Advanced Microeconomics

Lecture 7: Screening and Optimal Mechanism Design

Yihu Hou

Renmin University of China

Spring 2026

Outline

1. From auctions to screening
2. Two-type screening: IC, IR, and information rent
3. Continuous types: monotonicity, envelope, and distortion
4. Myerson's optimal auction as the auction analog of screening

From Lecture 6 to Lecture 7

Lecture 6: auction language

- allocation rule $x_i(v)$,
- payment rule $t_i(v)$,
- IC and IR,
- revenue equivalence and VCG.

Lecture 7: screening language

- quantity or quality $q(\theta)$,
- payment $p(\theta)$,
- IC and IR again,
- rent-distortion tradeoff in menu design.

Question there

When bidders privately know their values, how does an auction mechanism map types into allocations and payments under IC and IR?

Question now

How should a principal design a menu when the agent privately knows how much they value quality?

Same backbone

Both are mechanism design problems. The difference is the **feasibility set**: auctions allocate scarce objects across agents, while screening studies how one buyer selects from a menu of contracts.

Screening as menu design

Economic story

- The principal does not observe type θ.
- The principal offers a menu of contracts (q, p).
- The agent chooses the contract that maximizes
$$\theta q - p.$$

Interpretation

- q can be quantity, quality, or probability of service.
- p can be price, fee, or transfer.

Direct-mechanism form

We can rewrite the same problem as:

$$\hat{\theta} \mapsto \big(q(\hat{\theta}),\ p(\hat{\theta})\big).$$

If the true type is θ, utility from reporting $\hat{\theta}$ is

$$U(\hat{\theta}; \theta) = \theta q(\hat{\theta}) - p(\hat{\theta}).$$

Constraints

- **IC**: truthful reporting beats misreporting.
- **IR**: truthful participation beats opting out.

Why higher types care more about quality

Gain from upgrading

Compare a low contract (q_L, p_L) and a high contract (q_H, p_H) with

$$\Delta q = q_H - q_L > 0, \qquad \Delta p = p_H - p_L.$$

The gain for type θ from *upgrading* from the low contract to the high contract is

$$\Delta U(\theta) = \theta \Delta q - \Delta p.$$

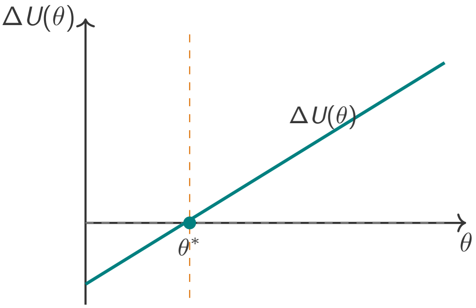

Single-crossing

Because $\Delta q > 0$, the function $\Delta U(\theta)$ is increasing in θ.
So if a lower type is willing to upgrade, a higher type is *even more* willing to upgrade.

Two-type screening setup

Environment

- Type is either $\underline{\theta}$ or $\overline{\theta}$, with $\overline{\theta} > \underline{\theta}$.
- The high type occurs with probability α.
- The principal offers two contracts:

$$(q_L, p_L), \qquad (q_H, p_H).$$

- Agent utility is quasilinear:

$$u(\theta) = \theta q - p.$$

Principal's problem

With production cost $c(q)$, expected profit is

$$\Pi = \alpha\big[p_H - c(q_H)\big] + (1-\alpha)\big[p_L - c(q_L)\big].$$

Goal

Choose the menu so that:

- each type selects the intended contract,
- participation is voluntary,
- expected profit is maximized.

The four constraints

Individual rationality

$$\underline{\theta} q_L - p_L \geq 0 \qquad \text{(IR-L)}$$
$$\overline{\theta} q_H - p_H \geq 0 \qquad \text{(IR-H)}$$

Meaning

Each type weakly prefers participating to opting out.

Incentive compatibility

$$\underline{\theta} q_L - p_L \geq \underline{\theta} q_H - p_H \qquad \text{(IC-L)}$$
$$\overline{\theta} q_H - p_H \geq \overline{\theta} q_L - p_L \qquad \text{(IC-H)}$$

Meaning

Each type prefers the contract intended for that type.

Which constraints bind in the canonical solution?

