Columbia B9136 April 24, 2025

Lecture 10 — Bayesian Mechanism Design and Prophet Inequalities

Lecturer: Will Ma *Scribe: Chengyue He, Haixiang Lan*

In this lecture, we will study the introduction to (Bayesian, single-parameter) mechanism design.

1 Introduction

- agents $t \in [T]$ with valuations $V_t \geq 0$
- one item, goal is to give to highest valuation

Previous prophet inequality setting:

- sequentially observe V_t, decide to give item or not;
- the valuations are drawn from known independent distributions;
- agents being truthful (they disclose their true valuation).

The above can be achieved by using threshold τ. Even if valuations are unknown, we can charge a price of τ. From prophet inequality, it implies that selling the item to the first agent willing to pay fixed price τ ensures the winning agent's valuation is $\geq \frac{1}{2}\mathbb{E}[\max_t V_t]$ in expectation. In this sense, valuations being known was not important, because posting price τ achieves the same outcome as knowing valuations and accepting the first valuation above τ (for the "welfare" objective of maximizing the winning agent's valuation).

In the offline setting, if all agents are present at the same time, we are able to achieve $\max_t V_t$ even if valuations are unknown and there is no distributional information. This is known as Welfare-maximizing prior-free mechanism (VCG mechanism/second-price mechanism).

However for revenue maximization, we need Bayesian priors to do anything meaningful in terms of mechanism design. We now formally define the mechanism design framework, which can capture an arbitrary auction with multiple self-interested agents whose valuations are private information.

2 Preliminaries in Mechanism Design

In single-parameter mechanism design problem, there is a seller with one indivisible item and a set of T agents (or buyers, bidders). Each agent $t \in [T]$ has a private valuation about the item that is summarized by just one number v_t. The seller want to sell the item to at most one of the agent. We make the following assumptions:

1. The seller has commitment power and publicly announces the selling protocol.
2. The agents are strategic and play their best-response strategy to the given protocol.

Revelation principle:

- Any protocol can be described by *direct mechanism*:
    - input: buyer reports V_t;

  - output: outcome (payment, whether buyer wins the item), possibly determined randomly.

- Direct mechanism must promise buyer that they are incentivized to report the truth.

2.1 Deterministic Valuations and Deterministic Mechanism

- $\boldsymbol{v} = (v_t)_{t=1}^T \in \mathbb{R}_{\geq 0}^T$: Valuation profile of the agents.
- $\boldsymbol{v}_{-t} = (v_{t'})_{t' \neq t}$: Valuation profile of other agents except t.
- A deterministic mechanism is given by a pair (X, P) where

$$\begin{array}{rcl} X : \mathbb{R}_{\geq 0}^T & \to & \left\{ x \in \{0,1\}^T : \sum_{t=1}^T x_t \leq 1 \right\} \\ \boldsymbol{v} & \to & X(\boldsymbol{v}) = (x_t)_{t=1}^T \end{array}$$

  is called an *allocation rule* such that $X_t(\boldsymbol{v}) = x_t$ indicates whether t wins the item, and

$$\begin{array}{rcl} P : \mathbb{R}_{\geq 0}^T & \to & \mathbb{R}_{\geq 0}^T \\ \boldsymbol{v} & \to & P(\boldsymbol{v}) = (p_t)_{t=1}^T \end{array}$$

  is called a *payment rule* such that $P_t(\boldsymbol{v}) = p_t$ states the amount of money paid by t.

Ideally, a mechanism should prevent agents from manipulation.

Definition 1 (Dominant-Strategy Incentive Compatible)**.** A deterministic mechanism (X, P) is *dominant-strategy incentive compatible (DSIC)* if for every agent $t \in [T]$, it has

$$\underbrace{v_t \cdot X_t(v_t, \boldsymbol{v}_{-t}) - P_t(v_t, \boldsymbol{v}_{-t})}_{\text{utility from telling truth } v_t} \geq \underbrace{v_t \cdot X_t(v_{t'}, \boldsymbol{v}_{-t}) - P_t(v_{t'}, \boldsymbol{v}_{-t})}_{\text{utility from telling lie } v_{t'}} \tag{1}$$

for all $v_t, v_t' \in \mathbb{R}_{\geq 0}$ and $\boldsymbol{v}_{-t} \in \mathbb{R}_{\geq 0}^{T-1}$. In other words, being truthful is a dominant strategy for every agent, no matter what others' valuations are.

Another nice property of a mechanism is to encourage every agent to participate.

Definition 2 (Dominant-Strategy Individual Rational)**.** A deterministic mechanism (X, P) is *dominant-strategy individual rational (DSIR)* if for every agent $t \in [T]$, it has

$$v_t \cdot X_t(v_t, \boldsymbol{v}_{-t}) - P_t(v_t, \boldsymbol{v}_{-t}) \geq 0 \tag{2}$$

for all $v_t \in \mathbb{R}_{\geq 0}$ and $\boldsymbol{v}_{-t} \in \mathbb{R}_{\geq 0}^{T-1}$. In other words, participating in the auction does not harm any agent (so that they would not leave the auction).

As a famous example, first-price auction[1] is DSIR but not DSIC. In contrast, second-price auction[2] is both DSIR and DSIC.

---

[1] The agent with the highest bid (reported valuation) wins the item, and only the winner pays her bid.

[2] Also known as the Vickrey mechanism: The agent with the highest bid (reported valuation) wins the item, and only the winner pays the second-highest bid.

