Foundations Multidimensional Mechanism Design

Debasis Mishra
Indian Statistical Institute, Delhi

January 2016

Outline

- What is (multidimensional) mechanism design?
- Lessons from one-dimensional mechanism design.
- Description of incentive compatible mechanisms.
- Revenue equivalence result.

The Setting

A finite set of agents and each agent has some **private information**.

A planner/mechanism designer takes two decision based on the (reported) private information:

- an *allocation* decision
- a *payment* decision

Once the decisions are taken, agent realizes utility, *quasilinear* in payments:

Utility from allocation - payment

Examples - Private and Public Good Allocation

- Single object sale: who to allocate and how much price to charge
- Multiple objects sale: who to allocate which object and how much to charge each agent
- Opening a public facility (a bridge or library or university): whether to open it or where to open it and how much to tax/subsidize each agent.

Reverse Engineering a Game

Mechanism designer does not know the private information of agents.

If he knew the private information, he knows what decisions he would have taken.

Needs to elicit information by providing incentives - design *mechanisms* or a *game-form*.

Design the game-form in a manner such that equilibrium of the resulting game coincides with what the designer wanted to do.

A One-Dimensional Example

A time slot (on a server) of 1 unit needs to be allocated to an agent.

The agent has value v for the slot and if he receives only $\alpha \in [0, 1]$ amount of time, then he gets a utility of αv and if he is asked to pay p, then he gets a net utility of

$$\alpha v - p.$$

The designer does not know how much value v agent has, but knows that it is some non-negative real number.

A Reserve Price Mechanism

The designer announces a *reserve price* r and asks agent his value.

If the reported value is more than r, then the entire slot is given to the agent. Else, no time is given to the agent.

If no time is given to the agent, then the agent pays zero. Else, he pays the reserve price.

The agent cannot increase his utility by reporting a false value - the mechanism is *incentive compatible*.

What is a mechanism?

Complicated to define a general mechanism.

But consider the following kind of mechanisms (direct mechanisms).

What is a mechanism?

Complicated to define a general mechanism.

But consider the following kind of mechanisms (direct mechanisms).

A **mechanism** is a pair of functions (f, p), where

$$f : V \rightarrow [0, 1]$$

and

$$p : V \rightarrow \mathbb{R},$$

where V is the set of all possible values - assume it to be an interval, say $[0, \beta]$.

Incentive Compatibility

Definition

A mechanism (f, p) *is* **incentive compatible** *if for every* $v, v' \in V$*, we have*

$$f(v) \cdot v - p(v) \geq f(v') \cdot v - p(v').$$

Incentive Compatibility

Definition

A mechanism (f, p) *is* **incentive compatible** *if for every* $v, v' \in V$*, we have*

$$f(v) \cdot v - p(v) \geq f(v') \cdot v - p(v').$$

Truth-telling is a *weakly dominant strategy*.

Revelation Principle. If there is a mechanism with a non-truthful "equilibrium", then there is a direct mechanism with a truthful equilibrium achieving the same outcome.

Question

What kind of mechanisms are incentive compatible?

Rewriting Incentive Compatibility

Given a mechanism $M \equiv (f, p)$, denote the net utility of truth-telling of an agent with value v as:

$$\mathcal{U}^M(v) := f(v) \cdot v - p(v).$$

Incentive compatibility is equivalent to requiring: for all $v, v' \in V$,

$$\begin{aligned} \mathcal{U}^M(v) &\geq f(v') \cdot v - p(v') \\ &= f(v') \cdot (v - v') + f(v') \cdot v' - p(v') \\ &= \mathcal{U}^M(v') + f(v') \cdot (v - v'). \end{aligned}$$

Convexity of Net Utility

Lemma

If $M \equiv (f, p)$ is incentive compatible, then $\mathcal{U}^M$ is a convex function.

Convexity of Net Utility

Lemma
If $M \equiv (f, p)$ is incentive compatible, then $\mathcal{U}^M$ is a convex function.

