

Lecture 11: Mechanism Design and Social Choice

*Lecturer: Yishay Mansour**Scribe: Dean Agron , Avi Sedaka , Shai Embon*

11.1 Mechanism Design

11.1.1 Motivation

So far we've seen various types of games, and investigated how they work. In this lecture we'll ask the opposite question: We want to reach a certain result. How should we plan a game that will lead to this result? This area is called **Mechanism Design**. In the second part of the lecture we will show that designing such a mechanism, even under simple assumptions, is sometimes impossible.

At any game, the rules (mechanism, institute) by which the game is taking place have a profound effect on both the way the players (agents) play, and on the decisions eventually taken. Examples:

- Auction – Are the bids given by sealed envelopes, or is it a public auction, with bids given orally? The method of offering the bids effects what participants learn on each other, the bids they give, and eventually the result of the auction.
- Scheduling – The scheduling-policy, effects which jobs are sent, where and when. the players are either machines or jobs.
- Public project (new highway, new library, park, defense system, etc.) – The way we spread the costs of the project across the society, effect the decision of whether the project is undertaken or not.

The aim is to plan a mechanism that guarantees certain objectives, based on the following assumptions:

- The agents have a strategic behavior. That is, they have a benefit function they want to maximize
- The agents have **private information** known only to them, that effect their decision.

For Example, Negotiation between a buyer and a seller.

Each side has a value of the deal (private information)

The seller will claim that his value is higher than it's real value (Increase the price)

The buyer will claim that his value is lower than it's real value (decrease the price)

Sample Questions:

- Plan a mechanism in which the negotiation takes place, that will cause an *efficient* merchandizing. That is, that a successful trade will take place whenever the buyer valuation exceeds the valuation of the seller.
- Is there an efficient mechanism where both the buyer and the seller agree to participate voluntarily?
- Elections – How do we plan fair elections. That is, how do we cause voters to vote by their real preferences, without lying?

Building blocks:

- Usually there is a central institution that governs the game
- Messages (For example, envelopes in auction) – These are the means of interaction with the agents. In the literature, also called types, strategies, or states.
- Payments – Optionally, we can tax or subsidize agents, in order to create incentives for a desired behavior.

Negative results:

We will show that in some cases, it is impossible to achieve all the following goals:

- Efficient result (maximum overall benefit)
- Voluntary participation
- Budget balancing (of tax/subsidy payments)

11.1.2 Formal Model

- Agents – $N = \{1..n\}$
- Decisions - $d \in D$
- Private information for agent i – $\theta_i \in \Theta_i$

- Value function – $v_i : D \times \Theta_i \longrightarrow \mathbb{R}$ – The benefit of player i with private information θ_i from decision $d \in D$
- Preferences – $v_i(d, \theta_i) > v_i(\hat{d}, \theta_i)$ means that agent i prefer decision d over decision \hat{d}

Example 1 – Public project

The society want to decide whether or not to build a certain public project

- Cost of project – c (cost for participant - $\frac{c}{n}$)
- Decision – $D = \{0, 1\}$ (1–do / 0–don't do the project)
- Benefit of agent i from doing the project – θ_i .
- Benefit of agent i –

$$v_i(d, \theta_i) = d(\theta_i - \frac{c}{n}) = \begin{cases} 0 & d=0 \\ \theta_i - \frac{c}{n} & d=1 \end{cases}$$

Example 2 – Allocating a private product

An indivisible good is to be allocated to one member of the society. For instance, an enterprise is to be privatized. We have $N = \{1..n\}$ potential buyers. We want to give the product to one of them.

- Decision – $D = \{1..n\}$
- Benefit of agent i if achieving the product – θ_i
- Value of agent i –

$$v_i(d, \theta_i) = I(d = i) \cdot \theta_i$$

- Later, we will decrease the payment from v_i to establish the Utility function.

11.1.3 Decision rules and efficiency

The society would like to make a decision that maximizes the overall benefit of it's members. For example:

- For the public project, we will decide to build the project if the sum of benefits over all agents is more then it's cost
- For the private product, we might want to give it to the agent that have the maximal benefit from it

We will now phrase it formally.

From now on, we denote by $\Theta = \Theta_1 \times \dots \times \Theta_n$ the joint type space.

