $f(\theta)$ is the shortest $s-t$ path in $G(\theta)$, i.e., $f(\theta) = p_1$. So $\sum_{j \neq i} v_j(f(\theta), \theta_j) = -cost(p_1 - \{i\})$.

To understand the second expression in the definition of $t_i(\theta)$ note that for each $p \in D(G \setminus \{i\})$, so for each $s-t$ path p in $G \setminus \{i\}$, we have

$$-\sum_{j \neq i} v_j(p, \theta_j) = \sum_{j \in p - \{i\}} \theta_j = \sum_{j \in p} \theta_j,$$

since the edge i does not belong to the path p. So $-\max_{p \in D(G \setminus \{i\})} \sum_{j \neq i} v_j(p, \theta_j)$ equals the length of the shortest $s-t$ path in $(G \setminus \{i\})(\theta_{-i})$, i.e., it equals $cost(p_2)$. □

So given θ and the above definitions of the paths p_1 and p_2 the central authority pays to each player i whose edge is used the amount $cost(p_2) - cost(p_1 - \{i\})$. The final utility of such a player is then $-\theta_i + cost(p_2) - cost(p_1 - \{i\})$, i.e., $cost(p_2) - cost(p_1)$. So by the choice of p_1 and p_2 it is positive. No payments are made to the other players and their final utilities are 0.

Consider an example. Take the communication network depicted in Figure 10.1.

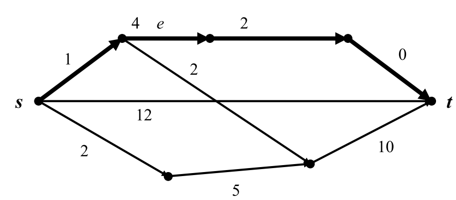

Figure 10.1: A communication network

This network has nine edges, so it corresponds to a decision problem with nine players. We assume that each player submitted the depicted length of the edge. Consider the player who owns the edge e, of length 4. To compute the payment he receives we need to determine the shortest $s-t$ path and the shortest $s-t$ path that does not include the edge e. The first path is the upper path, depicted in Figure 10.1 in bold. It contains the edge e and has the length 7. The second path is simply the edge connecting s and t and its length is 12. So, assuming that the players submit the costs truthfully, according to Note 49 player e receives the payment $12 - (7 - 4) = 9$ and his final utility is $9 - 4 = 5$.

10.4 Green and Laffont result

Until now we studied only one class of incentive compatible direct mechanisms, namely Groves mechanisms. Are there any other ones? J. Green and J.-J. Laffont showed that when the decision rule is efficient, under a natural assumption no other incentive compatible direct mechanisms exist. To formulate the relevant result we introduce the following notion.

Given a decision problem $(D, \Theta_1, \ldots, \Theta_n, v_1, \ldots, v_n, f)$, we call the utility function v_i ***complete*** if

$$\{v \mid v : D \rightarrow \mathbb{R}\} = \{v_i(\cdot, \theta_i) \mid \theta_i \in \Theta_i\},$$

that is, if each function $v : D \rightarrow \mathbb{R}$ is of the form $v_i(\cdot, \theta_i)$ for some $\theta_i \in \Theta_i$.

Theorem 50 *Consider a decision problem* $(D, \Theta_1, \ldots, \Theta_n, v_1, \ldots, v_n, f)$ *with an efficient decision rule* f*. Suppose that each utility function* v_i *is complete. Then each incentive compatible direct mechanism is a Groves mechanism.*

To prove it first observe that each direct mechanism originating from a decision problem $(D, \Theta_1, \ldots, \Theta_n, v_1, \ldots, v_n, f)$ can be written in a 'Groves-like' way, by putting

$$t_i(\theta) := \sum_{j \neq i} v_j(f(\theta), \theta_j) + h_i(\theta),$$

where each function h_i is defined on Θ and *not* on Θ_{-i}, as in the Groves mechanisms.

