CS599: Algorithm Design in Strategic Settings
Fall 2012
Lecture 3: Mechanism Design Preliminaries

Instructor: Shaddin Dughmi

Administrivia

- HW out soon (monday), due in two weeks
- Office hours next week rescheduled
- Email list
- Announcements on class page

Outline

Outline

Rationality

Some of you asked for a formalization of rationality…

Definition

A utility function on choice set A is a map $u : A \to \mathbb{R}$.

Definition

When choice set A is a family of lotteries over some other choice set B, a utility function $u : A \to \mathbb{R}$ is a Von-Neumann Morgenstern utility function if there is a utility function $v : B \to \mathbb{R}$ over B such that $u(a) = \mathrm{E}_{b \sim a}[v(b)]$.

Rationality

Some of you asked for a formalization of rationality…

Definition

A utility function on choice set A is a map $u : A \to \mathbb{R}$.

Definition

When choice set A is a family of lotteries over some other choice set B, a utility function $u : A \to \mathbb{R}$ is a Von-Neumann Morgenstern utility function if there is a utility function $v : B \to \mathbb{R}$ over B such that $u(a) = \mathrm{E}_{b \sim a}[v(b)]$.

We assume agents are equipped with VNM utility functions over (distributions over) outcomes of a game / mechanism, and moreover they act to maximize (expected) utility.

Definition

A rational agent always chooses the element of his choice set maximizing his (expected) utility.

Arguments in Favor of Nash Equilibrium

- MWG has a nice discussion
- Favorite arguments: self-enforcing agreement, stable social convention

Outline

Single-item Allocation

- n players
- Player i's private data (type): $v_i \in \mathbb{R}_+$
- Outcome: choice of a winning player, and payment from each player
- Utility of a player for an outcome is his value for the outcome if he wins, less payment

Objectives: Revenue, welfare.

Single-item Allocation

First Price Auction

1. Collect bids
2. Give to highest bidder
3. Charge him his bid

Single-item Allocation

Second-price (Vickrey) Auction

1. Collect bids
2. Give to highest bidder
3. Charge second highest bid

Example: Public Project

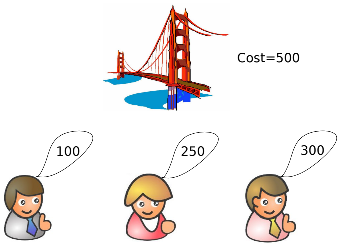

- n players
- Player i's private data (type): $v_i \in \mathbb{R}_+$
- Outcome: choice of whether or not to build, and payment from each player covering the cost of the project if built
- Utility of a player for an outcome is his value for the project if built, less his payment

Goal: Build if sum of values exceeds cost

Shortest Path Procurement

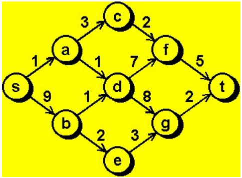

- Players are edges in a network, with designated source/sink
- Player i's private data (type): cost $c_i \in \mathbb{R}_+$
- Outcome: choice of s-t shortest path to buy, and payment to each player
- Utility of a player for an outcome is his payment, less his cost if chosen.

Goal: buy path with lowest total cost (welfare), or buy a path subject to a known budget, ...

Example: Voting

- n players
- m candidates
- Player i's private data (type): total preference order on candidates
- Outcome: choice of winning candidate

Goal: ??

Outline

Recall: Incomplete Information Game

A game of strict incomplete information is a tuple (N, A, T, u), where

- N is a finite set of players. Denote $n = |N|$ and $N = \{1, \dots, n\}$.
- $A = A_1 \times \dots A_n$, where A_i is the set of actions of player i. Each $\vec{a} = (a_1, \dots, a_n) \in A$ is called an action profile.
- $T = T_1 \times \dots T_n$, where T_i is the set of types of player i. Each $\vec{t} = (t_1, \dots, t_n) \in T$ is called an type profile.
- $u = (u_1, \dots u_n)$, where $u_i : T_i \times A \to \mathbb{R}$ is the utility function of player i.

For a Bayesian game, add a common prior $\mathcal{D}$ on types.