Definition – Decision rule

$$d : \Theta \longrightarrow D$$

Definition – Efficiency

A decision d is efficient if it maximizes the overall benefit:

$$\forall \theta \in \Theta, d' \in D : \sum_i v_i(d(\theta), \theta_i) \geq \sum_i v_i(d', \theta_i)$$

For example:

- The public project should be done only if the overall benefit is higher than the cost:
 $d = 1 \Leftrightarrow \sum_i v_i(1, \theta_i) \geq \sum_i v_i(0, \theta_i) \Leftrightarrow \sum_i \theta_i - c \geq 0 \Leftrightarrow \sum_i \theta_i \geq c$
- Private product – The efficient decision is $d = \text{maxarg}_i\{\theta_i\}$

Payments

The definition of efficiency as the basis for decision making has a problem: We base it on a private information, which we don't actually have. We could ask the agents to give us their private value, but the answers will often be dishonest or deceptive.

For instance, in the public project, an agent with $\theta_i < \frac{c}{n}$ has an incentive to underreport his value, and claim he has no profit from the project, and hence try to manipulate the decision to his own benefit. In the same way, an agent with $\theta_i > \frac{c}{n}$ has an incentive to overreport his value (meaning $\theta_i = c$). This could result in wrong decision. Other mechanisms of decision aimed to bypass this problem, such as voting and decide whether to build the project by the majority's vote, could also result in a decision which is not efficient.

The question is, can we create incentives for the agents to reveal their real private information? Many times the answer is yes. We can balance their interests by putting taxes (to reduce the profit) or by subsidizing agents (to reduce the loss). That is done by a **payment function**, which indicates how much each agent receives: for player i ,

$$t_i : \Theta_i \longrightarrow \mathbb{R}$$

For all the players,

$$t : \Theta \longrightarrow \mathbb{R}^n$$

Definition – Social Choice Function

A social choice function is a pair of decision and payment:

$$f(\theta) = (d(\theta), t(\theta))$$

From now on, we will assume $\hat{\theta}$ to be the value declared by the agents, while θ is their true private value. Decision - $d(\hat{\theta})$ Payment - $t(\hat{\theta})$

The utility that agent i receives is based on it's benefit from the decision taken, plus the payment he receives from the mechanism (both are based on it's declared value):

$$u_i(\hat{\theta}, \theta_i, d, t) = v_i(d(\hat{\theta}), \theta_i) + t_i(\hat{\theta})$$

A utility function at this form is called Quasi-linear.

Definition – Balanced and Feasible payments

The payments are:

- Feasible, if $\forall \theta : \sum_i t(\theta_i) \leq 0$
- Balanced, if $\forall \theta : \sum_i t(\theta_i) = 0$

We want the payments to be balanced. If the payments are feasible and not balanced, we have a surplus we have to waste, or return to outside source. We can not return the change to the society, because the only way to do it is by changing the payment functions.

Balance is an important property if we wish (d, t) to be efficient rather than just d . If payments are feasible and not balanced, then there is some loss to the society relative to an efficient decision with no transfers.

Definition – Mechanism (M, g)

- Messages of agents: $(m_1, \dots, m_n) \in M$, while $M = M_1 \times \dots \times M_n$

- Outcome function: The outcome function $g : M \longrightarrow D \times \mathbb{R}^n$ is defined as $g(m) = (g_d(m), g_{t,1}(m), \dots, g_{t,n}(m))$

While:

- $g_d : M \longrightarrow D$ is the decision function
- $g_{t,i} : M \longrightarrow \mathbb{R}$ is the transfer payment function of agent i

For example, First price auction:

- Benefit of agent i , if he achieves the product, is m_i
- Outcome:

$$g_d(m) = \max \arg_i \{m_i\}$$

$$g_{t,i}(m) = \begin{cases} 0, & \text{No win} \\ -m_i, & \text{Win} \end{cases}$$

Mechanism Design and Dominant strategies

Definition – Dominant strategy

A strategy $m_i \in M_i$ is a *dominant strategy* for $\theta_i \in \Theta_i$ if it is superior to all other player strategies, regardless of other players strategies: formally,

$$\forall m_i, m'_i \in M : v_i(g_d(m_{-i}, m_i), \theta_i) + g_{t,i}(m_{-i}, m_i) \geq v_i(g_d(m_{-i}, m'_i), \theta_i) + g_{t,i}(m_{-i}, m'_i)$$

For example: Confess at the prisoner dilemma.

A dominant strategy is a compelling property of a mechanism, if it exists. It allows us a better prediction of which strategies will be employed by the agents. However, because it is such a strong property, it exists only at narrow space of problems.

