Auction Theory Foundations

Game Theory · Module 6 — Mechanism Design

Node 6.3

Position in Knowledge Graph

System: Game Theory **Structural Domain:** Mechanism Design **Node:** 6.3

Auctions are the canonical applied domain of mechanism design, the place where the abstract machinery of social choice functions, incentive compatibility, and the revelation principle meets a concrete economic institution with thousands of years of history and trillions of dollars of annual throughput. This node develops the **independent private values (IPV) model**, the four classical single-item auction formats (first-price sealed-bid, second-price sealed-bid / Vickrey, English ascending, Dutch descending), and the **Revenue Equivalence Theorem** (Myerson 1981, Riley–Samuelson 1981) — the remarkable result that all standard auctions yield the same expected revenue under IPV with risk-neutral bidders. The node establishes the Myerson monotonicity and envelope conditions that characterize incentive-compatible auction mechanisms, forming the bridge from the abstract revelation principle (6.2) to the concrete revenue-optimization program (6.4).

Formal Definition

Independent Private Values (IPV) model. A single indivisible good is sold to one of n bidders. Each bidder i has a private **valuation** $v_i \in [\underline{v}, \bar{v}]$ drawn independently from a distribution F_i with density $f_i > 0$. Bidder i's utility from winning at price p is $v_i - p$; utility from losing is 0. Bidders are risk-neutral.

An **auction mechanism** is a direct mechanism (x, t) where: - $x_i(v) \in [0, 1]$ is bidder i's probability of winning given the reported profile $v = (v_1, \ldots, v_n)$, with $\sum_i x_i(v) \leq 1$. - $t_i(v) \in \mathbb{R}$ is the transfer (payment) from bidder i to the seller.

Define the **interim allocation** and **interim payment**:

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

Bidder i's **interim expected utility** from reporting $\hat{v}_i$ when true value is v_i:

$$U_i(\hat{v}_i; v_i) = v_i \cdot Q_i(\hat{v}_i) - T_i(\hat{v}_i).$$

Bayesian IC requires $U_i(v_i; v_i) \geq U_i(\hat{v}_i; v_i)$ for all $v_i, \hat{v}_i$. **Individual rationality (IR)** requires $U_i(v_i; v_i) \geq 0$ for all v_i.

The four classical auctions:

1. **First-price sealed-bid (FPSB).** Each bidder submits a sealed bid b_i. Highest bidder wins, pays their bid. $x_i(b) = \mathbf{1}\{b_i > \max_{j \neq i} b_j\}$, $t_i(b) = b_i \cdot x_i(b)$.
2. **Second-price sealed-bid (Vickrey).** Each bidder submits a sealed bid. Highest bidder wins, pays the second-highest bid. $t_i(b) = \max_{j \neq i} b_j \cdot x_i(b)$. DSIC: bidding $b_i = v_i$ is dominant.

3. **English ascending auction.** Price rises continuously; bidders drop out when price exceeds their value. Last remaining bidder wins at the price where the second-to-last dropped out. Strategically equivalent to second-price under IPV.
4. **Dutch descending auction.** Price falls from a high starting point; the first bidder to "stop" the clock wins and pays the current price. Strategically equivalent to first-price under IPV.

Intuition

The fundamental tension in auctions is between **efficiency** (allocate the good to the bidder who values it most) and **revenue** (extract as much surplus as possible from the winner). Second-price and English auctions are efficient and DSIC under IPV — truth-telling is dominant because your bid determines *whether* you win but not *what you pay*. First-price and Dutch auctions are efficient in equilibrium but not DSIC — bidders shade their bids below true value because their bid determines both allocation and payment. The shading is a strategic response to the price-determination rule.

The Revenue Equivalence Theorem says that despite these different strategic structures, all four auctions yield the same expected revenue under the IPV model with symmetric risk-neutral bidders. The intuition: in any IC mechanism, once you fix the allocation rule (highest-value bidder wins), the envelope condition pins down the payment rule up to a constant. If the lowest type pays zero (the standard IR condition), the constant is fixed, and revenue is identical across all mechanisms with the same efficient allocation. The differences between auctions show up only when the assumptions break: risk aversion, asymmetry, interdependent values, or collusion.