Step 1: IC already implies monotonicity

Adding IC-H and IC-L gives

$$(\overline{\theta} - \underline{\theta})(q_H - q_L) \geq 0,$$

so

$$q_H \geq q_L.$$

Step 2: IC pins down the rent gap

Let

$$U_L = \underline{\theta} q_L - p_L, \qquad U_H = \overline{\theta} q_H - p_H.$$

Then IC-H implies

$$U_H - U_L \geq (\overline{\theta} - \underline{\theta}) q_L,$$

while IC-L implies

$$U_H - U_L \leq (\overline{\theta} - \underline{\theta}) q_H.$$

Profit falls when rents rise

For a fixed allocation pair (q_L, q_H),

$$\Pi = \alpha\left[\overline{\theta} q_H - c(q_H) - U_H\right] + (1-\alpha)\left[\underline{\theta} q_L - c(q_L) - U_L\right].$$

So the principal wants the smallest rents that still satisfy IC and IR.

Canonical binding constraints

The minimum feasible rents are

$$U_L = 0, \qquad U_H = (\overline{\theta} - \underline{\theta}) q_L.$$

Set the low type's rent to its minimum $U_L = 0$, then choose the smallest feasible rent gap. Hence IR-L and IC-H bind. Moreover, IC-H together with IR-L already implies IR-H.

Payments and information rent

Payments implied by the binding constraints

$$p_L = \underline{\theta} q_L$$

$$p_H = \overline{\theta} q_H - (\overline{\theta} - \underline{\theta}) q_L$$

Utilities

$$U_L = \underline{\theta} q_L - p_L = 0$$

$$U_H = \overline{\theta} q_H - p_H = (\overline{\theta} - \underline{\theta}) q_L$$

Information rent

The high type gets positive rent because otherwise it would imitate the low contract.

$$U_H = (\overline{\theta} - \underline{\theta}) q_L.$$

The more generous the low contract is, the more rent must be left to the high type.

Key asymmetry

The low type's rent is driven down to zero. The high type's rent is what remains after the principal tries to prevent imitation.

Pricing intuition in the (q, p) diagram

Step 1: price the first q_L units

The low contract must satisfy

$$p_L \leq \underline{\theta} q_L$$

by IR-L. Since profit rises with p_L, the principal chooses

$$p_L = \underline{\theta} q_L.$$

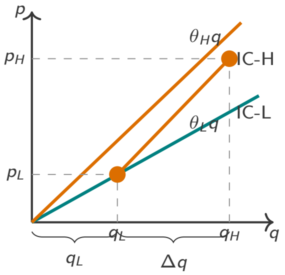

Step 2: price the incremental units

Let

$$\Delta q = q_H - q_L, \qquad \Delta p = p_H - p_L.$$

To make only the high type upgrade,

$$\underline{\theta}\Delta q \leq \Delta p \leq \overline{\theta}\Delta q.$$

The lower bound is IC-L; the upper bound is IC-H. Revenue maximization chooses the upper bound:

$$\Delta p = \overline{\theta}\Delta q.$$

Where the downward distortion comes from

Substitute the binding constraints into expected profit:

$$\Pi = \alpha\left[\overline{\theta}q_H - c(q_H)\right] + (1-\alpha)\left[\underline{\theta}q_L - c(q_L)\right] - \alpha(\overline{\theta} - \underline{\theta})q_L.$$

First two terms

These are the usual surplus terms:

- serve the high type at quality q_H,
- serve the low type at quality q_L.

Last term

$$\alpha(\overline{\theta} - \underline{\theta})q_L$$

is the extra information-rent cost created by offering quality in the low contract.

effective marginal cost of $q_L = (1-\alpha)c'(q_L) + \alpha(\overline{\theta} - \underline{\theta})$.

Main logic

Raising q_L helps the low type, but it also makes imitation more tempting for every high type. That is why the low contract is distorted downward in the optimum.

First-best benchmark and the high type

First-best allocations

If type were observable, then for each $j \in \{L, H\}$,

$$q_j^{FB} \in \text{argmax}_q \ \theta_j q - c(q), \qquad c'(q_j^{FB}) = \theta_j$$

whenever the optimum is interior.

High-type allocation

The second-best high-type allocation solves

$$\max_{q_H} \ \alpha\left[\overline{\theta} q_H - c(q_H)\right].$$

So it satisfies the same FOC:

$$c'(q_H^{SB}) = \overline{\theta}.$$

Hence

$$q_H^{SB} = q_H^{FB}.$$

No distortion at the top

Serving the high type more does not relax any incentive problem above it. So the high type stays at first best.