Recall: Incomplete Information Game

A game of strict incomplete information is a tuple (N, A, T, u), where

- N is a finite set of players. Denote $n = |N|$ and $N = \{1, \ldots, n\}$.
- $A = A_1 \times \ldots A_n$, where A_i is the set of actions of player i. Each $\vec{a} = (a_1, \ldots, a_n) \in A$ is called an action profile.
- $T = T_1 \times \ldots T_n$, where T_i is the set of types of player i. Each $\vec{t} = (t_1, \ldots, t_n) \in T$ is called an type profile.
- $u = (u_1, \ldots u_n)$, where $u_i : T_i \times A \to \mathbb{R}$ is the utility function of player i.

For a Bayesian game, add a common prior $\mathcal{D}$ on types.

Example: Vickrey Auction

- $A_i = \mathbb{R}$ is the set of possible bids of player i.
- $T_i = \mathbb{R}$ is the set of possible values for the item.
- For $v_i \in T_i$ and $b \in A$, we have $u_i(v_i, b) = v_i - b_{-i}$ if $b_i > b_{-i}$, otherwise 0.

Strategies in Incomplete Information Games

- Strategies of player i
  - Pure strategy $s_i : T_i \to A_i$: a choice of action $a_i \in A_i$ for every type $t_i \in T_i$.
    - Example: Truthtelling is a strategy in the Vickrey Auction
    - Example: Bidding half your value is also a strategy
  - Mixed strategy: a choice of distribution over actions A_i for each type $t_i \in T_i$
    - Won't really use... all our applications will involve pure strategies

Strategies in Incomplete Information Games

- Strategies of player i
  - Pure strategy $s_i : T_i \rightarrow A_i$: a choice of action $a_i \in A_i$ for every type $t_i \in T_i$.
    - Example: Truthtelling is a strategy in the Vickrey Auction
    - Example: Bidding half your value is also a strategy
  - Mixed strategy: a choice of distribution over actions A_i for each type $t_i \in T_i$
    - Won't really use... all our applications will involve pure strategies

Note

In a strategy, player decides how to act based only on his private info (his type), and NOT on others' private info nor their actions.

Equilibria

$s_i : T_i \to A_i$ is a dominant strategy for player i if, for all $t_i \in T_i$ and $a_{-i} \in A_{-i}$ and $a'_i \in A_i$,

$$u_i(t_i, (s_i(t_i), a_{-i})) \geq u_i(t_i, (a'_i, a_{-i}))$$

Equivalently: $s_i(t_i)$ is a best response to $s_{-i}(t_{-i})$ for all t_i, t_{-i} and s_{-i}.

Illustration: Vickrey Auction

Vickrey Auction

Consider a Vickrey Auction with incomplete information.

Illustration: Vickrey Auction

Vickrey Auction

Consider a Vickrey Auction with incomplete information.

Claim

The truth-telling strategy is dominant for each player.

Bayes-Nash Equilibrium

As before, a strategy s_i for player i is a map from T_i to A_i. Now, we define the extension of Nash equilibrium to this setting.

A pure Bayes-Nash Equilibrium of a Bayesian Game of incomplete information is a set of strategies $s_1, \ldots, s_n$, where $s_i : T_i \to A_i$, such that for all i, $t_i \in T_i$, $a'_i \in A_i$ we have

$$\mathop{\mathbf{E}}_{t_{-i} \sim \mathcal{D} | t_i} u_i(t_i, s(t)) \geq \mathop{\mathbf{E}}_{t_{-i} \sim \mathcal{D} | t_i} u_i(t_i, (a'_i, s_{-i}(t_{-i})))$$

where the expectation is over t_{-i} drawn from p after conditioning on t_i.

- Note: Every dominant strategy equilibrium is also a Bayes-Nash Equilibrium
- But, unlike DSE, BNE is guaranteed to exist.

Example: First Price Auction

Example: First Price Auction

- $A_i = T_i = [0, 1]$
- $u_i(v_i, b) = v_i - b_i$ if $b_i > b_j$ for all $j \neq i$, otherwise 0.
- $\mathcal{D}$ draws each $v_i \in T_i$ independently from $[0, 1]$.

Show that the strategies $b_i(v_i) = v_i/2$ form a Bayes-Nash equilibrium.