2.2 Randomized Mechanism

A mechanism can be randomized. In particular, a *randomized mechanism* is given by a pair $(\tilde{X}, \tilde{P})$ such that

$$\tilde{X} : \mathbb{R}^T_{\geq 0} \to \left\{ x \in [0,1]^T : \sum_{t=1}^{T} x_t \leq 1 \right\}$$

where $\tilde{X}_t(\boldsymbol{v}) = x_t \in [0,1]$ is the *winning probability* of t, and

$$\tilde{P} : \mathbb{R}^T_{\geq 0} \to \Delta(\mathbb{R}^T_{\geq 0}),$$

where $\tilde{P}_t(\boldsymbol{v}) = p_t$ is the *payment* of agent t. Equivalently, randomized mechanism can be seen as a distribution over deterministic mechanisms $X \sim \tilde{X}, P \sim \tilde{P}$. We can still define DSIC (resp., DSIR) properties for a randomized mechanism, where one needs to let (1) (resp., (2)) hold in expectation under $\tilde{X}, \tilde{P}$. More specifically:

Definition 3 (Dominant-Strategy Incentive Compatible)**.** A randomized mechanism $(\tilde{X}, \tilde{P})$ is *dominant-strategy incentive compatible (DSIC)* if for every agent $t \in [T]$, it has

$$\underbrace{\mathbb{E}_{\tilde{X},\tilde{P}} \left[v_t \cdot X_t(v_t, \boldsymbol{v}_{-t}) - P_t(v_t, \boldsymbol{v}_{-t}) \right]}_{\text{utility from telling truth } v_t} \geq \underbrace{\mathbb{E}_{\tilde{X},\tilde{P}} \left[v_t \cdot X_t(v_{t'}, \boldsymbol{v}_{-t}) - P_t(v_{t'}, \boldsymbol{v}_{-t}) \right]}_{\text{utility from telling lie } v_{t'}} \tag{3}$$

for all $v_t, v'_t \in \mathbb{R}_{\geq 0}$ and $\boldsymbol{v}_{-t} \in \mathbb{R}^{T-1}_{\geq 0}$.

Definition 4 (Dominant-Strategy Individual Rational)**.** A randomized mechanism $(\tilde{X}, \tilde{P})$ is *dominant-strategy individual rational (DSIR)* if for every agent $t \in [T]$, it has

$$\mathbb{E}_{\tilde{X},\tilde{P}} \left[v_t \cdot X_t(v_t, \boldsymbol{v}_{-t}) - P_t(v_t, \boldsymbol{v}_{-t}) \right] \geq 0 \tag{4}$$

for all $v_t \in \mathbb{R}_{\geq 0}$ and $\boldsymbol{v}_{-t} \in \mathbb{R}^{T-1}_{\geq 0}$.

2.3 Bayesian Mechanism Design

Bayesian mechanism design assumes that the buyers know the valuation distributions of others and that others will report truthfully. Generally, Bayesian mechanism design deals with the following environment:

- $\boldsymbol{V} = (V_t)_{t=1}^T$: Valuation profile, where $V_t \in \Delta(\mathbb{R}_{\geq 0})$.
- $\boldsymbol{V}_{-t} = (V_{t'})_{t' \neq t}$: Valuations of others.
- For each $t \in [T]$, V_t is drawn independently from known distribution F_t.
- A randomized mechanism $(\tilde{X}, \tilde{P})$.

Now, we extend incentive compatibility and individual rationality to Bayesian setting:

Definition 5 (Bayesian Incentive Compatible)**.** A (randomized) mechanism $(\tilde{X}, \tilde{P})$ is *Bayesian incentive compatible (BIC)* if for every agent $t \in [T]$, it has

$$\underbrace{\mathbb{E}_{\tilde{X},\tilde{P},\boldsymbol{V}_{-t}} [v_t \cdot \tilde{X}_t(v_t, \boldsymbol{V}_{-t}) - \tilde{P}_t(v_t, \boldsymbol{V}_{-t})]}_{\text{utility from telling truth } v_t} \geq \underbrace{\mathbb{E}_{\tilde{X},\tilde{P},\boldsymbol{V}_{-t}} [v_t \cdot \tilde{X}_t(v_{t'}, \boldsymbol{V}_{-t}) - \tilde{P}_t(v_{t'}, \boldsymbol{V}_{-t})]}_{\text{utility from telling lie } v_{t'}} \tag{5}$$

for all $v_t \in \mathbb{R}_{\geq 0}$ and $v'_t \in \mathbb{R}_{\geq 0}$. In other words, if the agent is certain about her type and assume other agent behaves as the belief, being truthful is the best-response in average for every agent.

Definition 6 (Bayesian Individual Rational)**.** A (randomized) mechanism $(\tilde{X}, \tilde{P})$ is *Bayesian individual rational (BIR)* if for every agent $t \in [T]$, it has

$$\mathbb{E}_{\tilde{X},\tilde{P},\boldsymbol{V}_{-t}}[v_t \cdot \tilde{X}_t(v_t, \boldsymbol{V}_{-t}) - \tilde{P}_t(v_t, \boldsymbol{V}_{-t})] \geq 0 \tag{6}$$

for all $v_t \in \mathbb{R}_{\geq 0}$. In other words, assuming other agent behaves as the belief, participating in the auction (in average) does not harm any agent.

To simplify the notation, we define $x_t(\cdot) := \mathbb{E}_{\tilde{X},\boldsymbol{V}_{-t}}\left[\tilde{X}_t(\cdot, \boldsymbol{V}_{-t})\right]$ and $p_t(\cdot) := \mathbb{E}_{\tilde{P},\boldsymbol{V}_{-t}}\left[\tilde{X}_t(\cdot, \boldsymbol{V}_{-t})\right]$. Note that $x_t(\cdot)$ describes the interim winning probability of agent t upon reporting $\cdot$ and $p_t(\cdot)$ describes the interim payment of agent t upon reporting $\cdot$.

We remark that BIC (resp., BIR) is a weaker property compared to DSIC (resp., DSIR), as a DSIC (resp., DSIR) mechanism requires that every sample from $\boldsymbol{V}$ satisfies (1) (resp., (2)), and thus in average it also satisfies (5) (resp., (6)).

