

A brief introduction to Mechanism Design Theory

Antonio Nicolò

Scuola Galileiana - Università degli Studi di Padova

19-20 Marzo 2012

Example 1

A society is deciding on whether or not to build a public project at a fixed cost c . For example the project might be a public swimming pool, a public library, a park, a defense system, or any of many public goods. The cost of the public project has to be equally divided among members of the society. Here the set of decisions is $A = \{0, 1\}$ where $a = 0$ representing not building the project and $a = 1$ representing building the project. For each agent $i \in N$, let $\alpha_i \in \mathbb{R}_+$ be the (monetary) *valuation* that i assigns to the public good if it is produced. By requiring that α_i be independent of the set of users we are implicitly assuming that there is no rivalry in the consumption of the public good. A *profile* $\alpha = (\alpha_i)_{i \in N} \in \mathbb{R}_+^N$ is a vector of valuations, one for each agent. Agent i 's preferences on the set of allocations A depend on i 's valuation $\alpha_i \in \mathbb{R}_+$ and are represented by the utility function $u_i : A \times \mathbb{R}_+ \rightarrow \mathbb{R}$, where

$$u_i(a, \alpha_i) = a\left(\alpha_i - \frac{c}{n}\right).$$

Example 2

A public good is *excludable* whenever a subset of agents (called *non-users*) can be excluded from its use, even when $a = 1$. The set of agents that are not excluded are called *users*. A generic subset of users will be denoted by S . For each agent $i \in N$, let $\alpha_i \in \mathbb{R}_+$ be the (monetary) *valuation* that i assigns to the public good if it is produced and i is a user. By requiring that α_i be independent of the set of users we are implicitly assuming that there is no rivalry in the consumption of the public good. A *profile* $\alpha = (\alpha_i)_{i \in N} \in \mathbb{R}_+^N$ is a vector of valuations, one for each agent. For each subset of agents $S \subseteq N$, let $\mathbf{1}_S : N \rightarrow \{0, 1\}$ be the indicator function where for all $i \in N$,

$$\mathbf{1}_S^i = \begin{cases} 1 & \text{if } i \in S \\ 0 & \text{if } i \notin S. \end{cases}$$

Example 2 cont.

Let $p = (p_i)_{i \in N} \in \mathbb{R}^N$ be a vector of *prices* (or contributions). The set of agents N has to decide whether or not to provide the public good ($x \in X$), its set of users $S \in 2^N$, and the vector of contributions $p \in \mathbb{R}^N$. An *allocation* is a triple $(x, S, p) \in X \times 2^N \times \mathbb{R}^N$ with the property that $x = 0$ implies $S = \emptyset$. Denote by $A \equiv \{(x, S, p) \in X \times 2^N \times \mathbb{R}^N \mid x = 0 \text{ implies } S = \emptyset\}$ the set of all alternatives. Agent i 's preferences on the set A depend on i 's valuation $\alpha_i \in \mathbb{R}_+$ and are represented by the utility function $v_i : A \times \mathbb{R}_+ \rightarrow \mathbb{R}$, where for each $(x, S, p, \alpha_i) \in A \times \mathbb{R}_+$,

$$u_i(x, S, p, \alpha_i) = \mathbf{1}_S^i \cdot x \cdot \alpha_i - p_i.$$

Example 3

An indivisible good is to be allocated to one member of society. For instance, the rights to an exclusive license are to be allocated. Here

$A = \left\{ a \in \{0, 1\}^n \text{ s.t. } \sum_{i=1}^n a_i = 1 \right\}$, where $a_i = 1$ indicates that the good

has been allocated to agent i . Let $\alpha_i \in \mathbb{R}_+$ be the (monetary) *valuation* that i assigns to the object. Society here has to decide how to allocate the object. If the object is allocated by means of a first-price auction then the utility of the winner is $v_i(a, \alpha_i, b_i) = a_i \alpha_i - b_i$ where $b_i \geq 0$ is agent i 's bid and the utility of the other agents is zero.

Example 4

A list of n patients is waiting for a kidney transplant. Let $\Omega = \{\omega_1, \dots, \omega_k\}$ denote the set of kidneys available for transplantation. To describe the situation of a patient who remains in dialysis, we denote by ω_0 the null kidney. We say that a patient who remains in dialysis is matched with ω_0 .

The probability of graft survival when kidney $\omega \in \Omega$ is transplanted to patient i depends on objective medical criteria. Thus, the fitness of each available kidney to each patient is directly observable by doctors. We denote by $p_i(\omega) \in [0, 1)$ the probability of graft survival when kidney ω is transplanted to patient i .

We say that patient i and kidney ω are **incompatible**, if $p_i(\omega) = 0$. We say that patient i and kidney ω are **compatible** if $p_i(\omega) > 0$. The **probability of graft survival matrix** $\mathbf{P} \in \mathbb{M}_{n \times n}$ is defined by $\mathbf{P}_{i,j} \equiv p_j(\omega_i)$ for each $i, j \in N$. The matrix \mathbf{P} contains all the relevant medical information about patients and donors that is observable by doctors.

Example 4 cont.

Each patient i is equipped with a preference –a complete, reflexive, and transitive binary relation \succsim_i on $\Omega \cup \{\omega_0\}$. Patients' preferences over kidneys are based on the information contained in the probability of graft survival matrix \mathbf{P} . Specifically, for each patient i , the preference \succsim_i is consistent with \mathbf{P} if i) for each pair $\omega, \omega' \in \Omega$

$p_i(\omega) \geq p_i(\omega') \Leftrightarrow \omega \succsim_i \omega'$; ii) for each ω such that $p_i(\omega) = 0$,

$\omega_0 \succ_i \omega$. By (i), when comparing two kidneys, every patient prefers the kidney that yields the highest probability of graft survival. By (ii), every patient prefers to remain in dialysis rather than receiving an incompatible kidney.

Given a probability of graft survival matrix \mathbf{P} , $\mathcal{R}_i^{\mathbf{P}}$ is the set of preferences which are consistent with \mathbf{P} for patient i . Note that patients' preferences are not completely determined by \mathbf{P} . Specifically, \mathbf{P} does not determine the rank of ω_0 in each patient's preference. Hence, the minimum probability of graft survival such that a patient prefers to undergo a transplant rather than remaining in dialysis is private information of the patient.

A solution of this problem is an **assignment** $a = [(1, \omega), \dots, (n, \omega')]$ such that for each $i, j \in N$, $i \neq j$ and each $\omega, \omega' \in \Omega$, if $(i, \omega), (j, \omega') \in a$, then $\omega \neq \omega'$. An assignment is an allocation of the available kidneys to the patients. which never allocates the same kidney to two different patients, unless that kidney is the null kidney.

In this problem the society has to decide how to allocate available kidneys to the patients.