Outline

General Form

Mechanism Design Setting (Prior-free)

Given by a tuple $(N, \mathcal{X}, T, u)$, where

- N is a finite set of players. Denote $n = |N|$ and $N = \{1, \ldots, n\}$.
- $\mathcal{X}$ is a set of outcomes.
- $T = T_1 \times \ldots T_n$, where T_i is the set of types of player i. Each $\vec{t} = (t_1, \ldots, t_n) \in T$ is called an type profile.
- $u = (u_1, \ldots u_n)$, where $u_i : T_i \times \mathcal{X} \to \mathbb{R}$ is the utility function of player i.

General Form

Mechanism Design Setting (Prior-free)

Given by a tuple $(N, \mathcal{X}, T, u)$, where

- N is a finite set of players. Denote $n = |N|$ and $N = \{1, \ldots, n\}$.
- $\mathcal{X}$ is a set of outcomes.
- $T = T_1 \times \ldots T_n$, where T_i is the set of types of player i. Each $\vec{t} = (t_1, \ldots, t_n) \in T$ is called an type profile.
- $u = (u_1, \ldots u_n)$, where $u_i : T_i \times \mathcal{X} \to \mathbb{R}$ is the utility function of player i.

In a Bayesian setting, supplement with a distribution $\mathcal{D}$ over T

General Form

Mechanism Design Setting (Prior-free)

Given by a tuple $(N, \mathcal{X}, T, u)$, where

- N is a finite set of players. Denote $n = |N|$ and $N = \{1, \ldots, n\}$.
- $\mathcal{X}$ is a set of outcomes.
- $T = T_1 \times \ldots T_n$, where T_i is the set of types of player i. Each $\vec{t} = (t_1, \ldots, t_n) \in T$ is called an type profile.
- $u = (u_1, \ldots u_n)$, where $u_i : T_i \times \mathcal{X} \to \mathbb{R}$ is the utility function of player i.

In a Bayesian setting, supplement with a distribution $\mathcal{D}$ over T

Example: Single-item Allocation

- Outcome: choice $x \in \{e_1, \ldots, e_n\}$ of winning player, and payment $p_1, \ldots, p_n$ from each
- Type of player i: value $v_i \in \mathbb{R}_+$.
- $u_i(v_i, x) = v_i x_i - p_i$.

Social Choice Functions

A principal wants to communicate with players and aggregate their private data (types) into a choice of outcome. Such aggregation captured by

A social choice function $f : T \to \mathcal{X}$ is a map from type profiles to outcomes.

Social Choice Functions

A principal wants to communicate with players and aggregate their private data (types) into a choice of outcome. Such aggregation captured by

A social choice function $f : T \to \mathcal{X}$ is a map from type profiles to outcomes.

Choosing a Social Choice Function

- A particular social choice function in mind (e.g. majority voting, utilitarian allocation of a single item, etc).
- An objective function $o : T \times \mathcal{X} \to \mathbb{R}$, and want $f(T)$ to (approximately) maximize $o(T, f(T))$
  - Either worst case over T (Prior-free) or in expectation (Bayesian)

Example: Single-item Allocation

- Welfare objective: $welfare(v, (x, p)) = \sum_i v_i x_i$
- Revenue objective: $revenue(v, (x, p)) = \sum_i p_i$

Mechanisms

To perform such aggregation, the principal runs a protocol, known as a mechanism. Formally,

A mechanism is a pair (A,g), where

- $A = A_1 \times \ldots A_n$, where A_i is the set of possible actions (think messages, or bids) of player i in the protocol. $\mathcal{A}$ is the set of action profiles.
- $g : A \to \mathcal{X}$ is an outcome function

Mechanisms

To perform such aggregation, the principal runs a protocol, known as a mechanism. Formally,

A mechanism is a pair (A, g), where

- $A = A_1 \times \ldots A_n$, where A_i is the set of possible actions (think messages, or bids) of player i in the protocol. $\mathcal{A}$ is the set of action profiles.
- $g : A \to \mathcal{X}$ is an outcome function

The resulting game of mechanism design is a game of incomplete information where when players play $a \in A$, player i's utility is $u_i(t_i, g(a))$ when his type is t_i.

Example: First price auction

- $A_i = \mathbb{R}$
- $g(b_1, \ldots, b_n) = (x, p)$ where $x_{i^*} = 1$, $p_{i^*} = b_{i^*}$ for $i^* = \text{argmax}_i\, b_i$, and $x_i = p_i = 0$ for $i \neq i^*$.