2.4 Welfare, Revenue and Consumer Surplus

Different mechanism designers have different goals. Notice that the result of a fixed mechanism $(\tilde{X}, \tilde{P})$ can be decomposed as

$$\underbrace{\sum_{t=1}^{T} V_t\tilde{X}_t(\boldsymbol{V}) - \tilde{P}_t(\boldsymbol{V})}_{\text{Consumer Surplus}} + \underbrace{\sum_{t=1}^{T} \tilde{P}_t(\boldsymbol{V})}_{\text{Revenue}} = \underbrace{\sum_{t=1}^{T} V_t\tilde{X}_t(\boldsymbol{V})}_{\text{Social Welfare}}. \tag{7}$$

The first-price auction and second-price are designed to maximize social welfare, as one can see that $\max_t V_t \geq \max\{\sum_{t=1}^{T} V_t\tilde{X}_t(\boldsymbol{V})\}$ is achieved by these mechanisms. In this class, however, we stand on the seller's side. Our goal is to see a mechanism that maximizes revenue. In particular, we want to design a mechanism that needs to be both BIC and BIR, and maximizes the expected revenue $\mathbb{E}_{\tilde{X},\tilde{P},\boldsymbol{V}}[\sum_{t=1}^{T} \tilde{P}_t(\boldsymbol{V})]$.

3 Revenue Maximization Auction

In this section, we study the mechanism that maximizes revenue first proposed by Myerson (1981). We focus on the setting introduced in Section 2.3.

3.1 Finite Valuation Assumption

Like most of the proof we did in this class, to simplify the notation and get a better interpretation, we assume that all the possible valuations are

$$r_1 > \cdots > r_m \geq 0,$$

and the valuation distribution V_t is defined by $(\lambda_{t,j})_{j=1}^{m}$ where $\sum_{j=1}^{m} \lambda_{t,j} = 1$ for every $t \in [T]$.

3.2 Tie-breaking Rule

A deterministic mechanism (X, P) usually needs to specify a tie-breaking rule. For example, in first-price auction, an agent with the highest bid wins the item, but such agent may not be unique. To finalize the allocation and payment, we need to discuss how the indivisible good to be assigned when tie appears. Some classical tie-breaking rules are:

- Give the item to the winner t with the smallest index.
- Randomly choose one winner.

These rules can also be captured by some perturbation rule (small changes in agents' valuations such that the winner will be always unique) on the valuations reported by the agents. Notice that a specific tie-breaking rule may influence the DSIC property of a mechanism. To avoid lengthy analysis, in most of the cases, we assume that there is no tie in the auction. One may also always apply the first tie-breaking rule and check every time when a tie appears. The benefit of the first tie-breaking rule over the second is that it yields a deterministic mechanism.

3.3 Upper Bound Analysis

Now, let us assume that a BIC + BIR mechanism exists. We will figure out an upper bound on the revenue that the seller can achieve. Starting from here, we denote $x_{t,j} := x_t(r_j)$, $p_{t,j} := p_t(r_j)$ and $x_j := x(r_j)$, $p_j := p(r_j)$ if we ignore the index t of the agent.

3.3.1 Single Agent Analysis

We start with a single agent t (and thus we omit t), the BIC constraints become:

$$\begin{aligned} & r_j x_j - p_j \geq r_j x_{j'} - p_{j'}, \forall j, j' \in [m] \\ \Rightarrow & r_j(x_j - x_{j'}) \geq p_j - p_{j'}, \forall j, j' \in [m]. \end{aligned} \tag{8}$$

Monotone Allocation Rule: Exchange j and j' in (8), we obtain

$$\begin{aligned} & r_{j'} x_{j'} - p_{j'} \geq r_{j'} x_j - p_j, \forall j, j' \in [m] \\ \Rightarrow & p_j - p_{j'} \geq r_{j'}(x_j - x_{j'}), \forall j, j' \in [m] \end{aligned} \tag{9}$$

From (8) and (9) we have

$$\begin{aligned} & r_j(x_j - x_{j'}) \geq r_{j'}(x_j - x_{j'}) \quad \forall j, j' \in [m], \\ \Rightarrow & (r_j - r_{j'})(x_j - x_{j'}) \geq 0, \quad \forall j, j' \in [m]. \end{aligned}$$

which means that if $r_j > r_{j'}$, then $x_j \geq x_{j'}$. This is a necessary condition for a BIC mechanism. We call this condition *weakly monotone* on interim winning probability.

We do not know the optimal $x_1, \ldots, x_m$ so far. However, given $1 \geq x_1 \geq \cdots \geq x_m \geq 0$ (caveat: Here $x_1 = x_2 = \ldots = x_m = 1$ can possibly feasible. Do not confuse this with the constraint $\sum_{t \in [T]} x_t \leq 1$.), how does the seller maximize the payment while satisfying BIR? Remember that the BIR constraints are

$$r_j x_j - p_j \geq 0, \quad \forall j \in [m]. \tag{10}$$

We observe that by BIC and BIR constraints,

$$\begin{aligned} & p_m \leq r_m x_m, \\ & p_{m-1} \leq p_m + r_{m-1}(x_{m-1} - x_m) \leq r_m x_m + r_{m-1}(x_{m-1} - x_m) \\ & p_{m-2} \leq p_{m-1} + r_{m-2}(x_{m-2} - x_{m-1}) \leq r_m x_m + r_{m-1}(x_{m-1} - x_m) + r_{m-2}(x_{m-2} - x_{m-1}) \end{aligned}$$

where the second line comes from $r_{m-1}x_{m-1} - p_{m-1} \geq r_{m-1}x_m - p_m$. Inductively, one can show that (setting $x_{m+1} = 0$):

$$p_i \leq \sum_{j=i}^{m} r_j(x_j - x_{j+1}), \quad \forall i \in [m] \tag{11}$$

Then, the expected revenue from the seller is

$$\begin{aligned}
\sum_{i=1}^{m} \lambda_i p_i &\leq \sum_{i=1}^{m} \left[\lambda_i \sum_{j=i}^{m} r_j(x_j - x_{j+1}) \right] \\
&= \sum_{j=1}^{m} (x_j - x_{j+1}) \underbrace{r_j \sum_{i=1}^{j} \lambda_i}_{=:R_j} \\
&= \sum_{j=1}^{m} (x_j - x_{j+1}) R_j \\
&\leq \max_{j \in [m]} R_j
\end{aligned} \tag{12}$$

