Optimal Auction Design

Game Theory · Module 6 — Mechanism Design

Node 6.4

Position in Knowledge Graph

System: Game Theory **Structural Domain:** Mechanism Design **Node:** 6.4

This node develops Myerson's (1981) solution to the problem of optimal auction design: given a seller with one item and n bidders with independently drawn private values, what mechanism maximizes expected revenue? The answer is strikingly clean. Revenue maximization reduces to allocating the item to the bidder with the highest **virtual valuation** — a transformed value that accounts for the information rent the seller must surrender to higher types — subject to a **reserve price** that excludes bidders whose virtual valuation is negative. The optimal auction is a second-price auction with a reserve price (in the symmetric regular case), or more generally a modified Vickrey auction with bidder-specific entry thresholds. This is the crown jewel of single-item mechanism design: a closed-form solution to a seemingly intractable optimization over the space of all incentive-compatible mechanisms.

Formal Definition

Environment. One seller, one indivisible item, n bidders. Bidder i has type $v_i \sim F_i$ on $[\underline{v}_i, \bar{v}_i]$ with density $f_i > 0$, independently across i. Quasi-linear utilities: winner gets $v_i - t_i$, loser gets $-t_i$ (typically $t_i = 0$ for losers). Seller values the item at $v_0 \geq 0$.

By the revelation principle (6.2), restrict to direct BIC mechanisms (Q, T) where $Q_i(v_i)$ is the interim win probability and $T_i(v_i)$ is the interim payment.

The **virtual valuation** of bidder i with type v_i is:

$$\psi_i(v_i) \;=\; v_i - \frac{1 - F_i(v_i)}{f_i(v_i)}.$$

A distribution F_i is **regular** if $\psi_i(v_i)$ is strictly increasing in v_i.

Myerson's optimal auction. Allocate the item to the bidder i^* with the highest virtual valuation, provided $\psi_{i^*}(v_{i^*}) \geq v_0$. If all virtual valuations are below v_0, the seller retains the item. Formally:

$$x_i(v) \;=\; \begin{cases} 1 & \text{if } \psi_i(v_i) > \max\big(v_0, \max_{j \neq i} \psi_j(v_j)\big), \\ 0 & \text{otherwise.} \end{cases}$$

Payments are determined by the envelope formula (6.3):

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

The **optimal reserve price** for bidder i is $r_i^* = \psi_i^{-1}(v_0)$, the type at which the virtual valuation equals the seller's reservation value.

Intuition

The seller faces a tradeoff: extracting more revenue from high-value bidders (by raising the price) versus losing sales when bidders turn out to have low values. The virtual valuation $\psi_i(v_i) = v_i - \frac{1-F_i(v_i)}{f_i(v_i)}$ captures this tradeoff. The term $\frac{1-F_i(v_i)}{f_i(v_i)}$ is the **information rent** — the surplus that a bidder of type v_i must be given to prevent them from mimicking a lower type. Each dollar of valuation above the reserve "costs" the seller an information-rent dollar that must be paid to all higher types. Virtual valuation nets out this cost.

The seller's problem is thus transformed: instead of maximizing revenue directly (which involves the complex IC constraints), maximize the expected virtual surplus:

$$\max_{(x,t)} \ \mathbb{E}\Big[\sum_i \psi_i(v_i) \cdot x_i(v)\Big].$$

This is a pointwise maximization problem, solved by allocating to the highest $\psi_i(v_i)$ — which is why the optimal auction allocates by virtual valuation, not by actual valuation.

The reserve price emerges because the seller should prefer to keep the item rather than sell it to a bidder whose virtual valuation is below v_0. Even if a bidder values the item above the seller's physical value, the information rent makes selling unprofitable if the virtual valuation is too low. The optimal reserve **excludes some willing buyers** — a deliberate sacrifice of efficiency for revenue.

In the symmetric regular case ($F_i = F$ for all i, ψ increasing), the optimal auction reduces to a second-price auction with reserve $r^* = \psi^{-1}(v_0)$. All the mechanism-design machinery collapses to a single number: the reserve price.

Structural Role

6.4 is the optimization node of Module 6 — it takes the feasibility conditions from 6.3 (monotonicity, envelope, IR) and maximizes the seller's objective over them. The structural sequence is:

- 6.1 defines IC and the social-choice framework.
- 6.2 reduces to direct mechanisms.
- 6.3 characterizes the feasible set (monotone allocation + envelope payment).
- **6.4 solves the optimization** over this feasible set.