Implementation of Social Choice Functions

We say a mechanism (A, g) implements social choice function $f : T \to \mathcal{X}$ in dominant-strategy [Bayes-Nash] equilibrium if there is a strategy profile $s = (s_1, \dots, s_n)$ with $s_i : T_i \to A_i$ such that

- $s_i : T_i \to A_i$ is a dominant-strategy [Bayes-Nash] equilibrium in the resulting incomplete information game
- $g(s_1(t_1), s_2(t_2), \dots, s_n(t_n)) = f(t_1, t_2, ..., t_n)$ for all $t \in T$

Example: First price, two players, i.i.d $U[0, 1]$

Implements in BNE the following social choice function: give the item to the player with the highest value and charges him half his value.

Example: Vickrey Auction

Implements in DSE the following social choice function: give the item to the player with the highest value and charges him the second highest value.

The Task of Mechanism Design

Task of Mechanism Design (Take 1)

Given a notion of a "good" social choice function from T to X, find

- A mechanism
  - An action space $A = (A_1, \ldots, A_n)$,
  - an outcome function $g : A \to \mathcal{X}$,
- an equilibrium $(s_1, \ldots, s_n)$ of the resulting game of mechanism design

such that the social choice function $f(t_1, \ldots, t_n) = g(s_1(t_1), \ldots, s_n(t_n))$ is "good."

The Task of Mechanism Design

Task of Mechanism Design (Take 1)

Given a notion of a "good" social choice function from T to X, find

- A mechanism
  - An action space $A = (A_1, \ldots, A_n)$,
  - an outcome function $g : A \to \mathcal{X}$,
- an equilibrium $(s_1, \ldots, s_n)$ of the resulting game of mechanism design

such that the social choice function $f(t_1, \ldots, t_n) = g(s_1(t_1), \ldots, s_n(t_n))$ is "good."

Problem

This seems like a complicated, multivariate search problem.

The Task of Mechanism Design

Task of Mechanism Design (Take 1)

Given a notion of a "good" social choice function from T to X, find

- A mechanism
  - An action space $A = (A_1, \ldots, A_n)$,
  - an outcome function $g : A \to \mathcal{X}$,
- an equilibrium $(s_1, \ldots, s_n)$ of the resulting game of mechanism design

such that the social choice function $f(t_1, \ldots, t_n) = g(s_1(t_1), \ldots, s_n(t_n))$ is "good."

Problem

This seems like a complicated, multivariate search problem.

Luckily

The revelation principle reduces the search space to just $g : T \to \mathcal{X}$.

Outline

Incentive-Compatibility

Direct Revelation

A mechanism (A, g) is a direct revelation mechanism if $A_i = T_i$ for all i.

i.e. in a direct revelation mechanism, players simultaneously report types (not necessarily truthfully) to the mechanism. Such mechanisms can simply be described via the function $g : T \to \mathcal{X}$.

Incentive-Compatibility

A direct-revelation mechanism is dominant-strategy [Bayesian] incentive-compatible (aka truthful) if the truth-telling is a dominant-strategy [Bayes-Nash] equilibrium in the resulting incomplete-information game.

Note: A direct revelation incentive-compatible mechanism implements its outcome function $g : T \to \mathcal{X}$, by definition.

The social choice function IS the mechanism!!

Examples

Vickrey Auction

Direct revelation mechanism, dominant-strategy incentive-compatible.

First Price Auction

Direct revelation mechanism, not Bayesian incentive compatible.

Example: Posted price

The auction that simply posts a fixed price to players in sequence until one accepts is not direct revelation.

Revelation Principle

Revelation Principle

If there is a mechanism implementing social choice function f in dominant-strategy [Bayes-Nash] equilibrium, then there is a direct revelation, dominant-strategy [Bayesian] incentive-compatible mechanism implementing f.

Revelation Principle

Revelation Principle

If there is a mechanism implementing social choice function f in dominant-strategy [Bayes-Nash] equilibrium, then there is a direct revelation, dominant-strategy [Bayesian] incentive-compatible mechanism implementing f.

