Mechanism Design - Introduction

- An MDP allowed us to model a sequential decision-making problem with uncertainty over future inputs, which we model using a known stochastic process

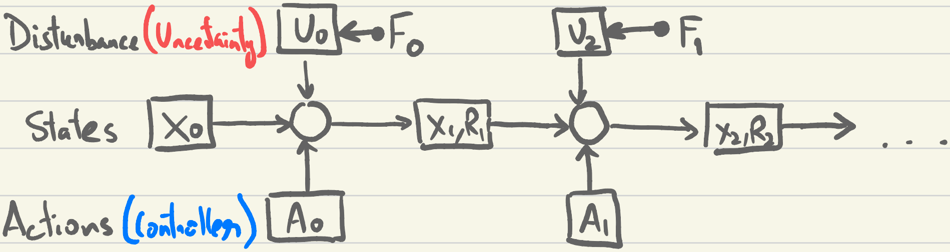

- Mechanism design typically deals with static (i.e., one-shot) decisions in settings where inputs are private to strategic agents, who participate in a way that maximizes their selfish interests.

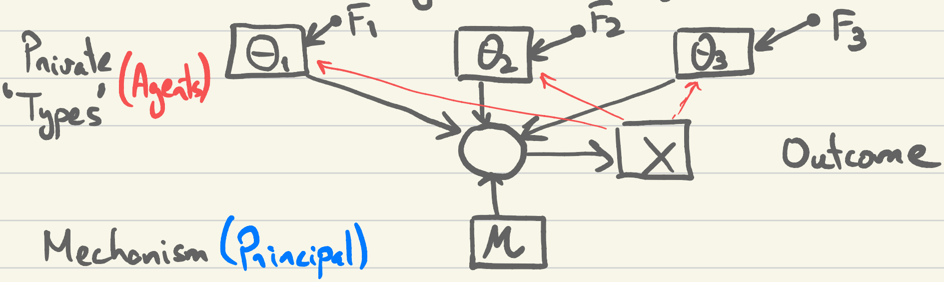

- Can be thought of as a three-stage process

① a principal commits to some mechanism (i.e., a protocol for communicating and making decisions)

② next, agents with private types participate in a way which maximizes their personal utilities

③ Finally the principal chooses an outcome (or allocation) at which point the principal and each agent realizes their own utilities (as a fn of outcome and types)

Mechanisms and Solution Concepts

- **Setting**
  - n agents $[n] = \{1, 2, \ldots, n\}$
  - $\mathcal{X} \equiv$ set of feasible allocations ← assume finite
  - Each agent i has a private type $\theta_i \in \Theta$
    - (Bayesian setting) $\theta_i \sim F_i$, independent, known
      'Independent private valuations'
    - θ_i may also be correlated in more complex models
  - Each agent i has utility fn $u_i : \Theta \times \mathcal{X} \to \mathbb{R}$
    we will denote this as $u_i(x \mid \theta)$
  - Agent i knows θ_i, and all the F_i
    Principal knows all F_i

- **Mechanism**: Comprises of 2 components
  1) Collection of Action Sets $\{A_i\}_{i \in [n]}$, one for each agent
  2) Outcome fn $\phi : A_1 \times A_2 \times \ldots \times A_n \to \mathcal{X}$

- Agent i is assumed to commit to only playing actions in A_i.
  Agent i strategy $\sigma_i : \Theta \to A_i$; Strategy profile $\underline{\sigma} = \{\sigma_i\}$

- **Solution Concept (Equilibrium)** - 'Behavioral' assumptions on how agents choose strategies
  - Dominant Strategy Equilibrium (DSE): A strategy profile $\underline{\sigma}$ where $\forall i$, σ_i is (weakly) dominant compared to any other $\hat{\sigma}_i$, irrespective of $\sigma_{-i} = \{\sigma_j\}_{j \neq i}$
  - Nash Equilibrium - Strategy profile $\underline{\sigma}$ where $\forall i$, σ_i is a best response to σ_{-i}
  - Bayes-Nash Equilibrium (BNE) - Given common prior $\{F_i\}$, $\underline{\sigma}$ is a BNE if $\forall i$, σ_i is a best response on average, i.e., over $\sigma_{-i}(\theta_{-i})$ where $\theta_{-i} \sim F_{-i}$

Eg - Allocating with agent → principal monetary transfers

- **Setting** - n agents (buyers) + 1 special agent (seller)
  - $\mathcal{X}$ ≡ allocation space
  - Agent i has private type θ_i, which determines her value function $v : \mathcal{X} \times \Theta \to \mathbb{R}$, where $v(x|\theta_i)$ denotes the value she assigns to an allocation x given her type θ_i
  - Buyer i can also be charged a payment p_i, resulting in a total utility $v(x_i|\theta_i) - p_i$.
    Usually, we also assume the buyer can opt out (with 0 utility). This gives the quasilinear utility model
    $$u_i(x_i, p_i|\theta_i) = (v_i(x_i|\theta_i) - p_i)^+$$