This is a remarkable result. It means that the choice between first-price and second-price is not a revenue question under IPV — it is a question of robustness, computational simplicity, susceptibility to collusion, and other extra-mechanism-design considerations. The English auction is preferred in practice for reasons outside the model: it is transparent, allows bidders to learn about opponents' valuations (irrelevant under IPV but valuable under common values), and is culturally familiar.

Structural Role

6.3 is the core applied node of Module 6. It converts the abstract revelation-principle reduction (6.2) into concrete characterizations of auction mechanisms:

- The **Myerson monotonicity + envelope** conditions (developed here) characterize the entire space of BIC auction mechanisms. This is what 6.4 optimizes over.
- The **Revenue Equivalence Theorem** is the structural baseline against which optimal auction design (6.4) is measured: the optimal auction typically differs from the standard auctions by introducing a **reserve price**, which breaks revenue equivalence by excluding low-value bidders.
- The IPV model is the workhorse of economic theory and the starting point for extensions to common values, affiliated values, multi-item auctions, and combinatorial auctions.

The node also introduces the structural analogy between auctions and poker tournaments (via ICM), which is developed in the Poker Anchor.

Mathematical Development

Myerson's characterization of BIC mechanisms

Theorem (Myerson 1981). A direct mechanism (Q, T) is BIC and IR iff:

1. $Q_i(v_i)$ is weakly increasing in v_i for each i (**monotonicity**).

2. The interim payment satisfies the **envelope formula**:

$$T_i(v_i) = v_i Q_i(v_i) - \int_{\underline{v}}^{v_i} Q_i(s)\,ds - U_i(\underline{v})$$

where $U_i(\underline{v}) \geq 0$ is the information rent of the lowest type (**IR pins this to 0** in the optimal mechanism).

Proof sketch. BIC requires $U_i(v_i) \equiv v_i Q_i(v_i) - T_i(v_i)$ to satisfy $U_i(v_i) \geq U_i(\hat{v}_i; v_i) = v_i Q_i(\hat{v}_i) - T_i(\hat{v}_i)$ for all $\hat{v}_i$. By revealed preference, for any $v_i > v_i'$:

$$U_i(v_i) \geq v_i Q_i(v_i') - T_i(v_i') = U_i(v_i') + (v_i - v_i')Q_i(v_i'),$$

and symmetrically $U_i(v_i') \geq U_i(v_i) - (v_i - v_i')Q_i(v_i)$. Combining:

$$(v_i - v_i')Q_i(v_i) \;\geq\; U_i(v_i) - U_i(v_i') \;\geq\; (v_i - v_i')Q_i(v_i').$$

Since $v_i > v_i'$, we get $Q_i(v_i) \geq Q_i(v_i')$: monotonicity. Taking $v_i - v_i' \to 0$, we get $U_i'(v_i) = Q_i(v_i)$ almost everywhere (envelope theorem). Integrating:

$$U_i(v_i) = U_i(\underline{v}) + \int_{\underline{v}}^{v_i} Q_i(s)\,ds.$$

Substituting $U_i(v_i) = v_i Q_i(v_i) - T_i(v_i)$ yields the payment formula. ■

Revenue Equivalence Theorem

Theorem (Myerson 1981, Riley–Samuelson 1981). Any two BIC mechanisms that allocate the good to the same bidder for every type profile and give the lowest type the same expected utility yield the same expected revenue.

Proof. From the envelope formula, expected revenue from bidder i is

$$\mathbb{E}[T_i(v_i)] = \mathbb{E}[v_i Q_i(v_i)] - \int_{\underline{v}}^{\bar{v}} \int_{\underline{v}}^{v_i} Q_i(s)\,ds\, f_i(v_i)\,dv_i - U_i(\underline{v}).$$

Integration by parts on the double integral yields

$$\mathbb{E}[T_i(v_i)] = \mathbb{E}\Big[\Big(v_i - \frac{1 - F_i(v_i)}{f_i(v_i)}\Big)Q_i(v_i)\Big] - U_i(\underline{v}).$$