Why the low type is distorted downward

Low-type allocation

The low-type allocation solves

$$\max_{q_L} \left[(1-\alpha)\underline{\theta} - \alpha(\overline{\theta} - \underline{\theta})\right] q_L - (1-\alpha)c(q_L).$$

If the low type is served

The FOC is

$$c'(q_L^{SB}) = \underline{\theta} - \frac{\alpha}{1-\alpha}(\overline{\theta} - \underline{\theta}) < \underline{\theta}.$$

Since c' is increasing,

$$q_L^{SB} < q_L^{FB}.$$

Interpretation

Relative to first best, the low contract is cut back because raising q_L also raises the information rent that must be left to the high type.

When the low type is excluded

Exclusion condition

If

$$\underline{\theta} - \frac{\alpha}{1-\alpha}(\overline{\theta} - \underline{\theta}) \leq c'(0),$$

then the optimum sets

$$q_L^{SB} = 0.$$

So exclusion is more likely when $\overline{\theta} - \underline{\theta}$ is large or when α is high.

Canonical result

The high type is efficient. The low type is distorted downward, and may be excluded entirely when the rent cost is too large.

What the two-type solution is really saying

Economic interpretation

- The principal would like to extract all surplus.
- Private information prevents full extraction.
- The downward distortion is an allocative inefficiency created by incentive constraints.
- Lowering q_L is valuable because it relaxes the high type's temptation to imitate.

Basic versus premium menu

The basic contract is intentionally degraded so that high types do not find it too tempting.

What should you remember?

- low type: zero rent,
- high type: positive information rent,
- low contract: distorted downward,
- top contract: no distortion in the canonical model.

Bridge to the continuum case

Once the type space becomes continuous, the same logic reappears as:

$$\text{monotonicity} + \text{envelope} + \text{payment identity.}$$

From two types to a continuum

Environment

Now let

$$\theta \in [\underline{\theta}, \overline{\theta}]$$

with distribution F and density f.
The principal chooses a direct mechanism

$$\hat{\theta} \mapsto \big(q(\hat{\theta}),\ p(\hat{\theta})\big).$$

Utility

If the true type is θ, reporting $\hat{\theta}$ yields

$$U(\hat{\theta}; \theta) = \theta q(\hat{\theta}) - p(\hat{\theta}).$$

Truthful utility is

$$U(\theta) = U(\theta; \theta).$$

Question

What restrictions does truthful IC impose on the shape of $q(\theta)$ and on the payment schedule $p(\theta)$?

IC implies monotone allocation

Take two types $\theta > \hat{\theta}$

Truthful IC gives two inequalities:

$$\theta q(\theta) - p(\theta) \geq \theta q(\hat{\theta}) - p(\hat{\theta}),$$
$$\hat{\theta} q(\hat{\theta}) - p(\hat{\theta}) \geq \hat{\theta} q(\theta) - p(\theta).$$

Add them

$$(\theta - \hat{\theta})\left[q(\theta) - q(\hat{\theta})\right] \geq 0.$$

Since $\theta - \hat{\theta} > 0$, we obtain

$$q(\theta) \geq q(\hat{\theta}).$$

Interpretation

Higher types must receive weakly more quality. If the allocation ran backward somewhere, nearby types would want to swap reports.

Continuous single-crossing

Same condition as in the two-type model

For two reports $r > \hat{r}$,

$$\big[\theta q(r) - p(r)\big] - \big[\theta q(\hat{r}) - p(\hat{r})\big]$$

has derivative $q(r) - q(\hat{r}) \geq 0$ with respect to θ. So higher types find the higher report relatively more attractive. This is exactly the continuous analogue of

$$\Delta U(\theta) = \theta \Delta q - \Delta p$$

from the two-type case.
More generally, single-crossing is the increasing-differences condition

$$\frac{\partial^2 u(q,\theta)}{\partial q\, \partial \theta} \geq 0.$$

Here $u(q,\theta) = \theta q$, so the cross-partial is simply 1.

Why we did not name it in auctions

Lecture 6 used the same structure implicitly:

$$v x_i(\hat{v}) - p_i(\hat{v})$$

is also linear in the true type v. So the same single-crossing property was already built into the auction environment even though we did not label it separately.