Convention
- $p_i > 0 \Rightarrow$ buyer i pays seller
- $p_i < 0 \Rightarrow$ buyer i receives money
- No two buyers can pay each other

- Some examples of allocation spaces
  1) Single item - $\mathcal{X} = [0,1]^n$, and the allocation vector $\underline{x}$ satisfies $\sum_{i=1}^n x_i \le 1$
     - $v(x_i|\theta_i) = \theta_i x_i$

Note: Here x_i ≡ either fraction of item (divisible) or probability of gettin item

  2) Multi (k) items - $\mathcal{X} = [0,1]^k$, $v_i(x_i|\theta_i) = \sum_{j=1}^k \theta_{ij} x_{ij}$

Here we assume an agent has additive values over items

  3) Combinatorial auction ≡ M items, $\mathcal{X}$ ≡ Dist^n over $2^{|M|}$ subsets

This allows us to model complex complement / substitute valuations

Eg. Given $S \subseteq [M]$, $v(S|\theta_i) = \max_{j \in S} \theta_{ij}$ (pure substitutes / unit demand)
or $v(S|\theta_i) = \theta_i \mathbb{1}\{S_i \subseteq S\}$ (pure complements / single minded)

  4) Unrestricted preferences ≡ $\mathcal{X}$ = Partition of $[M]$

Can model externalities - agent i's utility depends on agent j's alloc^n

---

Mechanism Seller specifies 2 functions (based on all buyer actions)

1) Allocation Function - $\underline{x} : A_1 \times \dots \times A_n \to \mathcal{X}^n$, x_i ≡ alloc^n of agent i
2) Payment Function - $\underline{p} : A_1 \times \dots \times A_n \to \mathbb{R}^n$, p_i ≡ payment of agent i

From Mechanisms to Optimization

- A mechanism $M = (\{A_i\}, \Phi)$ is said to implement an allocation fn $f : \Theta^n \to X$ under a DSE / BNE if $f(\theta_1, \ldots, \theta_n) = \Phi(\sigma_1(\theta_1), \sigma_2(\theta_2), \ldots, \sigma_n(\theta_n))$ $\forall \{\theta_i\}$ under a DSE/BNE strategy profile $\underline{\sigma}$

- A Direct Revelation (DR) mechanism is one where $A_i = \Theta$, ie., agents are 'asked to reveal' their types

- A DR mechanism is
  - Dominant Strategy Incentive Compatible (DSIC) if revealing θ_i (truthtelling) is a DSE
  - Bayesian Incentive Compatible (BIC) if truthtelling is a BNE

Thm (Revelation Principle) - Any mechanism $M = (\{A_i\}, \Phi)$ which implements an allocation fn $f : \Theta^n \to X$ under a DSE / BNE can be emulated (ie, same f implemented) by a DR mechanism which is DSIC/BIC

Pf - Simulation argument!

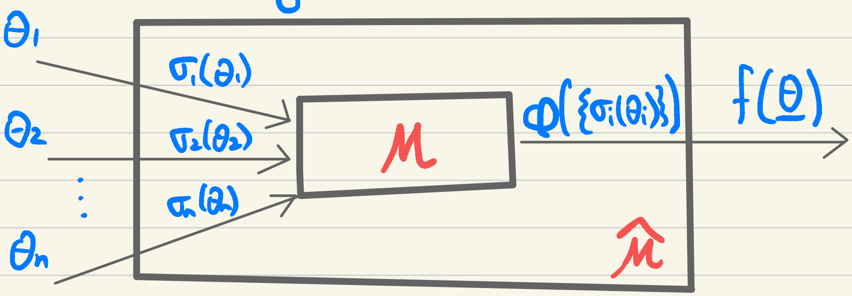

Because $\sigma_i(\theta_i)$ is a DSE/BNE, if we offer to 'simulate' $\forall i$ σ_i when we get θ_i as input, then agents are 'incentivized' to report θ_i $\Rightarrow$ Overall mechanism is DSIC/BIC

DSIC|BIC Auctions

- Using the revelation principle, we can now describe mechanisms for the allocation with transfers setting
- Recall - Agent i has private type $\theta_i \in \Theta$, gets allocation x_i, payment p_i
  - Value fn $v_i(x|\theta_i)$, utility $u_i(x|\theta_i) = (v_i(x|\theta_i) - p_i)^+$
- By revelation principle, $A_i = \Theta$, ie, agents asked to report values

- Mechanism $M \equiv$ 1) Ask agents to report types $\{t_i\} \in \Theta^n$
  2) Set allocation rule $\underline{x} = x(\underline{t})$
  Payment rule $\underline{p} = p(\underline{t})$
  3) Agent i gets utility $u_i(\underline{t}) = (v_i(x_i|\theta_i) - p_i)^+$
- Notation - For any i, vector $t \equiv (t_i, t_{-i})$, where $t_{-i} = \{t_j\}_{j \neq i}$
  - Agent i 'follows a truthtelling strategy' if she reports $t_i = \theta_i$