If Q_i is the same across two mechanisms and $U_i(\underline{v}) = 0$ in both, then expected revenue is the same. ■

Equilibrium bidding in the first-price auction (symmetric IPV)

With n symmetric bidders, $v_i \sim F$ iid, the symmetric BNE bid function $\beta(v)$ satisfies the first-order condition from maximizing expected utility $(v - b)F(\beta^{-1}(b))^{n-1}$:

$$\beta(v) = v - \frac{\int_{\underline{v}}^{v} F(s)^{n-1} ds}{F(v)^{n-1}}.$$

For $F = U[0,1]$, this simplifies to $\beta(v) = \frac{n-1}{n} v$ — each bidder shades by v/n.

Verification of revenue equivalence. Expected revenue in the first-price auction is $\mathbb{E}[\beta(v_{(1)})]$ where $v_{(1)}$ is the highest order statistic. In the second-price auction, expected revenue is $\mathbb{E}[v_{(2)}]$. For $U[0,1]$:

$$\mathbb{E}[\beta(v_{(1)})] = \frac{n-1}{n} \cdot \frac{n}{n+1} = \frac{n-1}{n+1}, \qquad \mathbb{E}[v_{(2)}] = \frac{n-1}{n+1}.$$

Equal, confirming the theorem.

Strategic equivalences

Under IPV with risk-neutral bidders: - First-price sealed-bid $\equiv$ Dutch (same strategy space and payoffs: in Dutch, stopping at price b is strategically identical to bidding b in FPSB). - Second-price sealed-bid $\equiv$ English ascending (last bidder wins at the dropout price of the second-to-last, which is the second-highest value under dominant-strategy bidding).

Under common values or risk aversion, these equivalences break.

All-pay auctions and war of attrition

An **all-pay auction** requires every bidder to pay their bid, not just the winner. The symmetric BNE bid function for n bidders with $v_i \sim U[0,1]$ is $\beta(v) = \frac{n-1}{n}v^n$. Expected revenue equals $\frac{n-1}{n+1}$, confirming revenue equivalence with standard auctions (all bidders pay, but they bid less).

The all-pay auction is the formal model of contests, lobbying, and R&D races: effort is sunk regardless of who wins. It also models poker pots pre-flop: every player who enters pays the blind/ante, regardless of whether they win the pot. The structural mapping is direct.

Common values and the winner's curse

When bidders' valuations are correlated — the **common-value** or **affiliated-value** model — the winner tends to be the bidder who overestimates the item's value the most. This is the **winner's curse**: winning is bad news about your estimate. Rational bidders shade their bids further to compensate, and the strategic analysis is substantially more complex. Milgrom and Weber (1982) showed that in the affiliated-values model, the English auction generates strictly higher expected revenue than the second-price sealed-bid auction, which generates strictly higher revenue than the first-price auction. Revenue equivalence fails because the English auction reveals information during the auction process, reducing the winner's curse and allowing more aggressive bidding.

Worked Example

Revenue equivalence with three uniform bidders

$n = 3$, $v_i \sim U[0,1]$ iid. Expected revenue:

$$R = \frac{n-1}{n+1} = \frac{2}{4} = 0.5.$$

First-price. Bid function $\beta(v) = \frac{2}{3}v$. CDF of highest value $v_{(1)}$ is $F_{(1)}(v) = v^3$, density $3v^2$. Expected revenue:

$$\mathbb{E}\left[\tfrac{2}{3}v_{(1)}\right] = \tfrac{2}{3}\int_0^1 v \cdot 3v^2\,dv = \tfrac{2}{3}\cdot\tfrac{3}{4} = \tfrac{1}{2}.$$

Second-price. Expected value of second-highest order statistic with $n = 3$, $U[0,1]$:

$$\mathbb{E}[v_{(2)}] = \frac{2}{4} = \frac{1}{2}.$$

Revenue equivalence confirmed.

