Advanced Microeconomics

Lecture 6: Auctions and Mechanism Design

Yihu Hou

Renmin University of China

Spring 2026

Outline

1. Why auctions are the clean entry point to mechanism design
2. Auction formats: first price, second price, English, and Dutch
3. The mechanism design language: direct mechanisms, IC, and IR
4. Revenue equivalence
5. VCG as externality pricing

From information economics to mechanism design

Lecture 5

We treated private information as a feature of the environment:

- types are privately observed,
- beliefs matter,
- signals and menus can reveal information.

Question then

Given the game, what equilibrium behavior should we expect?

Lecture 6

We now put the designer at the center:

- choose a rule,
- anticipate strategic reporting or bidding,
- make the desired behavior optimal.

Question now

Given an objective, what game or rule should we design?

Mechanism design

Mechanism design is often called “reverse game theory”: instead of solving a given game, we choose the game so that equilibrium behavior has the property we want.

Why auctions are a useful starting point

Clean private information

Each bidder has a private valuation:

$$v_i \in V_i.$$

The seller does not know the valuations, but wants to allocate one object and collect payments.

Observable actions

Bids are simple messages:

$$b_i = \text{a number submitted by bidder } i.$$

Clear allocation and payment

Every auction rule specifies:

- who gets the object,
- how much each bidder pays,
- how these depend on submitted bids.

Clear objectives

Common objectives include:

- efficiency: allocate to the highest valuation,
- revenue: maximize expected payment to the seller.

The baseline environment

Private-value single-object auction

- one indivisible object,
- n bidders,
- bidder i has private value v_i,
- standard auction benchmark: only the winner pays.

If bidder i wins and pays p_i, then

$$u_i = v_i - p_i.$$

If bidder i loses, then

$$u_i = 0.$$

Special case used for calculations

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

independently across bidders.

Why start here?

This keeps the first benchmark close to real auctions: you pay only when you get the object. Later, when we write a general mechanism, payments can be more flexible.

Private value, common value, and affiliated signals

Private value

- each bidder values the object personally,
- others' signals do not change my valuation.

Examples: collectibles, idiosyncratic use value.

Common value

- the object has one true value,
- bidders observe noisy signals.

Examples: oil leases, mineral rights.

Affiliated signals

- signals are positively correlated,
- others' bids reveal information about my estimate.

This is closer to many real auctions.

Scope today

Today we impose three separate benchmark assumptions:

- **Private values**: bidder i knows v_i, and v_i is not revised based on others' signals or on the fact of winning.
- **Quasilinear utility**: money enters utility linearly, so bidders are risk-neutral over monetary transfers.
- **One indivisible object**: the outcome is just win or lose an object with deterministic value v_i.

Why the winner's curse belongs to common-value auctions

Common-value logic

Suppose the object has one true value w, but bidders only observe noisy signals:

$$s_i = w + \varepsilon_i.$$

If I win, that usually means my signal was the most optimistic one in the room.

Why winning is bad news

Conditional on winning, my estimate is biased upward:

$$\mathbb{E}[w \mid s_i, \text{ win}] < s_i.$$

So bidding as if s_i were the true value tends to make the winner overpay.

Private-value benchmark

With private values, bidder i's valuation is just v_i.

- Winning tells me others value the object less.
- But it does **not** lower my own valuation.

So the main issue is strategic bidding, not adverse information revealed by winning.

Winner's curse

The curse is not "the winner paid a high price." The curse is that winning itself reveals bad news about the object's value when values are common or signals are affiliated.

Four common auction formats

Sealed-bid formats

- **First-price**: highest bidder wins and pays own bid.
- **Second-price**: highest bidder wins and pays the second-highest bid.

Strategic distinction

In a first-price auction, my bid affects both my probability of winning and the price I pay if I win.

Open formats

- **English**: price rises until only one bidder remains.
- **Dutch**: price falls until one bidder accepts.

Strategic distinction

Open formats reveal some information through the timing of exit or acceptance.

