Mechanism Design (with and) without Money

"Algorithmic Game Theory", Ch. 9-10
"A Course in Game Theory", Ch. 10
"A Primer in Social Choice Theory", Ch. 5

davide.grossi@uni.lu
Individual and Collective Reasoning Group

UNIVERSITÉ DU LUXEMBOURG

Outline

- Social Choice Theory
- Implementation Theory
- Mechanism Design With Money
- Mechanism Design Without Money

Social Choice Theory

davide.grossi@uni.lu
Individual and Collective Reasoning Group

Collective decision-making

- Elections
- Auctions
- Program Committees

Social Choice

A Social Choice structure is a quadruple:

$$\mathfrak{S} = \langle \mathtt{Agn}, \mathtt{Iss}, \mathtt{Prf}, \mathtt{Sc} \rangle$$

s.t:

- $\mathtt{Agn}$ is a finite set of agents such that $1 \leq |\mathtt{Agn}|$;
- $\mathtt{Iss}$ is a finite set of issues such that $3 \leq |\mathtt{Iss}|$;
- $\mathtt{Prf}$ is the set of all preference profiles, i.e., $|\mathtt{Agn}|$-tuples $\mathfrak{p} = (\preceq_i)_{i \in \mathtt{Agn}}$ where each $\preceq_i$ is a total order over $\mathtt{Iss}$;
- $\mathtt{Sc}$ is a function taking each $\mathfrak{p} \in \mathtt{Prf}$ to an element in $\mathtt{Iss}$,i.e.:

$$\mathtt{Sc} : \mathtt{Prf} \longrightarrow \mathtt{Iss}$$

.

Input	Function	Output
$\preceq_1$ ... $\preceq_n$	$\dashrightarrow$	a

What *classes* of Social Choice functions are *possible*?

Incentive-compatibility

A social choice function $\mathtt{Sc}$ can be strategically manipulated by agent i if for some profile $\mathfrak{p} = (\preceq_1, \ldots, \preceq_n)$, there exists another profile $\mathfrak{p}' = (\preceq_{-i}, \preceq'_i)$ s.t.

$$\mathtt{Sc}(\mathfrak{p}) \preceq_i \mathtt{Sc}(\mathfrak{p}')$$

for $\mathtt{Sc}(\mathfrak{p}) \neq \mathtt{Sc}(\mathfrak{p}')$. A function $\mathtt{Sc}$ is *incentive compatible* (or strategy-proof) if it cannot be manipulated.

- *An agent can force a different alternative which he prefers by misrepresenting his preferences*
- Majority voting on a set of issues with 2 elements is strategy-proof

Why incentive compatibility?

- If we want to construct a social choice function as an *algorithm* we have, first of all, to elicit the preferences of the agents
- Preferences are private
- Incentive compatibility guarantees that our algorithm elicits the right information

Dictatorship

A social choice function $\mathtt{Sc}$ is a dictatorship if there exists an agent i s.t. $\forall \mathfrak{p} \in \mathtt{Prf}$, and $\forall b \neq a \in \mathtt{Iss}$:

$$b \preceq_i a \quad \Rightarrow \quad \mathtt{Sc}(\mathfrak{p}) = a$$

A function $\mathtt{Sc}$ is non-dictatorial if there is no dictator.

- Is a dictatorship incentive compatible?

Gibbard-Satterthwaite Theorem ('73, '75)

- A. Gibbard, "Manipulation of voting schemes: a general result", *Econometrica*, Vol. 41, No. 4 (1973), pp. 587–601
- M. A. Satterthwaite, "Strategy-proofness and Arrow's Conditions: Existence and Correspondence Theorems for Voting Procedures and Social Welfare Functions", *Journal of Economic Theory* 10 (April 1975), 187–217

Gibbard-Satterthwaite Theorem ('73, '75)

Let `Iss` > 2. If a social choice function `Sc` is:

1. onto `Iss` (aka non-imposition) and
2. is incentive compatible

then it is a dictatorship.

- How does this jeopardize the possibility of finding algorithms for collective decision-making?