Definition

A Social Choice Function $f = (d, t)$ is *implemented in dominant strategies* by a mechanism (M, g) if:

- There are functions $m_i : \Theta_i \longrightarrow M_i$ such that for any agent i with strategy $\theta_i \in \Theta$, $m_i(\theta_i)$ is a dominant strategy
- $\forall \theta \in \Theta : g(m(\theta)) = f(\theta)$

11.1.4 Direct Mechanism and the Revelation theorem

We'll now show that if there is a dominant strategy, there is no need to use the complicated (M, g) mechanism. There is an equivalent and much simpler method, called *Direct Mechanism (DM)*

Definition

A social choice function f can be viewed as a mechanism, where $M_i = \Theta_i$, hence $m_i(\theta_i) = \theta_i$, and $g = f$. That is called *Direct Mechanism (DM)*

Definition

A DM $f(d, t)$ is *dominant strategy incentive compatible (DS - IC)* if for each agent i with type $\theta_i \in \Theta_i$, $\theta_i \in M_i$ is a dominant strategy for $\theta_i \in \Theta_i$

A social choice function f is *strategy proof* if it is dominant strategy incentive compatible.

In other words, saying that DM is dominant strategy incentive compatible means that telling the truth is the dominant strategy.

Theorem 11.1 *Revelation principle for dominant strategies*

If a mechanism (M, g) implements a social choice function $f = (d, t)$ in dominant strategies, then the DM f is dominant strategy incentive compatible

Proof:

Since (M, g) implements f in dominant strategies, then there are $m(\theta) = (m_1(\theta), \dots, m_n(\theta))$ such that:

$$\forall \theta : g(m(\theta)) = f(\theta)$$

And also:

$$\forall m_i, \hat{m}_i = m_i(\theta_i)$$

Because $m_i(\theta_i)$ is a dominant strategy, then so is θ_i □

Negative Results:

We will see later that there is no voting method between 3 or more candidates that guarantees that the voters will tell their true preferences.

Theorem 11.2 – *Grove method*

If:

- d is an efficient decision rule

- There is a function $h_i : \Theta_{-i} \rightarrow \mathbb{R}$ such that:

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

then (d, t) is dominant strategy incentive compatible

Remark: Under some additional conditions, this is the only way to define dominant strategy incentive compatible.

Proof:

By way of contradiction, suppose d is an efficient decision rule, and that for each agent i there are h_i that define together $t : \Theta \rightarrow \mathbb{R}^n$, But (d, t) is not dominant strategy incentive compatible.

Then, there is an agent i , and states θ and θ'_i such that it is more beneficial to agent i to report θ'_i rather than θ_i . Namely,

$$v_i(d(\theta_{-i}, \theta'_i), \theta_i) + t_i(\theta_{-i}, \theta'_i) > v_i(d(\theta), \theta_i) + t_i(\theta)$$

We'll expand t_i explicitly, and write d' instead of $d(\theta_{-i}, \theta'_i)$ for simplicity:

$$v_i(d', \theta_i) + h_i(\theta_{-i}) + \sum_{j \neq i} v_j(d', \theta_j) > v_i(d(\theta), \theta_i) + h_i(\theta_{-i}) + \sum_{j \neq i} v_j(d(\theta), \theta_j)$$

Therefore:

$$\sum_{j=1}^n v_j(d', \theta_j) > \sum_{j=1}^n v_j(d(\theta_{-i}, \theta_i), \theta_j)$$

And d' contradicts the efficiency of $d(\theta)$. The conclusion is that (d, t) is dominant strategy incentive compatible. □

Clark Mechanism

Clark suggested a specific implementation of h_i

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

which is actually the decision that would have been taken if agent i had not participated.

Therefore, t_i is defined as follows:

$$t_i = \sum_{j \neq i} v_j(d, \theta_j) - \max_{d \in D} \left\{ \sum_{j \neq i} v_j(d, \theta_j) \right\}$$

Properties of Clark mechanism:

- The price is non-positive. That means that the agents pay to the mechanism, and assures *feasible payments*.
- If agent i does not effect the decision, then $t_i(\theta) = 0$.
- $t_i(\theta)$ can be thought of as the marginal loss for the other agents.

Example 1 – Second price auction

We give the products to the agent with the highest bid. How much should we charge for it?

- $d(\theta) = \text{maxarg}\{\theta_i\}$
- $v_i(d, \theta_i) = \begin{cases} 0 & d \neq i \\ \theta_i & d = i \end{cases}$
- For the agent i with the highest value:

$$\begin{aligned} - \sum_{j \neq i} v_j(d(\theta), \theta_j) &= 0 \\ - h_i &= -v_{2^{nd}}(d_{2^{nd}}, \theta_{2^{nd}}) = -\theta_{2^{nd}} \\ - \Rightarrow t_i(\theta) &= -\theta_{2^{nd}} \end{aligned}$$

- The participation of all other agents does not effect the decision, i.e. $t_j(\theta) = 0$

Explanation: The winner is the agent with the highest bid, and he pays the 2nd highest offer. No other payments are done.