The optimal-auction result has deep connections:

- **Revenue equivalence (6.3)** is the *unconstrained* case: without a reserve, all efficient auctions yield the same revenue. The optimal reserve breaks efficiency (not all bidders participate) and breaks revenue equivalence (the optimal auction outperforms standard auctions).
- **Implementation theory (6.5)** considers whether the optimal mechanism can be implemented in stronger solution concepts (Nash, SPE).
- **Matching (6.6)** is the multi-sided analogue: instead of one seller optimizing, two sides of a market are matched, and the design question is stability rather than revenue.

Mathematical Development

Deriving the revenue formula

From 6.3's envelope formula and the integration-by-parts identity:

$$\text{Revenue} \ = \ \sum_i \mathbb{E}[T_i(v_i)] \ = \ \sum_i \mathbb{E}[\psi_i(v_i) \cdot Q_i(v_i)] - \sum_i U_i(\underline{v}_i).$$

Setting $U_i(\underline{v}_i) = 0$ (optimal for the seller: give no rent to the lowest type):

$$\text{Revenue} \;=\; \mathbb{E}\Big[\sum_i \psi_i(v_i) \cdot x_i(v)\Big].$$

This is the **Myerson revenue formula**. The seller maximizes expected virtual surplus, pointwise in v. If ψ_i is increasing (regular), the allocation $x_i^*(v) = \mathbf{1}\{\psi_i(v_i) > \max(v_0, \max_{j \neq i} \psi_j(v_j))\}$ is monotone in v_i (higher v_i implies higher ψ_i implies higher win probability), so the IC constraint is satisfied automatically. The optimal mechanism is well-defined.

Computing the optimal reserve (symmetric case)

Symmetric bidders: $F_i = F$, $f_i = f$, $\psi_i = \psi$ for all i. Seller value $v_0 = 0$. Optimal reserve r^* satisfies $\psi(r^*) = 0$:

$$r^* - \frac{1 - F(r^*)}{f(r^*)} = 0 \quad \Longrightarrow \quad r^* = \frac{1 - F(r^*)}{f(r^*)}.$$

For $F = U[0,1]$: $\psi(v) = 2v - 1$, so $r^* = 1/2$. The optimal auction is a second-price auction with reserve 1/2.

Revenue gain from the reserve. The reserve price is the single most important design variable. It transforms an efficient mechanism (second-price auction without reserve, which gives the item to the highest bidder at the second-highest price) into a revenue-maximizing mechanism by excluding low-value bidders. The excluded bidders would have paid something, so the reserve sacrifices efficiency for revenue — but the revenue gain exceeds the efficiency loss from the seller's perspective.

Revenue comparison. Without reserve ($n = 2, U[0,1]$): $R_0 = \mathbb{E}[v_{(2)}] = 1/3$. With optimal reserve $r^* = 1/2$:

$$R^* = \int_0^1 \int_0^1 T_1(v_1, v_2) + T_2(v_1, v_2)\, dv_1\, dv_2.$$

Under the second-price + reserve mechanism: if both bid above 1/2, winner pays max(second bid, 1/2). If only one bids above 1/2, they win and pay 1/2. If neither bids above 1/2, no sale. Computing:

$$R^* = 2 \int_{1/2}^1 \Big[\int_0^{v_1} \max(v_2, 1/2)\, dv_2\Big] dv_1 = \frac{5}{12} \approx 0.417.$$

Improvement: from $1/3 \approx 0.333$ to $5/12 \approx 0.417$, a 25% revenue increase from a single design parameter.

Irregular distributions and ironing

If ψ_i is not monotone, the pointwise maximization may produce a non-monotone allocation (violating IC). Myerson's solution is **ironing**: replace ψ_i with its monotone upper envelope $\bar{\psi}_i$, obtained by "pooling" types in regions where ψ_i is decreasing. Formally, $\bar{\psi}_i$ is the convex conjugate of the convex hull of the function $v \mapsto \int_{\underline{v}}^{v} \psi_i(s) f_i(s)\, ds$. The ironed virtual valuation is monotone, and the optimal auction allocates by ironed virtual valuation.