- All previous examples have three common features. The society

- All previous examples have three common features. The society
- *has to take a decision which affects all the individuals; to this end, it*

- All previous examples have three common features. The society
- *has to take a decision which affects all the individuals*; to this end, it
- *has to decide which is a correct decision to take*. Finally, in order to take the correct decision, the society

- All previous examples have three common features. The society
- *has to take a decision which affects all the individuals*; to this end, it
- *has to decide which is a correct decision to take*. Finally, in order to take the correct decision, the society
- *has to elicit some piece of information from its members, which are private information*.

- All previous examples have three common features. The society
- *has to take a decision which affects all the individuals*; to this end, it
- *has to decide which is a correct decision to take*. Finally, in order to take the correct decision, the society
- *has to elicit some piece of information* from its members, which are private information.
- Individuals have to collectively take a decision and therefore to aggregate their preferences by means of a procedure which has to satisfy some normative requirement in order to be "correct".

- All previous examples have three common features. The society
- *has to take a decision which affects all the individuals*; to this end, it
- *has to decide which is a correct decision to take*. Finally, in order to take the correct decision, the society
- *has to elicit some piece of information* from its members, which are private information.
- Individuals have to collectively take a decision and therefore to aggregate their preferences by means of a procedure which has to satisfy some normative requirement in order to be "correct".
- For instance in case of Example 1 (a pure binary public good) a society may desire to always choose an *efficient* alternative. In case of a public good an efficient decision is such that $a = 1$ if and only if $\sum_{i=1}^n \alpha_i \geq c$, that is the public good is built if and only if the sum of monetary valuations of all individuals is equal or larger than the cost of the public good.

- All previous examples have three common features. The society
- *has to take a decision which affects all the individuals*; to this end, it
- *has to decide which is a correct decision to take*. Finally, in order to take the correct decision, the society
- *has to elicit some piece of information* from its members, which are private information.
- Individuals have to collectively take a decision and therefore to aggregate their preferences by means of a procedure which has to satisfy some normative requirement in order to be "correct".
- For instance in case of Example 1 (a pure binary public good) a society may desire to always choose an *efficient* alternative. In case of a public good an efficient decision is such that $a = 1$ if and only if $\sum_{i=1}^n \alpha_i \geq c$, that is the public good is built if and only if the sum of monetary valuations of all individuals is equal or larger than the cost of the public good.
- Therefore in order to select an efficient alternative the decision maker has to elicit each agent's monetary valuation α_i ; which is private information.

Agents, Alternatives and Preferences

- Let consider a society formed by N agents with $N = \{1, 2, \dots, n\}$ and $n \geq 2$.

Agents, Alternatives and Preferences

- Let consider a society formed by N agents with $N = \{1, 2, \dots, n\}$ and $n \geq 2$.
- Agents have to jointly take a decision which is going to affect all the members of the society. Let A be the set of decision or social alternatives. Depending on the problem A can be a finite set of alternatives, a subset of \mathbb{R}^k with $k \geq 1$, a metric space or even a more abstract space.

Agents, Alternatives and Preferences

- Let consider a society formed by N agents with $N = \{1, 2, \dots, n\}$ and $n \geq 2$.
- Agents have to jointly take a decision which is going to affect all the members of the society. Let A be the set of decision or social alternatives. Depending on the problem A can be a finite set of alternatives, a subset of \mathbb{R}^k with $k \geq 1$, a metric space or even a more abstract space.
- Each agent $i \in N$ has preferences which are represented by a binary relation R_i complete and transitive on A . For any pair $x, y \in A$, xR_iy (xP_iy) means that agent i weakly (strongly) prefers x to y .

Agents, Alternatives and Preferences

- Let consider a society formed by N agents with $N = \{1, 2, \dots, n\}$ and $n \geq 2$.
- Agents have to jointly take a decision which is going to affect all the members of the society. Let A be the set of decision or social alternatives. Depending on the problem A can be a finite set of alternatives, a subset of \mathbb{R}^k with $k \geq 1$, a metric space or even a more abstract space.
- Each agent $i \in N$ has preferences which are represented by a binary relation R_i complete and transitive on A . For any pair $x, y \in A$, xR_iy (xP_iy) means that agent i weakly (strongly) prefers x to y .
- Let \mathcal{R} be the set of preferences for each agent $i \in N$. A profile of preference $R = (R_1, \dots, R_n)$ is a list of n preferences, one each agent $i \in N$.

Gibbard-Satterthwaite theorem

A society has to take a decision, that is to pick an alternative in the set A . To take a "correct" decision it needs to elicit information about individuals' preferences.

A social choice function $f : \mathcal{R}^n \rightarrow A$ selects for each preference profile an alternative, $f(R) \in A$. However, it is not necessarily the case that each individual has the incentive to truthfully report his or her own preferences.

Example 1 Cont.

Suppose $N = \{1, 2, 3\}$ and each agent $i \in N$ is informed about the other agents' valuations. Suppose that $c = 1$ and $\alpha_1 = 0.6$, $\alpha_2 = 0.3$ and $\alpha_3 = 0.2$ and the social choice function is efficient, that is for each preference profile it takes the efficient decision. If agent 3 truthfully reports her valuation she gets a negative utility since her contribution $\frac{1}{3}$ is larger than her valuation. If agent 3 misreports her valuation announcing $\alpha'_3 = 0.01$ for instance, she gets zero utility since the public good is not built.

Consider any preference profile $R \in \mathcal{R}^n$. Let $(R'_i, R_{-i}) = (R_1, \dots, R_{-i}, R'_i, R_{+i}, \dots, R_n)$.

Definition

Agent $i \in N$ manipulates the social choice function $f : \mathcal{R}^n \rightarrow A$ if there exists a profile $R \in \mathcal{R}^n$ and a preference $R'_i \neq R_i$ such that

$$f(R'_i, R_{-i}) P_i f(R_i, R_{-i}).$$

A social choice function is not manipulable (strategy-proof) if there is no agent $i \in N$ who can manipulate it.

Example

Let $N = \{1, 2\}$ and $A = \{x, y, z\}$. Assume that agents' preferences over alternatives are strict. Suppose that each agent has a veto power. The worst individual preferences of any individual is never selected. If both agents consider the same alternative as the worst one, then the top of agent 2's alternative is selected. Consider the following profile P and preference P'_1 for agent 1 :

P_1	P_2	P'_1
x	y	x
y	x	z
z	z	y

It is easy to check that $f(P) = \{y\}$ while $f(P'_1, P_2) = \{x\}$ and therefore at profile P agent 1 has incentive to manipulate by reporting preference P'_1 .