Remind: Convex functions are continuous in its interior and differentiable almost everywhere.

So a mechanism $M \equiv (f, p)$ is incentive compatible if the function $\mathcal{U}^M$ satisfies for all $v, v' \in V$,

$$\mathcal{U}^M(v) \geq \mathcal{U}^M(v') + f(v') \cdot (v - v').$$

Careful Look

For every $v, v' \in V$, we need to satisfy incentive constraints $v \longrightarrow v'$ and $v' \longrightarrow v$:

$$\mathcal{U}^M(v) \geq \mathcal{U}^M(v') + f(v') \cdot (v - v').$$
$$\mathcal{U}^M(v') \geq \mathcal{U}^M(v) + f(v) \cdot (v' - v).$$

Hence, we have

$$f(v) \cdot (v - v') \geq \mathcal{U}^M(v) - \mathcal{U}^M(v') \geq f(v') \cdot (v - v')$$

As $v' \longrightarrow v$, we see that $f(v)$ is the derivative of $\mathcal{U}^M$ at v (if it is differentiable at v).

Allocation is Subgradient of Net Utility

For convex functions, $f(v)$ will be termed a *subgradient* of $\mathcal{U}^M$ at v: they are not decreasing functions.

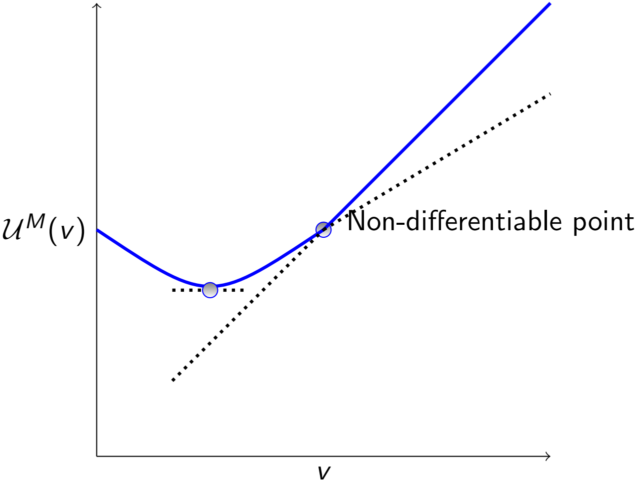

Figure: A convex function and its subgradients

Fundamental Theorem of Convex Analysis (Calculus)

Every convex function can be written as a definite integral of its subgradient.

Fundamental Theorem of Convex Analysis (Calculus)

Every convex function can be written as a definite integral of its subgradient.

For every $v \in [0, \beta]$,

$$\mathcal{U}^M(v) = \mathcal{U}^M(0) + \int_0^v f(x)dx.$$

Fundamental Theorem of Convex Analysis (Calculus)

Every convex function can be written as a definite integral of its subgradient.

For every $v \in [0, \beta]$,

$$\mathcal{U}^M(v) = \mathcal{U}^M(0) + \int_0^v f(x)dx.$$

Alternatively,

$$p(v) = p(0) + f(v) \cdot v - \int_0^v f(x)dx.$$

Two Necessary Conditions for Incentive Compatibility

f is non-decreasing.

For every $v \in [0, \beta]$,

$$\mathcal{U}^M(v) = \mathcal{U}^M(0) + \int_0^v f(x)dx.$$

Two Necessary Conditions for Incentive Compatibility

f is non-decreasing.

For every $v \in [0, \beta]$,

$$\mathcal{U}^M(v) = \mathcal{U}^M(0) + \int_0^v f(x)dx.$$

They are also sufficient.

