Classic Mechanism Design (II)

David C. Parkes

Division of Engineering and Applied Science,

Harvard University

CS 286r–Spring 2002

Direct Revelation Mechanisms (DRM)

In a DRM, $\mathcal{M} = (\Theta, g)$, the strategy space, $S = \Theta$, and an agent simply reports a type to the mechanism, with outcome rule, $g : \Theta \rightarrow \mathcal{O}$.

Def. [**incentive-compatible**] A DRM is (Bayes-)Nash *incentive-compatible* if truth-revelation is a (Bayes-)Nash equilibrium, *i.e.* $s_i^*(\theta_i) = \theta_i$, for all $\theta \in \Theta$.

Def. [**strategyproof**] A DRM is *strategyproof* if truth-revelation is a dominant strategy eq., for all $\theta \in \Theta$.

Note: The SCF implemented by an *incentive-compatible* DRM is precisely the outcome function, $g(\theta)$.

"$g(\theta)$ is truthfully implementable..."

The Revelation Principle

[Gibbard 73;Green & Laffont 77, Myerson 79]

Thm. For any mechanism, $\mathcal{M}$, there is a direct, IC mechanism with the same outcome.

“the computations that go on within the mind of any bidder in the nondirect mechanism are shifted to become part of the mechanism in the direct mechanism.” [McAfee&McMillan, 87]

Consider:

–the IC direct-revelation implementation of a first-price sealed-bid auction

–the IC direct-revelation implementation of an English (ascending-price) auction.

Proof

Consider mechanism, $\mathcal{M} = (S, g)$, that implements SCF, $f(\theta)$, in a dominant strategy equilibrium. In otherwords, $g(s^*(\theta)) = f(\theta)$, for all $\theta \in \Theta$, where s^* is a dominant strategy eq.

Construct direct mechanism, $\mathcal{M}' = (\Theta, f(\theta))$. By contradicton, suppose:

$$\exists \theta_i' \neq \theta_i \text{ s.t. } u_i(f(\theta_i', \theta_{-i}), \theta_i) > u_i(f(\theta_i, \theta_{-i}), \theta_i)$$

for some $\theta_i' \neq \theta_i$, some θ_{-i}. But, because $f(\theta) = g(s^*(\theta))$, this implies that

$$u_i(g(s_i^*(\theta_i'), s_{-i}^*(\theta_{-i})), \theta_i) > u_i(g(s^*(\theta_i), s_{-i}^*(\theta_{-i})), \theta_i)$$

which contradicts the strategyproofness of s^* in mechanism, $\mathcal{M}$.

Theoretical Implications

- **Focus goals.** If $\mathcal{M}$ is the only DRM that implements outcome function $k(\theta)$ with properties $\mathcal{P}$ then *any* mechanism must implement the same transfers as $\mathcal{M}$.
- **Impossibility.** If *no* DRM, $\mathcal{M}$, can implement SCF, $f(\theta)$, with properties $\mathcal{P}$, then no mechanism can implement SCF $f(\theta)$.

A modeler can limit the search for an optimal mechanism to the class of direct, IC mechanisms. Useful, because the number of mechanisms is huge.

Math Programming Approach

[Myerson, "Optimal auction design", Math. Oper. Res., 1981]

Consider a single-item allocation problem with N agents, let $p_i(v_1, \ldots, v_N)$ denote the expected payment from agent i to the mechanism and $x_i(v_1, \ldots, v_N)$ denote the probability with which agent i is allocated the item. Let v_0 denote the value of the seller.

$$\max_{\{p_i, x_i\}} \sum_{i=1}^{N} E\left[p_i(v_1, \ldots, v_N)\right] - \sum_{i=1}^{N} E\left[x_i(v_1, \ldots, v_N)\right] v_0$$

$$\text{s.t.} \quad \sum_{i=1}^{N} x_i(v_1, \ldots, v_N) \leq 1, \quad \forall v \qquad \text{(feas)}$$

$$Eu_i(v_i) \geq Eu_i(\hat{v}_i), \quad \forall \hat{v}_i \neq v_i, \forall v_i, \forall i \qquad \text{(IC)}$$

$$Eu_i(v_i) \geq 0, \quad \forall v_i \qquad \text{(IR)}$$

where

$$Eu_i(\hat{v}_i) = E_{v_{-i}}\left[x_i(v_1, \ldots, \hat{v}_i, \ldots, v_N) v_i\right] - E\left[p_i(v_1, \ldots, \hat{v}_i, \ldots, v_N)\right]$$

Practical Implications?

- Incentive-compatibility is “free”
  – any outcome implemented by mechanism, $\mathcal{M}$, can be implemented by incentive-compatible mechanism, $\mathcal{M}'$.
- “Fancy” mechanisms are unneccessary
  – any outcome implemented by a mechanism with complex strategy space, S, can be implemented by a DRM.

But, few procedures in practical use are direct & IC, perhaps their are some unmodeled costs, computational problems?

Gibbard-Satterthwaite Impossibility

[Arrow 51, Gibbard & Satterthwaite 73, 75]

Consider SCF, $f(\theta)$, and an outcome space $\mathcal{O}$. Let $R_f \subseteq \mathcal{O}$ denote the *range* of f, i.e. $R_f = \{o \in \mathcal{O} \, : \, \exists \theta \in \Theta \, \mathrm{s.t.} \, o = f(\theta)\}$.

Let $o_i^* \in \mathcal{O}$ denote the outcome that maximizes the value, $u_i(o, \theta_i)$, over $o \in R_f$.