Implementation

Implementation

An Implementation Problem for the Social Choice structure $\mathfrak{S}$ is a structure:

$$(\mathfrak{S}, \mathcal{G})$$

where $\mathcal{G}$ is a set of strategic game forms $G = (\mathtt{Agn}, \mathtt{Str}, \mathtt{g})$ s.t:

- $\mathtt{Agn}$ is the finite set of agents of $\mathfrak{S}$;
- $\mathtt{Str}$ is the set $\prod_{i \in \mathtt{Agn}} S_i$ of all strategy profiles, where S_i is the set of strategies of agent i;
- $\mathtt{g}$ is the outcome function of the game: $\mathtt{g} : \mathtt{Str} \longrightarrow \mathtt{Iss}$

Given $(\mathfrak{S}, \mathcal{G})$, find a game $G \in \mathcal{G}$ and a solution concept S s.t.:

$$\mathtt{g}(S(G, \mathfrak{p})) = \mathtt{Sc}(\mathfrak{p})$$

If such a G exists then $\mathtt{Sc}$ is S*-implementable*.

Truthful implementation

A Direct Implementation Problem for the Social Choice structure $\mathfrak{S}$ is a structure:

$$(\mathfrak{S}, \mathcal{G})$$

where $\mathcal{G}$ is a set of strategic game forms, aka *direct revelation mechanisms*, $G = (\texttt{Agn}, \texttt{Str}, \texttt{g})$ s.t:

- $\texttt{Agn}$ is the finite set of agents of $\mathfrak{S}$;
- $\texttt{Str} = \texttt{Prf}$, i.e., the set of strategy profiles is the set of preference profiles;
- $\texttt{g}$ is the outcome function of the game: $\texttt{g} : \texttt{Str} \longrightarrow \texttt{Iss}$

Agents play the game by declaring their preferences!

Truthful implementation

A Direct Implementation Problem for the Social Choice structure $\mathfrak{S}$ is a structure:

$$(\mathfrak{S}, \mathcal{G})$$

where $\mathcal{G}$ is a set of strategic game forms, aka *direct revelation mechanisms*, $G = (\mathtt{Agn}, \mathtt{Str}, \mathtt{g})$ s.t:

- $\mathtt{Agn}$ is the finite set of agents of $\mathfrak{S}$;
- $\mathtt{Str} = \mathtt{Prf}$, i.e., the set of strategy profiles is the set of preference profiles;
- $\mathtt{g}$ is the outcome function of the game: $\mathtt{g} : \mathtt{Str} \longrightarrow \mathtt{Iss}$

Given a Direct Implementation Problem $(\mathfrak{S}, \mathcal{G})$, find a *direct revelation mechanism* $G \in \mathcal{G}$ and a solution concept S s.t.:

$$\begin{aligned} \mathtt{g}(S(G, \mathfrak{p})) &= \mathtt{Sc}(\mathfrak{p}) \\ S(G, \mathfrak{p}) &= \mathfrak{p} \end{aligned}$$

If such a G exists, then $\mathtt{Sc}$ is *truthfully S-implementable*

Revelation in *Dominant Strategies*

A *dominant strategy equilibrium* of a strategic game $(G, \mathfrak{p})$ with $G = (\mathtt{Agn}, \mathtt{Str}, \mathtt{g})$, is a strategy profile $s^* \in \mathtt{Str}$ s.t. $\forall i \in \mathtt{Agn}$, and $\forall s \in \mathtt{Str}$:

$$\mathtt{g}((s_{-i}, s_i)) \preceq_i \mathtt{g}((s_{-i}, s_i^*))$$

with $\preceq_i$ being the i-th projection of $\mathfrak{p}$. A social choice function is said to be *DSE-implementable* if it is implementable w.r.t. dominant strategy equilibrium.

Theorem (Revelation Principle). Given an Implementation problem $(\mathfrak{S}, \mathcal{G})$, if there exists a game form $G \in \mathcal{G}$ DSE-implementing $\mathtt{Sc}$, then there exists a direct revelation mechanism G^d s.t. $\forall \mathfrak{p} \in \mathtt{Prf}$:

$$\mathrm{DSE}(G^d, \mathfrak{p}) = \mathfrak{p}$$

that is, $\mathtt{Sc}$ is *truthfully DSE-implementable.*

DSE & Strategy-proofness

Fact. If $\mathtt{Sc}$ is *truthfully DSE-implementable* then it is incentive compatible, i.e.:

$$\forall \mathfrak{p} : \; DSE(G, \mathfrak{p}) = \mathfrak{p} \quad \Rightarrow \quad \forall i, \forall \mathfrak{p}, \mathfrak{p}' : \; \mathtt{Sc}((\mathfrak{p}_{-i}, \mathfrak{p}'_i)) \preceq_i \mathtt{Sc}((\mathfrak{p}_{-i}, \mathfrak{p}_i))$$

where G is a direct revelation mechanism, and $\preceq_i$ is the i^{th} projection of $\mathfrak{p}$.