- $M = (x,p)$ is dominant strategy incentive compatible (DSIC) if $\forall i \in [n], \forall t_{-i} \in \Theta^{n-1}, \forall \theta_i \in \Theta, \forall t_i \in \Theta$,

$$v_i(x_i(\theta_i, t_{-i})|\theta_i) - p_i(\theta_i, t_{-i}) \geq v_i(x_i(t_i, t_{-i})|\theta_i) - p_i(t_i, t_{-i})$$

  i.e, truthtelling is (weakly) dominant no matter what others report

- Suppose $\theta_i \sim F_i$, independently for some known prior $\{F_i\}_{i \in [n]}$
  $M = (x,p)$ is Bayesian incentive compatible (BIC) if $\forall i \in [n], \forall \theta_i \in \Theta, \forall t_i \in \Theta$

$$\mathbb{E}_{\theta_{-i}}[v(X_i(\theta_i, \theta_{-i})|\theta_i) - p_i(\theta_i, \theta_{-i})] \geq \mathbb{E}_{\theta_{-i}}[v(X_i(t_i, \theta_{-i})|\theta_i) - p_i(t_i, \theta_{-i})]$$

  i.e., assuming others tell the truth, then truthtelling is (weakly) dominant in expectation (ie, truthtelling is a BNE)

- $M = (x,p)$ is ex-ante / ex-interim / ex-post individually rational (IR) if

$$\text{ex-ante: } \mathbb{E}_{\theta}[v_i(x_i(\theta)|\theta_i) - p_i(\theta)] \geq 0$$

$$\text{ex-interim: } \mathbb{E}_{\theta_{-i}}[v_i(x_i(\theta_i, \theta_{-i})|\theta_i) - p_i(\theta_i, \theta_{-i})] \geq 0$$

$$\text{ex-post: } v_i(x_i(\theta)|\theta_i) - p_i(\theta) \geq 0$$

LP formulation of single-item auction

- Objectives -

  1) Revenue (or seller surplus) - $R(x,p) = \sum_i p_i$

  2) Social Welfare $W(x,p) = \sum_i \underbrace{\left(v_i(x_i|\theta_i) - p_i\right)}_{\text{buyer i's surplus}} + \underbrace{\sum_i p_i}_{\text{seller surplus}} = \sum_i v_i(x_i|\theta_i)$

We get different LPs depending on objective + solution concept

- Welfare Max under DSE - variables $\{x_i(\underline{\theta}),\ p_i(\underline{\theta})\}\ \forall \underline{\theta} \in \Theta^n$

Suppose true type vector is $\hat{\underline{\theta}}$

ie -

$$\max \sum_{i\in[n]} x_i(\hat{\underline{\theta}})\,\hat{\theta}_i$$

(DSIC)	s.t. $x_i(\theta_i,\theta_{-i})\,\theta_i - p_i(\theta_i,\theta_{-i}) \geq x_i(t,\theta_{-i})\theta_i - p_i(t,\theta_{-i})$	$\forall i, \theta_{-i}, \theta_i, t$
(EP-IR)	$x_i(\theta_i,\theta_{-i})\,\theta_i - p_i(\theta_i,\theta_{-i}) \geq 0$	$\forall i, \theta_i, \theta_{-i}$
(Implementability)	$\sum_{i\in[n]} x_i(\underline{\theta}) \leq 1$	$\forall \underline{\theta}$
	$x_i(\underline{\theta}) \geq 0,\ p_i(\underline{\theta}) \geq 0$	$\forall i, \underline{\theta}$

- Revenue Max under BNE - Suppose $\theta_i \sim F_i$, $\perp\!\!\!\perp \Rightarrow F(\underline{\theta}) = \prod_{i=1}^{n} F_i(\theta_i)$

$$\max \sum_{\underline{\theta}\in\Theta^n} F(\underline{\theta})\left[\sum_{i\in[n]} p_i(\underline{\theta})\right]$$

(BIC)	s.t. $\sum_{\theta_{-i}} F(\theta_{-i})\left(x_i(\theta_i,\theta_{-i})\theta_i - p_i(\theta_i,\theta_{-i})\right) \geq \sum_{\theta_{-i}} F(\theta_{-i})\left(x_i(t,\theta_{-i})\theta_i - p_i(t,\theta_{-i})\right)$	$\forall i, \theta_i, t$
(EI-IR)	$\sum_{\theta_{-i}} F(\theta_{-i})\left(x_i(\theta_i,\theta_{-i})\theta_i - p_i(\theta_i,\theta_{-i})\right) \geq 0$	$\forall i, \theta_i$
(Implementability)	$\sum_{i\in[n]} x_i(\underline{\theta}) \leq 1$	$\forall \underline{\theta}$
	$x_i(\underline{\theta}) \geq 0,\ p_i(\underline{\theta}) \geq 0$	$\forall i, \underline{\theta}$

- Notes

  1) # of vars $= \overset{x,p}{2} \times \overset{i\in[n]}{n} \times |\Theta|^n$ $\leftarrow \underline{\theta} = (\theta_1, \theta_2, \ldots, \theta_n)$

  2) Allocation/payment rules are complex menus mapping every type profile vector ($|\Theta|^n$ in number) to an allocation/payment for each agent.