ORIE 4154/5154: Pricing and Market Design **Fall 2026**

Lecture 22 — Simple versus Optimal Auctions: Revenue Equivalence and Bulow–Klemperer

Lecturer: Sid Banerjee *Date: November 10, 2026 (draft)*

22.1 Can a new payment rule improve the same allocation?

Lecture 18 used VCG to implement efficient allocation. Lectures 19–20 then characterized the larger family of truthful single-parameter rules, and Lecture 21 chose one of those rules to maximize revenue. We now compare what can be gained by changing payments, by changing allocations, and by attracting additional buyers.

The payment identity first gives a precise revenue-equivalence statement. Its implication for design is useful: once a truthful allocation and its normalization are fixed, there is no extra freedom to raise revenue merely by rewriting the payment rule.

22.2 Revenue equivalence from the payment identity

Theorem 22.1 (DSIC payment and revenue equivalence). *Fix a single-parameter environment with quasilinear agents and type domains starting at zero. If two DSIC mechanisms have the same allocation rule $x_i(b)$ for every agent and report profile, and the same payments $p_i(0, b_{-i})$, then their conditional expected payments coincide at every report profile. Their expected revenues therefore coincide under any joint distribution for which the expectations are finite.*

Proof: For either mechanism, Lecture 20 gives

$$p_i(v, b_{-i}) = v x_i(v, b_{-i}) - \int_0^v x_i(z, b_{-i})\, dz + p_i(0, b_{-i}).$$

The right side is identical in the two mechanisms. Thus their payments agree conditional on reports, averaging any internal randomization. Taking expectations over values and summing over agents preserves equality. ■

This pointwise DSIC statement needs neither independence nor regularity. For deterministic mechanisms the actual payment functions agree. For randomized mechanisms it is the conditional expected payments that agree, as in Lecture 20.

In particular, the efficient single-item allocation with normalized DSIC payments already gives the second-price rule. To improve its revenue while maintaining the same normalization, a seller must change something beyond the payment formula—for example, use a reserve that changes which types receive the item. Truthful mechanisms with different allocation rules need not earn the same revenue.

Remark 22.2 (Interim version). *For independent private values and risk-neutral agents, define* $q_i(v) = \mathbb{E}_{V_{-i}}[x_i(v, V_{-i})]$ *and* $t_i(v) = \mathbb{E}_{V_{-i}}[p_i(v, V_{-i})]$. *The interim identity from Lecture 20 is*

$$t_i(v) = vq_i(v) - \int_0^v q_i(z)\,dz - u_i^{\text{int}}(0). \tag{22.1}$$

Consequently, two BIC mechanisms with the same interim allocations and lowest-type utilities have the same interim expected payments and total expected revenue. BIC only compares expected utility when the other agents are truthful, whereas DSIC holds against every fixed report profile. The same-allocation and normalization hypotheses remain essential in this broader formulation.

22.3 The simple benchmark: Vickrey revenue and market thickness

Consider n iid buyers with continuous CDF F and a no-reserve second-price auction. It allocates to the highest value. Agent i with value v wins exactly when every rival value is lower, up to probability-zero ties, so its interim allocation is

$$q_i(v) = F(v)^{n-1}.$$

For Uniform$[0, 1]$ buyers, this is v^{n-1}, and (22.1) with zero lowest-type utility yields

$$t_i(v) = v^n - \int_0^v z^{n-1}\,dz = \frac{n-1}{n}v^n. \tag{22.2}$$

Integrating over each buyer's value gives

$$\rho_n^{\text{SP}} = n\int_0^1 t_i(v)\,dv = \frac{n-1}{n+1}. \tag{22.3}$$

This also gives zero revenue at $n = 1$, when the competing bid is defined as zero.