Where revenue equivalence fails: risk-averse bidders

If bidders have CARA utility $u(w) = 1 - e^{-\alpha w}$, $\alpha > 0$, they overbid in the first-price auction (less shading, because losing is more painful in utility terms). Expected revenue in FPSB exceeds second-price. In second-price, the dominant strategy is still $b_i = v_i$ regardless of risk attitudes (payment does not depend on own bid), so second-price revenue is unchanged. First-price revenue rises with risk aversion.

Poker Anchor

Concept mapping

Formal object	Poker instantiation
Single item for sale	The tournament prize pool
Bidders with private valuations	Players with private skill levels and bankrolls
Buy-in / entry fee	The "bid" to enter the auction for the prize pool
Allocation rule $x_i(v)$	Probability of finishing in each prize position
Revenue equivalence	Under certain conditions, different tournament structures yield the same total prize-pool extraction
Reserve price	Minimum buy-in threshold that excludes some potential entrants

Concrete situation: ICM (Independent Chip Model) and auction-like tournament dynamics

In a poker tournament, chips do not convert linearly to money. The **Independent Chip Model** computes each player's dollar equity as a function of chip stacks and the remaining prize pool:

$$\text{ICM}_i(c_1, \ldots, c_n; \text{payouts}) = \sum_{k=1}^{n} \text{Prize}_k \cdot P(\text{player } i \text{ finishes } k\text{th}),$$

where the finishing probabilities are computed from chip proportions. The key structural feature: **ICM equity is concave in chip count**. Doubling your chips does not double your equity. Losing half your chips costs more equity than gaining the same amount adds.

This concavity creates **auction-like dynamics at the bubble**. At the bubble (one elimination from the money), every chip risked has asymmetric equity impact: losing a pot costs more in ICM dollars than winning the same pot gains. The bubble is analogous to a **risk-averse first-price auction**: each player's "bid" (chips committed to a pot) carries an ICM cost that is convex in the amount risked, just as a risk-averse bidder's utility is concave in the payment. The result is the same: **players overbid conservatively**, tightening their ranges to avoid elimination.

The ICM concavity is why tournament strategy deviates from cash-game strategy. In a cash game, $v_i - p$ is linear in p (risk-neutral). In a tournament, the utility of winning chips c is approximately $U(c) \approx c^{\alpha}$ with $0 < \alpha < 1$ (concave). This maps directly to the risk-aversion extension of auction theory: risk-averse bidders shade less in first-price auctions (conserve rather than gamble), and tournament players near the bubble tighten their ranges for the same structural reason.

Revenue equivalence breaks at the bubble. In cash games, different bet-sizing strategies that implement the same allocation (same hands win) are revenue-equivalent. In tournaments, the non-linear chip-to-dollar mapping means that bet sizing matters beyond allocation: larger pots create larger ICM distortions, so two betting strategies that implement the same "winner" profile can have different ICM equity consequences. This is the tournament analogue of revenue equivalence failing under risk aversion.

Concrete situation: satellite tournaments as nested auctions

A satellite tournament (winner gets a seat in a larger tournament) is literally a multi-unit auction: k seats are "sold" to $n > k$ bidders via a tournament mechanism. The buy-in is the bid, the item is the seat, and the allocation rule is determined by finishing position. In a winner-take-all satellite with one seat, the tournament is a first-price common-value auction: each player's "value" for the seat includes both their skill edge and their bankroll constraints, and the "payment" is the risk of busting.

What this understanding buys you

ICM awareness is risk-aversion awareness. When you hear "play tighter at the bubble," the formal content is: ICM concavity makes you effectively risk-averse, so you shade your bids (tighten your ranges) just as a risk-averse first-price bidder shades their bid. Understanding this formally lets you quantify *how much* tighter: you can compute the ICM-adjusted EV of a shove and compare it to the chip-EV, and the gap is the "risk premium" induced by the tournament structure.

Tournament structures are mechanism designs. The tournament operator choosing between a flat payout structure (top 15% paid) and a top-heavy structure (winner-take-most) is making a mechanism-design decision. Flat payouts increase ICM concavity at the bubble (more players affected), producing tighter play. Top-heavy payouts reduce bubble effects but increase variance. The revenue equivalence theorem tells you these different structures produce the same *expected* prize-pool allocation under strong assumptions — and the ICM deviations from those assumptions are where the design choices actually bite.