Lemma 51 *For each incentive compatible direct mechanism*

$$(D \times \mathbb{R}^n, \Theta_1, \ldots, \Theta_n, u_1, \ldots, u_n, (f, t)),$$

given the above representation, for all $i \in \{1, \ldots, n\}$

$$f(\theta_i, \theta_{-i}) = f(\theta'_i, \theta_{-i}) \text{ implies } h_i(\theta_i, \theta_{-i}) = h_i(\theta'_i, \theta_{-i}).$$

Proof. Fix $i \in \{1, \ldots, n\}$. We have

$$u_i((f,t)(\theta_i, \theta_{-i}), \theta_i) = \sum_{j=1}^{n} v_j(f(\theta_i, \theta_{-i})), \theta_j) + h_i(\theta_i, \theta_{-i})$$

and

$$u_i((f,t)(\theta'_i, \theta_{-i}), \theta_i) = \sum_{j=1}^{n} v_j(f(\theta'_i, \theta_{-i})), \theta_j) + h_i(\theta'_i, \theta_{-i}),$$

so, on the account of the incentive compatibility, $f(\theta_i, \theta_{-i}) = f(\theta'_i, \theta_{-i})$ implies $h_i(\theta_i, \theta_{-i}) \geq h_i(\theta'_i, \theta_{-i})$. By symmetry $h_i(\theta'_i, \theta_{-i}) \geq h_i(\theta_i, \theta_{-i})$, as well. $\square$

Proof of Theorem 50.

Consider an incentive compatible direct mechanism

$$(D \times \mathbb{R}^n, \Theta_1, \ldots, \Theta_n, u_1, \ldots, u_n, (f,t))$$

and its 'Groves-like' representation with the functions $h_1, \ldots, h_n$. We need to show that no function h_i depends on its ith argument. Suppose otherwise. Then for some i, θ and θ'_i

$$h_i(\theta_i, \theta_{-i}) > h_i(\theta'_i, \theta_{-i}).$$

Choose an arbitrary ϵ from the open interval $(0, h_i(\theta_i, \theta_{-i}) - h_i(\theta'_i, \theta_{-i}))$ and consider the following function $v : D \to \mathbb{R}$:

$$v(d) := \begin{cases} \epsilon - \sum_{j \neq i} v_j(d, \theta_j) & \text{if } d = f(\theta'_i, \theta_{-i}) \\ -\sum_{j \neq i} v_j(d, \theta_j) & \text{otherwise} \end{cases}$$

By the completeness of v_i for some $\theta''_i \in \Theta_i$

$$v(d) = v_i(d, \theta''_i)$$

for all $d \in D$.

Since $h_i(\theta_i, \theta_{-i}) > h_i(\theta'_i, \theta_{-i})$, by Lemma 51 $f(\theta_i, \theta_{-i}) \neq f(\theta'_i, \theta_{-i})$, so by the definition of v

$$v_i(f(\theta_i, \theta_{-i}), \theta''_i) + \sum_{j \neq i} v_j(f(\theta_i, \theta_{-i}), \theta_j) = 0. \tag{10.1}$$

Further, for each $d \in D$ the sum $v_i(d, \theta_i'') + \sum_{j \neq i} v_j(d, \theta_j)$ equals either 0 or ϵ. This means that by the efficiency of f

$$v_i(f(\theta_i'', \theta_{-i}), \theta_i'') + \sum_{j \neq i} v_j(f(\theta_i'', \theta_{-i}), \theta_j) = \epsilon. \tag{10.2}$$

Hence, by the definition of v we have $f(\theta_i'', \theta_{-i}) = f(\theta_i', \theta_{-i})$, and consequently by Lemma 51

$$h_i(\theta_i'', \theta_{-i}) = h_i(\theta_i', \theta_{-i}). \tag{10.3}$$

We have now by (10.1)

$$\begin{aligned} & u_i((f,t)(\theta_i, \theta_{-i}), \theta_i'') \\ = & v_i(f(\theta_i, \theta_{-i}), \theta_i'') + \sum_{j \neq i} v_j(f(\theta_i, \theta_{-i}), \theta_j) + h_i(\theta_i, \theta_{-i}) \\ = & h_i(\theta_i, \theta_{-i}). \end{aligned}$$