The asymmetric case

With asymmetric bidders ($F_1 \neq F_2$), the optimal auction is *not* a standard auction format. It allocates to the highest virtual valuation, which may not be the highest actual valuation — the auction can be biased toward the "weaker" bidder (the one with higher information rent). For example, with $v_1 \sim U[0,1]$ and $v_2 \sim U[0,2]$, the virtual valuations are $\psi_1(v) = 2v - 1$ and $\psi_2(v) = 2v - 2$. Bidder 2 is disadvantaged despite having potentially higher values, because their information rent is larger.

This implies that **the optimal auction discriminates against stronger bidders** — a counterintuitive result that has deep implications for policy (affirmative action in procurement, handicaps in contests).

Connection to monopoly pricing

Myerson's optimal auction generalizes classical monopoly pricing. A monopolist facing a single buyer with value $v \sim F$ sets price p^* to maximize $p(1-F(p))$. The first-order condition is $1 - F(p) - pf(p) = 0$, i.e., $p - \frac{1-F(p)}{f(p)} = 0$, i.e., $\psi(p) = 0$. So the monopoly price equals the optimal reserve. With $n > 1$ bidders, competition partially substitutes for the reserve: even without a reserve, bidders bid against each other. The optimal reserve adds an additional revenue extraction on top of competition.

For $n = 1$ (monopoly): optimal revenue $= \max_p p(1-F(p))$. For $n \to \infty$: the highest value converges to $\bar{v}$, and the optimal reserve becomes irrelevant (the second-highest value approaches $\bar{v}$, so revenue approaches the full surplus). The reserve matters most when n is small — precisely when competition is weakest and the seller's market power is greatest.

Budget balance and efficiency

The optimal auction is not budget-balanced in the multi-seller extension and not efficient (some trades are blocked by the reserve). The **Myerson–Satterthwaite theorem** (1983) shows that for bilateral trade (one buyer, one seller) with overlapping value distributions, no mechanism is simultaneously BIC, IR for both parties, efficient, and budget-balanced. This is the fundamental impossibility of efficient trade under asymmetric information, and the optimal auction is the revenue-maximizing response to it.

The Myerson–Satterthwaite result is the mechanism-design analogue of the "lemons" problem (Akerlof 1970): when information is asymmetric, some mutually beneficial trades are impossible because the informed party's willingness to trade reveals adverse information. In auctions, this manifests as the reserve price: the seller withholds the good from low-type bidders even though trade would be efficient, because the information rent required to distinguish high types from low types exceeds the gains from trading with low types.

Worked Example

Optimal auction with two uniform bidders

$n = 2$, $v_i \sim U[0, 1]$, seller value $v_0 = 0$.

Virtual valuation: $\psi(v) = 2v - 1$. Regular (increasing).

Optimal reserve: $\psi(r^*) = 0 \Rightarrow r^* = 1/2$.

Optimal mechanism: second-price auction with reserve $r^* = 1/2$.

Expected revenue computation. Three cases: 1. Both $v_1, v_2 \geq 1/2$: probability $1/4$. Revenue $= \mathbb{E}[\max(v_{(2)}, 1/2) \mid v_1, v_2 \geq 1/2]$. Since $v_{(2)}|v_1, v_2 \geq 1/2$ has CDF $G(x) = (2x-1)^2$ on $[1/2, 1]$, $\mathbb{E}[v_{(2)}|v_1, v_2 \geq 1/2] = 2/3$. Revenue from this case: $(1/4)(2/3) = 1/6$. 2. Exactly one $v_i \geq 1/2$: probability $1/2$. Revenue $= 1/2$ (the reserve). Revenue from this case: $(1/2)(1/2) = 1/4$. 3. Both $v_1, v_2 < 1/2$: probability $1/4$. No sale. Revenue: 0.

Total: $1/6 + 1/4 = 5/12 \approx 0.417$.

Comparison. Without reserve: revenue $= 1/3 \approx 0.333$. Gain from optimal reserve: $5/12 - 1/3 = 1/12 \approx 0.083$, a 25% improvement.