Here R_j can be interpreted as the expected revenue from posting price r_j for the agent because $R_j = r_j \mathbb{P}(V \geq r_j)$. The first inequality is the implication of the BIC and BIR constraint on p_i. The first equality is a integration by part technique. The last inequality is immediate by noting that $x_1 \geq \ldots \geq x_m \geq 0$ and $\|(x_1 - x_2, ..., x_m - x_{m+1})\|_1 \leq 1$. The last inequality can be attained by setting $j^* = \text{argmax}_j R_j$ and setting $x_j = 1$ for $j \leq j^*$ and $x_j = 0$ for $j > j^*$. Note that (12) is the upper bound of the revenue, and it is attained by simply a posted price r_{j^*}. In this way, we not only solve (1) the payment rule: for valuations $j \leq j^*$, the price is always r_{j^*} and for $j > j^*$, the price is always 0; but also (2) solve the optimal $x_1, \ldots, x_m$: accepts when valuations satisfy $j \leq j^*$ and rejects when valuatiosn satisfy $j > j^*$, which answers the initial question. In conclusion: the best mechanism for selling to a single agent is **a posted price**.

3.3.2 Multiple Agents Analysis

From the analysis on a single agent, we can extend the framework to multiple agents. We know that given BIC and BIR constraints, the revenue is bounded by

$$\sum_{t=1}^{T} \sum_{j=1}^{m} \lambda_{t,j} p_{t,j} \leq \sum_{t=1}^{T} \sum_{j=1}^{m} R_{t,j}(x_{t,j} - x_{t,j+1}) = \sum_{t=1}^{T} \sum_{j=1}^{m} x_{t,j}(R_{t,j} - R_{t,j-1}) \tag{13}$$

where $x_{t,m+1} = 0$ and $R_{t,0} = 0$. Here, we define

Definition 7 (Virtual Valuation)**.** Fix agent $t \in [T]$, the virtual valuation of t is a transformed function $\phi_t(\cdot) : \mathbb{R}_{\geq 0} \to \mathbb{R}$ such that for true valuation r_j[3],

$$\phi_t(r_j) = \frac{R_{t,j} - R_{t,j-1}}{\lambda_{t,j}}. \tag{14}$$

---

[3] For j such that $\lambda_{t,j} = 0$, this is undefined and irrelevant in the discussion.

Continuing the analysis in (13), we have

$$\begin{aligned}
\sum_{t=1}^{T}\sum_{j=1}^{m}\lambda_{t,j}p_{t,j} &\le \sum_{t=1}^{T}\sum_{j=1}^{m} x_{t,j}(R_{t,j}-R_{t,j-1}) \\
&= \sum_{t=1}^{T}\sum_{j=1}^{m} x_{t,j}\lambda_{t,j}\frac{R_{t,j}-R_{t,j-1}}{\lambda_{t,j}} \\
&= \sum_{t=1}^{T}\sum_{j=1}^{m} x_{t,j}\lambda_{t,j}\phi_t(r_j) \\
&= \sum_{t=1}^{T}\sum_{j=1}^{m} \mathbb{E}_{\boldsymbol{V}_{-t},\tilde{X}}[\tilde{X}(r_j,\boldsymbol{V}_{-t})]\cdot\phi_t(r_j)\cdot\mathbb{P}_{V_t}(V_t=r_j) \\
&= \mathbb{E}_{\boldsymbol{V},\tilde{X}}\left[\sum_{t=1}^{T}\sum_{j=1}^{m}\tilde{X}(V_t,\boldsymbol{V}_{-t})\cdot\phi_t(r_j)\cdot\mathbf{1}(V_t=r_j)\right] && \text{(by independence)} \\
&= \mathbb{E}_{\boldsymbol{V},\tilde{X}}\left[\sum_{t=1}^{T}\tilde{X}(V_t,\boldsymbol{V}_{-t})\cdot\phi_t(V_t)\right] \\
&\le \mathbb{E}_{\boldsymbol{V}}\left[[\max_{t\in[T]}\phi_t(V_t)]^+\right]. && (15)
\end{aligned}$$

Here, $\mathbb{E}_{\boldsymbol{V}}[\phi_t(V_t)]$ is called the *expected virtual surplus*. We remark that, until this point, we do not know that the expected revenue can match expected virtual surplus; while the former depends on the mechanism, the latter does not. Our next goal is to derive a mechanism whose expected revenue matches the expected virtual surplus.

We start from the easy things: the last inequality can be attained by setting $t^* := \text{argmax}_t\phi_t(V_t)$ and (1) if $\phi_{t^*}(V_{t^*}) \ge 0$, we set $X_{t^*}(\boldsymbol{V}) = 1$ and $X_t(\boldsymbol{V}) = 0$ for $t \ne t^*$; (2) otherwise, we set $X_t(\boldsymbol{V}) = 0$ for all t.

We have not checked an important assumption: the monotonicity of x_{tj} with respect to j (for all t). Only when x_{tj} has the monotonic property can we construct the payment rule above serving as an upper bound. We will guarantee this monotone property by the following lemma:

Lemma 8. *For any fixed t, if $\phi_t(r_1) \ge \phi_t(r_2) \ge \ldots \ge \phi_t(r_m)$, then $x_{t1} \ge x_{t2} \ge \ldots \ge x_{tm}$.*

Proof. For any t, we have

$$X_t(\boldsymbol{V}) = \mathbf{1}\left\{t = \text{argmax}_{t'}\phi_{t'}(V_{t'})\right\}.$$

We first fix $\boldsymbol{V}_{-t}$:

$$\begin{aligned}
X_t(r_j,\boldsymbol{V}_{-t}) &= \mathbf{1}\left\{\phi_t(r_j) \ge \phi_{t'}(V_{t'}), \forall t'\right\}, \\
X_t(r_i,\boldsymbol{V}_{-t}) &= \mathbf{1}\left\{\phi_t(r_i) \ge \phi_{t'}(V_{t'}), \forall t'\right\}.
\end{aligned}$$

If $\phi_t(r_j) \ge \phi_t(r_i)$ then we have the implication

$$\phi_t(r_i) \ge \phi_{t'}(V_{t'}), \forall t' \Rightarrow \phi_t(r_j) \ge \phi_{t'}(V_{t'}), \forall t'.$$

Hence,

$$X_t(r_j,\boldsymbol{V}_{-t}) \ge X_t(r_i,\boldsymbol{V}_{-t}).$$

By taking expectation with respect to $\boldsymbol{V}_{-t}$, we have $x_{t,j} \ge x_{t,i}$. □

We remark that the above lemma are based on the assumption that $\phi_t(r_1) \geq \phi_t(r_2) \geq \ldots \geq \phi_t(r_m)$, This may not be true even if $r_1 \geq r_2 \geq \ldots \geq r_m$. We will proceed with this assumption for now, and later work out a modified definition of $\phi_t(\cdot)$ to guarantee the monotone property without loss of optimality.