A useful equivalence map

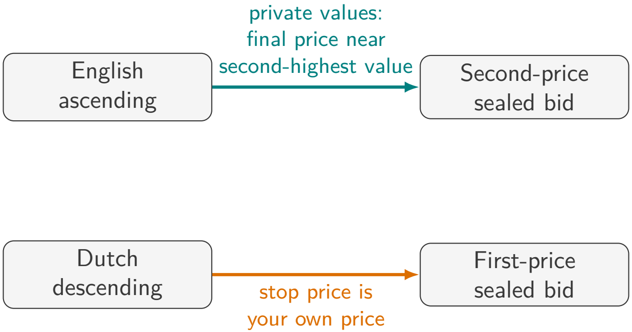

Interpretation

English is close to second-price because the winner pays approximately the value of the last competitor. Dutch is strategically close to first-price because accepting the clock price is like submitting the winning bid.

First-price auction: bid shading

Rule

Highest bid wins and pays own bid.

Tradeoff

If a bidder with value v submits bid b:

- a higher b raises the probability of winning,
- but lowers the surplus $v - b$ conditional on winning.

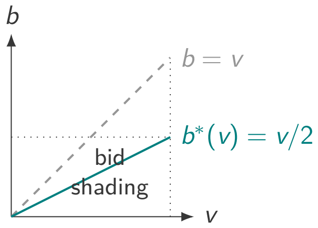

Recall from Lecture 5

With two bidders and $v_i \sim U[0, 1]$, the symmetric BNE is $b^*(v) = v/2$.

First-price equilibrium: guess a linear strategy

Guess a symmetric linear strategy

Suppose every opponent bids

$$b(v) = \beta v, \qquad 0 < \beta < 1.$$

If bidder i with value v deviates and submits bid b, then bidder i wins exactly when every opponent's value is below b/β. For $b \leq \beta$,

$$\Pr(\text{win} \mid b) = \left(\frac{b}{\beta}\right)^{n-1}.$$

So expected payoff from bid b is

$$U(b \mid v) = (v - b)\left(\frac{b}{\beta}\right)^{n-1}.$$

First-price equilibrium: verify the guess

Best response and verification

Ignoring the constant $\beta^{-(n-1)}$, type v solves

$$\max_b\, (v-b)b^{n-1}.$$

The first-order condition is

$$(n-1)(v-b)b^{n-2} - b^{n-1} = 0 \quad \Longrightarrow \quad b^*(v) = \frac{n-1}{n}v.$$

So the guessed linear form is verified by

$$\beta^* = \frac{n-1}{n}.$$

Remark

For class, the guess-and-verify route is the cleanest. More generally, one can start from an arbitrary increasing rule $b = \beta(v)$ and solve a differential equation.

Second-price sealed-bid auction

Rule

- every bidder submits a sealed bid b_i,
- highest bid wins,
- winner pays the second-highest bid,
- losers pay zero.

Notation

Let

$$m_i = \max_{j \neq i} b_j$$

be the highest competing bid faced by bidder i.

Key observation

If bidder i wins, the payment is m_i.
It does **not** depend on the exact level of b_i, as long as $b_i > m_i$.

Dominant strategy claim

For every bidder i, bidding truthfully

$$b_i(v_i) = v_i$$

is a weakly dominant strategy.

Why truth-telling is dominant: the threshold view

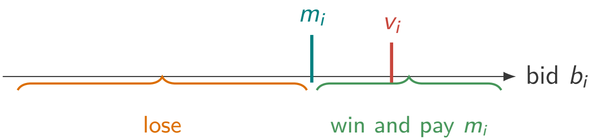

If $v_i > m_i$

Winning is profitable, because payoff is

$$v_i - m_i > 0.$$

Bidding $b_i = v_i$ wins exactly in this case. Underbidding below m_i can only turn a profitable win into a loss.

Message

When the object is worth more to you than the threshold price, you want to cross the threshold; bidding your value does that without changing the price you pay.

Why truth-telling is dominant: the other case

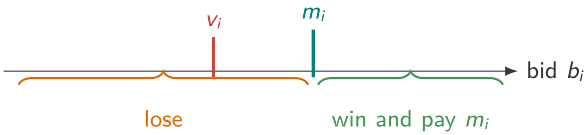

If $v_i < m_i$

Winning is unprofitable, because payoff would be

$$v_i - m_i < 0.$$

Bidding $b_i = v_i$ loses exactly in this case. Overbidding above m_i can only turn a zero payoff into a negative payoff.

Dominant strategy

Truthful bidding is weakly best no matter what m_i is. This is stronger than a Bayesian equilibrium argument.

First-price vs. second-price: same allocation, different incentives

First-price

winner pays own bid.