Myerson Characterization

Theorem (Myerson)

A mechanism $M \equiv (f, p)$ *is incentive compatible if and only if*

- **Monotonicity.** *for all* $v, v' \in [0, \beta]$ *with* $v > v'$*, we have* $f(v) \geq f(v')$.
- **Payoff Equivalence.** *for all* $v \in [0, \beta]$*,*

$$\mathcal{U}^M(v) = \mathcal{U}^M(0) + \int_0^v f(x)dx.$$

Proof of Sufficiency

Remind: incentive compatibility requires us to show for all $v, v' \in [0, \beta]$,

$$\mathcal{U}^M(v) \geq \mathcal{U}^M(v') + f(v') \cdot (v - v').$$

Proof of Sufficiency

Remind: incentive compatibility requires us to show for all $v, v' \in [0, \beta]$,

$$\mathcal{U}^M(v) \geq \mathcal{U}^M(v') + f(v') \cdot (v - v').$$

$$\mathcal{U}^M(v) - \mathcal{U}^M(v') = \int_{v'}^{v} f(x)dx \geq f(v') \cdot (v - v').$$

Why is it Useful?

Monotonicity is an easy condition to check.

Payoff (revenue) equivalence is very powerful: Payoff (from truth-telling) at every type is uniquely determined upto an additive constant by the allocation rule.

If we take two incentive compatible mechanisms (f, p) and (f, p'), they can only differ in the payment at the lowest value.

If we can set the payment at the lowest value, then the every incentive compatible mechanism can be written in terms of f - search space is the space of all non-decreasing f.

Implementation

Since a mechanism is almost uniquely pinned down by an allocation rule, we frame the question slightly differently.

Definition

An allocation rule f is **implementable** *if there exists a payment rule p such that (f, p) is incentive compatible.*

An immediate corollary of Myerson's theorem is that f is implementable if and only if it is non-decreasing - p that implements f is almost uniquely determined by the revenue/payoff equivalence formula.

Beyond this Example

Techniques developed here works for almost all one-dimensional mechanism design problems.

Techniques extend to multiple agents case also (more later).

It forms the core of optimization done in mechanism design - expected revenue maximizing mechanism, budget-balanced mechanism, fair mechanism etc.

They are also starting point for analyzing multidimensional mechanism design problems.

Multidimensional Mechanism Design

Private information of agents is no longer a single real number.

- **Dichotomous Preference.** Agent has some feasible dates of travel - travel gives agent a value only his feasible day but gives zero value on other days.

  Value and *feasible dates* are private information.

- **Multi object auction.** Multiple objects for sale. Agent has value for each bundle of objects.

- **Locating a public good.** A public good (bridge) can be located at many places. Each location gives a different value to the agent.

The One Agent Model

A finite set of alternatives $A = \{a, b, c, \ldots\}$ - set of objects, set of locations etc.

Private information of agent is a value *vector* $v \in \mathbb{R}^{|A|}$.

A **mechanism** $M \equiv (f, p)$ is

$$f : V \rightarrow \mathcal{L}(A), \quad p : V \rightarrow \mathbb{R},$$

where $\mathcal{L}(A)$ are the set of all lotteries over A and V is the set of all possible values.

The One Agent Model Contd.

An allocation rule is deterministic if $f_a(v) \in \{0, 1\}$ for all a and for all $v \in V$.

If the agent has value v and reports v', then he gets a net utility of

$$f(v') \cdot v - p(v').$$

Usual Examples

A represents the set of all objects and agent can be given one of the objects. $v(a)$ represents the value for object a.

A represents the set of all *bundles* of objects and agent can be given any bundle. $v(a)$ represents the value for bundle a.

A represents the set of all possible locations of the public good - only one location will be chosen. $v(a)$ represents the value for location a.

Example - Dichotomous Preference

A is the set of all travel dates.

Private information consists of (S, x), where $S \subseteq A$ are the feasible travel dates and $x \in \mathbb{R}_+$ is the value of travel.

The value vector $v \in \mathbb{R}^{|A|}$ will consist of a vector where $v(a) = x$ if $a \in S$ and $v(a) = 0$ otherwise.