- DSE implementation is strictly related to the problem of preference elicitation

No DSE implementation!

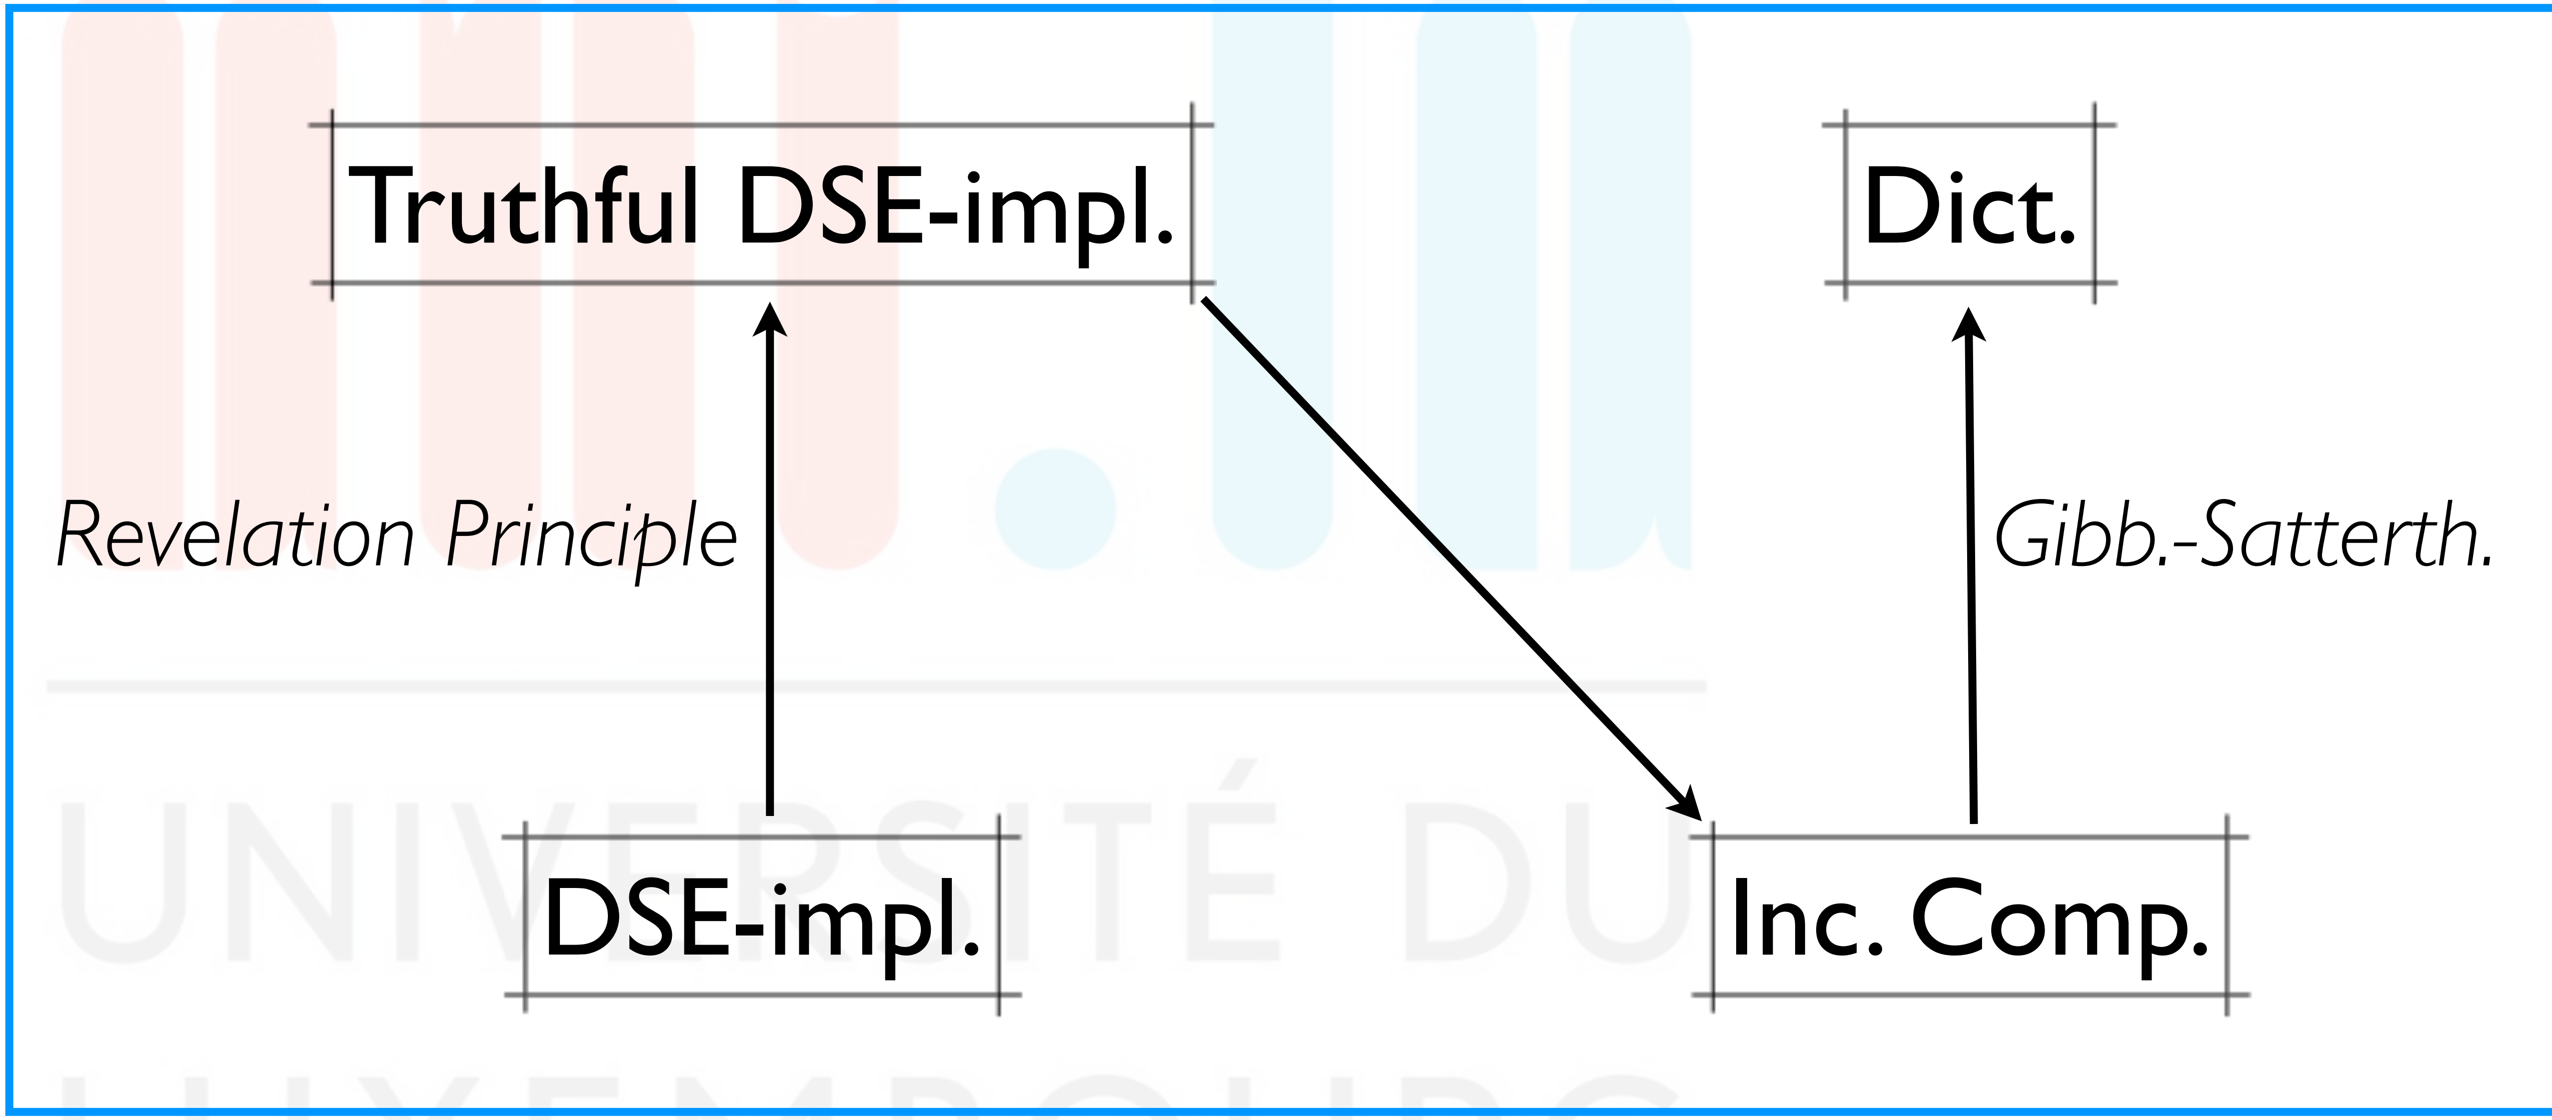

How is MD possible?

- No non-trivial incentive compatible social choice function!
- No DSE implementations of non-trivial social choice functions!

- *Under what conditions* **can we prove the existence of non-trivial incentive compatible social choice functions become by** *designing mechanisms*?

With Money

The virtues of money

Total orders $\preceq_i \subseteq \texttt{Iss} \times \texttt{Iss}$ vs. *valuation functions* $v_i : \texttt{Iss} \longrightarrow \mathbb{R}$

1. Preference intensity can be *measured* and interpersonal comparisons become possible
2. The unit of measure of preference intensity is transferable. Payments become possible:

$$u_i(a) = v_i(a) - p$$

Auctions

An auction is a Social Choice structure:

$$\mathfrak{S} = \langle \mathtt{Agn}, \mathtt{Iss}, \mathtt{Prf}, \mathtt{Sc} \rangle$$

s.t:

- $\mathtt{Agn}$ is a finite set of agents such that $1 \leq |\mathtt{Agn}|$;
- $\mathtt{Iss} := \mathtt{Agn} \times \mathbb{R}$;
- $\mathtt{Prf}$ is the set of all valuations of the auctioned item, i.e., $|\mathtt{Agn}|$-tuples $\mathfrak{p} = (w_i)_{i \in \mathtt{Agn}}$ where each $w_i' \in \mathbb{R}$;
- $\mathtt{Sc}$ is a function $\mathtt{Sc} : \mathtt{Prf} \longrightarrow \mathtt{Iss}$

- Valuations determine total preorders on Iss
- Sc picks a winner and establishes a payment

What's the price?

1. Highest bidder wins, and no payment?
2. Highest bidder wins and pays the bid?

- W. Vickrey, *"Counterspeculation, auctions and competitive sealed tenders"*, Journal of Finance, 8-37, 1961

Vickrey Auction

Define $\texttt{Sc}$ as follows. Let the winner be the agent i with the highest declared valuation w_i an let i pay the second highest declared valuation $p = \max_{j \neq i} w_j$.

Theorem. Let $u(w_i)$ denote the utility of i if i bids w_i. For any profile of declared valuations $(w_1, \ldots, w_n)$, and valuation w'_i, it holds that $u(w'_i) \leq u(w_i)$.

- Possibility of an incentive compatible mechanism under a specific subclass of total orders!
- Mechanism design aims at the *generalization* of such possibility: between Vickrey and Gibbard-Satterthwaite.

Input	Function	Output
w_1 … w_n		(i, p)

Incentive compatible functions are implementable (with Money)! The possibility is proven by the existence of the Vickrey auction!