- no dominant strategy in general,
- equilibrium bid depends on n and F,
- bidders shade below value.

Second-price

winner pays the highest losing bid.

- truthful bidding is dominant,
- does not require knowing n or F for the basic incentive result,
- winner still gets positive surplus when $v_i > m_i$.

Same efficient allocation under standard assumptions

Both allocate the object to the highest valuation in equilibrium, but they implement that allocation through different incentive logic.

$n = 2$, uniform example: expected revenue

First-price

With $b^*(v) = v/2$, the winning bid is

$$\frac{1}{2}\max\{v_1, v_2\}.$$

$$\Pr(\max\{v_1, v_2\} \leq z) = z^2 \quad \Rightarrow \quad f_{\max}(z) = 2z.$$

Hence

$$\mathbb{E}[\max\{v_1, v_2\}] = \int_0^1 z \cdot 2z \, dz = \frac{2}{3},$$

expected revenue is

$$\frac{1}{2} \cdot \frac{2}{3} = \frac{1}{3}.$$

Second-price

Truthful bidding implies the winner pays

$$\min\{v_1, v_2\}.$$

$$\Pr(\min\{v_1, v_2\} > z) = (1 - z)^2 \quad \Rightarrow \quad f_{\min}(z) = 2(1 - z).$$

Hence

$$\mathbb{E}[\min\{v_1, v_2\}] = \int_0^1 z \cdot 2(1 - z) \, dz = \frac{1}{3},$$

expected revenue is

$$\frac{1}{3}.$$

Question

Is this equality just a coincidence of the two-bidder uniform example, or does it reflect a deeper principle?

Auction rules as mechanisms

General mechanism

Each bidder sends a message $m_i \in M_i$. The mechanism chooses:

$$x_i(m_1, \ldots, m_n) \in [0, 1],$$

the probability that bidder i gets the object, and

$$p_i(m_1, \ldots, m_n),$$

the payment made by bidder i.

Utility

For true value v_i,

$$u_i(m; v_i) = v_i x_i(m) - p_i(m).$$

Feasibility

For a single object,

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

Unifying move

A mechanism is fully described by its **message spaces**, **allocation rule**, and **payment rule**. First-price, second-price, English, and Dutch auctions differ only in these three objects.

A concrete bridge to direct mechanisms

Original indirect description

In the two-bidder uniform first-price auction, equilibrium bidding is

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

Players submit bids, the highest bid wins, and the winner pays their own bid.

Equivalent direct description

Ask each bidder to report a value $\hat{v}_i$. Internally convert it to the bid

$$b_i = \frac{\hat{v}_i}{2}.$$

Because dividing by 2 preserves ranking, the highest report wins.

Resulting direct rule

The direct mechanism is

$$x_i(\hat{v}) = \begin{cases} 1, & \hat{v}_i = \max_j \hat{v}_j, \\ 0, & \text{otherwise}, \end{cases}$$

and

$$p_i(\hat{v}) = \begin{cases} \hat{v}_i/2, & \text{if } i \text{ wins}, \\ 0, & \text{otherwise}. \end{cases}$$

Teaching point

A direct mechanism can ask for values and internally convert them into the messages used in the original auction.

Indirect versus direct mechanisms

Indirect mechanism

Agents send arbitrary messages:

$$m_i \in M_i.$$

The mechanism chooses

$$x_i(m), \qquad p_i(m).$$

A strategy is a rule

$$\sigma_i : V_i \to M_i$$

mapping a true type into a message.

Direct mechanism

The message space is the type space itself:

$$\hat{v}_i \in V_i.$$

The mechanism chooses

$$x_i(\hat{v}), \qquad p_i(\hat{v}).$$

A strategy is now a reporting rule

$$\rho_i : V_i \to V_i.$$

Truth-telling means $\rho_i(v_i) = v_i$; a deviation is a **misreport** that **mimics** another type.

Key distinction

IC means no type wants to mimic any other report. The strategy space is still there; direct mechanisms just make the messages interpretable as types.

The revelation principle

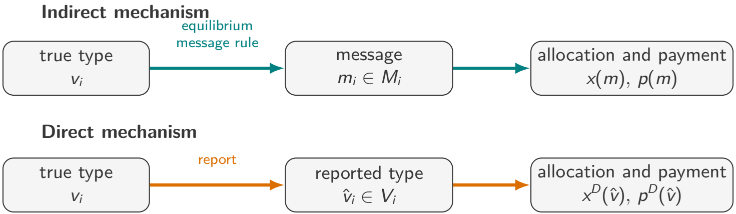

Statement

If an outcome can be implemented by some mechanism in equilibrium, then there is a direct mechanism in which each agent reports a type and truthful reporting implements the same outcome.