Graphical Illustration

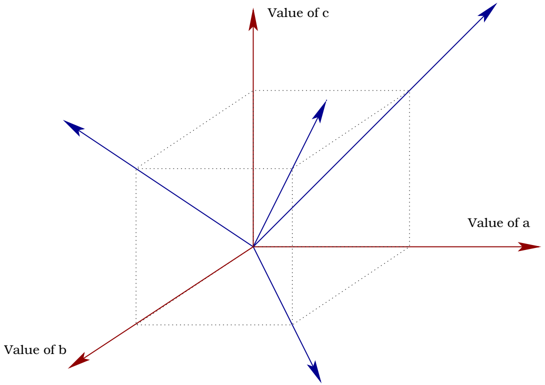

Figure: Dichotomous values

Example - Scheduling

A is the dates of a month.

A firm (agent) needs delivery of a product in that month. The firm has an "ideal" date - any date away from it gives him less value.

Note: Ideal date is not known and none of the values are known.

Value vector is *single-peaked* - values away from ideal date are increasingly worse.

Example - Scheduling (contd.)

Suppose three dates $a \succ b \succ c$.

Then, any value v can have four possibilities:

$$v(a) \geq v(b) \geq v(c)$$
$$v(c) \geq v(b) \geq v(a)$$
$$v(b) \geq v(a) \geq v(c)$$
$$v(b) \geq v(c) \geq v(a).$$

What is not allowed:

$$v(a) > v(c) > v(b)$$
$$v(c) > v(a) > v(b).$$

Graphical Illustration

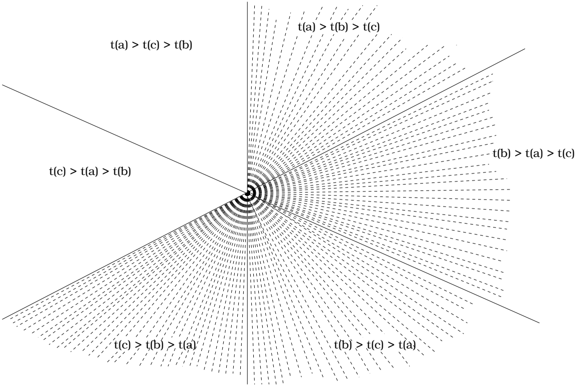

Figure: Single Peaked Values

Value Space

The domain of values play an important role in mechanism design - defines the set of incentive constraints we need to worry about.

Mechanism designer has complete knowledge of domain of values.

The structure of domain of values determine the structure of incentive compatible mechanisms.

An Example

Two alternatives $\{a, b\}$. Domain of values is the colored region (except the boundaries).

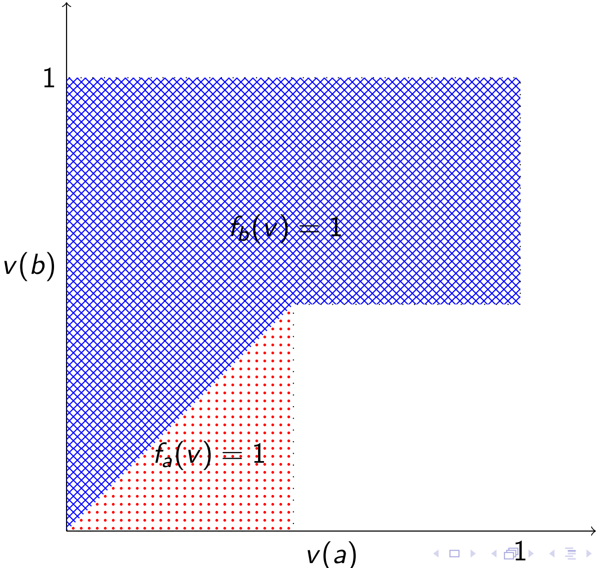

A Mechanism for this Example

f as shown and $p(v) = 0$ for all v.

Is this incentive compatible?

- In red region, $v(a) > v(b)$ and $f_a(v) = 1$ - so agent will not manipulate.
- In blue region, above 45-degree line, $v(b) > v(a)$ and $f_b(v) = 1$ - so agent will not manipulate.
- In blue region, below 45-degree line, $v(b) < v(a)$ and $f_b(v) = 1$ - so agent can manipulate to red region.