Example 2 – Building graph pathes

Given a graph $G = (V, E)$ with weights \vec{c} on the each $e \in E$.

- Each $e \in E$ is an agent
- The cost of each agent is c_e (benefit is $-c_e$)
- The objective: to choose a minimal cost path
- $t_e(\vec{c}) = (\text{Cost of shortest path} - c_e) - (\text{shortest path without } e)$
 $t_e(\vec{c}) = (\text{shortest path without } e) - (\text{cost of shortest path with } e)$

Budget balancing vs. Incentive Compatible

It is not enough to require efficiency of decision. The society should also aspire to balanced payments. Transfer functions that are not balanced, cause waste, and it can be considerable. An example for it can be seen in the public project case:

Consider two agents. Project cost is $c = \frac{3}{2}$, and $\Theta_1 = \Theta_2 = \mathbb{R}$.

According to Clark, there are h_1 and h_2 . We'll check what we can conclude from the feasibility of the payments.

1. Assume $\theta_1 = \theta_2 = 1$,
 $d(1, 1) = 1$, and therefore:

$$\begin{aligned} 0 &\geq t_1(1, 1) + t_2(1, 1) = h_1(1) + 1 + h_2(1) + 1 \\ &\Rightarrow -2 \geq h_1(1) + h_2(1) \end{aligned}$$

2. Assume $\theta_1 = \theta_2 = 0$,
 $d(0, 0) = 0$, and therefore:

$$0 \geq t_1(0, 0) + t_2(0, 0) \geq h_1(0) + h_2(0)$$

From both inequalities we can conclude that there must be either $-1 \geq h_1(1) + h_2(0)$ or $-1 \geq h_1(0) + h_2(1)$. Since $d(1, 0) = d(0, 1) = 0$, in one of the cases $(0, 1)$ or $(1, 0)$ the payments are negative. That means that: a. to have a dominant strategy incentive compatible mechanism with efficient decision rule, one cannot satisfy balance. b. In some cases, there are payments without a project.

11.2 Social Choice

Social choice is the general problem of mapping a set of multiple individual preferences to a single social preference, that will best reflect the aggregate individual preferences. Common examples for such scenarios are public elections, voting, etc. While in mechanism design we have shown how to select a mechanism that will exhibit desired behavior, for social choice we shall show that given very simple and reasonable requirements, no such mechanism exists. (An important difference is that there are no payments to the players.)

We shall show two theorems for two slightly different social choice scenarios: Arrow's Impossibility Theorem and the Gibbard-Satterthwaite Theorem.

11.2.1 Arrow's Impossibility Theorem

Arrow's Impossibility Theorem deals with social ordering. Given a set of alternatives $\mathcal{A} = \{A, B, C, \dots\}$, a *transitive preference* is a ranking of the alternatives from top to bottom, with ties allowed. Given a set of individuals (a society, so to speak), a *social preference function* is a function associating any tuple of personal transitive preferences, one per individual, with a single transitive preference called the *social preference*.

Definition A **Transitive Preference** is an ordering of \mathcal{A} with ties allowed.

For example: $A > B > C = D = E > F > G$.

Definition Given alternatives \mathcal{A} and a set of individuals \mathcal{N} , a **Social Profile** is an association of a transitive preference per individual.

We will typically denote individual by numbers $1 \dots N$, in which case a social profile will be represented by an N -tuple of transitive preferences. We will also represent a profile by a matrix, where each column is the transitive preference of single individual, ranked from top to bottom. For example:

1	2	3	4	5	6
A	A	B,C	C	A,B,D	A,B,C,D
C	B	-	A,B	-	-
B	D	A	-	-	-
D	C	D	D	C	-

Definition A **Social Preference Function** is a function associating each profile with a transitive preference, called the social preference.

1	2	3	4	5	6	→	Social
A	A	B,C	C	A,B,D	A,B,C,D		A
C	B	-	A,B	-	-		C
B	D	A	-	-	-		B,D
D	C	D	D	C	-		-

Definition A social preference function respects **Unanimity** if the social preference strictly prefers α over β whenever all of the individuals strictly prefer α over β .

Definition A social preference function respects **I.I.A. (Independence of Irrelevant Alternatives)** if the relative social ranking (higher, lower or indifferent) of α and β depends only on their relative ranking in the profile (and not on the ranking of other alternatives).