Why truth-telling works in the constructed direct mechanism

$$x^D(\hat{v}) = x(m^*(\hat{v})), \qquad p^D(\hat{v}) = p(m^*(\hat{v})).$$

Step 1: truthful reporting reproduces the old equilibrium path

If type v_i reports $\hat{v}_i = v_i$, the direct mechanism feeds the same message $m_i^*(v_i)$ into the old mechanism. So the allocation and payment are exactly the same as in the original equilibrium.

Step 2: a misreport reproduces an old deviation

If the same type reports $\hat{v}_i \neq v_i$, the direct mechanism instead feeds in $m_i^*(\hat{v}_i)$. That gives exactly the payoff type v_i would obtain in the original indirect mechanism by deviating to that message while the other players keep following their equilibrium message rules.

Conclusion

Because the original equilibrium message $m_i^*(v_i)$ is already a best response for type v_i, truthful reporting weakly beats every misreport $\hat{v}_i$ in the direct mechanism.

Incentive compatibility and individual rationality

Direct utility

If bidder i has true value v_i but reports $\hat{v}_i$, expected utility is

$$U_i(\hat{v}_i; v_i) = \mathbb{E}_{v_{-i}}[v_i x_i(\hat{v}_i, v_{-i}) - p_i(\hat{v}_i, v_{-i})] .$$

Truthful payoff

$$U_i(v_i; v_i).$$

Incentive compatibility

Truthful reporting is optimal:

$$U_i(v_i; v_i) \geq U_i(\hat{v}_i; v_i) \quad \forall \hat{v}_i, \; v_i.$$

Individual rationality

Participation is weakly better than opting out:

$$U_i(v_i; v_i) \geq 0 \quad \forall v_i.$$

Economic interpretation

IC controls lying. IR controls participation.

First-price auction as a direct truthful mechanism

Direct rule induced by the equilibrium bid function

In the two-bidder uniform benchmark, the equilibrium first-price bid is

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

So a direct version asks for reports $\hat{v}_i$, gives the object to the highest report, and charges

$$p_i(\hat{v}) = \begin{cases} \hat{v}_i/2, & \text{if } i \text{ wins,} \\ 0, & \text{otherwise.} \end{cases}$$

Interpretation

The direct mechanism is truthful only because it has already built the first-price shading rule into the payment formula.

Why the direct first-price version is truthful

Deviation calculation

If bidder i has true value v_i but reports $\hat{v}_i$, expected utility is

$$U_i(\hat{v}_i; v_i) = \left(v_i - \frac{\hat{v}_i}{2}\right) \Pr(\hat{v}_i > v_j).$$

With $v_j \sim U[0,1]$, this becomes

$$U_i(\hat{v}_i; v_i) = \left(v_i - \frac{\hat{v}_i}{2}\right) \hat{v}_i.$$

Best report

Differentiate with respect to $\hat{v}_i$:

$$\frac{\partial U_i(\hat{v}_i; v_i)}{\partial \hat{v}_i} = v_i - \hat{v}_i.$$

So the optimal report is

$$\hat{v}_i = v_i.$$

Second-price auction as a direct truthful mechanism

Allocation

If reported values are $\hat{v} = (\hat{v}_1, \ldots, \hat{v}_n)$,

$$x_i(\hat{v}) = \begin{cases} 1, & \hat{v}_i = \max_j \hat{v}_j, \\ 0, & \text{otherwise.} \end{cases}$$

Ignore ties or break them randomly.

Payment

If bidder i wins,

$$p_i(\hat{v}) = \max_{j \neq i} \hat{v}_j.$$

If bidder i loses,

$$p_i(\hat{v}) = 0.$$

IC and efficiency

The second-price auction is IC in dominant strategies and allocates the object to the highest reported value. Under truth-telling, this is the highest true value.