This simplifies the task of mechanism design

Task of Mechanism Design (Take 2)

Given a notion of a "good" social choice function from T to X, find such a function $f : T \to X$ such that truth-telling is an equilibrium in the following mechanism:

- Solicit reports $\widetilde{t_i} \in T_i$ from each player i (simultaneous, sealed bid)
- Choose outcome $f(\widetilde{t_1}, \ldots, \widetilde{t_n})$

Example

2 players, with values i.i.d uniform from $[0, 1]$, facing the first-price auction.

First-price Auction

1. Solicit bids b_1, b_2
2. Give item to highest bidder, charging him his bid

Recall

The strategies where each player reports half their value are in BNE. In other words, when player 1 knows his value v_1, and faces player 2 who is bidding uniformly from $[0, 1/2]$, he maximizes his expected utility $(v_1 - b_1).2b_1$ by bidding $b_1 = v_1/2$. And vice versa.

Example

2 players, with values i.i.d uniform from $[0, 1]$, facing the first-price auction.

First-price Auction

1. Solicit bids b_1, b_2
2. Give item to highest bidder, charging him his bid

Recall

The strategies where each player reports half their value are in BNE. In other words, when player 1 knows his value v_1, and faces player 2 who is bidding uniformly from $[0, 1/2]$, he maximizes his expected utility $(v_1 - b_1).2b_1$ by bidding $b_1 = v_1/2$. And vice versa.

Therefore . . .

the first price auction implements in BNE the social choice function which gives the item to the highest bidder, and charges him half his bid

Example

Modified First-price Auction

1. Solicit bids b_1, b_2
2. Give item to highest bidder, charging him half his bid
   - Equivalently, simulate a first price auction where bidders bid $b_1/2, b_2/2$

Claim

Truth-telling is a BNE in the modified first-price auction.

Therefore, the modified auction implements the same social-choice function in equilibrium, but is truthful.

Example

Modified First-price Auction

1. Solicit bids b_1, b_2
2. Give item to highest bidder, charging him half his bid
   - Equivalently, simulate a first price auction where bidders bid $b_1/2, b_2/2$

Claim

Truth-telling is a BNE in the modified first-price auction.

Therefore, the modified auction implements the same social-choice function in equilibrium, but is truthful.

Proof

Assume player 2 bids truthfully. Player 1 faces a (simulated) first price auction where his own bid is halved before participating, and player 2 bids uniformly from $[0, 1/2]$. To respond optimally in the simulation, he bids $b_1 = v_1$ and lets the mechanism halve his bid on his behalf.

Proof (Bayesian Setting)

Consider mechanism (A, g), with BNE strategies $s_i : T_i \rightarrow A_i$.

- Implements $f(t_1, \ldots, t_n) = g(s_1(t_1), \ldots, s_n(t_n))$ in BNE
- For all i and t_i, action $s_i(t_i)$ maximizes player i's expected utility when other players are playing $s_{-i}(t_{-i})$ for $t_{-i} \sim \mathcal{D}|t_i$.

Proof (Bayesian Setting)

Consider mechanism (A, g), with BNE strategies $s_i : T_i \to A_i$.

- Implements $f(t_1, \ldots, t_n) = g(s_1(t_1), \ldots, s_n(t_n))$ in BNE
- For all i and t_i, action $s_i(t_i)$ maximizes player i's expected utility when other players are playing $s_{-i}(t_{-i})$ for $t_{-i} \sim \mathcal{D}|t_i$.

Modified Mechanism

1. Solicit reported types $\widetilde{t}_1, \ldots, \widetilde{t}_n$
2. Choose outcome $f(\widetilde{t}_1, \ldots, \widetilde{t}_n) = g(s_1(\widetilde{t}_1), \ldots, s_n(\widetilde{t}_n))$
    - Equivalently, simulate (A, g) when players play $s_i(t_i)$

Proof (Bayesian Setting)

Consider mechanism (A, g), with BNE strategies $s_i : T_i \to A_i$.

- Implements $f(t_1, \ldots, t_n) = g(s_1(t_1), \ldots, s_n(t_n))$ in BNE
- For all i and t_i, action $s_i(t_i)$ maximizes player i's expected utility when other players are playing $s_{-i}(t_{-i})$ for $t_{-i} \sim \mathcal{D}|t_i$.