Definition Given a social preference function, an individual n is a **Dictator** if the social preference strictly prefers α over β whenever n strictly prefers α over β . If a dictator exists, the social preference function is a **Dictatorship**.

Unless specifically stated, relationships (such as prefer, above, first, last, etc) are non-strict. We are now ready to present Arrow's Impossibility Theorem.

Theorem 11.3 Arrow's Impossibility Theorem For 3 or more alternatives, any social preference function that respects transitivity, unanimity and I.I.A., is a dictatorship.

Proof: Assume a social preference function meeting the conditions set in the theorem. Let $B \in \mathcal{A}$ be chosen arbitrarily.

Claim 11.4 For any profile where B is always either strictly first or strictly last in the ranking (for all individuals), the social preference must place B either strictly first or strictly last as well.

Proof: Assume to the contrary that the social preference does not place B strictly first or strictly last. Then there exist two alternatives A and C (different from each other and B) such that in the social preference, $A \geq B$ and $B \geq C$. Since all individuals place B strictly first or strictly last, moving C strictly above A for an individual is possible without changing the relative preference between A and B or B and C for this individual. This is depicted in Figure 11.1. Due to I.I.A., we conclude that moving C strictly above A for all individuals should not change the relative social preference between A and B , and between B and C , so we still have $A \geq B$ and $B \geq C$, which implies $A \geq C$ due to transitivity. But this contradicts unanimity, because now all individuals strictly prefer C over A .

Therefore, the social preference must place B strictly first or strictly last. \square




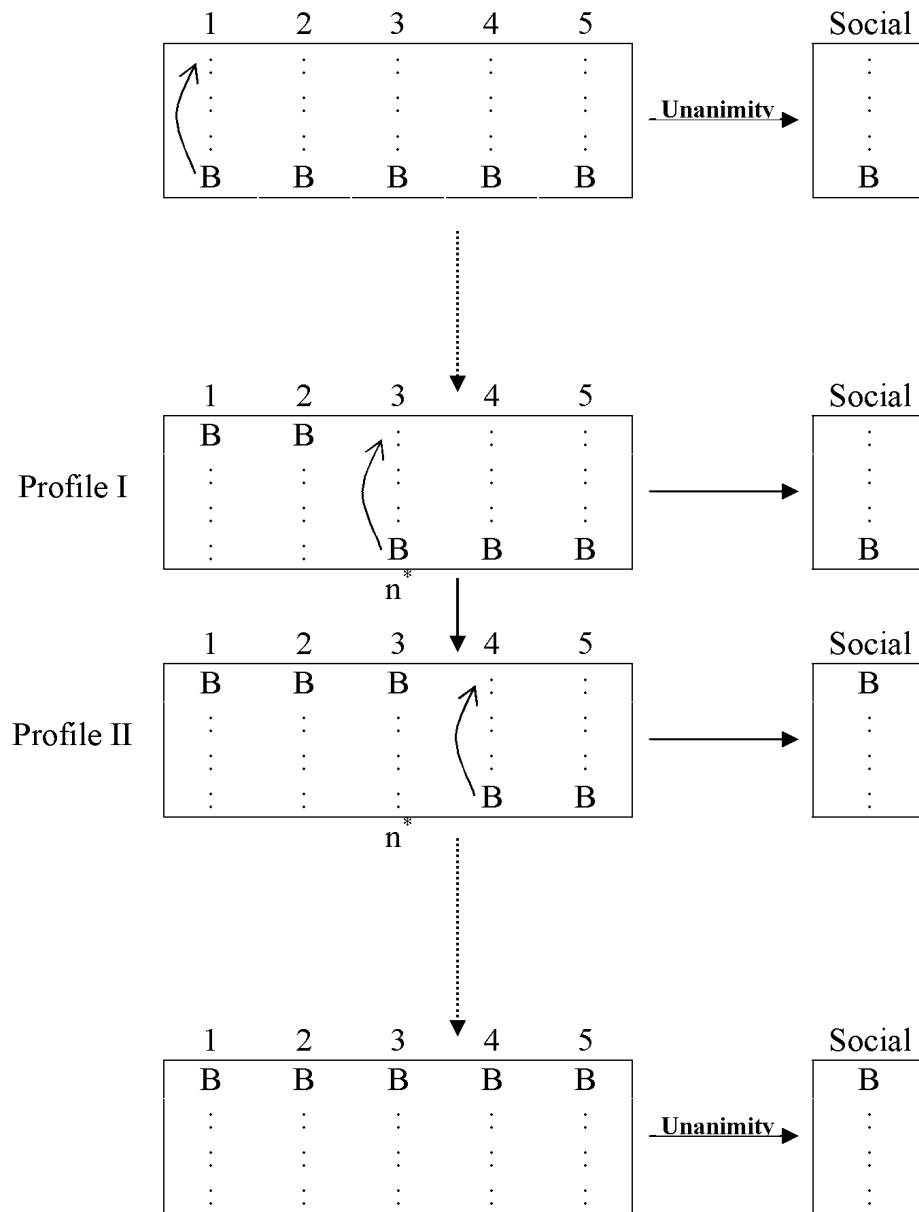
Figure 11.1: Moving C strictly above A without changing relative preference to B