3.4 Attaining Expected Virtual Surplus

We show that indeed there exists a BIC + BIR mechanism that achieves the revenue as in (15), thus all inequalities in the last section can be attained as equalities. Again, we first show the result holds for a single agent.

Lemma 9. *Given* $1 \geq x_1 \geq \cdots \geq x_m \geq 0$. *Then, setting* $(p_j)_{j=1}^m$ *as equality* (11) *induces a mechanism that satisfies BIC and BIR.*

Proof. We first show BIC. Fix a true valuation j and a fake valuation i. If $j < i$, we can obtain (according to (11))

$$\begin{aligned} p_{i-1} - p_i &= r_{i-1}(x_{i-1} - x_i) \leq r_j(x_{i-1} - x_i), \\ p_{i-2} - p_{i-1} &= r_{i-2}(x_{i-2} - x_{i-1}) \leq r_j(x_{i-2} - x_{i-1}), \\ &\cdots \\ p_j - p_{j+1} &= r_j(x_j - x_{j+1}), \end{aligned}$$

where the inequalities hold since $x_1 \geq \cdots \geq x_m$. Summing up the inequalities we obtain $p_j - p_i \leq r_j(x_j - x_i)$, which verifies (8).

Similarly, if $j > i$,

$$\begin{aligned} p_{j-1} - p_j &= r_{j-1}(x_{j-1} - x_j) \geq r_j(x_{j-1} - x_j), \\ p_{j-2} - p_{j-1} &= r_{j-2}(x_{j-2} - x_{j-1}) \geq r_j(x_{j-2} - x_{j-1}), \\ &\cdots \\ p_i - p_{i+1} &= r_i(x_i - x_{i+1}) \geq r_j(x_i - x_{i+1}). \end{aligned}$$

Summing up the inequalities we obtain $p_i - p_j \geq r_j(x_i - x_j)$, which also verifies (8). Therefore, for every pair $i \neq j \in [m]$, the BIC condition holds.

For BIR, first we observe that by definition

$$r_m x_m - p_m = 0,$$

and for $j < m$,

$$r_j x_j - p_j \geq r_j x_m - p_m \geq r_m x_m - p_m = 0,$$

where the first inequality is by BIC. □

In Section 3.3.1, we have discussed that by selling the single agent the item with price $\max_{j\in[m]} R_j$, such mechanism achieves maximum (expected) revenue. This is a posted price allocation rule, thus weakly monotone (because bidding higher gets more). Therefore, this mechanism also verifies the condition in Lemma 9. Therefore, combine the payment rule given in Lemma 9, this mechanism is BIC + BIR.

For multiple agents, we consider the mechanism with an allocation rule that assigns the item to whoever achieves $\max_{t\in[T]}[\phi_t(v_t)]^+$[4], and *assume the allocation rule is weakly monotone*, then we

---

[4] We note that the mechanism acts after seeing the realization $\boldsymbol{v}$.

obtain the payment rule that satisfies BIC + BIR, by applying the payment rule from Lemma 9 for each agent. We omit the algebra here.

However, monotonicity does not always hold. We will see a (counter-)example in Example 11.

Example 10. In this example, we illustrate the concepts introduced so far in the lecture.

Suppose that there are two agents where $V_1 = \text{uniform}\{2,3\}$ and $V_2 = \text{uniform}\{3,4\}$. Thus we have $r_1 = 4, r_2 = 3, r_3 = 2$. We can then compute the virtual valuations, which are the slopes of the linear functions in Figure 1. There are four realizations of the valuation vector $\boldsymbol{v}$, each with probability $\frac{1}{4}$. If we always assign the item to the agent with highest virtual valuation, then the allocation rule is given by According to the allocation rule, we can compute

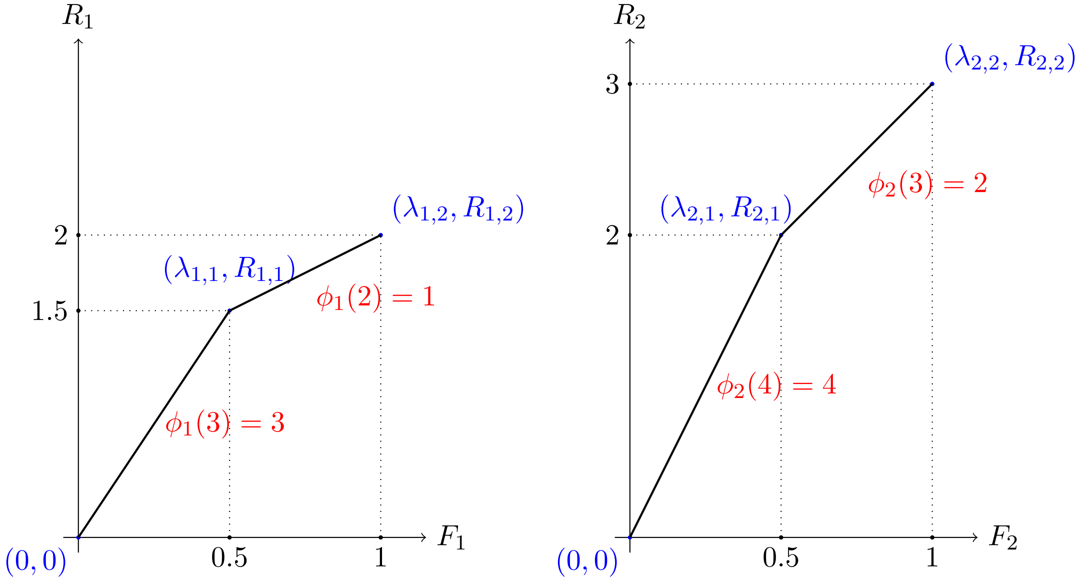

Figure 1: Revenue curve in Example 10.