The cost of efficiency

The optimal auction is not efficient: when both bidders have values in $(0, 1/2)$, neither wins, but both value the item above 0. The seller withholds the good for revenue. This is the fundamental tradeoff of mechanism design: efficiency (total surplus maximization) and revenue (seller surplus maximization) are generically in tension. VCG maximizes efficiency; Myerson maximizes revenue. Only in the case of no reserve ($v_0 = 0$ and $r^* = 0$) do they coincide, and that requires $\psi(\underline{v}) \geq 0$, which fails whenever $\underline{v} < \frac{1-F(\underline{v})}{f(\underline{v})}$.

Poker Anchor

Concept mapping

Formal object	Poker instantiation
Virtual valuation $\psi_i(v_i)$	ICM-adjusted expected value of a hand (chip EV corrected for tournament equity)
Reserve price r^*	The minimum hand strength at which entering a pot is +EV in ICM terms
Information rent	The equity you "leak" to opponents by having to play your strong hands consistently
Optimal reserve excluding low types	Folding hands that are +chipEV but -ICMEV
Seller's revenue	Tournament operator's rake / the house edge in the mechanism

Concrete situation: ICM-adjusted push/fold as optimal-auction-like behavior

At the tournament bubble with 4 players remaining and 3 spots paid, consider the short stack (10 BB) on the button deciding whether to shove all-in. In chip EV, any hand with positive expectation against the calling ranges should be shoved. But the ICM adjustment acts exactly like a virtual-valuation transformation.

Let the short stack's hand strength be v (equity against a random hand). The chip-EV of shoving is proportional to v. But the ICM-EV is proportional to $\psi(v) = v - \rho(v)$, where $\rho(v)$ is the "ICM tax" — the non-linear penalty for risking elimination. The optimal shoving threshold r^* satisfies $\psi(r^*) = 0$: you shove only when the ICM-adjusted value is positive.

This is *exactly* Myerson's optimal reserve. The tournament structure (payouts, stack sizes) determines the "distribution" and "information rent"; the player computes the virtual valuation and shoves only above the threshold. The tournament operator, by choosing the payout structure, is effectively choosing the reserve price of the mechanism. A top-heavy payout structure lowers the ICM tax (less to lose by busting relative to what you gain by winning), which lowers r^* and produces more aggression. A flat payout structure raises the ICM tax and raises r^*, producing tighter play.

Concrete situation: the "Myerson reserve" in high-stakes cash

In a high-stakes cash game, the effective "reserve price" is the minimum pot size at which it is worth committing significant equity. With deep stacks (300BB+), players can construct complex multistreet lines, and the "reserve" drops (even marginal edges are worth pursuing because the implied odds are large). With short stacks (30BB), the reserve rises (you need a stronger hand to commit because the risk-reward ratio is fixed). This is the cash-game analogue of the reserve-price optimization, with stack depth playing the role of the seller's outside option.

What this understanding buys you

The optimal reserve is the central strategic parameter. In both Myerson's auction theory and tournament poker, the single most important design/strategy variable is the threshold below which you do not participate. Getting this threshold right — whether it is a reserve price in an auction, a shoving threshold in a tournament, or a minimum hand strength for 3-betting in cash — is more valuable than any amount of post-flop sophistication. This is Myerson's deep insight: the entire revenue-optimization problem reduces to choosing one number (the reserve), and everything else follows from the envelope formula.

ICM is virtual valuation. When you run an ICM calculator and it tells you a hand is "too weak to shove," it is computing the virtual valuation and finding it negative. The structural content is identical: the concavity of the ICM function creates an information-rent wedge between chip value and dollar value, and the optimal play accounts for this wedge exactly as Myerson's optimal auction accounts for the virtual-valuation wedge.

Payout structure is mechanism design. The tournament operator choosing the payout curve is Myerson's seller choosing the reserve. Flatter payouts extract more "revenue" (more rake-equivalent conservatism from players) at the cost of less action. Top-heavy payouts produce more action (lower effective reserve) at the cost of less revenue from the conservatism channel. Understanding this tradeoff is understanding optimal auction design.

The "optimal bluff threshold" is an optimal reserve. In river betting, the bettor chooses which hands to bet (for value and as bluffs) and which to check. The threshold between "bluff" and "give up" is structurally identical to the optimal reserve: hands below the threshold have negative virtual value (the information rent of including them in the betting range exceeds the pot equity they capture), and the optimal threshold sets virtual value to zero. A solver computing the optimal bluffing frequency on the river is solving Myerson's optimal auction problem in miniature — the bettor is the "seller" of the pot, the caller is the "buyer," and the bluff threshold is the reserve price.