Interim objects behind IC

Interim objects

For bidder i, define the interim allocation probability and interim expected payment by

$$x_i(v_i) = \mathbb{E}_{v_{-i}}[x_i(v_i, v_{-i})] , \qquad p_i(v_i) = \mathbb{E}_{v_{-i}}[p_i(v_i, v_{-i})] .$$

Truthful interim utility is then

$$U_i(v_i) = v_i x_i(v_i) - p_i(v_i).$$

Two IC inequalities

For any two reports v and $\hat{v}$, truthful IC implies

$$U_i(v) \geq v\, x_i(\hat{v}) - p_i(\hat{v}),$$

$$U_i(\hat{v}) \geq \hat{v}\, x_i(v) - p_i(v).$$

Next step

These two inequalities are the starting point for monotonicity and the envelope formula.

Allocation monotonicity is a necessary IC condition

Start from the two IC inequalities

Add

$$U_i(v) \geq v\,x_i(\hat{v}) - p_i(\hat{v})$$

and

$$U_i(\hat{v}) \geq \hat{v}\,x_i(v) - p_i(v)$$

after substituting

$$U_i(v) = v\,x_i(v) - p_i(v), \qquad U_i(\hat{v}) = \hat{v}\,x_i(\hat{v}) - p_i(\hat{v}).$$

This gives

$$(v - \hat{v})\big(x_i(v) - x_i(\hat{v})\big) \geq 0.$$

What does monotonicity mean?

If $v > \hat{v}$, then $x_i(v) \geq x_i(\hat{v})$. So allocation monotonicity is a **necessary** condition for truthful implementation in any single-dimensional direct mechanism.

Envelope theorem: why $U_i'(v) = x_i(v)$

General envelope logic

If
$$W(t) = \max_a f(a, t),$$
then, under regularity conditions,
$$W'(t) = \frac{\partial f(a, t)}{\partial t}\Big|_{a=a^*(t)}.$$
The indirect effect through the optimal choice $a^*(t)$ drops out to first order.

Apply it to truthful reporting

Define
$$f_i(\hat{v}, v) = v\,x_i(\hat{v}) - p_i(\hat{v}).$$
Truthful IC means
$$U_i(v) = \max_{\hat{v}} f_i(\hat{v}, v), \qquad \hat{v}^*(v) = v.$$
This is the general envelope formula with parameter $t = v$ and choice variable $a = \hat{v}$. Hence
$$U_i'(v) = \frac{\partial f_i(\hat{v}, v)}{\partial v}\Big|_{\hat{v}=v} = x_i(v)$$
at differentiability points.

Another route to $U_i'(v) = x_i(v)$

No type wants to mimic its left neighbor

Under IC, for any small $h > 0$,

$$U_i(v) \geq v\,x_i(v-h) - p_i(v-h) = U_i(v-h) + h\,x_i(v-h),$$

because type v does not mimic report $v-h$. Likewise,

$$U_i(v-h) \geq (v-h)x_i(v) - p_i(v) = U_i(v) - h\,x_i(v),$$

because type $v-h$ does not mimic report v.

Squeeze the difference quotient

Rearranging gives

$$x_i(v-h) \leq \frac{U_i(v) - U_i(v-h)}{h} \leq x_i(v).$$

So whenever $x_i(\cdot)$ is continuous at v, the middle term converges to $x_i(v)$. This gives the same conclusion:

$$U_i'(v) = x_i(v).$$

Envelope formula

Integrate the derivative

From the envelope condition,

$$U_i(v) = U_i(0) + \int_0^v x_i(t)\, dt.$$

IR implies $U_i(0) \geq 0$. In the standard auction benchmark, type 0 also never wins and pays 0, so IR pins down

$$U_i(0) = 0.$$

Therefore,

$$U_i(v) = \int_0^v x_i(t)\, dt.$$

Meaning

The utility of type v is the accumulated area under the interim allocation curve. This is the information rent left to that type.

Payment identity and sufficiency

Recover the payment rule

Since

$$U_i(v) = v\,x_i(v) - p_i(v),$$

we get

$$p_i(v) = v\,x_i(v) - U_i(0) - \int_0^v x_i(t)\,dt.$$

So once the allocation rule and boundary utility are fixed, the payment schedule is pinned down.

Necessary or sufficient?

Monotonicity is necessary for truthful IC. In the standard single-dimensional quasilinear setting, a monotone allocation together with this payment formula is also sufficient for truthful implementation.