In turn, by (10.2) and (10.3),

$$\begin{aligned} & u_i((f,t)(\theta_i'', \theta_{-i}), \theta_i'') \\ = & v_i(f(\theta_i'', \theta_{-i}), \theta_i'') + \sum_{j \neq i} v_j(f(\theta_i'', \theta_{-i}), \theta_j) + h_i(\theta_i'', \theta_{-i}) \\ = & \epsilon + h_i(\theta_i', \theta_{-i}). \end{aligned}$$

But by the choice of ϵ we have $h_i(\theta_i, \theta_{-i}) > \epsilon + h_i(\theta_i', \theta_{-i})$, so

$$u_i((f,t)(\theta_i, \theta_{-i}), \theta_i'') > u_i((f,t)(\theta_i'', \theta_{-i}), \theta_i''),$$

which contradicts the incentive compatibility for the joint type $(\theta_i'', \theta_{-i})$. $\square$

Chapter 11

Pre-Bayesian Games

Mechanism design, as introduced in the previous chapter, can be explained in game-theoretic terms using pre-Bayesian games In strategic games, after each player selected his strategy, each player knows the payoff of *every other player*. This is not the case in pre-Bayesian games in which each player has a private type on which he can condition his strategy. This distinguishing feature of pre-Bayesian games explains why they form a class of ***games with incomplete information***. Formally, they are defined as follows.

Assume a set $\{1, \ldots, n\}$ of players, where $n > 1$. A ***pre-Bayesian game*** for n players consists of

- a non-empty set A_i of ***actions***,
- a non-empty set Θ_i of ***types***,
- a ***payoff function*** $p_i : A_1 \times \ldots \times A_n \times \Theta_i \rightarrow \mathbb{R}$,

for each player i.

Let $A := A_1 \times \ldots \times A_n$. In a pre-Bayesian game Nature (an external agent) moves first and provides each player i with a type $\theta_i \in \Theta_i$. Each player knows only his type. Subsequently the players simultaneously select their actions. The payoff function of each player now depends on his type, so after all players selected their actions, each player knows his payoff but does not know the payoffs of the other players. Note that given a pre-Bayesian game, every joint type $\theta \in \Theta$ uniquely determines a strategic game, to which we refer below as a θ-game.

A ***strategy*** for player i in a pre-Bayesian game is a function $s_i : \Theta_i \rightarrow A_i$. A strategy $s_i(\cdot)$ for player i is called

- ***best response*** to the joint strategy $s_{-i}(\cdot)$ of the opponents of i if for all $a_i \in A_i$ and $\theta \in \Theta$

$$p_i(s_i(\theta_i), s_{-i}(\theta_{-i}), \theta_i) \geq p_i(a_i, s_{-i}(\theta_{-i}), \theta_i),$$

- ***dominant*** if for all $a \in A$ and $\theta_i \in \Theta_i$

$$p_i(s_i(\theta_i), a_{-i}, \theta_i) \geq p_i(a_i, a_{-i}, \theta_i),$$

Then a joint strategy $s(\cdot)$ is called an ***ex-post equilibrium*** if each $s_i(\cdot)$ is a best response to $s_{-i}(\cdot)$. Alternatively, $s(\cdot) := (s_1(\cdot), \ldots, s_n(\cdot))$ is an ex-post equilibrium if

$$\forall \theta \in \Theta \ \forall i \in \{1, \ldots, n\} \ \forall a_i \in A_i \ p_i(s_i(\theta_i), s_{-i}(\theta_{-i}), \theta_i) \geq p_i(a_i, s_{-i}(\theta_{-i}), \theta_i),$$

where $s_{-i}(\theta_{-i})$ is an abbreviation for the sequence of actions $(s_j(\theta_j))_{j \neq i}$.

So $s(\cdot)$ is an ex-post equilibrium iff for every joint type $\theta \in \Theta$ the sequence of actions $(s_1(\theta_1), \ldots, s_n(\theta_n))$ is a Nash-equilibrium in the corresponding θ-game. Further, $s_i(\cdot)$ is a dominant strategy of player i iff for every type $\theta_i \in \Theta_i$, $s_i(\theta_i)$ is a dominant strategy of player i in every (θ_i, θ_{-i})-game.