Let r_f denote the range of the social choice function f . Let $m(R_i, B) \equiv \{x \in B \mid x R_i y \text{ for all } y \in B\}$ denote the set of maximal alternatives in B according to preference R_i . Note that if preference is strict and the set A is finite $\#m(P_i, B) = 1$ for all non empty $B \subseteq A$.

Definition

The social choice function $f : \mathcal{R}^n \rightarrow A$ is dictatorial if there exists an agent $i \in N$ such that for all profiles $R \in \mathcal{R}^n$, $f(R) \in m(R_i, r_f)$.

Before stating and proving the Gibbard-Satterthwaite theorem we need a further piece of notation.

Fixed preferences P_1 of agent 1 agent 2's option set (relative to f) is

$$o_2(P_1) \equiv \{x \in A \mid \exists P_2 \in \mathcal{P}^2 \text{ such that } f(P_1, P_2) = x\}.$$

Theorem

Let $f : \mathcal{R}^n \rightarrow A$ a social choice function with $\#r_f > 2$. If f is not manipulable, then f is dictatorial.

Proof. We prove the theorem for $n = 2$. By induction the proof can be extended to $n \geq 2$. (but it is out of the scope of these introductory notes). We proceed by proving some easy lemmata.

Lemma 1 For all $P \in \mathcal{P}^2$, $f(P) = m(P_2, o_2(P_1))$.

Proof. Suppose $z = f(P)$ and $x = m(P_2, o_2(P_1))$. Since $x \in o_2(P_1)$, there exists P'_2 such that $x = f(P_1, P'_2)$. If $x \neq z$, then xP_2z and therefore $f(P_1, P'_2)P_2f(P_1, P_2)$ and agent 2 manipulate f at P by declaring P'_2 .

Lemma 2 For all $P_1 \in P$, $m(P_1, r_f) \in o_2(P_1)$.

Proof. Let $x = m(P_1, r_f)$. Since $x \in r_f$, then $\exists \bar{P}$ such that $f(\bar{P}) = x$. Let $z = f(P_1, \bar{P}_2)$. If $x = z$ then $m(P_1, r_f) \in o_2(P_1)$. If $x \neq z$ then $f(\bar{P}_1, \bar{P}_2)P_1f(P_1, \bar{P}_2)$ and therefore f is manipulable by agent 1 at (P_1, \bar{P}_2) by reporting \bar{P}_1 .

Lemma 3 For all $x \in r_f$ if $x = m(P_1, r_f) = m(P_2, r_f)$, then $f(P) = x$

Proof. Suppose $x \in r_f$ and $x = m(P_1, r_f) = m(P_2, r_f)$. By Lemma 2 $x \in o_2(P_1) \cap o_1(P_2)$. Hence $x = m(P_1, o_1(P_2)) = m(P_2, o_2(P_1))$. By Lemma 1 $f(P) = x$.

Lemma 4 For all $x \in r_f$, if $x = m(P_1, r_f) = m(P'_1, r_f)$ then $o_2(P_1) = o_2(P'_1)$.

Proof. Suppose $x = m(P_1, r_f) = m(P'_1, r_f) \in r_f$ and there exists $z \in o_2(P_1) \setminus o_2(P'_1)$. By Lemma 2, $x \in o_2(P_1) \cap o_2(P'_1)$. Let $\bar{P}_2 \in \mathcal{P}$ be any preference such that $z\bar{P}_2x\bar{P}_2y$ for all $y \in r_f \setminus \{x, z\}$ (\bar{P}_2 exists since we assume universal domain and $\#r_f \geq 3$). It follows that, (i) $f(P_1, \bar{P}_2) = z$, since by Lemma 1 $z = m(\bar{P}_2, o_2(P_1)) = f(P_1, \bar{P}_2)$ and (ii) $f(P'_1, \bar{P}_2) = x$, since $z \notin o_2(P'_1)$, $x \in o_2(P'_1)$ and by definition of \bar{P}_2 and Lemma 1, $x = m(\bar{P}_2, o_2(P'_1))$. Hence, $f(P'_1, \bar{P}_2)P_1f(P_1, \bar{P}_2)$ and f is manipulable by agent 1 at profile (P_1, \bar{P}_2) by reporting P'_1 .

Lemma 5 For all $P_1 \in \mathcal{P}$, either $\#o_2(P_1) = 1$ or $o_2(P_1) = r_f$.

Proof. Suppose $\bar{P}_1 \in \mathcal{P}$ and $x, y, z \in r_f$ such that $x, y \in o_2(\bar{P}_1)$ and $z \notin o_2(\bar{P}_1)$. Without loss of generality we can assume that $x = m(\bar{P}_1, A)$ and $z\bar{P}_1y$; in fact if it is not the case, we can modify \bar{P}_1 due to the fact that by Lemma 2, the best alternative in the range belongs to the option set and by Lemma 4 the option set only depends on best alternative in the range. Let $\bar{P}_2 \in \mathcal{P}$ be such that $z\bar{P}_2y\bar{P}_2w$ for all $w \in r_f \setminus \{y, z\}$ (\bar{P}_2 exists since we assume universal domain and $\#r_f \geq 3$). Hence, $f(\bar{P}_1, \bar{P}_2) = y$ because $z \notin o_2(\bar{P}_1)$ and, by Lemma 1 and the fact that $y \in o_2(\bar{P}_1)$, $f(\bar{P}_1, \bar{P}_2) = m(\bar{P}_2, o_2(\bar{P}_1))$. Consider another preference $P'_1 \in \mathcal{P}$ such that $z = m(P'_1, r_f)$. By Lemma 2, $z \in o_2(P'_1)$. By Lemma 1, $f(P'_1, \bar{P}_2) = z$. It follows that, $f(P'_1, \bar{P}_2)\bar{P}_1f(\bar{P}_1, \bar{P}_2)$; agent 1 manipulates f at (\bar{P}_1, \bar{P}_2) by reporting P'_1 .

Lemma 6 Either $\#o_2(P_1) = 1$ for all $P_1 \in \mathcal{P}$, or $o_2(P_1) = r_f$ for all $P_1 \in \mathcal{P}$.

Proof. Suppose not. By Lemma 5 there exist $\hat{P}_1, \bar{P}_1 \in \mathcal{P}$ such that $o_2(\hat{P}_1) = \{x\}$ and $o_2(\bar{P}_1) = r_f$. Since $\#r_f \geq 3$, by Lemma 4 (the option set only depends on the best alternative in the range) we can assume that $x\bar{P}_1z$ for some $z \in r_f$. Consider any preference $P_2 \in \mathcal{P}$ such that $m(P_2, r_f) = z$. Then, (i) $f(\bar{P}_1, P_2) = z$, by Lemma 1, and (ii) $f(\hat{P}_1, P_2) = x$ because $o_2(\hat{P}_1) = \{x\}$. It follows that, $f(\hat{P}_1, P_2) \bar{P}_1 f(\bar{P}_1, P_2)$; agent 1 manipulates f at profile (\bar{P}_1, P_2) by reporting \hat{P}_1 .