The order statistics make the economics visible. Write $V_{(1)} \geq V_{(2)} \geq \cdots \geq V_{(n)}$. The highest value determines efficient welfare; the second-highest value determines the winner's payment. Thus random Vickrey revenue is $R_n^{\text{SP}} = V_{(2)}$ for $n \geq 2$. In the uniform model,

$$\mathbb{E}[V_{(1)}] = \int_0^1 (1 - z^n)\,dz = \frac{n}{n+1}, \qquad \mathbb{E}[V_{(2)}] = \frac{n-1}{n+1}.$$

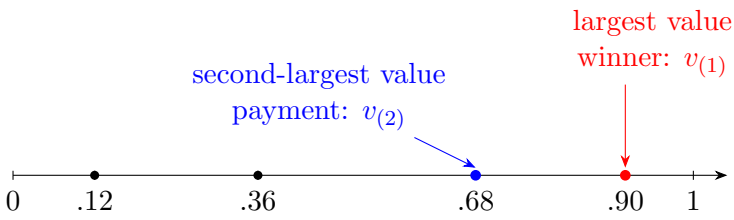

As more iid uniform buyers enter, expected revenue and efficient welfare both approach 1, and their difference is $1/(n+1)$. Competition reduces how much surplus is left with the winning buyer in this model.

22.4 Why reserves are consistent with revenue equivalence

A reserve changes the allocation rule. Under a second-price auction with reserve r, the interim winning probability is

$$q_i^r(v) = \begin{cases} 0, & v < r, \\ F(v)^{n-1}, & v \geq r. \end{cases}$$

Thus a reserve auction and a no-reserve auction do not meet the same-allocation hypothesis. Revenue equivalence places no obstacle to the improvement from 1/3 to 5/12 for two uniform buyers in Lecture 21.

A reserve extracts more when only one buyer has a sufficiently high value, but can lose sales when all values are low. For iid regular values, the virtual-value calculation identifies the cutoff balancing those effects. Setting it requires knowledge of F; the no-reserve second-price rule itself requires no distributional input.

22.5 Bulow–Klemperer: one more bidder

We use the continuous, full-support, finite-mean model of Lecture 21: values have support $[0, \bar{v}]$, possibly unbounded, with $F(0) = 0$ and a positive continuous density in the interior. Buyers are risk neutral, values are iid, the seller's cost is zero, and IR and payment normalization are as stated in that lecture. Let ρ_n^* denote the optimal expected revenue with n buyers, and ρ_n^{SP} the expected revenue of a no-reserve second-price auction.

Theorem 22.3 (Bulow–Klemperer, smooth private-value version). *If F is regular, then for every $n \geq 1$,*

$$\boxed{\rho_{n+1}^{\mathrm{SP}} \geq \rho_n^*.} \tag{22.4}$$

Here the optimal auction is the monopoly-reserve auction from Lecture 21, which is also optimal among BIC, interim-IR mechanisms in this model.

This compares a simple auction with $n+1$ iid buyers against the best auction with n buyers. It is not a claim that reserves never help, or that every possible additional buyer creates this benefit. We give the full comparison argument.

22.5.1 The key lemma: optimize subject to always selling

Lemma 22.4 (Always-allocate benchmark). *With iid regular values under the assumptions above, the no-reserve second-price auction maximizes expected revenue among normalized DSIC mechanisms that always allocate the item to exactly one buyer at every report profile.*

Proof: For an always-allocate rule, $x_i(v) \geq 0$ and $\sum_i x_i(v) = 1$. Thus

$$\sum_i \phi(v_i) x_i(v) \leq \max_i \phi(v_i)$$

at every profile. This remains true when all the virtual values are negative: always allocating requires choosing one of them. Regularity and identical distributions imply that a highest-valued buyer also has a highest virtual value. The second-price auction always serves such a buyer, so it attains this upper bound pointwise. It is normalized DSIC, and the virtual-surplus identity converts the comparison into a revenue comparison. ■

The absence of a maximum with zero on the right is crucial. The unrestricted optimal auction can withhold the item when virtual values are negative; the always-allocate benchmark cannot.

22.5.2 A comparison auction that preserves the old revenue

Proof (Proof of the Bulow–Klemperer theorem): Define an auxiliary auction $\mathcal{A}$ with $n+1$ buyers:

1. Run the optimal n-buyer reserve auction using only buyers $1, \ldots, n$.
2. If that auction does not sell the item, give it to buyer $n+1$ for free.

The first n buyers face exactly their old allocation and payments. The last buyer's allocation depends only on the first n reports, and its payment is always zero, so its own report cannot affect its payoff. Therefore $\mathcal{A}$ is normalized DSIC and ex post IR.