Manipulating in Both Dimensions

This manipulation requires agent to change his values on *a and b*.

Incentive constraints where agent only changes *either a* or *b* value are satisfied.

Moral of the story: *Incentive constraints in both dimensions have to be taken care of - not one dimension at a time*.

Incentive Constraints

The second example is again with $A = \{a, b\}$. Consider an allocation rule f which either picks a or picks b (no other lottery) in some domain V.

Define

$$V_a^f := \{v \in V : f_a(v) = 1\}$$

and

$$V_b^f := \{v \in V : f_b(v) = 1\}.$$

If (f, p) is an incentive compatible mechanism and $v, v' \in V$ such that $f(v) = f(v')$, we have $p(v) = p(v')$.

Hence, there are two numbers q_a and q_b such that for all $v \in V_a^f$, we have $p(v) = q_a$ and for all $v \in V_b^f$, we have $p(v) = q_b$.

Incentive Constraints

Two sets of incentive constraints. For every $v \in V_a^f$ and every $v' \in V_b^f$, we must have

$$v(a) - q_a \geq v(b) - q_b$$
$$v'(b) - q_b \geq v'(a) - q_a.$$

Rewriting this, we note that

$$\inf_{v \in V_a^f} \left[v(a) - v(b) \right] \geq q_a - q_b$$
$$\inf_{v' \in V_b^f} \left[v'(b) - v'(a) \right] \geq q_b - q_a.$$

A Necessary Condition

A necessary condition on f can be immediately found by adding these incentive constraints.

$$\inf_{v \in V_a^f} \left[v(a) - v(b) \right] + \inf_{v' \in V_b^f} \left[v'(b) - v'(a) \right] \geq 0.$$

Now, consider an arbitrary mechanism (f, p). Suppose f satisfies this condition.

Then, setting $q_a = 0$ and

$$q_b = \inf_{v' \in V_b^f} \left[v'(b) - v'(a) \right],$$

we can construct an incentive compatible mechanism

Sufficiency

We need to ensure

$$\inf_{v \in V_a^f} \left[v(a) - v(b) \right] \geq q_a - q_b$$

$$\inf_{v' \in V_b^f} \left[v'(b) - v'(a) \right] \geq q_b - q_a.$$

Sufficiency

We need to ensure

$$\inf_{v \in V_a^f} \left[v(a) - v(b) \right] \geq q_a - q_b$$
$$\inf_{v' \in V_b^f} \left[v'(b) - v'(a) \right] \geq q_b - q_a.$$

Second inequality is satisfied trivially. For first, note

$$q_a - q_b = - \inf_{v' \in V_b^f} \left[v'(b) - v'(a) \right] \leq \inf_{v \in V_a^f} \left[v(a) - v(b) \right],$$

as desired by the first inequality.

Observations

Rewriting: For every $v \in V_a^f$ and every $v' \in V_b^f$, we have

$$v(a) - v(b) \geq v'(a) - v'(b),$$

Different from Myerson monotonicity, but related.

We look at *differences of values* between the two alternatives.

To see this, suppose this condition is not true:

$$v(b) - v(a) > v'(b) - v'(a).$$

But $f_b(v') = 1$ and $f_a(v') = 1$.

Questions

What are necessary and sufficient conditions for f to be implementable?

Does payoff/revenue equivalence still hold? What is the analogue of Myerson's theorem?

These questions map out the search space of incentive compatible mechanisms.

Incentive Compatibility

Definition

A mechanism $M \equiv (f, p)$ *is* **incentive compatible** *if for every* $v, v' \in V$*, we have*

$$v \cdot f(v) - p(v) \geq v \cdot f(v') - p(v'),$$

or equivalently,

$$\mathcal{U}^M(v) \geq \mathcal{U}^M(v') + (v - v') \cdot f(v').$$

An allocation rule f *is* **implementable** *if there exist a payment rule* p *such that* (f, p) *is incentive compatible.*

Extending Monotonicity

Definition

An allocation rule f is **monotone** *if for every $v, v' \in V$, we have*

$$(v - v') \cdot \big(f(v) - f(v')\big) \geq 0.$$

Why this extension?