Now it is easy to show that $f : \mathcal{P}^2 \rightarrow A$ is dictatorial. By Lemma 6, either $\#o_2(P_1) = 1$ for all $P_1 \in \mathcal{P}$ or $o_2(P_1) = r_f$ for all $P_1 \in \mathcal{P}$. If $\#o_2(P_1) = 1$ for all $P_1 \in \mathcal{P}$, by Lemma 2, $o_2(P_1) = m(P_1, r_f)$. Hence, $f(P_1, P_2) = m(P_1, r_f)$ for all $(P_1, P_2) \in \mathcal{P}^2$; that is, agent 1 is a dictator.. If $o_2(P_1) = r_f$ for all $P_1 \in \mathcal{P}$, by Lemma 1, $f(P_1, P_2) = m(P_2, o_2(P_1))$ for all $P_2 \in \mathcal{P}$. By assumption $m(P_2, o_2(P_1)) = m(P_2, r_f)$ for all $P_2 \in \mathcal{P}$; that is agent 2 is a dictator.

Domain Restrictions

- The Gibbard-Satterthwaite theorem provides a strong negative result on the possibility of implementing desirable social choice function.

Domain Restrictions

- The Gibbard-Satterthwaite theorem provides a strong negative result on the possibility of implementing desirable social choice function.
- We devote some time in studying how to overcome this negative results.

Domain Restrictions

- The Gibbard-Satterthwaite theorem provides a strong negative result on the possibility of implementing desirable social choice function.
- We devote some time in studying how to overcome this negative results.
- The impossibility theorem comes from two sources: first the requirement of universal domain, second the notion of manipulability.

Domain Restrictions

- The Gibbard-Satterthwaite theorem provides a strong negative result on the possibility of implementing desirable social choice function.
- We devote some time in studying how to overcome this negative results.
- The impossibility theorem comes from two sources: first the requirement of universal domain, second the notion of manipulability.
- We first focus on the first aspect: in many interesting problems we can have information over agents' preferences that can be used to restrict the agents' preference domain.

Domain Restrictions

- The Gibbard-Satterthwaite theorem provides a strong negative result on the possibility of implementing desirable social choice function.
- We devote some time in studying how to overcome this negative results.
- The impossibility theorem comes from two sources: first the requirement of universal domain, second the notion of manipulability.
- We first focus on the first aspect: in many interesting problems we can have information over agents' preferences that can be used to restrict the agents' preference domain.
- The fact that *any preference* is admissible give agents' a large strategy space and allow them to easily manipulate by reporting untruthful information about their preference.

Domain Restrictions

- The Gibbard-Satterthwaite theorem provides a strong negative result on the possibility of implementing desirable social choice function.
- We devote some time in studying how to overcome this negative results.
- The impossibility theorem comes from two sources: first the requirement of universal domain, second the notion of manipulability.
- We first focus on the first aspect: in many interesting problems we can have information over agents' preferences that can be used to restrict the agents' preference domain.
- The fact that *any preference* is admissible give agents' a large strategy space and allow them to easily manipulate by reporting untruthful information about their preference.
- **By restricting agents' strategy space (the preference domain) we can restore positive results.**

Definition

A social choice function $f : \mathcal{D}^n \rightarrow \mathcal{R}$ is efficient if for all $R \in \mathcal{R}^n$ and for each pair of alternatives $x, y \in A$, if $x P_i y$ for all $i \in N$ then $f(R) \neq y$.

Definition

A social choice function $f : \mathcal{D}^n \rightarrow \mathcal{R}$ is anonymous if for all $\sigma : N \rightarrow N$ and $R \in \mathcal{D}^n$, $f(R_1, \dots, R_n) = f(R_{\sigma(1)}, \dots, R_{\sigma(n)})$.

A milder requirement than efficiency is unanimity. A social choice function $f : \mathcal{D}^n \rightarrow A$ is unanimous if for all $R \in \mathcal{D}^n$ such that $\bigcap_{i \in N} m(R_i, A) \neq \emptyset$, $f(R) \in \bigcap_{i \in N} m(R_i, A)$.

Single-Peaked domain

- We consider social choice problems where there is a "natural" order on the set of alternatives A that is unanimously accepted by all agents.

Single-Peaked domain

- We consider social choice problems where there is a "natural" order on the set of alternatives A that is unanimously accepted by all agents.
- The location a public good (an hospital), of a party in the left-wing right-wing ideological space, the choice of a temperature in a room, etc.

Single-Peaked domain

- We consider social choice problems where there is a "natural" order on the set of alternatives A that is unanimously accepted by all agents.
- The location a public good (an hospital), of a party in the left-wing right-wing ideological space, the choice of a temperature in a room, etc.
- Everybody agrees that a room with $68F^\circ$ is colder than a room with $74F^\circ$, but individuals may disagree on which is their preferred temperature in the room.

Single-Peaked domain

- We consider social choice problems where there is a "natural" order on the set of alternatives A that is unanimously accepted by all agents.
- The location a public good (an hospital), of a party in the left-wing right-wing ideological space, the choice of a temperature in a room, etc.
- Everybody agrees that a room with $68F^\circ$ is colder than a room with $74F^\circ$, but individuals may disagree on which is their preferred temperature in the room.
- Moreover if the ideal temperature in a room for agent i is $74F^\circ$ then he has to prefer temperature $72F^\circ$ to $68F^\circ$.

Single-Peaked domain

- We consider social choice problems where there is a "natural" order on the set of alternatives A that is unanimously accepted by all agents.
- The location a public good (an hospital), of a party in the left-wing right-wing ideological space, the choice of a temperature in a room, etc.
- Everybody agrees that a room with $68F^\circ$ is colder than a room with $74F^\circ$, but individuals may disagree on which is their preferred temperature in the room.
- Moreover if the ideal temperature in a room for agent i is $74F^\circ$ then he has to prefer temperature $72F^\circ$ to $68F^\circ$.
- We consider problems where each agent has a preferred alternative (is top alternative) denoted $t(R_i)$ and those alternatives which are farer from the ideal one are worst. To simplify we can restrict our attention on the interval $[0, 1] \subset \mathbb{R}$ and the natural order on $[0, 1]$ is the order $>$ of real numbers.

Definition

A preference $R_i \in \mathcal{R}$ is single-peaked if

(1) there exists a unique top alternative $t(R_i)$: $t(R_i) P_i y$ for all $y \in [0, 1] \setminus \{t(R_i)\}$.

(2) for each pair of alternatives $x, y \in [0, 1]$ such that $y < x \leq t(R_i)$ or $t(R_i) \leq x < y$, $x P_i y$.