Claim 11.5 *There exists an individual n^* and a specific profile such that n^* can swing B from the strictly last position in the social preference to the strictly first position by changing his preference.*

Proof: We observe an arbitrary profile where B is strictly last for all individuals. Due to unanimity, B must be strictly last in the social preference. Now let the individuals, from 1 to N , move B from the strictly last position to the strictly first position successively. Due to the previous claim, in any stage B must be strictly first or strictly last in the social preference. Because it starts strictly last, and must end strictly first, there must be an individual whose change causes B to move from the former position to the latter. We denote this individual by n^* . Denote by profile I the profile just before n^* changes his preference, and by profile II the profile just after the change. Profile I is the profile mentioned in the claim, and n^* is the individual. This is depicted in Figure 11.2. □

Note that n^* will have this behavior for any profile where all individuals $i < n^*$ place B strictly first and all individuals $i \geq n^*$ place B strictly last. The reason is that the (strict) relative preferences between B and any other alternative in such a profile and in profile I are identical, so this must hold in the social preference, and thus B must still be strictly last in any such profile. The same is true for the changed profile and profile II, where B must be strictly first. Therefore the choice of n^* is only dependent on B , and the order of the players. Later we will show that n^* is also independent of the order of the players. Therefore we can denote $n^* = n(B)$.

Claim 11.6 *n^* is a dictator for any pair of alternatives A and C that does not include B .*

Figure 11.2: Existence of n^*




Figure 11.3: Profile IV

Proof: Recall that by I.I.A., the social preference between A and C depend only on the individuals preference between A and C. Given any profile III where for n^* , $A > C$, create profile IV from profiles II and III by:

1. Start with profile II
2. Have n^* move A strictly above B, without changing the relative preferences among all alternatives other than A.
3. Have all other individuals change the preference between A and C to be identical to profile III, while B remains in it's profile II position.

In profile IV, the relative preference between A and B is identical to their relative preference in profile I for all individuals, and due to I.I.A. we must have $A > B$ in the social preference for profile IV (because in profile I , B is strictly last). The relative preference between C and B is in profile IV is identical to that of profile II for all individuals, thus (I.I.A) we must have $B > C$ in the social preference for profile IV (because in profile II B is strictly first). This is depicted in figure 11.3.

Therefore we must have $A > C$ in the social preference for profile IV. But the relative preferences between A and C in profiles III and IV are identical, so we must also have $A > C$ in the social preference for profile III. This is true for any profile III with n^* strictly preferring A over C, thus n^* is a dictator for A and C. □

Claim 11.7 n^* is a dictator for any pair of alternatives A and B.

Choose a third alternative C . By the same construction above, there exists $n(C)$ that is a dictator for any pair of alternatives exclusive of C , such as A and B . But $n(B)$ definitely effects the relative social preference of A and B , so he is the only possible dictator for A and B , thus $n(C) = n(B)$. Therefore $n(B)$ is also a dictator for A and B .

We have shown that there is a single individual that is dictator for any pair of alternatives, thus the social preference function is a dictatorship. \square

Claim 11.8 *The identity of n^* is independent of the order of the individuals.*

Proof: For contradiction assume that using a different order we get $k^* \neq n^*$. Consider a profile where k^* has $A > C$ and n^* has $A < C$. By Claim 11.6 n^* (or k^*) is a dictator for A and C . But only one can hold. Therefore $n^* = k^*$. \square

11.2.2 Gibbard-Satterthwaite Theorem

We shall now deal with an even simpler problem. The general scenario is similar to the one described previously. The difference is that we will only be interested in a single "most-desired" alternative, instead of an entire social preference. Instead of looking for a social preference function, we are looking for an election mechanism.

Definition An **Election Mechanism** is a function mapping each social profile to a single alternative (the **elected alternative**).