- A necessary condition: derived by adding incentive constraints

$$\mathcal{U}^M(v) \geq \mathcal{U}^M(v') + (v - v') \cdot f(v')$$
$$\mathcal{U}^M(v') \geq \mathcal{U}^M(v) + (v' - v) \cdot f(v).$$

- Reduces to Myerson monotonicity if v is one-dimensional.
- If f is *deterministic* (only 0/1 probabilities), the condition reduces to

$$v(f(v)) - v(f(v')) \geq v'(f(v)) - v'(f(v')).$$

Myerson Extended

Theorem

Suppose $V \subseteq \mathbb{R}^{|A|}$ is convex. A mechanism $M \equiv (f, p)$ is incentive compatible if and only if

(a) *f is monotone,*

(b) *for every $v, v' \in V$,*

$$\mathcal{U}^M(v') = \mathcal{U}^M(v) + \int_0^1 \psi^{v,v'}(z)dz,$$

where $\psi^{v,v'}(z) = (v' - v) \cdot f(v + z(v' - v))$ for all $z \in [0, 1]$.

Necessity

The necessity of these conditions:

- f monotone follows by adding incentive constraints.
- Second condition follows by looking at the line segment joining v and v' (a one-dimensional space) and applying Myerson.
- To be able to apply incentive constraints along all the points on the line segment, we need convexity of V.

Sufficiency

Sufficiency is established by showing that

- ψ is non-decreasing if f is monotone.
- Then,

$$
\begin{aligned}
&\mathcal{U}^M(v') - \mathcal{U}^M(v) - (v' - v) \cdot f(v) \\
&= \int_0^1 \psi^{v,v'}(z)dz - (v' - v) \cdot f(v) \\
&\geq \psi^{v,v'}(0) - (v' - v) \cdot f(v) \\
&= 0,
\end{aligned}
$$

Good News: Payoff Equivalence

If we fix at $\mathcal{U}^M$ some v_0, it fixes the $\mathcal{U}^M(v)$ for all v via f.

So, if we have (f, p) and (f, p'), then for all $v, v' \in V$,

$$p(v) - p(v') = p'(v) - p'(v').$$

Good News: Payoff Equivalence

If we fix at $\mathcal{U}^M$ some v_0, it fixes the $\mathcal{U}^M(v)$ for all v via f.

So, if we have (f, p) and (f, p'), then for all $v, v' \in V$,

$$p(v) - p(v') = p'(v) - p'(v').$$

However, the revenue equivalence formula here is stronger - it needs to hold for every pair of values v, v' as compared to v, v_0 for some v_0 and for all v.

The Implementability Question

Since payoff equivalence holds, can we characterize implementable allocation rules?

Conjecture: If V is convex, f is implementable if and only if it is monotone.

The Implementability Question

Since payoff equivalence holds, can we characterize implementable allocation rules?

Conjecture: If V is convex, f is implementable if and only if it is monotone.

The conjecture is *false* in general.

The reason is: fixing $\mathcal{U}^M(v_0)$ for some v_0 and finding $\mathcal{U}^M(v)$ for all v using the payoff equivalence formula **does not** imply that payoff equivalence formula is true for all v, v'.

Stronger Conditions

Consider a sequence of values $v^1, \ldots, v^k$. Consider the incentive constraints along the cycle: $v_1 \to v^2 \to v^3 \to \ldots \to v^k \to v^1$:

$$
\begin{aligned}
v^1 \cdot f(v^1) - p(v^1) &\geq v^1 \cdot f(v^2) - p(v^2) \\
v^2 \cdot f(v^2) - p(v^2) &\geq v^2 \cdot f(v^3) - p(v^3) \\
\ldots &\geq \ldots \\
\ldots &\geq \ldots \\
v^{k-1} \cdot f(v^{k-1}) - p(v^{k-1}) &\geq v^{k-1} \cdot f(v^k) - p(v^k) \\
v^k \cdot f(v^k) - p(v^k) &\geq v^k \cdot f(v^1) - p(v^1).
\end{aligned}
$$

Adding these payments cancel and you get a necessary condition for f.