Envelope formula: local step

Envelope step

Define interim utility

$$U(\hat{\theta}; \theta) = \theta q(\hat{\theta}) - p(\hat{\theta}),$$

and let truthful utility be

$$U(\theta) = U(\theta; \theta).$$

Under truthful IC, type θ solves

$$\max_{\hat{\theta}} \; U(\hat{\theta}; \theta).$$

So the envelope theorem gives

$$U'(\theta) = \left. \frac{\partial U(\hat{\theta}; \theta)}{\partial \theta} \right|_{\hat{\theta} = \theta} = q(\theta).$$

Meaning

IC implies that the slope of the truthful rent function equals the allocation rule.

Envelope formula: integral form

Integrate the slope condition

Since

$$U'(\theta) = q(\theta),$$

we obtain

$$U(\theta) = U(\underline{\theta}) + \int_{\underline{\theta}}^{\theta} q(t)\, dt.$$

Meaning

Once $q(\theta)$ is fixed, the entire rent schedule is pinned down up to the boundary term $U(\underline{\theta})$.

Bottom IR as the boundary condition

Canonical boundary term

In the canonical solution,

$$U(\underline{\theta}) = 0.$$

Why this binds

Because the principal maximizes profit, any positive bottom rent is a pure giveaway. Raising every payment by the same constant preserves IC, so the optimum drives this constant down to zero.

Why this matters

The envelope formula tells us that IC leaves only one free constant. Bottom IR pins down that constant, so the payment schedule is fully determined once $q(\theta)$ is chosen.

Payment identity

Start from the envelope representation

$$U(\theta) = \theta q(\theta) - p(\theta) = U(\underline{\theta}) + \int_{\underline{\theta}}^{\theta} q(t)\ dt.$$

Rearrange:

$$p(\theta) = \theta q(\theta) - U(\underline{\theta}) - \int_{\underline{\theta}}^{\theta} q(t)\ dt.$$

If bottom IR binds

$$p(\theta) = \theta q(\theta) - \int_{\underline{\theta}}^{\theta} q(t)\ dt.$$

Same message as in auctions

Payments are not an independent free variable once IC is imposed. The allocation rule determines rents, and rents determine payments.

Rewriting the principal's objective

Start from expected profit

$$\Pi = \int_{\underline{\theta}}^{\overline{\theta}} \left[p(\theta) - c(q(\theta))\right] f(\theta)\ d\theta.$$

Substitute the payment identity:

$$\Pi = \int_{\underline{\theta}}^{\overline{\theta}} \left[\theta q(\theta) - U(\underline{\theta}) - \int_{\underline{\theta}}^{\theta} q(t)\ dt - c(q(\theta))\right] f(\theta)\ d\theta.$$

The only nontrivial term is

$$\int_{\underline{\theta}}^{\overline{\theta}} \int_{\underline{\theta}}^{\theta} q(t) f(\theta)\ dt\ d\theta.$$

The Fubini step

Swap the order of integration

By Fubini,

$$\int_{\underline{\theta}}^{\overline{\theta}} \int_{\underline{\theta}}^{\theta} q(t) f(\theta)\; dt\; d\theta = \int_{\underline{\theta}}^{\overline{\theta}} \int_{t}^{\overline{\theta}} q(t) f(\theta)\; d\theta\; dt$$

$$= \int_{\underline{\theta}}^{\overline{\theta}} \big[1 - F(t)\big] q(t)\; dt.$$

Renaming t back to θ yields the compact formula below.

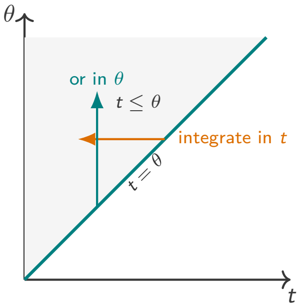

$$\underline{\theta} \le t \le \theta \le \overline{\theta}.$$

Key takeaway

The rent term is weighted by $1 - F(\theta)$, because improving the contract at type θ leaves extra rent to every higher type.

