Complete Information Implementation Theory - Maskin's Theorem and Beyond

Workshop, IPM, Tehran

Arunava Sen, Indian Statistical Institute

November 19,20, 2014

The King Solomon Problem: A Motivating Example

- Two women, referred to as 1 and 2 both claim to be the mother of a child.
- King Solomon has to decide whether (i) to give the child to 1 (outcome a) (ii) to give the child to 2 (outcome b) or (iii) to cut the child in half (outcome c).
- Two "states of the world", θ and ϕ. In state θ, 1 is the real mother while 2 is the impostor; the reverse is true in state ϕ.
- Preferences of the two women over the outcomes $\{a, b, c\}$ depend on the state.

King Solomon contd.

- Preferences of the two women:

State θ		State ϕ	
1	2	1	2
a	b	a	b
b	c	c	a
c	a	b	c

Table : Preferences in states θ and ϕ

- King Solomon's objectives are specified by a social choice function (scf) $f : \{\theta, \phi\} \to \{a, b, c\}$ where $f(\theta) = a$ and $f(\phi) = b$.
- King Solomon does not know which state of the world has occurred.

King Solomon contd.

- He devises a "mechanism" of the following kind. Both women are asked to reveal the state (i.e. the identity of the true mother). If both women agree that state is θ, outcome *a* is enforced; if both agree that it is state ϕ, outcome *b* is enforced; if they disagree outcome *c* is enforced.
- Shown below where 1's messages are shown along the rows and 2's along the columns.

	θ	ϕ
θ	*a*	*c*
ϕ	*c*	*b*

Table : King Solomon's mechanism

- Does this work?
- No!

King Solomon contd.

- In θ, the unique pure strategy Nash equilibrium of this game is for both women to announce ϕ leading to outcome b. Similarly, the equilibrium in state ϕ is for both women to announce θ leading to the outcome a. This mechanism gives the baby to wrong mother in each state!
- QUESTION: Does there exist a better mechanism?

A General Formulation

- $A = \{a, b, c, \ldots\}$: set of outcomes or alternatives.
- $I = \{1, 2, \ldots N\}$: set of agents or players.
- $\Theta = \{\theta, \phi, \psi, \ldots\}$: set of states.
- $R_i(\theta)$: preference ordering of agent i of the elements of A in state θ, i.e $R_i(\theta)$ is a complete, reflexive and antisymmetric binary relation defined on the elements of A.
- A Social Choice Correspondence (scc) F associates a non-empty subset of A denoted by $F(\theta)$ with every state $\theta \in \Theta$. The SCC specifies the objectives of the planner/principal/mechanism designer.
- A *Social Choice Function* (scf) is a singleton-valued scc.

Mechanisms and Implementation

- A *mechanism* G is an $N+1$ tuple $(M_1, M_2 \ldots, M_N; g)$ where M_i, $i = 1, 2 \ldots, N$ is the message set of agent i and g is a mapping $g : M_1 \times \ldots \times M_N \to A$.
- For all $\theta \in \Theta$, the pair (G, θ) constitutes a game in normal form. We let $\mathrm{NE}(G, \theta)$ denote the set of pure strategy Nash equilibria in (G, θ), i.e.
  $\bar{m} \in \mathrm{NE}(G, \theta) \Rightarrow g(\bar{m}) R_i(\theta) g(m_i, \bar{m}_{-i})$ for all $m_i \in M_i$, $i \in I$.
- The mechanism $G \equiv (M_1, M_2, \ldots, M_N)$ implements the scc F if

$$g(\mathrm{NE}(G, \theta)) = F(\theta) \text{ for all } \theta \in \Theta.$$

- We require all Nash equilibria to be optimal according to F. Restrict attention to pure strategy equilibria.
- QUESTION: What are the scc's that can be implemented?