- Let $\mathcal{SP} \subset \mathcal{R}$ the set of single-peaked preferences on $[0, 1]$.

Definition

A preference $R_i \in \mathcal{R}$ is single-peaked if

(1) there exists a unique top alternative $t(R_i)$: $t(R_i) P_i y$ for all $y \in [0, 1] \setminus \{t(R_i)\}$.

(2) for each pair of alternatives $x, y \in [0, 1]$ such that $y < x \leq t(R_i)$ or $t(R_i) \leq x < y$, $x P_i y$.

- Let $\mathcal{SP} \subset \mathcal{R}$ the set of single-peaked preferences on $[0, 1]$.
- Note that it may be that for some $R_i \in \mathcal{SP}$, we have $y P_i x$ and $|t(R_i) - x| < |t(R_i) - y|$; nevertheless x and y must be placed on different sides of the ideal point $t(R_i)$ (single-peaked preferences are not necessarily symmetric!).

Definition

A preference $R_i \in \mathcal{R}$ is single-peaked if

(1) there exists a unique top alternative $t(R_i)$: $t(R_i)P_i y$ for all $y \in [0, 1] \setminus \{t(R_i)\}$.

(2) for each pair of alternatives $x, y \in [0, 1]$ such that $y < x \leq t(R_i)$ or $t(R_i) \leq x < y$, $xP_i y$.

- Let $\mathcal{SP} \subset \mathcal{R}$ the set of single-peaked preferences on $[0, 1]$.
- Note that it may be that for some $R_i \in \mathcal{SP}$, we have $yP_i x$ and $|t(R_i) - x| < |t(R_i) - y|$; nevertheless x and y must be placed on different sides of the ideal point $t(R_i)$ (single-peaked preferences are not necessarily symmetric!).
- A social choice function $f : \mathcal{SP}^n \rightarrow [0, 1]$ is a voting system if for all pairs of profiles $R, R' \in \mathcal{UM}^n$ such that $t(R_i) = t(R'_i)$ for all $i \in N$, then $f(R) = f(R')$; that is a voting system only takes into account the vector of ideal points in selecting the social alternatives (top-only property).

- Moulin (1980) characterizes the set of strategy-proof voting systems in the single-peaked preference domain.

- Moulin (1980) characterizes the set of strategy-proof voting systems in the single-peaked preference domain.
- Surprisingly the set is large and there are many rules which are not dictatorial and satisfy other desirable normative properties! All strategy-proof rules are generalization of the median voter rule.

- Moulin (1980) characterizes the set of strategy-proof voting systems in the single-peaked preference domain.
- Surprisingly the set is large and there are many rules which are not dictatorial and satisfy other desirable normative properties! All strategy-proof rules are generalization of the median voter rule.
- Suppose for sake of simplicity that n is odd and each agent has single-peaked preference on $[0, 1]$.

- A social choice function $f : \mathcal{SP}^n \rightarrow [0, 1]$ is a median voter rule if for all $R \in \mathcal{SP}^n$, selects the median among the n top alternatives of the agent; that is $f(R) = \text{med}\{t(R_1), \dots, t(R_n)\}$.

- A social choice function $f : \mathcal{SP}^n \rightarrow [0, 1]$ is a median voter rule if for all $R \in \mathcal{SP}^n$, selects the median among the n top alternatives of the agent; that is $f(R) = \text{med}\{t(R_1), \dots, t(R_n)\}$.
- Consider now a "generalized" median voter n with, $n - 1$ fictitious (or phantom) voters : $\frac{n-1}{2}$ votes for alternative 0 and the other $\frac{n-1}{2}$ votes for alternative 1 (we are assuming that n is odd). The median of the n top alternatives of the n voters and the median among the n top alternative and the $n - 1$ phantoms coincide, that is for all $R = (R_1, \dots, R_n) \in \mathcal{UM}^n$,

$$f(R) = \text{med}\{t(R_1), \dots, t(R_n), \underbrace{0, \dots, 0}_{\frac{n-1}{2} \text{ votes}}, \underbrace{1, \dots, 1}_{\frac{n-1}{2} \text{ votes}}\} = \text{med}\{t(R_1), \dots, t(R_n)\}$$

- Consider now $n - 1$ phantom voters with top anywhere in the interval $[0, 1]$ and for n both odd or even. Let $v_{n-1}, \dots, v_1 \in [0, 1]$ denote the $n - 1$ fictitious votes. Given $v_{n-1}, \dots, v_1 \in [0, 1]$ we define the social choice function $f : \mathcal{UM}^n \rightarrow [0, 1]$ such that for all $R \in \mathcal{SP}^n$,

$$f(R) = \text{med}\{t(R_1), \dots, t(R_n), v_{n-1}, \dots, v_1\}.$$

- Consider now $n - 1$ phantom voters with top anywhere in the interval $[0, 1]$ and for n both odd or even. Let $v_{n-1}, \dots, v_1 \in [0, 1]$ denote the $n - 1$ fictitious votes. Given $v_{n-1}, \dots, v_1 \in [0, 1]$ we define the social choice function $f : \mathcal{UM}^n \rightarrow [0, 1]$ such that for all $R \in \mathcal{SP}^n$,

$$f(R) = \text{med}\{t(R_1), \dots, t(R_n), v_{n-1}, \dots, v_1\}.$$

- This is a family of social choice function (each function for each possible vector of $n - 1$ fictitious votes). A social choice function is a strategy-proof, anonymous and efficiency voting system if and only if it belongs to this family.

Theorem

(Moulin, 1980) A social choice function $f : \mathcal{SP}^n \rightarrow [0, 1]$ is a strategy-proof, efficient and anonymous voting system if and only if there exist $n - 1$ phantom voters $0 \leq v_{n-1} \leq \dots \leq v_1 \leq 1$ such that for all $R \in \mathcal{SP}^n$,

$$f(R) = \text{med}\{t(R_1), \dots, t(R_n), v_{n-1}, \dots, v_1\}.$$

If instead of $n - 1$ fictitious votes we use $n + 1$ then the social choice function keeps all the previous properties except efficiency.

Theorem

(Moulin, 1980) A social choice function $f : \mathcal{SP}^n \rightarrow [0, 1]$ is a strategy-proof and anonymous voting system if and only if there exist $n + 1$ phantom voters $0 \leq v_n \leq \dots \leq v_0 \leq 1$ such that for all $R \in \mathcal{UM}^n$,

$$f(R) = \text{med}\{t(R_1), \dots, t(R_n), v_n, \dots, v_0\}.$$

- Moulin (1980) actually characterizes the (larger) class of strategy-proof voting systems (not necessarily anonymous) and actually the previous two results are corollaries of this result.