We also have the following immediate observation.

Note 52 (Dominant Strategy) *Consider a pre-Bayesian game G. Suppose that $s(\cdot)$ is a joint strategy such that each $s_i(\cdot)$ is a dominant strategy. Then it is an ex-post equilibrium of G.* □

Example 23 As an example of a pre-Bayesian game, suppose that

- $\Theta_1 = \{U, D\}$, $\Theta_2 = \{L, R\}$,
- $A_1 = A_2 = \{F, B\}$,

and consider the pre-Bayesian game uniquely determined by the following four θ-games. Here and below we marked the payoffs in Nash equilibria in these θ-games in bold.

L:

U	F	B
F	**2, 1**	2, 0
B	0, 1	2, 1

R:

U	F	B
F	2, 0	**2, 1**
B	0, 0	2, 1

D		F	B
	F	3, 1	2, 0
	B	**5, 1**	4, 1

	F	B
F	3, 0	2, 1
B	5, 0	**4, 1**

This shows that the strategies $s_1(\cdot)$ and $s_2(\cdot)$ such that

$$s_1(U) := F,\ s_1(D) := B,\ s_2(L) = F,\ s_2(R) = B$$

form here an ex-post equilibrium. □

However, there is a crucial difference between strategic games and pre-Bayesian games. We call a pre-Bayesian game ***finite*** if each set of actions and each set of types is finite. By the ***mixed extension*** of a finite pre-Bayesian game

$$(A_1, \ldots, A_n, \Theta_1, \ldots, \Theta_n, p_1, \ldots, p_n)$$

we mean below the pre-Bayesian game

$$(\Delta A_1, \ldots, \Delta A_n, \Theta_1, \ldots, \Theta_n, p_1, \ldots, p_n).$$

Example 24 Consider the following pre-Bayesian game:

- $\Theta_1 = \{U, B\}$, $\Theta_2 = \{L, R\}$,
- $A_1 = A_2 = \{C, D\}$,

L:

U		C	D
	C	2, 2	0, 0
	D	3, 0	**1, 1**

R:

	C	D
C	2, 1	0, 0
D	3, 0	**1, 2**

B		C	D
	C	**1, 2**	3, 0
	D	0, 0	2, 1

	C	D
C	**1, 1**	3, 0
D	0, 0	2, 2

Even though each θ-game has a Nash equilibrium, they are so 'positioned' that the pre-Bayesian game has no ex-post equilibrium. Even more, if we consider a mixed extension of this game, then the situation does not change. The reason is that no new Nash equilibria are then added to the original θ-games.

Indeed, each of these original θ-games is solved by IESDS and hence by the IESDMS Theorem 33(ii) has a unique Nash equilibrium. This shows that a mixed extension of a finite pre-Bayesian game does not need to have an ex-post equilibrium, which contrasts with the existence of Nash equilibria in mixed extensions of finite strategic games. □

This motivates the introduction of a new notion of an equilibrium. A strategy $s_i(\cdot)$ for player i is called ***safety-level best response*** to the joint strategy $s_{-i}(\cdot)$ of the opponents of i if for all strategies $s'_i(\cdot)$ of player i and all $\theta_i \in \Theta_i$

$$\min_{\theta_{-i} \in \Theta_{-i}} p_i(s_i(\theta_i), s_{-i}(\theta_{-i}), \theta_i) \geq \min_{\theta_{-i} \in \Theta_{-i}} p_i(s'_i(\theta_i), s_{-i}(\theta_{-i}), \theta_i).$$

Then a joint strategy $s(\cdot)$ is called a ***safety-level equilibrium*** if each $s_i(\cdot)$ is a safety-level best response to $s_{-i}(\cdot)$.

The following theorem was established by Monderer and Tennenholz.