Definition An election mechanism M that decides an election is **Unanimous** if it elects alternative A whenever all individuals rate A as strictly first.

Definition A mechanism M that decides an election is defined to be a **strategy proof** one, when:

The dominant strategy of each voter is voting in the order of his real preferences ("telling the truth"). Namely, if the voter prefers candidate A over B , his dominant strategy will be to rank A above B . In other words, it is worthy for every voter to "tell the truth".

Definition A mechanism M that decides an election is defined to be a **dictatorial** one, when:

There exist a dictator, namely a voter v such that if v votes for candidate A , then A will win the election regardless of the other voters' votes.

Definition A profile is defined as a set that includes the preferences of all voters.

We are now ready to present the central theorem of this section.

Theorem 11.9 (Gibbard-Satterthwaite Theorem) *An election mechanism for 3 or more alternatives which is:*

- *Unanimous*
 - *Strategy proof*
- is a dictatorship.*

This theorem will also be referred to as the **GS Theorem**. We precede the proof with a few lemmas.

Lemma 11.10 *Irrelevance Lemma*

Suppose that an alternative A is selected by the mechanism given a profile P . Then a modification of P which raises the ranking of an alternative X in the preference ranking of a single voter i , will cause either A or X to be selected by the mechanism.

Proof:

Suppose an alternative C , $C \neq X$, $C \neq A$ is elected. In P , A was elected, and the fact that i raised its vote for X caused C to be elected. There exist two cases:

1. If i prefers A to C , i would not raise his vote for X , even if the higher ranking of X is the truth for him, since then he causes C to be voted instead of A , in contradiction to strategy proof.
2. If i prefers C to A , i would never had reported his (maybe real) Low vote for X . He could gain profit from not reporting the truth, giving X a high vote, and thus C would be voted. Again, contradiction to strategy proof.

□

Lemma 11.11 *Unanimous last Lemma*

If an alternative B is ranked last by all voters, B is not elected.

Proof:

Assume B is elected. Then, for an alternative A , $A \neq B$, suppose that every voter, one at a time, raises his vote for A to the top of their priority. Then, by strategy proof, B is still elected, because otherwise for a voter interested in B to be elected, it would be profitable to him to vote for B last, thus not reporting the truth, and contradicting strategy proof.

But after all voters had raised their votes for A to the top, A must be elected, because of unanimity. Contradiction, thus the assumption that B was elected is not true, meaning B is not elected.

□

Proof of the GS Theorem:

Step 1

Begin with a profile P which is arbitrary in any sense, besides the ranking of B in it—every voter in P ranks B as his last vote. Thus, B is not elected, by "Unanimous last"

Lemma. One voter at a time (arrange the voters in some order), have B "jump" from the bottom to the top of the voter's preferences. At the end of this process, all voters vote for B first, and thus B is elected (unanimity). So, let r be the first voter for which the change of his vote will cause B to be elected. r is called the **pivot for B** .

In all of the following tables, '?' stands for unknown, '...' means the sequence written before and after the dots (in the same line) is kept in between them.

Profile 1

1	2	.	.	r	.	.	n
B	.	.	B	K	?	?	?
?	?	?
?	?	?
?	.	.	?	B	B	.	B

$X \neq B$ is elected.

Profile 2

1	2	.	.	r	.	.	n
B	.	.	B	B	?	.	?
?	?	?
?	?
?	.	.	.	?	B	.	B

B is elected.

Consider profile 2:

1. **If any voter $i > r$ changes his vote, B is still elected.** Otherwise i , who does not want B to be elected, would change his vote to prevent it, in contradiction to strategy proof.

2. **If any voter $i \leq r$ keeps B ranked first, and changes the ranking of the other alternatives, B is still elected.** Otherwise i , who prefers B to be elected, would prefer to submit his original ranking even if his real preferences is represented by the modified one. Once again, contradiction to strategy proof.

Conclusion 1: B is elected if the first r voters rank B first.

Consider profile 1:

1. **If any voter $i < r$ changes his vote, B is still not elected.** Otherwise it is profitable for i to do this change, such that B will be elected. Contradiction to strategy proof.

2. **If any voter $i \geq r$ keeps B ranked last, and changes the ranking of the other alternatives, B is still not elected.**

Otherwise it is profitable for i to submit his original ranking (to prevent B from being elected) even if his real preferences are represented by the modified one. Contradiction to strategy proof.

Conclusion 2: B is not elected if the voters r through n rank B last.

We will show that the voter r is a dictator.