- Moulin (1980) actually characterizes the (larger) class of strategy-proof voting systems (not necessarily anonymous) and actually the previous two results are corollaries of this result.
- To conclude it is worthy to stress that the assumption that f is a voting system is fundamental and it is not a consequence of the property of strategy-proofness. For instance consider the following social choice function $f : \mathcal{SP}^n \rightarrow [0, 1]$ such that for all $R \in \mathcal{SP}^n$,

$$f(R) = \begin{cases} 0 & \text{if } \#\{i \in N \mid 0R_i1\} \geq \#\{i \in N \mid 1P_i0\} \\ 1 & \text{if } \#\{i \in N \mid 0R_i1\} < \#\{i \in N \mid 1P_i0\}. \end{cases} \quad (1)$$

This social choice function f is strategy-proof and anonymous but it is not a voting system. In fact the function depends on how agents compare pairs of alternatives (the two extremes). Moreover f is not efficient neither unanimous.

Quasi linear preferences

- We consider now another preference domain restriction which plays an important role in many social decision problems.

Quasi linear preferences

- We consider now another preference domain restriction which plays an important role in many social decision problems.
- Consider any binary collective decision $x = \{0, 1\}$. For each agent $i \in N$, let $\alpha_i \in \mathbb{R}$ be agent i 's (monetary) *valuation*, that can be either positive or negative, if $x = 1$ and normalize to zero the valuation of the decision $x = 0$. Let $t = (t_i)_{i \in N} \in \mathbb{R}^N$ be a vector of *transfer* (if positive, or price if negative).

Quasi linear preferences

- We consider now another preference domain restriction which plays an important role in many social decision problems.
- Consider any binary collective decision $x = \{0, 1\}$. For each agent $i \in N$, let $\alpha_i \in \mathbb{R}$ be agent i 's (monetary) *valuation*, that can be either positive or negative, if $x = 1$ and normalize to zero the valuation of the decision $x = 0$. Let $t = (t_i)_{i \in N} \in \mathbb{R}^N$ be a vector of *transfer* (if positive, or price if negative).
- An *allocation* is a pair $(x, t) \in \{0, 1\} \times \mathbb{R}^N$. Let A be the set of allocations. Agent i 's preferences on the set of allocations A depend on i 's valuation $\alpha_i \in \mathbb{R}$ and are represented by the utility function $u_i : A \times \mathbb{R} \rightarrow \mathbb{R}$, where for each $(x, t, \alpha_i) \in A \times \mathbb{R}$,

$$u_i(x, t, \alpha_i) = v_i(\alpha_i, x) + t_i(\alpha) = x\alpha_i + t_i(\alpha).$$

- Since the society N will remain fixed, a profile $\alpha = (\alpha_i)_{i \in N} \in \mathbb{R}_+^N$ completely describes a *problem*.

- Since the society N will remain fixed, a profile $\alpha = (\alpha_i)_{i \in N} \in \mathbb{R}_+^N$ completely describes a *problem*.
- Writing $t_i(\alpha_i)$ we emphasize that the transfer t_i may in general depends on agents' valuations.

- Since the society N will remain fixed, a profile $\alpha = (\alpha_i)_{i \in N} \in \mathbb{R}_+^N$ completely describes a *problem*.
- Writing $t_i(\alpha_i)$ we emphasize that the transfer t_i may in general depends on agents' valuations.
- We write $v_i(\alpha_i, x) : \mathcal{R}_i \times X \rightarrow R$ to illustrate the utility of agent i when the decision x is taken without taking into consideration the transfer t_i ; hence in this simple setting we have $v_i(\alpha_i, 1) = \alpha_i$ and $v_i(\alpha_i, 0) = 0$.

- Example 5 above already showed that agents may have incentive to misreport their valuations in the particular case that the binary decision is the decision whether to provide a public good at a cost c , or not provide it.

- Example 5 above already showed that agents may have incentive to misreport their valuations in the particular case that the binary decision is the decision whether to provide a public good at a cost c , or not provide it.
- **Example 5** *Suppose $N = \{1, 2, 3\}$ and each agent $i \in N$ is informed about the other agents' valuations. Suppose that $c = 1$ and $\alpha_1 = 0.6$, $\alpha_2 = 0.3$ and $\alpha_3 = 0.2$ and the social choice function is efficient, that is for each preference profile it takes the efficient decision. If agent 3 truthfully reports her valuation she gets a negative utility since her contribution $\frac{1}{3}$ is larger than her valuation. If agent 3 misreports her valuation announcing $\alpha'_3 = 0.01$ for instance, she gets zero utility since the public good is not built.*

- If according to the rule, all agents pay the same amount, an agent who has a valuation lower than the price $\frac{c}{n}$ may have incentive to underreport his valuation, if, as it occurs in example 5, by underreporting he is able to switch the public good decision from providing to not providing the public good.

- If according to the rule, all agents pay the same amount, an agent who has a valuation lower than the price $\frac{c}{n}$ may have incentive to underreport his valuation, if, as it occurs in example 5, by underreporting he is able to switch the public good decision from providing to not providing the public good.
- It follows that it is fundamental how the transfer function is designed in order to provide incentives to agents to sincerely report their willingness to pay.

- Consider the binary collective decision $x = \{0, 1\}$. In a binary decision it may be desirable that the decision of choosing $x = 1$ be efficient, that is $x = 1$ if and only if $\sum_{i=1}^n \alpha_i \geq 0$.

Grove's schemes: Efficient Decisions

- Consider the binary collective decision $x = \{0, 1\}$. In a binary decision it may be desirable that the decision of choosing $x = 1$ be efficient, that is $x = 1$ if and only if $\sum_{i=1}^n \alpha_i \geq 0$.
- Under this approach the main goal of a public decision maker is to find a transfer function such that the decision x is efficient.

- Consider the binary collective decision $x = \{0, 1\}$. In a binary decision it may be desirable that the decision of choosing $x = 1$ be efficient, that is $x = 1$ if and only if $\sum_{i=1}^n \alpha_i \geq 0$.
- Under this approach the main goal of a public decision maker is to find a transfer function such that the decision x is efficient.
- The following proposition characterizes the class of transfer rules such that the social choice function $f(\alpha)$ which takes efficient decision at each profile α is strategy-proof.

Theorem

(Groves; Green and Laffont) i) If x is an efficient decision and there exists a function $h_i : \mathcal{R}_{-i} \rightarrow \mathbb{R}$ such that

$$t_i(\alpha_{-i}) = h_i(\alpha_{-i}) + \sum_{j \neq i} v_j(\alpha_j, x), \quad (2)$$

then the allocation rule $f(\alpha) = (x(\alpha), t_i(\alpha_{-i}))$ is strategy-proof.

ii) if x is an efficient decision and the rule f is strategy proof, then there exists a function $h_i : \mathcal{R}_{-i} \rightarrow \mathbb{R}$ such that $t_i(\alpha_{-i}) = h_i(\alpha_{-i}) + \sum_{j \neq i} v_j(\alpha_j, x)$.