V_2 V_1	$4, (\phi = 4)$	$3, (\phi = 2)$
$3, (\phi = 3)$	2 wins	1 wins
$2, (\phi = 1)$	2 wins	2 wins

Table 1: BIC + BIR allocation rule.

$$x_1(3) = \frac{1}{2}, x_1(2) = 0, x_2(4) = 1, x_2(3) = \frac{1}{2}.$$

By setting the payment to every agent according to (11), we obtain the payment rule

$$p_1(3) = 1.5, p_1(2) = 0, p_2(4) = 3.5, p_2(3) = 1.5.$$

Under this payment rule, the seller collects the expected revenue[5]

$$\mathbb{E}[P_1(V_1, V_2) + P_2(V_1, V_2)] = \frac{(1.5 + 3.5) + (1.5 + 1.5) + (0 + 3.5) + (0 + 1.5)}{4} = 3.25.$$

---

[5] The four events in the summation correspond to $(V_1, V_2) = (3,4), (3,3), (2,4), (2,3)$ in order.

V_2 V_1	$4, (\phi = 4)$	$3, (\phi = 2)$
$3, (\phi = 3)$	1.5	1.5
$2, (\phi = 1)$	0	0

V_2 V_1	$4, (\phi = 4)$	$3, (\phi = 2)$
$3, (\phi = 3)$	3.5	1.5
$2, (\phi = 1)$	3.5	1.5

Table 2: BIC + BIR payment rule: left for agent 1, right for agent 2.

On the other hand, the expected virtual surplus is

$$\mathbb{E}[[\max_{t \in \{1,2\}} \phi_t(V_t)]^+] = \frac{\max\{3,4\} + \max\{3,2\} + \max\{1,4\} + \max\{1,2\}}{4} = 3.25.$$

Therefore, the upper bound is attained.

It can be observed that this example matches the BIC + BIR mechanism, but is not DSIC/DSIR. For example, the loser (agent 1 in the case $v_1 = 3, v_2 = 4$) sometimes has to pay a positive payment.

3.5 Ironing to ensure monotonic virtual valuations

Example 11. In this example, we will see that the mechanism for assigning the item to the agent with the highest virtual valuation is not monotone. Suppose that there is only a single agent where $V = \text{uniform}\{1, 3, 4, 8\}$. Similarly, we can draw the revenue curve as we did in the previous example.

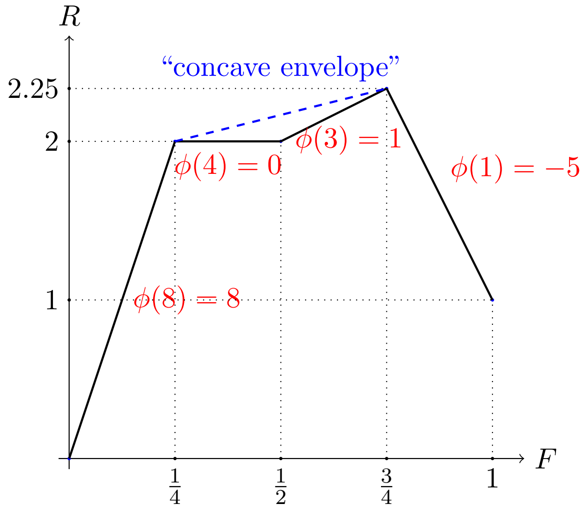

Figure 2: Revenue curve in Example 11.

Notice that the virtual valuation is not monotonic! That is, having lower valuation may result in higher allocation. One way to overcome is to "ironing" (make the revenue curve concave, or taking the upper envelope of the curve) to fix non-monotonicity. Informally, we illustrate the idea in Figure 2:

- Argument is quite geometric.
- Draw revenue curve formed by points $(0, R_0), (\lambda_1, R_1), (\lambda_1 + \lambda_2, R_2), \ldots, (1, R_m)$, denote the function induced by this curve by f.

- Slopes of line segments are virtual valuations $\phi_1, ..., \phi_m$ respectively.
- Take concave envelope $\to \bar{f}$, which is a continuous piecewise-linear, concave function from $[0,1]$ to $\mathbb{R}_{\geq 0}$.
- Let $\bar{R}_j = \bar{f}(\lambda_1 + \cdots + \lambda_j)$. Note that $\bar{R}_j \geq R_j$.
- Re-define $\bar{\phi}_j = \bar{f}'(\lambda_1 + \cdots + \lambda_j)$ where $\bar{f}'$ is the slope of $\bar{f}$ to the left. Note that by concavity we have $\bar{\phi}_1 \geq \ldots \geq \bar{\phi}_m$.

We have two key facts from the ironing procedure:

- $\bar{R}_j \geq R_j$ for all j, since upper concave envelope lies above the original curve.
- If $\bar{\phi}_j > \bar{\phi}_{j+1}$ then it must be $\bar{R}_j = R_j$. Or equivalently, if $\bar{R}_j > R_j$ then it must be $\bar{\phi}_j = \bar{\phi}_{j+1}$.