Modified Mechanism

1. Solicit reported types $\widetilde{t}_1, \ldots, \widetilde{t}_n$
2. Choose outcome $f(\widetilde{t}_1, \ldots, \widetilde{t}_n) = g(s_1(\widetilde{t}_1), \ldots, s_n(\widetilde{t}_n))$
    - Equivalently, simulate (A, g) when players play $s_i(t_i)$

- Assume all players other than i report truthfully
- When i's type is t_i, other players playing $s_{-i}(t_{-i})$ for $t_{-i} \sim \mathcal{D}|t_i$ in simulated mechanism
- As stated above, his best response in simulation is $s_i(t_i)$.
- Mechanism transforms his bid by applying s_i, so best to bid t_i.

Outline

Mechanism Design Impossibilities

The revelation principle reduces mechanism design to the design of direct-revelation, truthful mechanisms.

Mechanism Design Impossibilities

The revelation principle reduces mechanism design to the design of direct-revelation, truthful mechanisms.

Unfortunately...

Absent structure on the outcome space and utility functions, no reasonably good mechanisms exist even in simple settings.

Examples coming up: single-item allocation without payments, voting

Mechanism Design Impossibilities

The revelation principle reduces mechanism design to the design of direct-revelation, truthful mechanisms.

Unfortunately...

Absent structure on the outcome space and utility functions, no reasonably good mechanisms exist even in simple settings.

Examples coming up: single-item allocation without payments, voting

Luckily

The structure that enables much of mechanism design is assuming that the outcome space incorporates monetary payments, and player utilities are linear in these payments.

Single-item Allocation Without Money

Question

- Consider allocating a single item among n players, with private values (types) $v_1, \ldots, v_n \in \mathbb{R}_+$ for the item, without access to monetary payments.
- Restricted to mechanisms that implement their social choice function in dominant strategies.
- What is the smallest worst-case approximation ratio for social welfare of such a mechanism? Prove it.
- WLOG by revelation principle: restrict attention to dominant-strategy truthful mechanisms $f : \mathbb{R}_+^n \to \{1, \ldots, n\}$.

The worst-case approximation ratio of mechanism f for social welfare is defined as

$$\max_{v \in \mathbb{R}_+^n} \frac{\max_i v_i}{v_{f(v)}}$$

Single-item Allocation Without Money

Question

- Consider allocating a single item among n players, with private values (types) $v_1, \ldots, v_n \in \mathbb{R}_+$ for the item, without access to monetary payments.
- Restricted to mechanisms that implement their social choice function in dominant strategies.
- What is the smallest worst-case approximation ratio for social welfare of such a mechanism? Prove it.
- WLOG by revelation principle: restrict attention to dominant-strategy truthful mechanisms $f : \mathbb{R}_+^n \to \{1, \ldots, n\}$.

The smallest worst-case approximation ratio is n. No mechanism can guarantee better than $1/n$ fraction of the optimal social welfare in dominant strategy equilibrium!

Voting

Recall: voting

- n players
- m candidates
- Player i's private data (type): total preference order on candidates
- Outcome: choice of winning candidate

Theorem (Gibbard-Satterthwaite)

Assume the number of candidates C is at last 3. Consider a voting mechanism implementing allocation rule $f : \Sigma^n \to C$ in dominant strategies. Either f is a dictatorship or some candidate can never win in f.

Outline

Incorporating Payments

To make much of modern mechanism design possible, we assume that

- The set of outcomes has a particular structure: every outcome includes a payment to and from each player.
- Player utilities vary linearly with their payment.

Examples: Single-item allocation, public project, shortest path procurement
Non-examples: Single-item allocation without money, voting.

Quasilinear Utilities

The Quasi-linear Setting

Formally, $\mathcal{X} = \Omega \times \mathbb{R}^n$.

- Ω is the set of allocations
- For $(\omega, p_1, \ldots, p_n) \in \mathcal{X}$, p_i is the payment from (or to) player i.

and player $i's$ utility function $u_i : T_i \times \mathcal{X} \to \mathbb{R}$ takes the following form

$$u_i(t_i, (\omega, p_1, \ldots, p_n)) = v_i(t_i, \omega) - p_i$$

for some valuation function $v_i : T_i \times \Omega \to \mathbb{R}$.