The Information Structure

- Once a state is realized, agents are assumed to play Nash equilibrium i.e. the payoffs are common knowledge to all players except the mechanism designer.
- e.g both women know who the real mother is not King Solomon does not.
- Contrast with incomplete information.
- Two applications:
  1. The private information is not *verifiable* by an outside party (for instance, by a judge or arbiter).
  2. The mechanism has to be put in place before (in the chronological sense) the realization of the state - e.g. electoral procedures.
- Classical questions in equilibrium theory can also be addressed within the framework of this model. For example, can Walrasian equilibrium be attained when there is asymmetric information and a small number of agents? Leonid Hurwicz

Ideas behind Maskin's Mechanism

- For every θ and $a \in F(\theta)$, there must exist a message vector which is a Nash equilibrium under θ whose outcome is a.
- Assume w.l.o.g that this message vector is labeled (a, θ), i.e. when everybody sends the message (a, θ), the outcome is a.
- Any deviation by player i must lead to an outcome in the set $L(a, i, \theta) = \{b \in A | aR_i(\theta)b\}$.
- If $N \geq 3$, it is easy to identify the deviant and "punish" him (i.e. pick an outcome in $L(a, i, \theta)$).
- The mechanism designer has already ensured that $F(\theta) \subset g(\mathrm{NE}(G, \theta))$ for all $\theta \in \Theta$.
- Now he must try to ensure that there are no other equilibria in (G, θ).

Maskin's Mechanism contd.

- Other candidate equilibria (message vectors)
- Message vectors that are non-unanimous, i.e some agents send (a, θ), others (b, ϕ) and yet others (c, ψ) etc.
- Relatively easy to deal with because designer knows that such a message vector does not need to be an equilibrium in any state of the world.
- Try to "destroy" it as an equilibrium by allowing all agents to deviate and get any alternative in A.

Ideas behind Maskin's mechanism contd.

- Message vectors that are unanimous - Everyone sends the message (a, θ).
- True state is however ϕ.
- Designer must continue to behave as though the true state is θ - no way for him to distinguish this situation from the one where the true state is θ.
- Outcome will be a and deviations by player i must lead to an outcome in $L(a, i, \theta)$.
- Therefore if $L(a, i, \theta) \subset L(a, i, \phi)$, then (a, θ) will be an equilibrium under ϕ.
- Therefore, if F is implementable, it must be the case that $a \in F(\phi)$. Property of Maskin-Monotonicity.

Maskin's Theorem

- The scc F satisfies Maskin-Monotonicity (MM)if, for all $\theta, \phi \in \Theta$ and $a \in A$,

$$[a \in F(\theta) \text{ and } L(a,i,\theta) \subset L(a,i,\phi) \text{ for all } i \in I] \Rightarrow [a \in F(\phi)]$$

- Suppose a is F-optimal in state θ. Suppose also that for all agents, all alternatives that are worse than a in θ are also worse than a in ϕ. Then a is also F-optimal in ϕ.
- The scc F satisfies No Veto Power (NVP) if, for all $a \in A$ and $\theta, \phi \in \Theta$,

$$[\#\{i \in I | aR_i(\theta)b \text{ for all } b \in A\} \geq N-1] \Rightarrow [a \in F(\theta)]$$

- If at least $N-1$ agents rank an alternative as maximal in a state, then that alternative must be F-optimal in that state.
- Weak condition. Trivially satisfied in environments where there is a private good that everyone wants more of.

Maskin's Theorem

- Theorem (Maskin 1977) If F is implementable, then F satisfies MM. Assume $N \geq 3$. If F satisfies MM and NVP, then it is implementable.
- MM is necessary and "almost" sufficient for implementation.

Understanding Maskin Monotonicity

The following sccs are Maskin Monotonic.

1. The Pareto Efficient Correspondence.
2. The Walrasian Correspondence in exchange economies provided that all Walrasian allocations are interior.
3. The Individually Rational Correspondence in exchange economies.
4. The Pareto Efficient and Individually Rational Correspondence. In general, the intersection of two MM sccs is also MM, i.e if F and G are MM sccs and $F(\theta) \cap G(\theta) \neq \emptyset$ for all $\theta \in \Theta$, then $F \cap G$ also satisfies MM.
5. The dictatorship scc. There exists an agent, say i such that for all $\theta \in \Theta$, $d(\theta) = \{a \in A | aR_i(\theta)b \text{ for all } b \in A\}$.

Understanding Maskin Monotonicity contd.

The following are examples of sccs that violate MM.

1. The King Solomon scc.
2. Scoring methods, for example, the plurality rule. Let $A = \{a, b, c, d\}$ and $I = \{1, 2, 3, 4, 5\}$.