Why the payment formula is sufficient for IC

Utility from reporting r

Fix a true type v and suppose the mechanism uses

$$p_i(r) = r\,x_i(r) - U_i(0) - \int_0^r x_i(t)\,dt.$$

Then reporting r gives utility

$$U_i(r; v) = v\,x_i(r) - p_i(r) = U_i(0) + \int_0^r x_i(t)\,dt + (v - r)x_i(r).$$

Why no type wants to mimic any other type

If $r < v$, then

$$U_i(v; v) - U_i(r; v) = \int_r^v \big(x_i(t) - x_i(r)\big)\,dt \geq 0.$$

If $r > v$, then

$$U_i(v; v) - U_i(r; v) = \int_v^r \big(x_i(r) - x_i(t)\big)\,dt \geq 0.$$

So every type weakly prefers truthful reporting whenever $x_i(\cdot)$ is monotone.

Why IR also follows

Higher types get weakly more utility

Also,

$$U_i(v) = U_i(0) + \int_0^v x_i(t)\, dt \geq U_i(0).$$

So to conclude IR for all types, we only need the lowest type's IR condition

$$U_i(0) \geq 0.$$

If one additionally wants zero rent for the lowest type, then one sets $U_i(0) = 0$.

Takeaway

Monotonicity plus the payment identity gives truthful IC, and IR for the lowest type propagates upward to all higher types.

Envelope formula: a picture

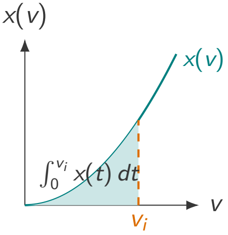

Reading the graph

If $U_i(0) = 0$, then type v_i's equilibrium utility is the area under the allocation probability curve:

$$U_i(v_i) = \int_0^{v_i} x_i(t)\,dt.$$

Payment follows

Because

$$U_i(v_i) = v_i x_i(v_i) - p_i(v_i),$$

the expected payment is

$$p_i(v_i) = v_i x_i(v_i) - U_i(v_i).$$

Revenue equivalence: setup

Environment

- n risk-neutral bidders,
- independent private values drawn from the same distribution on $[0, \bar{v}]$,
- two mechanisms have the same allocation rule,
- the lowest type receives the same expected utility, usually $U_i(0) = 0$.

Revenue equivalence theorem

Any two IC mechanisms with the same allocation rule and the same utility for the lowest type generate the same expected payments for every type, and hence the same expected revenue.

Auction implication

The theorem is not about four named auctions only. Any IC mechanism, however unusual, has the same expected revenue if it induces the same allocation rule and gives the same utility to the lowest type. The standard four auctions are just familiar examples under the benchmark assumptions.

Revenue equivalence: why the result is so strong

Step 1: same allocation

If two mechanisms have the same allocation rule, then they have the same interim allocation probability:

$$x_i(v_i).$$

Step 2: same rent

If both are IC and have the same boundary utility,

$$U_i(v_i) = U_i(0) + \int_0^{v_i} x_i(t)\, dt$$

is the same in both mechanisms.

Step 3: same payment

Expected payment satisfies

$$p_i(v_i) = v_i x_i(v_i) - U_i(v_i).$$

So if x_i and U_i are the same, p_i is the same.

Core intuition

The payment rule can look very different ex post, but IC plus the same allocation rule pins down the expected payment schedule.

Useful non-example

Suppose the highest report gets the object but the winner always pays only $0.01. This will generally not deliver the same revenue as standard auctions, because arbitrary reports need not allocate the object to the highest valuation. The real difference is that the **allocation rule in actual play** is no longer the same.

Revenue equivalence does not mean all auctions are the same

Start with a concrete realization

Suppose there are two bidders with values 3 and 5.

- In a second-price auction, truthful bids give revenue 3.
- In a first-price auction, equilibrium bids are 1.5 and 2.5, so revenue is 2.5.

Question for class: if revenue equivalence is true, why are these revenues different?

Answer

Revenue equivalence is about **expected** revenue, or equivalently expected payments by type. It does not say the seller collects the same amount in every realization.

What can break equivalence?

- risk aversion,
- different reserve prices or allocation rules,
- correlated or common values,
- budget constraints or entry costs,
- collusion.

Takeaway

Revenue equivalence is about expected payments, not equality realization by realization.

VCG: externality pricing for social choice

General social choice problem

There is a set of feasible outcomes O. Society wants to choose an outcome based on agents' reported values.

VCG logic

Choose the outcome maximizing reported total surplus, then make each agent pay the externality imposed on everyone else.