Virtual-value representation

Plug the Fubini step back in

$$\Pi = \int_{\underline{\theta}}^{\overline{\theta}} \left[\theta - \frac{1 - F(\theta)}{f(\theta)}\right] q(\theta) f(\theta)\ d\theta$$
$$- \int_{\underline{\theta}}^{\overline{\theta}} c(q(\theta)) f(\theta)\ d\theta - U(\underline{\theta}).$$

Screening virtual value

$$\psi(\theta) = \theta - \frac{1 - F(\theta)}{f(\theta)}.$$

Then the principal behaves as if type θ contributes virtual surplus

$$\psi(\theta) q(\theta) - c(q(\theta)).$$

Pointwise candidate

Ignoring monotonicity for a moment, choose

$$q^{SB}(\theta) \in \operatorname{argmax}_q\ \psi(\theta) q - c(q).$$

First-order condition

When the solution is interior,

$$c'(q^{SB}(\theta)) = \psi(\theta)$$

pointwise in θ.

When pointwise maximization is implementable

Monotonicity conditions

In this single-crossing environment, the pointwise candidate is increasing if

- $\psi(\theta)$ is increasing (regularity),
- $c(q)$ is convex so $c'(q)$ is increasing.

Then the first-order condition produces a monotone allocation rule.

Why this matters

In a single-dimensional quasilinear problem, an increasing allocation rule together with the payment identity is sufficient for truthful IC.

If the pointwise solution is not monotone

The pointwise maximizer cannot be implemented directly. The nonmonotone region must be replaced by an **ironed** or **pooled** allocation, which creates bunching of nearby types.

Why low types are distorted and the top is not

First-best benchmark

Ignoring IC, pointwise efficiency solves

$$\theta = c'(q^{FB}(\theta)).$$

Second-best screening

With IC,

$$\psi(\theta) = c'(q^{SB}(\theta)).$$

Since $\psi(\theta) < \theta$ for interior types,

$$q^{SB}(\theta) < q^{FB}(\theta).$$

Top type

At the upper endpoint,

$$1 - F(\bar{\theta}) = 0.$$

So

$$\psi(\bar{\theta}) = \bar{\theta}.$$

So the top type satisfies the same condition as first best.

No distortion at the top

Reducing service for the highest type relaxes no higher IC constraint, because there is no type above it. The distortion burden falls on lower types.

A concrete continuous example

Suppose

$$\theta \sim U[0,1], \qquad c(q) = \frac{q^2}{2}.$$

Then

$$\psi(\theta) = \theta - \frac{1-\theta}{1} = 2\theta - 1.$$

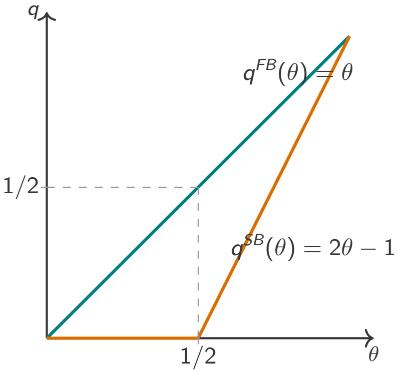

Solution

$$q^{FB}(\theta) = \theta$$
$$q^{SB}(\theta) = \max\{2\theta - 1, 0\}.$$

Reading the picture

- Types below $1/2$ are excluded.
- Types above $1/2$ are served, but below first best.
- At $\theta = 1$, there is no distortion: $q^{SB}(1) = q^{FB}(1) = 1$.

Why screening can be solved one consumer at a time

Screening with many consumers

A monopolist may face many consumers, not just one. But once the menu is posted, consumer A's choice does not change what can be offered to consumer B. So expected profit is additive across consumers, and we can analyze one representative consumer at a time.

Auctions with many bidders

An auction may also involve many agents, or multiple objects. The key difference is that allocations are linked by feasibility constraints:

$$\sum_{i=1}^{n} x_i(v) \leq 1$$

for a single object, and analogous capacity constraints with multiple objects.
Giving the object to bidder i changes what is feasible for bidder j.

Key difference

Both models may involve many agents. The real distinction is whether outcomes are coupled by scarcity. Screening often separates consumer by consumer; auctions must be solved jointly.

From revenue equivalence to Myerson

What screening taught us

In screening, once IC gives us the payment identity, profit can be rewritten as

$$\int \left[\psi(\theta)q(\theta) - c(q(\theta))\right] f(\theta)\ d\theta.$$

So the principal chooses the allocation rule to maximize profit.

What Lecture 6 did and did not say

Revenue equivalence said:

- if the allocation rule is fixed,
- then IC pins down rents and expected revenue.