We can show the benefit of the ironing procedure for the following lemma:

Lemma 12. *After the ironing procedure, our mechanism, say $\bar{X}$, by allocating to the agent with the highest non-negative ironed virtual valuation, is BIC and BIR. Moreover, it has the expected revenue that upper bounds any BIC and BIR mechanism $\tilde{X}$. Thus, the mechanism $\bar{X}$ is optimal.*

Proof. After the ironing procedure, we have monotonic ironing virtual evaluations, thus we have monotonic $x_{t,j}$ with respect to j, for any t. Hence, we have

$$\begin{aligned}
\sum_{t=1}^{T}\sum_{j=1}^{m} \lambda_{t,j} p_{t,j} \leq & \sum_{t=1}^{T}\sum_{j=1}^{m} R_{t,j}(x_{t,j} - x_{t,j+1}) \\
\leq & \sum_{t=1}^{T}\sum_{j=1}^{m} \bar{R}_{t,j}(x_{t,j} - x_{t,j+1}) \\
= & \sum_{t=1}^{T}\sum_{j=1}^{m} x_{t,j}(\bar{R}_{t,j} - \bar{R}_{t,j-1}) \\
= & \sum_{t=1}^{T}\sum_{j=1}^{m} x_{t,j}\lambda_{t,j}\frac{\bar{R}_{t,j} - \bar{R}_{t,j-1}}{\lambda_{t,j}} \\
= & \sum_{t=1}^{T}\sum_{j=1}^{m} x_{t,j}\lambda_{t,j}\bar{\phi}_t(r_j) \\
= & \mathbb{E}_{\boldsymbol{V},\bar{X}}\left[\sum_{t=1}^{T} \bar{X}(V_t, \boldsymbol{V}_{-t}) \cdot \bar{\phi}_t(V_t)\right] \\
\leq & \mathbb{E}_{\boldsymbol{V}}\left[[\max_{t\in[T]} \bar{\phi}_t(V_t)]^+\right]. \qquad \text{(Expected Nonnegative Ironed Virtual Surplus)}
\end{aligned}$$

The first inequality is achievable because now we have monotonic $x_{t,j}$ and we can utilize the payment rule from (11), which is BIC and BIR from Lemma 9. The only issue is to prove the attainability of the second inequality. This is by a clever argument: if $x_{t,j} - x_{t,j+1} > 0$, then by Lemma 8, $\bar{\phi}_t(r_j) > \bar{\phi}_t(r_{j+1})$. Thus, by the second key fact of the ironing procedure, $\bar{R}_{t,j} = R_{t,j}$. Therefore, it is always true that $R_{t,j}(x_{t,j} - x_{t,j+1}) = \bar{R}_{t,j}(x_{t,j} - x_{t,j+1})$. This inequality is essentially an equality. In conclusion, we have shown that our mechanism $\bar{X}$, which is BIC and BIR, has expected revenue with upper bounds that of t any BIC and BIR mechanism $\tilde{X}$, which means $\bar{X}$ is optimal. $\square$

Remark 13. Vickrey's mechanism maximizes the welfare (because the agent with highest valuation always wins), but does not maximize revenue. Myerson's mechanism does not maximize the welfare, but maximizes the seller's revenue. In the Example 10, Vickrey's mechanism only achieves revenue 2.5, whereas Myerson's mechanism achieves revenue 3.25.

To be more specific, for the second price auction, the expected revenue is less than the mechanisms before. The expected revenue is $\frac{3+3+2+2}{4} = 2.5$.

V_2 V_1	4	3
3	2 wins (pay 3)	1 wins (pay 3)
2	2 wins (pay 2)	2 wins (pay 2)

Table 3: Second Price Auction

4 From BIC/BIR to DSIC/DSIR

The above mechanism was only BIC/BIR but not DSIC/DSIR. We now try to construct a mechanism that is DSIC/DSIR, essentially by defining the payment rule in the same way (setting (11) to equality), but based on the actual allocation instead of the interim allocation. In particular, we define the allocation rule (which allocates to the highest non-negative ironed virtual valuation)

$$X_t(r_j, \boldsymbol{V}_{-t}) := \mathbf{1}(\bar{\phi}_t(r_j) > [\max_{t' \neq t} \bar{\phi}_{t'}(V_{t'})]^+)$$

and the payment rule

$$P_t(r_j, \boldsymbol{V}_{-t}) := \sum_{j'=j}^{m} r_{j'}(X_t(r_{j'}, \boldsymbol{V}_{-t}) - X_t(r_{j'+1}, \boldsymbol{V}_{-t})).$$

Lemma 14. *The above allocation and payment rule is DSIC and DSIR.*

Proof. For any fixed t and fixed $\boldsymbol{V}_{-t}$, denote $x_j = X_t(r_j, \boldsymbol{V}_{-t})$ and $p_j = P_t(r_j, \boldsymbol{V}_{-t})$. From Lemma 8, we have $x_1 \geq x_2 \geq \ldots \geq x_m$. If $j < i$, we can obtain (according to (11))

$$\begin{aligned} p_{i-1} - p_i &= r_{i-1}(x_{i-1} - x_i) \leq r_j(x_{i-1} - x_i), \\ p_{i-2} - p_{i-1} &= r_{i-2}(x_{i-2} - x_{i-1}) \leq r_j(x_{i-2} - x_{i-1}), \\ &\cdots \\ p_j - p_{j+1} &= r_j(x_j - x_{j+1}), \end{aligned}$$

where the inequalities hold since $x_1 \geq \cdots \geq x_m$. Summing up the inequalities we obtain $p_j - p_i \leq r_j(x_j - x_i)$, which verifies (8).