State θ					State ϕ				
1	2	3	4	5	1	2	3	4	5
a	a	b	c	d	a	a	b	b	b
b	b	c	b	b	b	b	c	c	c
c	c	d	d	c	c	c	d	d	d
d	d	a	a	a	d	d	a	a	a

Table : Preferences in states θ and ϕ

- $a = F(\theta)$ and $b = F(\phi)$. However $aR_i(\theta)x \rightarrow aR_i(\phi)x$ for all $x \in A$ and for all $i \in I$. Hence MM is violated.

Understanding Maskin Monotonicity contd.

3. The class of scfs satisfying MM over “large domains” is small. For instance, if one considers scfs defined over the domain of all strict orderings, the only ones which satisfy MM and the “full range” condition are the dictatorial ones. Over the domain of all orderings, only the constant scf satisfies MM.

- Overall conclusion: MM is a strong condition.

$N = 2$

- Extra condition required to ensure that equilibria can be sustained.
- Suppose agent 1 sends the message (a, θ) while 2 sends the message (b, ϕ) where $a \in F(\theta)$ and $b \in F(\phi)$. It could be that 1 is deviating unilaterally from the Nash equilibrium which supports (b, ϕ) or that 2 is deviating unilaterally from the Nash equilibrium which supports (a, θ).
- Resulting outcome must not upset either equilibrium, i.e. it must be an alternative in *both* $L(b, 1, \phi)$ and $L(a, 2, \theta)$.
- New necessary condition (which does not appear in the $N \geq 3$ case):
  for all $\theta, \phi \in \Theta$ and $a, b \in A$ such that $a \in F(\theta)$ and $b \in F(\phi)$
  $$L(b, 1, \phi) \cap L(a, 2, \theta) \neq \emptyset$$
- Other conditions are also required for implementation.

Subgame Perfect Implementation

- A mechanism is specified in extensive form - (finite) game tree, a player partition, an information partition and a mapping which associates elements of A with every terminal node of the tree.
- A mechanism Γ together with a state $\theta \in \Theta$ is a game in extensive form.
- Let $\mathrm{SPE}(\Gamma, \theta)$ denote the set of subgame pefect equilibrium outcomes of (Γ, θ).
- SCC F can be implemented if there exists a mechanism Γ such that $F(\theta) = \mathrm{SPE}(\Gamma, \theta)$ for all $\theta \in \Theta$.
- Scc that can be implemented in Nash equilibrium implies it can be implemented in SPNE (using Maskin)
- But can more be implemented?

Example

- $A = \{a, b, c\}$, $I = \{1, 2, 3\}$ and $\Theta = \{ \text{ all strict orderings over } A\}$.
- scf: $f(\theta) = \arg\max_{R_1(\theta)}\{a, \text{ majority winner over } \{b, c\}\}$.
- f is not monotonic.

State θ			State ϕ		
1	2	3	1	2	3
b	c	c	b	b	b
a	b	b	a	c	c
c	a	a	c	a	a

Table : Preferences in states θ and ϕ

- $f(\theta) = a$ and $L(a, i, \theta) \subset L(a, i, \phi)$ for all $i \in I$. However $f(\phi) = b$. Therefore f does not satisfy MM and is not implementable.

Example

- f can be implemented in SPNE.
- Extensive mechanism in Figure below.

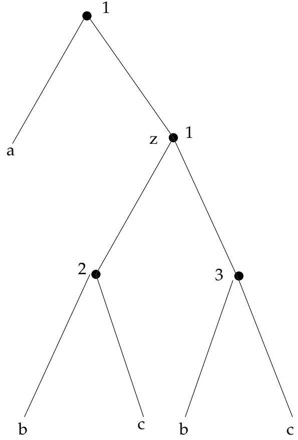

Figure : Tree implementing f

SPNE: A General Result

- How much does the class of implementable sccs expand when we consider SPNE rather than Nash implementation?
- Answer: "Quite a lot".
- Assume $N \geq 3$ and suppose the environment satisfies certain mild conditions ("economic domain"). Let F be any Pareto-efficient, interior scc. Then F can be implemented in SPNE. Furthermore, only two rounds are required.
- Result due to Moore-Repullo (1988) and Abreu-Sen (1989)

Discussion

- Large body of literature which shows that the scope for implementation increases very dramatically if
  - (i) the solution concept is “refined” from Nash to subgame-perfect Nash
  - (ii) iterated elimination of weakly dominated strategies
  - (iii) trembling-hand perfect equilibrium
  - (iv) randomization is allowed in the mechanism and the notion of implementation is weakened to concepts like “virtual implementation”
- Here we follow a behavioural approach - players are assumed to have a “slight” preference for honesty.