Cycle Monotonicity

For every allocation rule f and for every v, v', define

$$\ell^f(v, v') := v' \cdot (f(v') - f(v)).$$

Definition

*An allocation rule f is K-**cycle monotone** if for all $k \leq K$ and for all cycle of values $(v^1, \ldots, v^k, v^{k+1} \equiv v^1)$, we have*

$$\sum_{j=1}^{k} \ell^f(v^j, v^{j+1}) = 0.$$

*An allocation rule f is **cyclically monotone** if is K-cycle*

monotone for all K.

Note monotonicity is 2-cycle monotonicity.

Graph Interpretation

Construct a complete directed graph with set of nodes as V (which can be finite or infinite).

For every directed edge (v, v'), the length is $\ell^f(v, v')$.

Cycle monotonicity requires that every cycle of this directed graph has non-negative length.

The Rochet-Rockafellar Theorem

Theorem

An allocation rule f is implementable if and only if it is cyclically monotone.

The Rochet-Rockafellar Theorem

Theorem

An allocation rule f is implementable if and only if it is cyclically monotone.

Necessity is by adding incentive constraints along each cycle.

Sufficiency is by constructing an explicit payment rule.

Fix any value (node) v_0 and set $p(v_0) = 0$, then for every $v \neq v_0$, set $p(v) = s(v_0, v)$ - the shortest path from v_0 to v.

These shortest paths are finite because of no negative length cycle.

Payments

The construction hints at many possible payment rules for the same f (satisfying cycle monotonicity).

Revenue equivalence says that they should differ from each other by a constant.

Revenue equivalence holds if V is convex (earlier theorem). Anywhere else?

Revenue Equivalence

Definition

An allocation rule $f : V \to \mathcal{L}(A)$ *satisfies* **revenue equivalence** *if for every* p, p' *such that* (f, p) *and* (f, p') *are incentive compatible, we have*

$$p(v) - p'(v) = p(v') - p'(v') \ \ \forall \ v, v' \in V.$$

Revenue Equivalence

Definition

An allocation rule $f : V \rightarrow \mathcal{L}(A)$ *satisfies* **revenue equivalence** *if for every* p, p' *such that* (f, p) *and* (f, p') *are incentive compatible, we have*

$$p(v) - p'(v) = p(v') - p'(v') \;\; \forall \, v, v' \in V.$$

Theorem

An implementable (cyclically monotone) allocation rule $f : V \rightarrow \mathcal{L}(A)$ *satisfies revenue equivalence if and only if*

$$s(v, v') + s(v', v) = 0 \; \forall \; v, v' \in V.$$

Revenue Equivalence

Definition

An allocation rule $f : V \to \mathcal{L}(A)$ *satisfies* **revenue equivalence** *if for every* p, p' *such that* (f, p) *and* (f, p') *are incentive compatible, we have*

$$p(v) - p'(v) = p(v') - p'(v') \;\; \forall \; v, v' \in V.$$

Theorem

An implementable (cyclically monotone) allocation rule $f : V \to \mathcal{L}(A)$ *satisfies revenue equivalence if and only if*

$$s(v, v') + s(v', v) = 0 \; \forall \; v, v' \in V.$$

This condition is satisfied if (a) V is convex or (b) if f is deterministic and V is connected.

Deterministic Allocation Rules

Cycle monotonicity is a very cumbersome condition to use.

Deterministic Allocation Rules

Cycle monotonicity is a very cumbersome condition to use.

Theorem

Suppose f is a deterministic allocation rule. Then, f is implementable if and only if it $|A|$-cycle monotone.

Still not very useful because we need to verify *all* cycles of length $|A|$.