Step 2

profile 3

1	2	.	.	r	.	.	n
?	.	.	?	K	?	?	?
?	?	?
?	?	?
B	B	.	.	.	B	.	B

Raise K in profile 3 to the top position for all voters. K is now chosen, by unanimity. Now raise B to the top positions for voters $1 \dots r - 1$:

profile 4

1	2	.	r-1	r	.	.	n
B	.	.	B	K	.	.	K
K	K	.	K	?	.	.	?
.
?	?	.	?	B	.	.	B

Since K was chosen in profile 3, and the modifications from profile 3 to 4 are only raises K or B , it is inferred from the Irrelevance Lemma above that either K or B are chosen in profile 4. But in profile 4 B is not chosen, by Conclusion 2 above (r through n ranked B last). Thus **K is chosen in profile 4.**

Now raise B to the second position for the voter r :

profile 5

1	2	.	r-1	r	r+1	.	n
B	.	.	B	K	K	.	K
K	K	.	K	B	?	.	?
?	?
?	?	.	?	?	B	.	B

K is still elected, since otherwise r would not have reported the change even if this is a change in his real preferences, in contradiction to strategy proof.

Reconsider profile 3.

Lemma 11.12 *In profile 3, K is elected.*

Proof:

Start with profile 3, and assume $G \neq K$ is elected. Raise B to the top position for the first $r - 1$ voters. By Conclusion 2, B is not elected, and by the Irrelevance Lemma, G is still elected.

Now raise B to the second position in the voter r 's ranking.

profile 6

1	2	.	r-1	r	r+1	.	n
B	.	.	B	K	?	.	?
?	?	.	?	B	?	.	?
?	?
?	?	.	?	?	B	.	B

B is elected in profile 6.

Assume B is not elected in profile 6. By the Irrelevance Lemma, G is still elected. By Conclusion 1, if we raise alternative B in the vote of r one step up, to the top position, then B is elected. In profile 6, r prefers B over $G \neq B$, so he would profit from falsely reporting B above K , in contradiction to strategy proof.

Now, in profile 6, raise K to the second place in the votes of voters 1 through $r - 1$, and to the top position for voters $r + 1$ through n .

Profile 7

1	2	.	r-1	r	r+1	.	n
B	.	.	B	K	K	.	K
K	.	.	K	B	?	.	?
?	?
?	?	.	?	?	B	.	B

B is still elected in profile 7:

Assume B is not elected.

Voters 1 through $r - 1$ who want B to be elected will profit by not reporting the change (even if it is truthful). Voters $r + 1$ through n who want B not to be elected will profit by falsely reporting the change. Thus B is elected in profile 7. But profile 7 = profile 5, and we proved above that $K \neq B$, is elected in profile 5, in contradiction. Thus, the assumption that $G \neq K$ is elected in profile 3 is not true. Meaning, **K is elected in profile 3.**

□

Step 3

Lemma 11.13 Consider an arbitrary profile P where r ranks some alternative $K \neq B$ on top. Then K or B is elected.

Proof:

First, modify the profile by moving B to the bottom for all voters. We get profile 3, and we showed that K is elected. Now, one voter at a time, restore the profile by raising B to its original position. By the Irrelevance Lemma, **either K or B is elected.**

□

Lemma 11.14 Consider an arbitrary profile P where r ranks some alternative $K \neq B$ on top. Then K is elected.

Proof:

Consider:

Profile 8

1	2	.	r-1	r	r+1	.	n
B	.	.	B	B	A	.	A
?	?	.	?	?	?	.	?
?	?	.	?	?	?	.	?
C	C	C

where $C \neq B$ and $C \neq K$. Again, similarly to step 1, have C jump in the ranking of the voters one by one, starting from voter 1 until a pivot m , the first for whom the election will become C is found. Symmetrically to step 2, the top choice of m is selected in profile 8. But from Conclusion 1, we know that B is chosen in profile 8. Meaning, the top choice of m in profile 8 is B , meaning $m \leq r$. But a symmetric argument, starting with m and then finding r (Everything is done exactly the same, replacing m with r and B with C), will lead to the conclusion that $r \leq m$, and so $m = r$. So r , the pivot in respect to B , is also the pivot in respect to C . Using Lemma 10.10 for C instead of B , we obtain that K or C are elected in P . Thus, In P :

K or C are elected.

K or B are elected.

And we obtain:

K is elected in P .

□

For $K = B$, similar arguments show that r is a pivot for A , as well as C , and that B is elected. Hence, for each K , $K \neq B$ or $K = B$, if r ranks some alternative K on top, Then K is elected. Hence, r is a dictator.

□