Single-object special case

When the outcomes are simply "who gets the object," the Vickrey auction is exactly the **second-price auction**: the highest-value bidder wins and pays the second-highest bid.

Why this matters

Auctions, public projects, assignment, and other allocation problems look different on the surface, but they can all be studied in the common language of mechanism design.

Key idea

Vickrey equals the second-price auction in the single-object case. What VCG adds is the **externality-pricing** interpretation, and that is what generalizes.

General VCG environment

Environment

- a finite or feasible set of outcomes O,
- agent i has valuation $v_i(o)$ for outcome $o \in O$,
- agents report valuation functions $\hat{v}_i(\cdot)$,
- quasilinear utility:

$$u_i(o, p_i) = v_i(o) - p_i.$$

Efficient outcome rule

VCG chooses

$$o^*(\hat{v}) \in \text{argmax}_{o \in O} \sum_{j=1}^{n} \hat{v}_j(o).$$

What the allocation rule does

It maximizes reported total surplus. The payment rule will make truthful reports support this efficient choice.

VCG payment rule

Payment by agent i

$$p_i(\hat{v}) = \max_{o \in O} \sum_{j \neq i} \hat{v}_j(o) - \sum_{j \neq i} \hat{v}_j(o^*(\hat{v})).$$

First term

The best total value that everyone else could get if agent i were absent:

$$\max_{o \in O} \sum_{j \neq i} \hat{v}_j(o).$$

Second term

The total value everyone else gets under the chosen outcome:

$$\sum_{j \neq i} \hat{v}_j(o^*(\hat{v})).$$

Externality price

The difference is the welfare loss imposed on others by letting agent i participate and influence the outcome.

VCG payment as a diagram

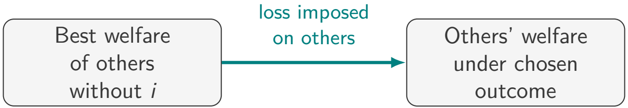

$$p_i = \max_o \sum_{j \neq i} \hat{v}_j(o) - \sum_{j \neq i} \hat{v}_j(o^*)$$

How to read this

If agent i changes the outcome in a way that makes others worse off, i pays for that loss. If i is not pivotal for others, the payment can be zero.

Same logic as second-price

In a single-object auction, the winner's presence prevents the second-highest bidder from getting the object. The loss to others is exactly the second-highest value.

Why VCG is truthful: payoff rewriting

Agent i's payoff under report profile $\hat{v}$

$$v_i(o^*) - p_i(\hat{v}) = v_i(o^*) + \sum_{j \neq i} \hat{v}_j(o^*) - \max_{o \in O} \sum_{j \neq i} \hat{v}_j(o).$$

Key observation

The final term is the best welfare others can obtain without agent i, so it does not depend on agent i's report.

Why VCG is truthful: best-response logic

Mechanism objective

VCG chooses

$$o^*(\hat{v}) \in \operatorname{argmax}_{o \in O} \left[\hat{v}_i(o) + \sum_{j \neq i} \hat{v}_j(o) \right].$$

Agent i's objective

From the payoff rewriting, the term

$$\max_{o \in O} \sum_{j \neq i} \hat{v}_j(o)$$

is fixed. So agent i wants the chosen outcome to maximize

$$v_i(o) + \sum_{j \neq i} \hat{v}_j(o).$$

If $\hat{v}_i = v_i$, the mechanism is maximizing exactly this same expression.

Alignment

Truthful reporting makes the agent's preferred outcome coincide with the efficient outcome.

Single-object VCG reduces to second price

If bidder i wins

Outcome $o = i$ means bidder i gets the object, so $o^* = i$. Without bidder i, the best other bidder gets the object:

$$\max_o \sum_{j \neq i} \hat{v}_j(o) = \max_{j \neq i} \hat{v}_j.$$

Under i winning, every other bidder gets zero object value:

$$\sum_{j \neq i} \hat{v}_j(o^*) = 0.$$

So

$$p_i = \max_{j \neq i} \hat{v}_j.$$

If bidder i loses

Suppose some other bidder wins. Then removing bidder i does not change who gets the object, so others' welfare is the same with or without i:

$$\max_o \sum_{j \neq i} \hat{v}_j(o) = \sum_{j \neq i} \hat{v}_j(o^*).$$

Hence

$$p_i = 0.$$

A losing bidder creates no externality on others.