Proof.

We only prove the sufficient part. Consider an efficient decision rule and suppose there exists $h_i : \mathbb{R}^{n-1} \rightarrow \mathbb{R}$ such that

$t_i(\alpha_{-i}) = h_i(\alpha_{-i}) + \sum_{j \neq i} v_j(\alpha_j, x)$. Take any α and suppose without loss of

generality loss that $x(\alpha) = 1$ is the efficient decision at α . By assumption

$\sum_{i=1}^n \alpha_i \geq 0$. Consider any agent i . If i reports α_i his utility is

$$u_i(x, t, \alpha) = v_i(x, \alpha_i) + t_i(\alpha_{-i}) = \alpha_i + h_i(\alpha_{-i}) + \sum_{j \neq i} \alpha_j.$$

If agent i misreports his $\hat{\alpha}_i$ but $\sum_{j \neq i} \alpha_j + \hat{\alpha}_i \geq 0$, then his deviation has no effect because the decision is the same and agent i gets the same transfer $t_i(\alpha_{-i})$, since it does not depend on his own report. The only way in which agent i can affect the decision is by reporting $\hat{\alpha}_i < 0$ such that $\sum_{j \neq i} \alpha_j + \hat{\alpha}_i < 0$. We write $u_i(x, t, \hat{\alpha}_i | \alpha_i)$ to indicate the utility of agent i who has willingness to pay α_i but he reports $\hat{\alpha}_i$. The utility of agent i when reporting $\hat{\alpha}_i$ with a (true) willingness to pay α_i is

$$u_i(x, t, \hat{\alpha}_i | \alpha_i) = h_i(\alpha_{-i}).$$

However by assumption the decision is efficient $\sum_{i=1}^n \alpha_i \geq 0$ but then $u_i(x, t, \alpha) \geq u_i(x, t, \hat{\alpha}_i | \alpha_i)$ and therefore to misreport is not profitable.

The Pivotal mechanism

Among the class of functions $h_i(\alpha_{-i})$, there is one simple version which has a very nice and intuitive interpretation. For all $i \in N$, let

$$h_i(\alpha_{-i}) = - \max_x \sum_{j \neq i} v_j(\alpha_j, x(\alpha)),$$

and therefore

$$t_i(\alpha_{-i}) = \sum_{j \neq i} v_j(\alpha_j, x(\alpha)) - \max_{x \in \{0,1\}} \sum_{j \neq i} v_j(\alpha_j, x).$$

- To understand, the above formula, suppose that $\sum_{j \neq i} \alpha_j < 0$. If $\alpha_i + \sum_{j \neq i} \alpha_j < 0$ then agent i is not *pivotal*, that is his report does not affect the social decision; the efficient decision is $x(\alpha) = 0$, $t_i(\alpha_{-i}) = 0$ and $u_i(x, t, \alpha_i) = 0$. If instead $\alpha_i + \sum_{j \neq i} \alpha_j > 0$ then the efficient decision is $x(\alpha) = 1$ and agent i is pivotal; in this case $t_i(\alpha_{-i}) = \sum_{j \neq i} \alpha_j < 0$ and $u_i(x, t, \alpha_i) = \alpha_i + \sum_{j \neq i} \alpha_j$.

- Notice that if $\alpha_i + \sum_{j \neq i} \alpha_j < 0$ and $\alpha_i > 0$ agent i could be tempted to overreport his willingness to pay to reverse the social decision. However if he reports $\hat{\alpha}_i$ such that $\hat{\alpha}_i + \sum_{j \neq i} \alpha_j > 0$ he gets $u_i(x, t, \hat{\alpha}_i | \alpha_i) = \alpha_i + \sum_{j \neq i} \alpha_j < 0 = u_i(x, t, \alpha_i)$.

- Notice that if $\alpha_i + \sum_{j \neq i} \alpha_j < 0$ and $\alpha_i > 0$ agent i could be tempted to overreport his willingness to pay to reverse the social decision.

However if he reports $\hat{\alpha}_i$ such that $\hat{\alpha}_i + \sum_{j \neq i} \alpha_j > 0$ he gets

$$u_i(x, t, \hat{\alpha}_i | \alpha_i) = \alpha_i + \sum_{j \neq i} \alpha_j < 0 = u_i(x, t, \alpha_i).$$

- At the same time, if the true α_i is such that $\alpha_i + \sum_{j \neq i} \alpha_j > 0$ agent i

prefers to report α_i , pay $t_i(\alpha_{-i}) = \sum_{j \neq i} \alpha_j < 0$ and get

$$u_i(x, t, \alpha_i) = \alpha_i + \sum_{j \neq i} \alpha_j > 0$$

than reporting $\hat{\alpha}_i$ such that $\hat{\alpha}_i + \sum_{j \neq i} \alpha_j < 0$ and get $u_i(x, t, \alpha_i) = 0$.

- Similarly, if $\sum_{j \neq i} \alpha_j > 0$ and $\alpha_i + \sum_{j \neq i} \alpha_j > 0$, agent i is not pivotal: we have $x(\alpha) = 1$ and $t_i(\alpha_{-i}) = 0$ and $u_i(x, t, \alpha_i) = \alpha_i$, while if $\alpha_i + \sum_{j \neq i} \alpha_j < 0$, then agent i is pivotal and $x(\alpha) = 0$; in this case $t_i(\alpha_{-i}) = -\sum_{j \neq i} \alpha_j < 0$ and $u_i(x, t, \alpha_i) = \alpha_i - \sum_{j \neq i} \alpha_j$. As before if the true willingness to pay is $\alpha_i < 0$ agent i could be tempted to underreport his willingness to pay to reverse the social decision. However if he reports $\hat{\alpha}_i$ such that $\hat{\alpha}_i + \sum_{j \neq i} \alpha_j < 0$ he gets $u_i(x, t, \hat{\alpha}_i | \alpha_i) = \alpha_i - \sum_{j \neq i} \alpha_j < 0 = u_i(x, t, \alpha_i)$.