Similarly, if $j > i$,

$$\begin{aligned} p_{j-1} - p_j &= r_{j-1}(x_{j-1} - x_j) \geq r_j(x_{j-1} - x_j), \\ p_{j-2} - p_{j-1} &= r_{j-2}(x_{j-2} - x_{j-1}) \geq r_j(x_{j-2} - x_{j-1}), \\ &\cdots \\ p_i - p_{i+1} &= r_i(x_i - x_{i+1}) \geq r_j(x_i - x_{i+1}). \end{aligned}$$

Summing up the inequalities we obtain $p_i - p_j \geq r_j(x_i - x_j)$, which also verifies (8). Therefore, for every pair $i \neq j \in [m]$, the DSIC condition holds: for all V_t, $\boldsymbol{V}_{-t}$, V_t' and all t,

$$\underbrace{V_t \cdot X_t(V_t, \boldsymbol{V}_{-t}) - P_t(V_t, \boldsymbol{V}_{-t})}_{\text{utility from telling truth } V_t} \geq \underbrace{V_t \cdot X_t(V_{t'}, \boldsymbol{V}_{-t}) - P_t(V_{t'}, \boldsymbol{V}_{-t})}_{\text{utility from telling lie } V_{t'}}$$

For DSIR, first we observe that by definition

$$r_m x_m - p_m = 0,$$

and for $j < m$,

$$r_j x_j - p_j \geq r_j x_m - p_m \geq r_m x_m - p_m = 0,$$

where the first inequality is by DSIC. Hence, for all V_t, $\boldsymbol{V}_{-t}$,

$$V_t \cdot X_t(V_t, \boldsymbol{V}_{-t}) - P_t(V_t, \boldsymbol{V}_{-t}) \geq 0.$$

$\square$

To interpret this rule, we define

$$j_t^{\text{crit}}(\boldsymbol{V}_{-t}) = \max\{j \in [m] : \bar{\phi}_t(r_j) > [\max_{t' \neq t} \bar{\phi}_{t'}(V_{t'})]^+\}$$

and $j_t^{\text{crit}}(\boldsymbol{V}_{-t}) = 0$ if the index set on right hand side is empty. Then, by definition

$$X_t(r_j, \boldsymbol{V}_{-t}) = \mathbf{1}[j \leq j_t^{\text{crit}}(\boldsymbol{V}_{-t})],$$

and

$$P_t(r_j, \boldsymbol{V}_{-t}) = r_{j_t^{\text{crit}}(\boldsymbol{V}_{-t})} \cdot X_t(r_j, \boldsymbol{V}_{-t}).$$

This can be interpreted in words: Winner pays the smallest valuation they could have reported and still won the auction. Consider each agent: it boils down to accepting or rejecting a price of $r_{j_t^{cirt}(\boldsymbol{V}_{-t})}$. This is called "taxation principle intuition". It can be shown that the above mechanism is DSIC + DSIR. We call this mechanism *Myerson's mechanism.*

The above mechanism can be illustrated still by the example above. The expected revenue is $\frac{4+3+3+3}{4} = 3.25$.

V_2 V_1	$4, (\phi = 4)$	$3, (\phi = 2)$
$3, (\phi = 3)$	2 wins (pay 4)	1 wins (pay 3)
$2, (\phi = 1)$	2 wins (pay 3)	2 wins (pay 3)

Table 4: DSIC + DSIR allocation rule.

Takeaway message: (1) randomization does not essentially improve the performance of the mechanism (deterministic versions are enough). (2) The above rule is DSIC and DSIR (will be proved later). Hence, they are BIC and BIR, and achieves the optimality among all BIC and BIR mechanisms. Hence, the above rule is optimal among all DSIC and DSIR mechanisms. We are now ready to show the DSIC and DSIR property.

5 Relation to Prophet Inequality

Recall that we have learned the online secretary problem:

- Known independent distributions $v_t \sim V_t, t \in [T]$.
- Agents arrive sequentially $t = 1, \dots, T$.
- We offer a price τ_t.
- If $V_t \geq \tau_t$, then we collect V_t and stop.

By prophet inequality, we can always collect at least $\frac{1}{2}\mathbb{E}[\max_t V_t]$.

In comparison, we can also define the online pricing problem:

- Known independent distributions $v_t \sim V_t, t \in [T]$.
- Agents arrive sequentially $t = 1, \dots, T$.
- We offer a price τ_t.
- If $V_t \geq \tau_t$, then we collect τ_t and stop.

It is not possible to collect $\frac{1}{2}\mathbb{E}[\max_t V_t]$ even if $T = 1$. Here, one may consider a different offline benchmark: optimal DSIC/DSIR, and this by definition is Myerson's mechanism, i.e. $\mathbb{E}[\max_t [\bar{\phi}_t(V_t)]^+]$.

Surprisingly, all the prophet inequality results hold if we apply this benchmark. Proof omitted. Intuition: Do prophet inequality for virtual valuation. See Chawla et al. (2010) for more.

Example 15. Consider the valuations are uniformly distributed on $\{1, 3, 4, 8\}$. If we model it as an online pricing problem, we consider the immediate revenue and the death probability:

- Offer price 8, immediate revenue 2, death probability 1/4,
- Offer price 4, immediate revenue 2, death probability 2/4 (bad decision),
- Offer price 3, immediate revenue 2.25, death probability 3/4,
- Offer price 1, immediate revenue 1, death probability 1 (bad decision).

If we model it as an accept/rejection problem. The virtual valuations are $\bar{\phi}(V)$ (a) $= 8$ w.p. $1/4$; (b) $= 1/2$ w.p. $1/2$ and (c) $= -5$ w.p. $1/4$. We set the thresholds as follows:

- Set threshold 8, immediate revenue 2, death probability 1/4,
- Set threshold 1/2, immediate revenue $2 + 1/2 \cdot 1/2 = 2.25$, death probability 3/4,
- Set threshold -5, immediate revenue $2 + 1/4 + (-5/4) = 1$, death probability 1 (bad choice).

6 Final Remarks

All results hold if it is possible to serve any subset $S \in \mathcal{F} \subseteq 2^{[T]}$ such that $\mathcal{F}$ is a downward-closed collection of feasible subsets. The maximum revenue can match "expected ironed virtual surplus"

$$\max_{S \in \mathcal{F}} \sum_{t \in S} \bar{\phi}_t(V_t)$$

Today, we have seen the special case when $\mathcal{F} = \{S \subseteq T : |S| \leq 1\}$.

Summary of today's lecture. The revenue-optimal mechanism among all randomized, BIC/BIR mechanisms is in fact deterministic and DSIC/DSIR, and satisfies the nice structure of Myerson's mechanism. In particular, if there is one buyer, then the best mechanism is a posted price.

References

S. Chawla, J. D. Hartline, D. L. Malec, and B. Sivan. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 311–320, 2010.

R. B. Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.