Conclusion

Only the winner is pivotal for others, and the winner pays the second-highest bid. That is exactly the Vickrey second-price auction.

A small VCG allocation example

Two possible outcomes

Suppose a department can assign one scarce seminar slot to project A or project B.

$$O = \{A, B\}.$$

Reported valuations are:

	A	B
1	8	2
2	4	7
3	0	8

Efficient choice

$$\sum_i \hat{v}_i(A) = 12, \qquad \sum_i \hat{v}_i(B) = 17.$$

VCG chooses B.

Payment in the VCG example

Recall the valuations

	A	B
1	8	2
2	4	7
3	0	8

$$\sum_i \hat{v}_i(A) = 12, \qquad \sum_i \hat{v}_i(B) = 17.$$

VCG chooses B.

Compute all three payments

agent i	$\max_o \sum_{j \neq i} \hat{v}_j(o)$	$\sum_{j \neq i} \hat{v}_j(B)$	p_i
1	$\max\{4,\ 15\} = 15$	15	0
2	$\max\{8,\ 10\} = 10$	10	0
3	$\max\{12,\ 9\} = 12$	9	3

Interpretation

Only agent 3 is pivotal: removing agent 3 changes the efficient outcome from B back to A, so agent 3 pays the welfare loss imposed on the others.

What VCG achieves and what it does not

What VCG achieves

- dominant-strategy truth-telling,
- efficient allocation based on reported values,
- a clean externality interpretation of payments,
- direct generalization of the second-price auction.

What VCG does not automatically solve

- revenue maximization,
- budget balance in public-good problems,
- collusion, false-name bidding, and weak cost recovery when agents are rarely pivotal,
- practical elicitation of complex valuation functions.

Where this returns

Public-good VCG is better treated later with externalities and public goods. Here the main purpose is to see the externality-pricing logic.

Why false-name bidding can break VCG

One identity

	A	B	p_i
1	8	2	0
2	4	7	0
3	0	8	3

Without bidder 3, totals are $A = 12$ and $B = 9$, so bidder 3 is pivotal.

Split into two identities

	A	B	p_i
1	8	2	0
2	4	7	0
$3a$	0	4	0
$3b$	0	4	0

Without $3a$ or $3b$ alone, totals are still $A = 12$ and $B = 13$, so neither fake identity is pivotal.

Broader lesson

VCG needs agents to be pivotal. With split identities or very large markets, each identity may impose almost no externality, so payments may fail to cover costs.

End of Lecture 6

- Auctions provide a clean entry point to mechanism design because they have private values, allocation rules, and payment rules.
- In a second-price auction, truthful bidding is weakly dominant because the bidder's own bid determines whether they cross the threshold, not the price conditional on winning.
- The revelation principle lets us focus on direct mechanisms with truthful reporting.
- IC and IR are the central constraints: IC controls misreporting, and IR controls participation.
- Revenue equivalence follows because IC pins down information rents once the allocation rule and boundary utility are fixed.
- VCG generalizes the second-price auction by making each agent pay the externality imposed on others.

Bridge to Lecture 7

What we did today

We used auctions to introduce the designer's language:

$$\text{allocation rule} + \text{payment rule} + \text{IC} + \text{IR}.$$

What comes next

Lecture 7 moves from auctions to screening:

- one principal designs a menu,
- agents privately know their types,
- different types self-select into different contracts,
- IC and IR determine which distortions and information rents are necessary.

Unifying theme

Auctions and screening look different on the surface, but both are mechanism design problems governed by the same IC and IR logic.

References

Selected references for Lecture 6

- Mas-Colell, A., Whinston, M. D., and Green, J. R. (1995). *Microeconomic Theory*. Oxford University Press, esp. Ch. 23.
- Jehle, G. A., and Reny, P. J. (2011). *Advanced Microeconomic Theory*. 3rd ed. Financial Times/Prentice Hall, esp. mechanism design and auction chapters.
- Vickrey, W. (1961). "Counterspeculation, Auctions, and Competitive Sealed Tenders." *Journal of Finance*.
- Clarke, E. H. (1971). "Multipart Pricing of Public Goods." *Public Choice*.
- Groves, T. (1973). "Incentives in Teams." *Econometrica*.
- Myerson, R. (1981). "Optimal Auction Design." *Mathematics of Operations Research*.