The construction always allocates the item and earns exactly the old auction's revenue on every realized profile of the first n buyers. Hence $\rho^{\mathcal{A}} = \rho_n^*$. The always-allocate lemma, now applied to $n+1$ iid buyers, gives

$$\rho_n^* = \rho^{\mathcal{A}} \leq \rho_{n+1}^{\mathrm{SP}},$$

which proves (22.4). ■

Both halves matter. Giving the extra buyer a free fallback does not itself create revenue; it turns the original auction into an always-allocate comparator. Regularity then makes second-price allocation optimal within that comparator class.

22.6 How much does the comparison mean for uniform buyers?

With two iid Uniform$[0, 1]$ buyers, the optimal reserve auction earns

$$\rho_2^* = \frac{5}{12} \approx .417.$$

With a third iid buyer, the no-reserve second-price auction earns

$$\rho_3^{\mathrm{SP}} = \frac{3-1}{3+1} = \frac{1}{2}.$$

Thus one extra buyer raises the simple auction's expected revenue above the old optimum. With one original buyer, the analogous comparison is $1/4 \leq 1/3$.

The inequality is in expectation. For example, with one original buyer of value .9, its optimal posted price .5 earns .5. If the additional buyer's value is .1, a no-reserve auction with the two buyers earns only .1. The theorem averages over independent draws of both values.

The uniform formulas also quantify how the reserve's value shrinks with market thickness:

$$\rho_n^* - \rho_n^{\mathrm{SP}} = \frac{2^{-n}}{n+1}.$$

For this model, distribution-specific reserve optimization has a rapidly decreasing incremental value as the number of comparable buyers grows.

22.7 Economic interpretation and scope

A no-reserve second-price auction is *prior independent*: its description does not require knowing F. The performance theorem still assumes iid regular values. Distinguish not needing a distribution to run the mechanism from not needing assumptions to analyze it.

The comparison points to two ways to raise revenue: improve the mechanism using better distributional information, or increase participation by attracting additional comparable buyers. Recruiting buyers has costs, which the theorem does not price. A new buyer with a different value distribution, correlated values, or a change in strategic participation would require a new comparison.

This returns to a recurring course theme: scale can make simple policies competitive with sophisticated optimization. Here the improvement is economic competition, rather than merely a reduction in statistical estimation error. The incentive guarantee remains essential at every market size.

22.8 Takeaways and what follows

1. The payment identity fixes normalized DSIC payments for a given allocation rule. Its interim version gives expected-payment equivalence when allocation probabilities and lowest-type utilities match.
2. Regularity turns highest-value allocation into highest-virtual-value allocation among mechanisms required to sell. This is the key comparison in the Bulow–Klemperer proof.
3. With iid regular buyers, one additional buyer lets a simple no-reserve auction match or exceed the optimal revenue with the original buyers.

Sometimes adding one competitor is more valuable than making the mechanism more sophisticated. The theorem gives this statement a precise meaning under its iid and regularity assumptions.

The characterization and revenue results of Lectures 19–22 used one private number per customer. Lecture 18's VCG theorem also applies to richer valuation functions, but revenue optimization becomes more complex with several products or menus. These questions lead to screening and bundling; Lecture 26 will return to the computational challenges of combinatorial allocation and VCG payments.

Recommended reading

Core.

Roughgarden, *Twenty Lectures on Algorithmic Game Theory*, Section 3.4 for payment uniqueness and Section 6.4 for the complete Bulow–Klemperer comparison. Karlin and Peres, *Game Theory, Alive*, Sections 14.4 and 14.10.1, provide the interim revenue-equivalence and competition results.

Course source.

The final pages of Banerjee's Spring 2017 *Optimal Revenue Mechanisms* use the same auxiliary auction and always-allocate argument. Vohra and Krishnamurthi, Section 5.1.5, discusses revenue comparisons and reserves.

Optional original paper.

Bulow and Klemperer, "Auctions Versus Negotiations," *American Economic Review* 86(1), 180–194, 1996; an earlier working-paper version is available from NBER. The proof here follows the shorter textbook and course-note route.