Implementation with Partially Honest Players

- Approach was initiated by Matsushima. Here we follow the approach of Dutta-Sen (2012).
- Idea: people strictly prefer to tell the truth *provided* their material payoffs are not affected.
- Let $G = (M, g)$ be a mechanism where for each $i \in N$, $M_i = \Theta \times S_i$ and S_i denotes the other components of the message space.
- For each i and $\theta \in \Theta$, let $T_i(\theta) = \theta \times S_i$. For any θ and i, we interpret $m_i \in T_i(\theta)$ as a *truthful* message as player i's reporting the true state of the world.

Partial Honesty contd.

- Extend a player's ordering over A to an ordering over the entire message space M as follows.
- True state θ. Players other than i send the message m_{-i}.
- Two possible messages $m_i, m_{-i} \in M_i$ where m_i is a truthful message. If i is *indifferent* between $g(m_i, m_{-i})$ and $g(m_i', m_{-i})$ according to $R_i(\theta)$, the she *strictly prefers* sending m_i; otherwise she prefers sending the message that is better according to $R_i(\theta)$.
- A player has a "lexicographic" ordering over M in each state θ. Dominant component is self-interest as given by $R_i(\theta)$; subordinate dominated component, which involves a preference for truth-telling.
- Let $\succeq_i^\theta$ denote i's ordering over M in state θ with its asymmetric component denoted by $\succ_i^\theta$.

Partial Honesty contd.

- Assumption A: There exists at least one partially honest individual and this fact is known to the planner. However, the identity of this individual may not be known to her
- Dutta-Sen (2012): Let $N \geq 3$ and assume that A holds. Then, every SCC satisfying NVP can be implemented.

Proof

- Pick a SCC F satisfying NVP. We use a mechanism very similar to that used to prove Maskin's Theorem.
- Let $M_i = \{(a_i, \theta_i, n_i, c_i) \in A \times \Theta \times \mathbb{N}_+ \times A \times A | a_i \in F(\theta_i)\}$, $i \in I$.

  The mapping $g : M_1 \times \ldots, M_N \rightarrow A$ is described as follows:

  (i) if $m_i = (a, \theta, ., ., .)$ for at least $N - 1$ players i, then $g(m) = a$.

  (iii) if (i) does not hold, then $g(m) = c_k$ where k is the lowest index in the set of agents who announce the highest integer , i.e $k = \arg\min\{i \in I | n_i \geq n_j \text{ for all } j \neq i\}$.

Proof contd.

- Suppose "true" state is ϕ and let $a \in F(\phi)$.
- m where $m_i = (a, \phi, ., .)$ is a Nash equilibrium under θ because no unilateral deviation can change the outcome - Each player is announcing the truth.
- To show that $NE(G, \succeq^\theta) \subseteq F(\theta)$.
- Consider arbitrary m. Suppose no more than $(N-1)$ individuals announce the same state of the world θ (where θ may be distinct from ϕ), the same $a \in F(\phi)$. Let the outcome be some b. Then, any one of $(N-1)$ individuals can deviate, precipitate (ii) and get any outcome she wants. Clearly, if the original announcement is to be a Nash equilibrium, then it must be the case that b is θ-maximal for $(N-1)$ individuals. By NVP $b \in F(\theta)$.

Proof contd.

- Suppose that *all* players unanimously announce θ, $b \in F(\theta)$, where $\theta \neq \phi$. Then, $g(m) = b$.
- Let i be a partially honest individual. i can deviate to the truthful announcement of $R\phi$, that is to some $\bar{m}_i(\phi) \in T_i(\phi)$. By construction $g(\bar{m}_i, m_{-i}) = b$. Since i is partially honest $g(\bar{m}_i, m_{-i}) \succ_i^{\theta} g(\bar{m}_i, m_{-i})$. Hence m is not a Nash equilibrium.