We say players have quasilinear utilities.

Example: Single-item Allocation

- $\Omega = \{e_1, \ldots, e_n\}$
- $u_i(t_i, (\omega, p_1, \ldots, p_n)) = t_i \omega_i - p_i$

Further simplification

Recall that, using the revelation principle, we got

Task of Mechanism Design (Take 2)

Given a notion of a "good" social choice function from T to X, find such a function $f : T \to X$ such that truth-telling is an equilibrium in the following mechanism:

- Solicit reports $\widetilde{t_i} \in T_i$ from each player i (simultaneous, sealed bid)
- Choose outcome $f(\widetilde{t_1}, \ldots, \widetilde{t_n})$

Further simplification

In quasilinear settings this breaks down further

Task of Mechanism Design in Quasilinear settings

Find a "good" allocation rule $f : T \to \Omega$ and payment rule $p : T \to \mathbb{R}^n$ such that the following mechanism is incentive-compatible:

- Solicit reports $\widetilde{t_i} \in T_i$ from each player i (simultaneous, sealed bid)
- Choose allocation $f(\widetilde{t})$
- Charge player i payment $p_i(\widetilde{t})$

We think of the mechanism as the pair (f, p).
Sometimes, we abuse notation and think of type t_i directly as the valuation $v_i : \Omega \to \mathbb{R}$.

Incentive-Compatibility

Incentive compatibility can be stated simply now

Incentive-compatibility (Dominant Strategy)

A mechanism (f, p) is dominant-strategy truthful if, for every player i, true type t_i, possible mis-report $\widetilde{t}_i$, and reported types t_{-i} of the others, we have

$$v_i(t_i, f(t)) - p_i(t) \geq v_i(t_i, f(\widetilde{t}_i, t_{-i})) - p_i(\widetilde{t}_i, t_{-i})$$

If (f, p) randomized, add expectation signs.

Incentive-Compatibility

Incentive compatibility can be stated simply now

Incentive-compatibility (Dominant Strategy)

A mechanism (f, p) is dominant-strategy truthful if, for every player i, true type t_i, possible mis-report $\widetilde{t}_i$, and reported types t_{-i} of the others, we have

$$v_i(t_i, f(t)) - p_i(t) \geq v_i(t_i, f(\widetilde{t}_i, t_{-i})) - p_i(\widetilde{t}_i, t_{-i})$$

If (f, p) randomized, add expectation signs.

Incentive-compatibility (Bayesian)

A mechanism (f, p) is Bayesian incentive compatible if, for every player i, true type t_i, possible mis-report $\widetilde{t}_i$, the following holds in expectation over $t_{-i} \sim D|t_i$

$$\mathbf{E}[v_i(t_i, f(t)) - p_i(t)] \geq \mathbf{E}[v_i(t_i, f(\widetilde{t}_i, t_{-i})) - p_i(\widetilde{t}_i, t_{-i})]$$

Examples

Vickrey Auction

- Allocation rule maps $b_1, \ldots, b_n$ to e_{i^*} for $i^* = \mathrm{argmax}_i \, b_i$
- Payment rule maps $b_1, \ldots, b_n$ to $p_1, \ldots, p_n$ where $p_{i^*} = b_{(2)}$, and $p_i = 0$ for $i \neq i^*$.

Dominant-strategy truthful.

First Price Auction

- Allocation rule maps $b_1, \ldots, b_n$ to e_{i^*} for $i^* = \mathrm{argmax}_i \, b_i$
- Payment rule maps $b_1, \ldots, b_n$ to $p_1, \ldots, p_n$ where $p_{i^*} = b_{(1)}$, and $p_i = 0$ for $i \neq i^*$.

For two players i.i.d $U[0, 1]$, players bidding half their value is a BNE.
Not Bayesian incentive compatible.

Examples

Modified First Price Auction

- Allocation rule maps $b_1, \ldots, b_n$ to e_{i^*} for $i^* = \text{argmax}_i\, b_i$
- Payment rule maps $b_1, \ldots, b_n$ to $p_1, \ldots, p_n$ where $p_{i^*} = b_{(1)}/2$, and $p_i = 0$ for $i \neq i^*$.

For two players i.i.d $U[0, 1]$, Bayesian incentive compatible.