- Similarly, if $\sum_{j \neq i} \alpha_j > 0$ and $\alpha_i + \sum_{j \neq i} \alpha_j > 0$, agent i is not pivotal: we have $x(\alpha) = 1$ and $t_i(\alpha_{-i}) = 0$ and $u_i(x, t, \alpha_i) = \alpha_i$, while if $\alpha_i + \sum_{j \neq i} \alpha_j < 0$, then agent i is pivotal and $x(\alpha) = 0$; in this case $t_i(\alpha_{-i}) = -\sum_{j \neq i} \alpha_j < 0$ and $u_i(x, t, \alpha_i) = \alpha_i - \sum_{j \neq i} \alpha_j$. As before if the true willingness to pay is $\alpha_i < 0$ agent i could be tempted to underreport his willingness to pay to reverse the social decision. However if he reports $\hat{\alpha}_i$ such that $\hat{\alpha}_i + \sum_{j \neq i} \alpha_j < 0$ he gets $u_i(x, t, \hat{\alpha}_i | \alpha_i) = \alpha_i - \sum_{j \neq i} \alpha_j < 0 = u_i(x, t, \alpha_i)$.
- Each individual i 's transfer function takes into account the marginal social impact (on other individuals) made by his report α_i . When looking at this social impact together with his own selfish utility, the individual has exactly the total social value in mind when deciding on a strategy. This leads to efficient decision making.

- A social choice function that always take an efficient decision is not necessarily efficient. Full efficiency in fact would require that resources are not wasted, that it not only the decision is efficient but the rule is also budget balanced. Let $C(d^f(\alpha))$ be the cost of the decision taken by social choice function f at the profile α .

Definition

A social choice rule f is budget balanced if for all $\alpha \in \mathbb{R}^n$,

$$-\sum_{i=1}^n t_i(\alpha) = C(d^f(\alpha)).$$

It is easy to see that full efficiency and strategy-proofness are not compatible.

Example

Suppose $N = \{1, 2\}$ Suppose that agents have to decide whether to provide or not a binary public good at a cost $c = 1$. Suppose $\alpha_i \in \mathbb{R}_+ \cup \{0\}$ for both $i = 1, 2$. Consider first the profile $\alpha_1 = 0.9$, $\alpha_2 = 0.9$. Since $\alpha_1 + \alpha_2 > 1$ then the efficient decision to take at this profile is to build the project. By budget balancedness $t_1(0.9, 0.9) + t_2(0.9, 0.9) = -1$. By 9 we have that if the allocation rule is strategy-proof then for both $i = 1, 2$

$$t_i(0.9) = h_i(0.9) + 0.9. \quad (3)$$

Therefore $-1 = h_1(0.9) + h_2(0.9) + 1.8$ and therefore $h_1(0.9) + h_2(0.9) = -2.8$. Similarly if $\alpha_1 = \alpha_2 = 0$, the efficient decision is to not provide the public good; budget balancedness implies $t_1(0, 0) + t_2(0, 0) = 0$ and therefore by 9 $0 = h_1(0) + h_2(0)$.

Now consider the profiles $\alpha_i = 0.9$ and $\alpha_{-i} = 0$. The efficient decision is to not provide the public good; budget balancedness implies $t_1(0.9, 0) + t_2(0.9, 0) = 0$ and $t_1(0, 0.9) + t_2(0, 0.9) = 0$ by 9
 $0 = h_1(0) + h_2(0.9) + 0.9 = h_1(0.9) + h_2(0) + 0.9$. However this system has no solution

$$h_1(0.9) + h_2(0.9) + 2.8 = 0$$

$$h_1(0) + h_2(0) = 0$$

$$h_1(0) + h_2(0.9) + 0.9 = 0$$

$$h_1(0.9) + h_2(0) + 0.9 = 0$$

- Another requirement that is relevant in many applications is that the social choice function guarantees at each profile a non negative utility to each agent. Sometimes this requirement is justified on a normative ground, sometimes can be considered as a feasibility constraint whenever agents who experience negative utility have the possibility to opt out and refuse to participate to the public decision process.

Definition

A social choice function is individual rational if for all $\alpha \in \mathbb{R}$ and for all $i \in N$, $u_i(f(\alpha)) \geq 0$.

Example

The properties of strategy-proofness, efficient decision and individually rationality are not compatible. Consider the previous setting with $N = \{1, 2\}$. Suppose that agents have to decide whether to provide or not a binary public good at a cost $c = 1$. Suppose $\alpha_i \in \mathbb{R}_+ \cup \{0\}$ for both $i = 1, 2$. If $\alpha_i = 2$ for both $i = 1, 2$ and the efficient decision is to provide the public good.. By feasibility the sum of individual contributions must be at least equal to 1 and therefore there exists at least one agent i who pays a positive amount. Suppose without loss of generality that $t_1(2, 2) > 0$. Consider now the profile $\alpha_2 = 2$ and $\alpha_1 = 0$. By efficiency the public good is still provided and by individual rationality, $t_1(0, 2) = 0$. However if the decision is efficient, then the public good is still provided and therefore strategy-proofness is violated.

The Second Price Auction

- The pivotal mechanism reduces to a well-known auction form in the context of the allocation of indivisible objects. In that context there is a simple auction that is dominant strategy incentive compatible.

The Second Price Auction

- The pivotal mechanism reduces to a well-known auction form in the context of the allocation of indivisible objects. In that context there is a simple auction that is dominant strategy incentive compatible.
- It turns out that this auction form, commonly referred to as a Vickrey auction, corresponds to the pivotal mechanism in this setting. Let us explore this relationship in the case of a single good.

The Second Price Auction

- The pivotal mechanism reduces to a well-known auction form in the context of the allocation of indivisible objects. In that context there is a simple auction that is dominant strategy incentive compatible.
- It turns out that this auction form, commonly referred to as a Vickrey auction, corresponds to the pivotal mechanism in this setting. Let us explore this relationship in the case of a single good.
- There are n individuals who each have a valuation μ_i for the object and $a \in \{1, \dots, n\} = A$ indicates the individual to whom the object is assigned. In that case, the efficient decision is such that $a(\mu_i) \in \arg \max_i \mu_i$. The pivotal mechanism takes an easy form. If $a(\mu) = i$ then $p_i(\mu) = -\max_{j \neq i} \mu_j$ and if $a(\mu) \neq i$ $p_i(\mu) = 0$.

- It follows that the object is allocated to the individual with the highest valuation and he pays an amount equal to the second highest valuation. No other payments are made.

- It follows that the object is allocated to the individual with the highest valuation and he pays an amount equal to the second highest valuation. No other payments are made.
- Now consider a second price (Vickrey) auction, where the high bidder is awarded the object and pays the second highest price. It is easy to see that bidding one's value is a dominant strategy in this auction, as it is the same reasoning as that behind the pivotal mechanism.

- It follows that the object is allocated to the individual with the highest valuation and he pays an amount equal to the second highest valuation. No other payments are made.
- Now consider a second price (Vickrey) auction, where the high bidder is awarded the object and pays the second highest price. It is easy to see that bidding one's value is a dominant strategy in this auction, as it is the same reasoning as that behind the pivotal mechanism.
- The pivotal mechanism and Vickrey auction implement the same social choice function.