That does *not* mean revenue cannot be improved. It means the seller must improve revenue by choosing a different allocation rule.

Myerson's question

Among all feasible auction allocation rules, which one maximizes expected revenue? The answer is the auction analog of screening: rewrite revenue using the envelope formula, obtain **virtual values**, and then choose the allocation rule accordingly. This is where reserve prices come from.

Myerson's revenue formula

Same envelope algebra as screening

For bidder i, let $x_i(v_i)$ be interim win probability. Under IC and IR,

$$u_i(v_i) = u_i(\underline{v}) + \int_{\underline{v}}^{v_i} x_i(t)\ dt, \qquad t_i(v_i) = v_i x_i(v_i) - u_i(v_i).$$

Take expectations:

$$\mathbb{E}[t_i(v_i)] = \int_{\underline{v}}^{\overline{v}} \left[v x_i(v) - u_i(\underline{v}) - \int_{\underline{v}}^{v} x_i(t)\ dt \right] f(v)\ dv.$$

Swap the integrals exactly as in screening:

$$\int_{\underline{v}}^{\overline{v}} \int_{\underline{v}}^{v} x_i(t) f(v)\ dt\ dv = \int_{\underline{v}}^{\overline{v}} [1 - F(t)] x_i(t)\ dt.$$

Revenue formula

Therefore

$$\mathbb{E}[t_i(v_i)] = \int_{\underline{v}}^{\overline{v}} \left[v - \frac{1 - F(v)}{f(v)} \right] x_i(v) f(v)\ dv - u_i(\underline{v}).$$

Define the virtual value

$$\phi(v) = v - \frac{1 - F(v)}{f(v)}.$$

Summing across bidders gives

$$\mathbb{E}\left[\sum_{i=1}^{n} t_i(v) \right] = \mathbb{E}\left[\sum_{i=1}^{n} \phi(v_i) x_i(v) \right] - \sum_{i=1}^{n} u_i(\underline{v}).$$

Same logic, different object

Screening replaces θ by $\psi(\theta)$. Optimal auction design replaces each bidder's value v_i by a virtual value $\phi(v_i)$. The envelope theorem is doing the same job in both problems.

Myerson's optimal allocation rule

Revenue objective

To maximize expected revenue, choose a feasible allocation rule that maximizes

$$\sum_{i=1}^{n} \phi(v_i)x_i(v).$$

Single-object case

Feasibility means

$$\sum_{i=1}^{n} x_i(v) \leq 1.$$

So the object goes to the bidder with the highest nonnegative virtual value.

Why this differs from efficiency

Efficiency chooses the highest *actual* value. Revenue maximization chooses the highest *virtual* value, and may refuse to sell even when bids are positive.

Regular case

If $\phi(v)$ is increasing, ranking bidders by $\phi(v)$ is the same as ranking by v, but now there is a reserve cutoff.

Reserve price interpretation

Regular i.i.d. environment

If $\phi(v)$ is increasing, then maximizing virtual value is equivalent to:

- ignore all bidders with $v < r$,
- among the rest, award the object to the highest bidder,
- choose payments to keep the mechanism IC.

Reserve price

The reserve r solves

$$\phi(r) = r - \frac{1 - F(r)}{f(r)} = 0.$$

So Myerson becomes:

- a second-price auction with reserve r, or
- any other IC mechanism implementing the same allocation rule.

Connection to Lecture 6

Revenue equivalence still applies *conditional on a fixed allocation rule*. Myerson improves revenue by changing the allocation rule itself.

Uniform example: Myerson = second price with reserve 1/2

Virtual value

If

$$v \sim U[0,1],$$

then

$$\phi(v) = v - \frac{1-v}{1} = 2v - 1.$$

So the reserve solves

$$\phi(r) = 0 \iff r = \frac{1}{2}.$$

How the rule works

- If bids are $(0.9, 0.6)$, sell to 0.9 and charge 0.6.
- If bids are $(0.9, 0.4)$, sell to 0.9 and charge 0.5.
- If bids are $(0.45, 0.30)$, do not sell.

Revenue versus efficiency

The last case is inefficient relative to pure surplus maximization, but it can be optimal for revenue because low-value sales generate no useful competition and depress expected payments.