Theorem 53 *Every mixed extension of a finite pre-Bayesian game has a safety-level equilibrium.* □

We now relate pre-Bayesian games to mechanism design. To this end we need one more notion. We say that a pre-Bayesian game is of a ***revelation-type*** if $A_i = \Theta_i$ for all $i \in \{1, \ldots, n\}$. So in a revelation-type pre-Bayesian game the strategies of a player are the functions on his set of types. A strategy for player i is called then ***truth-telling*** if it is the identity function $\pi_i(\cdot)$ on Θ_i.

Now mechanism design can be viewed as an instance of the revelation-type pre-Bayesian games. Indeed, we have the following immediate, yet revealing observation.

Theorem 54 *Given a direct mechanism*

$$(D \times \mathbb{R}^n, \Theta_1, \ldots, \Theta_n, u_1, \ldots, u_n, (f, t))$$

associate with it a revelation-type pre-Bayesian game, in which each payoff function p_i is defined by

$$p_i((\theta'_i, \theta_{-i}), \theta_i) := u_i((f, t)(\theta'_i, \theta_{-i}), \theta_i).$$

Then the mechanism is incentive compatible iff in the associated pre-Bayesian game for each player truth-telling is a dominant strategy.

By Groves Theorem 45 we conclude that in the pre-Bayesian game associated with a Groves mechanism, $(\pi_1(\cdot), \ldots, \pi_n(\cdot))$ is a dominant strategy ex-post equilibrium.

Bibliography

[1] K. R. Apt, F. Rossi, and K. B. Venable. Comparing the notions of optimality in CP-nets, strategic games and soft constraints. *Annals of Mathematics and Artificial Intelligence*, 52(1):25–54, 2008.

[2] K. Binmore. *Fun and Games: A Text on Game Theory.* D.C. Heath, 1991.

[3] K. Binmore. *Playing for Real: A Text on Game Theory.* Oxford University Press, Oxford, 2007.

[4] C. Boutilier, R. I. Brafman, C. Domshlak, H. H. Hoos, and D. Poole. CP-nets: A tool for representing and reasoning with conditional ceteris paribus preference statements. *J. Artif. Intell. Res. (JAIR)*, 21:135–191, 2004.

[5] D. Fudenberg and J. Tirole. *Game Theory.* MIT Press, Cambridge, Massachusetts, 1991.

[6] G. Jehle and P. Reny. *Advanced Microeconomic Theory.* Addison Wesley, Reading, Massachusetts, second edition, 2000.

[7] M. Kearns, M. Littman, and S. Singh. Graphical models for game theory. In *Proceedings of the 17th Conference in Uncertainty in Artificial Intelligence (UAI '01)*, pages 253–260. Morgan Kaufmann, 2001.

[8] A. Mas-Collel, M. D. Whinston, and J. R. Green. *Microeconomic Theory.* Oxford University Press, Oxford, 1995.

[9] R. B. Myerson. *Game Theory: Analysis of Conflict.* Harvard University Press, Cambridge, Massachusetts, 1991.

[10] J. F. Nash. Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences, USA*, 36:48–49, 1950.

[11] J. F. Nash. Non-cooperative games. *Annals of Mathematics*, 54:286–295, 1951.

[12] N. Nisan, T. Roughgarden, É. Tardos, and V. J. Vazirani, editors. *Algorithmic Game Theory*. Cambridge University Press, 2007.

[13] M. J. Osborne. *An Introduction to Game Theory*. Oxford University Press, Oxford, 2005.

[14] M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. The MIT Press, Cambridge, Massachusetts, 1994.

[15] H. Peters. *Game Theory: A Multi-Leveled Approach*. Springer, Berlin, 2008.

[16] K. Ritzberger. *Foundations of Non-cooperative Game Theory*. Oxford University Press, Oxford, 2002.

[17] S. L. Roux, P. Lescanne, and R. Vestergaard. Conversion/preference games. *CoRR*, abs/0811.0071, 2008.

[18] Y. Shoham and K. Leyton-Brown. *Essentials of Game Theory: A Concise, Multidisciplinary Introduction*. Morgan and Claypool Publishers, Princeton, 2008.

[19] S. Tijs. *Introduction to Game Theory*. Hindustan Book Agency, Gurgaon, India, 2003.

[20] J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1944.