Bid shading = range tightening. The formal mapping is: in a first-price auction, the equilibrium bid $\beta(v) < v$ represents shading below true value. In poker, the analogous shading is folding hands that would be +EV in chips but are -EV in ICM dollars. The amount of shading is determined by the same structural variables: the number of opponents, the distribution of their types (stacks), and the curvature of the value function (ICM concavity).

Cross-Links

- **6.1 Social choice and IC** — the IC definitions that constrain auction mechanisms.
- **6.2 Revelation principle** — reduces auction design to the study of direct BIC mechanisms.
- **6.4 Optimal auction design** — Myerson's program: optimize revenue over the space characterized here.
- **4.1 Bayesian games** — the IPV model is a Bayesian game with independent types.
- **1.5 Zero-sum and minimax** — cash poker is zero-sum; tournaments are not, due to ICM.
- **5.3 Trigger strategies** — repeated auction settings where collusion is sustained by trigger strategies.

Structural Compression

Invariant: In the IPV model with risk-neutral bidders, any two BIC mechanisms with the same allocation rule and the same lowest-type utility yield the same expected revenue. The Myerson monotonicity + envelope conditions characterize the full space of BIC auction mechanisms: once you specify a monotone allocation rule and pin down the lowest type's rent, the payment rule is determined.

Mechanism: Envelope theorem applied to BIC: the utility function of a truthful agent is convex in type, its derivative is the interim allocation probability, and integrating recovers the payment rule up to a constant. Revenue equivalence follows because the integral depends only on the allocation rule, not on the payment format.

Cross-System Echo: Revenue equivalence is the auction-theoretic analogue of **path independence of conservative forces** in physics. The "revenue" (work done by the mechanism) depends only on the "allocation" (final state, i.e., who gets the good) and not on the "path" (specific auction format). Just as a conservative force field has a potential function that determines work, a BIC mechanism has an envelope function that determines payments. The analogy extends: when the field is *not* conservative (risk aversion, interdependent values, collusion), path dependence returns, and the choice of auction format matters.

Exercises

1. Derive the symmetric BNE bid function for the first-price sealed-bid auction with n bidders and values drawn from $F(v) = v^2$ on $[0, 1]$. Compute expected revenue and verify revenue equivalence with the second-price auction.

2. Show that in the second-price auction, bidding $b_i = v_i$ is a weakly dominant strategy regardless of risk attitudes, number of bidders, or the distribution F. Why does the proof not extend to the first-price auction?

3. Consider a two-bidder first-price auction with $v_1 \sim U[0, 1]$ and $v_2 \sim U[0, 2]$ (asymmetric). Does a symmetric equilibrium exist? Find the equilibrium bid functions (they will differ across bidders) and compute expected revenue. Does revenue equivalence hold?

4. Compute the ICM equity of a player with 5000 chips in a 9-player tournament with equal starting stacks of 5000, top 3 paid (50%, 30%, 20%). Then compute the ICM equity after the player doubles up to 10000 while one opponent busts. Verify that the equity gain from doubling is less than twice the initial equity (concavity).

5. In a Dutch auction with three risk-neutral IPV bidders and $v_i \sim U[0, 1]$, at what price should a bidder with value v stop the clock? Show that this is strategically identical to the first-price sealed-bid equilibrium bid.

6. Construct an example where revenue equivalence fails due to common values (i.e., bidders' valuations are correlated). Compare expected revenue in first-price and second-price auctions in this setting. Which is higher?

7. **Argue, structurally,** why ICM concavity in tournament poker is the exact analogue of risk aversion in auction theory, and why this mapping implies that revenue equivalence (between different bet-sizing strategies with the same allocation) fails at the tournament bubble. Identify the structural feature of the chip-to-dollar conversion that generates the concavity.

Game Theory · Module 6 — Mechanism Design · Node 6.3
Concepts: mechanism-design, incentive-compatibility, bayesian-nash-equilibrium, expectation-value
Prerequisites: 6.1, 6.2
Enables: 6.4

hbar.university — own.your.cognition