The same rule implemented by first-price bids

Two bidders, uniform values

Keep the same environment:

$$v_1, v_2 \sim U[0,1], \qquad r = \frac{1}{2}.$$

Types below r stay out. For a type $v \geq r$, deviating to the bid of type $\hat{v}$ gives

$$u(\hat{v}; v) = \left[v - b(\hat{v})\right] \Pr(v_j \leq \hat{v}) = \left[v - b(\hat{v})\right] \hat{v}.$$

So in a symmetric increasing equilibrium,

$$u'(v) = v, \qquad u(r) = 0, \qquad u(v) = \frac{v^2 - r^2}{2}.$$

Using $u(v) = (v - b(v))v$,

$$b(v) = \frac{v}{2} + \frac{r^2}{2v} = \frac{v}{2} + \frac{1}{8v}, \qquad v \in \left[\frac{1}{2}, 1\right].$$

What changed?

Without reserve, the benchmark bid is

$$b(v) = \frac{v}{2}.$$

With reserve $1/2$, that formula is no longer correct:

- the cutoff type now satisfies $b(1/2) = 1/2$,
- the top type bids $b(1) = 5/8 > 1/2$,
- the whole bid schedule shifts up, not just truncates,
- because $b(v)$ is increasing, the highest value above $1/2$ still submits the highest bid above reserve.

So first price with reserve $1/2$ implements the same allocation rule as Myerson, even though payments are generated differently from second price.

VCG, revenue equivalence, and Myerson

VCG

- objective: actual surplus,
- allocation: efficient,
- payment: externality pricing.

Revenue equivalence

- fixes an allocation rule,
- then IC pins down rents,
- hence expected revenue is fixed.

Myerson

- objective: expected revenue,
- allocation: maximize virtual surplus,
- instrument: reserve or exclusion.

One sentence summary

VCG says how to price an efficient allocation. Revenue equivalence says pricing differences do not matter once allocation is fixed. Myerson says the seller should sometimes change the allocation rule itself.

Screening and optimal auction side by side

	Screening	**Optimal auction**
Allocation object	quality or quantity $q(\theta)$	winning probability $x_i(v)$
Rent derivative	$U'(\theta) = q(\theta)$	$u_i'(v_i) = x_i(v_i)$
Payment identity	$p = \theta q - U$	$t_i = v_i x_i - u_i$
Virtual object	$\psi(\theta) = \theta - \frac{1-F}{f}$	$\phi(v) = v - \frac{1-F}{f}$
Distortion instrument	lower $q(\theta)$ for low types	reserve price / no sale for low bids
Top-type result	no distortion at top	highest virtual value among active bidders wins

Core message

Auctions and screening are not two separate theories. They are two applications of the same IC-envelope-payment logic under different feasibility constraints.

End of Lecture 7

- Screening is mechanism design with a menu: the principal chooses allocation and payment so types self-select.
- In the two-type model, IR-L and IC-H are the key binding constraints in the canonical solution.
- Information rent is the price the principal pays for hidden information.
- In the continuous model, IC implies allocation monotonicity and the envelope formula, which pins down the payment identity.
- Low types are distorted because improving their contract raises the information rent left to higher types.
- Myerson's auction is the auction analog of screening: maximize virtual surplus, not actual surplus.

Bridge to Lecture 8

What we controlled here

The principal controlled:

- allocations,
- payments,
- participation incentives,
- reporting incentives.

What comes next

Lecture 8 shifts attention from **mechanism design** to **information design** and bargaining:

- change what agents know,
- change what beliefs they form,
- change how bargaining power translates into outcomes.

Continuity

The common theme is still strategic response under asymmetric information. The designer just moves from choosing transfers and menus to choosing information structures.

References

Selected references for Lecture 7

- Mas-Colell, A., Whinston, M. D., and Green, J. R. (1995). *Microeconomic Theory*. Oxford University Press, esp. Ch. 13 and Ch. 23.
- Jehle, G. A., and Reny, P. J. (2011). *Advanced Microeconomic Theory*. 3rd ed. Financial Times/Prentice Hall, esp. screening and mechanism design chapters.
- Bolton, P., and Dewatripont, M. (2005). *Contract Theory*. MIT Press.
- Myerson, R. (1981). "Optimal Auction Design." *Mathematics of Operations Research*.
- Mussa, M., and Rosen, S. (1978). "Monopoly and Product Quality." *Journal of Economic Theory*.