Simpler Conditions

Theorem

Suppose f is a deterministic allocation rule. Further, suppose V is either (a) convex or (b) contains all single-peaked value vectors. Then, f is implementable if and only if it is (2-cycle) monotone.

Simpler Conditions

Theorem

Suppose f is a deterministic allocation rule. Further, suppose V is either (a) convex or (b) contains all single-peaked value vectors. Then, f is implementable if and only if it is (2-cycle) monotone.

Theorem

Suppose f is a deterministic allocation rule. Further, suppose V contains all dichotomous value vectors. Then, f is implementable if and only if it is 3-cycle monotone.

Remind - Single Peaked Values

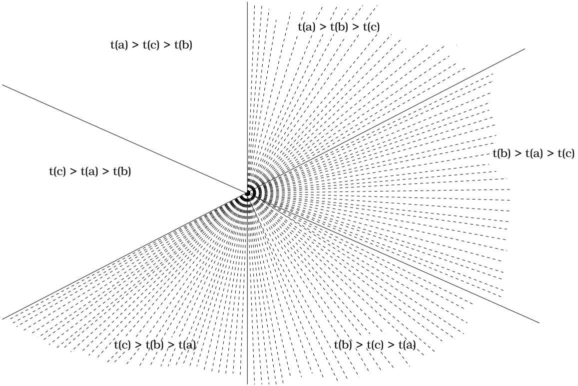

Figure: Single Peaked Values

Remind - Dichotomous Values

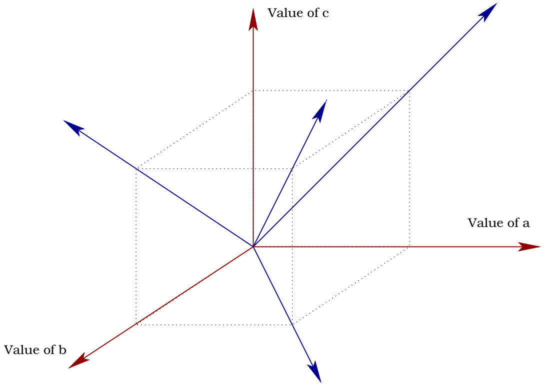

Figure: Dichotomous values

Extension to Many Agents

Suppose the set of agents is $N = \{1, \ldots, n\}$.

Every agent i has a space of values V_i and let $V = V_1 \times \ldots \times V_n$.

A mechanism is $(f, p_1, \ldots, p_n)$, where

$$f : V \to \mathcal{L}(A), \quad p_i : V \to \mathbb{R}, \ \forall\ i \in N.$$

Definition

A mechanism $(f, p_1, \ldots, p_n)$ *is* **dominant strategy incentive compatible (DSIC)** *if for every* $i \in N$*, for every* v_{-i}*, and for every* $v_i, v_i' \in V_i$*, we have*

$$f(v_i, v_{-i}) \cdot v_i - p_i(v_i, v_{-i}) \geq f(v_i', v_{-i}) \cdot v_i - p_i(v_i', v_{-i}).$$

How to Extend?

With more agents and DSIC, we fix an agent i and values of other agents at v_{-i}.

The incentive constraints are then defined for agent i once we fix it like this.

Hence, the problem is similar to a collection of "many" one-agent problems.

All the results need to be stated for all $i \in N$ and for all v_{-i}.

Beyond DSIC

DSIC is a very demanding solution concept - irrespective of what other agents report, an agent must have (weakly) dominant strategy to report his true value.

Weaker solution concepts possible: Bayesian IC - this requires defining what each agent believes other agents' values are.

Some *but not all* results can be extended with Bayesian IC.

Summary

The structure of set of incentive compatible mechanisms is complicated - unlike one dimensional problems.

However, monotonicity and revenue equivalent results of Myerson extend to many interesting multidimensional domains.

Optimizing over the set of all incentive compatible mechanisms becomes hard because of the complicated structure.

New ideas are needed that can utilize known structure to cleverly define and solve optimization problems.