Asymmetric optimal play is biased toward the underdog. In the asymmetric case, the optimal auction discriminates against the "strong" bidder. The poker analogue: when you are the favorite in a hand (you have the stronger range), the optimal play of your opponent is biased toward bluffing and semi-bluffing more aggressively — they must play as if their virtual valuation were shifted upward, just as the weaker bidder in an asymmetric auction is effectively given a handicap by the optimal mechanism.

Cross-Links

- **6.3 Auction theory foundations** — the IPV model, revenue equivalence, and Myerson's characterization that this node optimizes over.
- **6.2 Revelation principle** — the reduction to direct mechanisms that makes the optimization tractable.
- **6.1 Social choice and IC** — the IC and IR constraints.
- **6.5 Implementation theory** — can the optimal auction be implemented in Nash or SPE?
- **4.2 Bayesian updating** — bidders update beliefs about opponents' types.
- **Economics (pricing theory)** — Myerson's optimal auction is the auction analogue of monopoly pricing with price discrimination.

Structural Compression

Invariant: The revenue-maximizing auction allocates to the bidder with the highest virtual valuation above the seller's reserve, where virtual valuation subtracts the information rent from the true valuation. In the symmetric regular case, the optimal mechanism is a second-price auction with a reserve price solving $\psi(r^*) = v_0$.

Mechanism: Virtual valuation transforms the revenue-maximization problem from a constrained optimization over IC mechanisms into an unconstrained pointwise maximization of expected virtual surplus. The transformation absorbs the IC constraints into the definition of ψ, and the optimal allocation is immediate. Payments follow from the envelope formula.

Cross-System Echo: Virtual valuation is the auction-theoretic analogue of **marginal cost** in monopoly pricing. A monopolist facing demand curve $D(p)$ maximizes revenue by setting price where marginal revenue equals marginal cost — and marginal revenue is $p - D(p)/D'(p)$, which has the same "value minus inverse hazard rate" structure as virtual valuation. Myerson's result is the auction generalization of this classical monopoly-pricing insight, extended to multiple competing bidders and private information.

Exercises

1. Compute the optimal reserve price for a single-bidder auction (monopoly pricing) with $v \sim U[0, 1]$. Show that the optimal reserve is $r^* = 1/2$ and the expected revenue is $1/4$. Compare to the no-reserve expected revenue of 0 (the seller gets nothing without competition unless they set a reserve).

2. For n symmetric bidders with $v_i \sim U[0, 1]$, compute the optimal reserve and expected revenue as a function of n. Show that the optimal reserve $r^* = 1/2$ is independent of n, and that the revenue gain from the reserve vanishes as $n \to \infty$.

3. Verify that the virtual valuation for the exponential distribution $F(v) = 1 - e^{-v}$, $v \geq 0$, is $\psi(v) = v - 1$, and that this is regular. Compute the optimal reserve.

4. Construct an example of an irregular distribution where $\psi(v)$ is not monotone. Apply ironing to obtain the ironed virtual valuation $\bar{\psi}(v)$. Describe the qualitative difference between the optimal auction with and without ironing.

5. In a two-bidder asymmetric auction with $v_1 \sim U[0, 1]$ and $v_2 \sim U[0, 2]$, compute the optimal auction. Show that it can award the item to bidder 1 even when $v_2 > v_1$, and explain why this is revenue-maximizing despite being inefficient.

6. A tournament has 10 players, buy-in \$100, and three payout places: \$500, \$300, \$200. Compute the ICM equity of a player at the bubble (4 players remaining) with 40% of the chips. Compare to the "fair share" (25% of prize pool). Explain why ICM equity exceeds fair share and relate this to the virtual-valuation concept.

7. **Argue, structurally,** why the optimal reserve price is always strictly positive when the seller's value is zero and the support of values includes zero — even though this means some willing buyers are excluded and total welfare is reduced. Identify the structural feature of information rents that forces this tradeoff.

Game Theory · Module 6 — Mechanism Design · Node 6.4
Concepts: mechanism-design, constrained-optimization, incentive-compatibility
Prerequisites: 6.2, 6.3
Enables: 6.5

hbar.university — own.your.cognition