Without Money

Single-peaked preferences

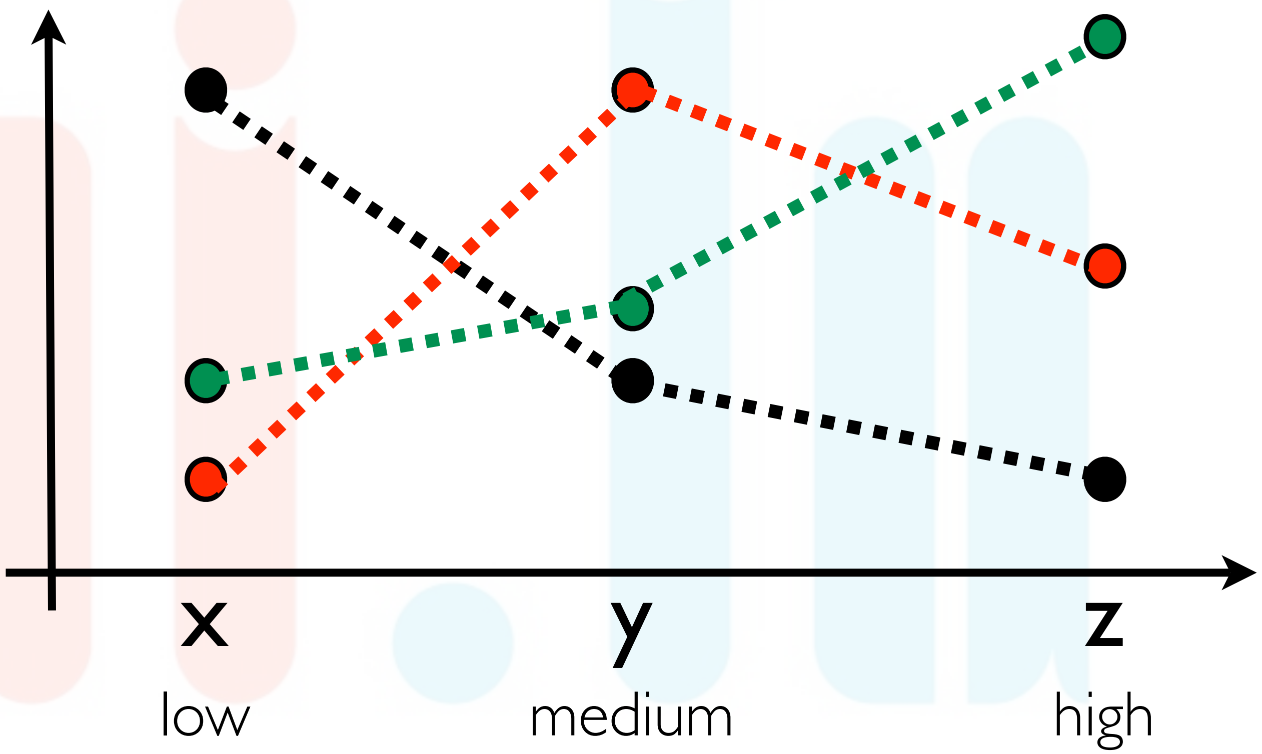

A preference profile $\mathfrak{p} = (\preceq_1, \ldots, \preceq_n)$ of total preorders on $\mathtt{Iss}$ is *single-peaked* if there exists a total order $\preceq^*$ on $\mathtt{Iss}$ s.t. $\forall i \in \mathtt{Agn}$:

$$y \preceq_i x \;\&\; B(x, y, z) \quad \Rightarrow \quad z \prec_i y$$

where B is the betweenness relation induced by $\preceq^*$.

Examples

- Laws and policies (from LEFT to RIGHT)
- Locations (from FAR to CLOSE)
- Dimensions (from SMALL to BIG)

- NB: no money is involved in such decision!
- NB: no interpersonal comparison needed!

Possibility of Strategy-Proofness

Theorem (Pairwise majority). Take an Implementation problem $(\mathfrak{S}, \mathcal{G})$ s.t. $|\texttt{Agn}|$ is odd. The direct revelation mechanism G with outcome function $\texttt{g}$ being *pairwise majority voting* is an incentive compatible social choice rule under DSE.

Theorem (Median voter). Take an Implementation problem $(\mathfrak{S}, \mathcal{G})$ s.t. $|\texttt{Agn}|$ is odd. The mechanism G where: i) agents declare their peak; ii) the outcome function $\texttt{g}$ selects the *median voter*'s peak is an incentive compatible social choice rule under DSE.

- They are equivalent mechanisms
- Both select the unique *Condorcet Winner*

Input	Function	Output
$\preceq_1$ ... $\preceq_n$	⟶	a

Single-peaked preferences are sufficient to yield the possibility of incentive-compatible social choice functions (e.g., pairwise majority, median voter rule)

... Moral of the Story

- MD moves from the acknowledgment of two related impossibilities:
  - NO non-dictatorial incentive compatible social choice functions;
  - NO DSE-implementations of non-dictatorial functions.
- MD is developed by *restricting* the type of allowed preferences.