Def. [Dictatorial] SCF $f(\theta)$ is dictatorial if there is an agent, h, s.t. $f(\theta) = o_h^*$, for all θ.

[Gibbard-Satterthwaite Impossibility] Suppose that the types include all possible strict orderings over $\mathcal{O}$. A SCF, $f(\theta)$, with $|R_f| > 2$, is implementable in dominant strategies (strategyproof) if and only if it is dictatorial.

Implications: collaborative filtering (Pennock et al.), web query aggregation (Kumar et al.), voting systems (Cranor).

Single-Peaked Preferences

[Moulin 80] [special case in which non dictatorial strategproof implementation is possible]

Def. [single-peaked] Suppose $\mathcal{O} \subseteq \mathbb{R}$, then preferences are single-peaked, if for every i, with utility, $u_i(k, \theta_i)$, there exists a *peak*, $p(\theta_i) \in \mathcal{O}$, s.t. for any $d, d' \in \mathcal{O}$, s.t. $p(\theta_i) \geq d > d'$, or $d' > d \geq p(\theta_i)$, then $u_i(d, \theta_i) > u_i(d', \theta_i)$.

Def. [median rule] Ask agents to declare their peaks, and select the median peak.

Thm. This "median rule" mechanism is Pareto efficient and strategy proof.

all orderings $\supset$ Single-peaked

$$\overset{\longrightarrow}{\text{easier}}$$

Introducing Transfers (Side-payments)

Define the outcome space, $\mathcal{O} = \mathcal{K} \times \mathbb{R}^N$, such that an outcome rule, $o = (k, t_1, \ldots, t_N)$, defines a *choice*, $k(s) \in \mathcal{K}$, and a transfer, $t_i(s) \in \mathbb{R}$ from agent i to the mechanism, given strategy profile $s \in S$.

Assume **quasilinear** preferences,

$$u_i(o, \theta_i) = v_i(k, \theta_i) - t_i$$

, with *valuation function*, $v_i(k, \theta_i)$ for agent i.

General/No-transfer $\supset$ Quasi-linear/Transfer

$$\overrightarrow{\text{easier}}$$

Budget Balance

Introduce constraints over the total transfers made from agents to the mechanism. Let $s^*(\theta)$ denote the equilibrium strategy of a mechanism.

- weak BB (or *feasible*)
  - *ex post:* $\sum_i t_i(s^*(\theta)) \geq 0$, for all θ
  - *ex ante:* $E_{\theta \in \Theta} \left[\sum_i t_i(s^*(\theta)) \right] \geq 0$
- strong BB
  - *ex post:* $\sum_i t_i(s^*(\theta)) = 0$, for all θ
  - *ex ante:* $E_{\theta \in \Theta} \left[\sum_i t_i(s^*(\theta)) = 0 \right]$

ex ante weak	$\supset$	ex post weak	$\downarrow$
$\cup$		$\cup$	
ex ante strong	$\supset$	ex post strong	
	$\longrightarrow$ harder		harder

Efficiency & Budget-balance Tension

[Hurwicz 75;Green & Laffont 79]

Def. [Efficiency] A choice rule, $k^* : \Theta \rightarrow \mathcal{K}$, is (*ex post*) efficient if for all $\theta \in \Theta$, $k^*(\theta)$ maximizes $\sum_{k \in \mathcal{K}} v_i(k, \theta_i)$.

Thm. [Green-Laffont Impossibility] If Θ allows *all* valuation functions from $\mathcal{K}$ to $\mathbb{R}$, then there no mechanism can implement an *efficient* and *ex post* budget-balanced SCF in dominant strategy.

Approaches: (a) restrict space of preferences; (b) drop budget-balance; (c) drop efficiency; (d) drop dominant strategy.

Groves Mechanisms

[Groves 73] **Drop:** budget-balance.

Def. A Groves mechanism, $\mathcal{M} = (\Theta, k, t_1, \ldots, t_N)$ is defined with *choice rule*,

$$k^*(\hat{\theta}) = \arg\max_{k \in \mathcal{K}} \sum_i v_i(k, \hat{\theta}_i)$$

, and *transfer rules*

$$t_i(\hat{\theta}) = h_i(\hat{\theta}_{-i}) - \sum_{j \neq i} v_j(k^*(\hat{\theta}), \hat{\theta}_j)$$

where $h_i(\cdot)$ is an (arbitrary) function that does not depend on the reported type, $\hat{\theta}_i$, of agent i.

Thm. [Groves 73] Groves mechanisms are strategyproof and efficient.

Proof. Agent i's utility for strategy $\hat{\theta}_i$, given $\hat{\theta}_{-i}$ from agents $j \neq i$, is:

$$u_i(\hat{\theta}_i) = v_i(k^*(\hat{\theta}), \theta_i) - t_i(\hat{\theta})$$
$$= v_i(k^*(\hat{\theta}), \theta_i) + \sum_{j \neq i} v_j(k^*(\hat{\theta}), \hat{\theta}_j) - h_i(\hat{\theta}_{-i})$$

Ignore $h_i(\hat{\theta}_{-i})$, and notice

$$k^*(\hat{\theta}) = \arg\max_{k \in \mathcal{K}} \sum_i v_i(k, \hat{\theta}_i)$$

$\Rightarrow$Strategyproofness, and efficiency immediately follow.

In fact, [Green&Laffont 77], Groves mechanisms are **unique**, in the sense that any mechanism that implements efficient choice, $k^*(\theta)$, in truthful dominant strategy must implement Groves transfers.