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

Lecture 21 — Optimal Auctions: Virtual Values and Reserves

Lecturer: Sid Banerjee *Date: November 5, 2026 (draft)*

21.1 From truthful payments to optimal revenue

Lecture 20 tells us how to pay for any monotone allocation rule. We now ask which such rule is best for a seller who knows the distribution of buyer values. The key is to rewrite expected payments entirely in terms of allocations. The resulting objective uses exactly the virtual values introduced in Lecture 2.

With one buyer and value $V \sim F$, a posted price p earns expected revenue

$$\rho(p) = p(1 - F(p)).$$

At a differentiable interior optimum, $1 - F(p) - pf(p) = 0$, equivalently

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

We will discover why the same equation sets an optimal auction reserve, even when there are several buyers.

21.2 Assumptions and the revenue benchmark

Values $V_1, \ldots, V_n$ are independent, with continuous CDFs F_i on supports $[0, \bar{v}_i]$; an endpoint $\bar{v}_i = \infty$ is allowed. Assume densities f_i are continuous and positive in the interior, $F_i(0) = 0$, and finite mean values. Buyers have risk-neutral quasilinear utilities, and the seller has zero cost for now. The integrals below use interior values; in the regular allocation rule, interpret endpoint virtual values by one-sided limits and use a consistent monotone boundary convention. Those endpoint events have probability zero.

A mechanism has a feasible allocation rule $x(v)$ and expected payments $p_i(v)$ conditional on the report profile, averaging any internal randomization as in Lecture 20. For a normalized DSIC mechanism,

$$p_i(v_i, v_{-i}) = v_i x_i(v_i, v_{-i}) - \int_0^{v_i} x_i(z, v_{-i})\, dz, \tag{21.1}$$

where $p_i(0, v_{-i}) = 0$. Under truthful reports the seller's random realized revenue is the sum of realized payments. Its expectation is

$$\rho = \mathbb{E}\left[\sum_{i=1}^{n} p_i(V)\right].$$

The expectations include values and, where relevant, internal randomization. The notation $p_i(v)$ itself is a deterministic conditional expected payment.

Starting the support at zero fixes the lowest-type boundary used below. These statements do not silently cover atoms, positive lower endpoints, or infinite expected values; those cases require their own boundary conventions.

21.3 Recall: virtual value is marginal revenue

Lecture 2 changed coordinates from a price to a served quantile. For one distribution, let $q = 1 - F(v)$ and $r(q) = qF^{-1}(1-q)$. Then

$$\frac{dv}{dq} = -\frac{1}{f(v)}, \qquad r'(q) = v - \frac{1-F(v)}{f(v)} = \phi(v).$$

Thus virtual value is the marginal revenue from increasing the served fraction of buyers. It adjusts a customer's value for the information rents created by serving additional types.

For buyer i, write

$$\phi_i(v) = v - \frac{1-F_i(v)}{f_i(v)}. \tag{21.2}$$

As in Lecture 2, F_i is *regular* when ϕ_i is nondecreasing, equivalently when its quantile revenue curve is concave. This is a distributional condition, distinct from the regularity of choice models in Lectures 16–17.

21.4 Revenue equals expected virtual surplus

Theorem 21.1 (The virtual-surplus identity). *Under the assumptions above, every normalized DSIC mechanism satisfies*

$$\boxed{\rho = \mathbb{E}\left[\sum_{i=1}^{n} \phi_i(V_i) x_i(V)\right].} \tag{21.3}$$

Regularity of the distributions is not required for this identity.

Proof: Fix agent i and the other values v_{-i}. Independence ensures that the conditional distribution of V_i remains F_i. Abbreviate $x(v) = x_i(v, v_{-i})$, $p(v) = p_i(v, v_{-i})$, $F = F_i$, and $f = f_i$. The payment identity gives

$$\mathbb{E}[p(V_i)] = \int_0^{\bar{v}_i} v x(v) f(v)\, dv - \int_0^{\bar{v}_i} \left(\int_0^v x(z)\, dz\right) f(v)\, dv. \tag{21.4}$$

In the double integral, the region is $0 \le z \le v \le \bar{v}_i$. Reversing the order of integration gives

$$\begin{aligned}\int_0^{\bar{v}_i} \int_0^{v} x(z) f(v)\, dz\, dv &= \int_0^{\bar{v}_i} x(z) \left(\int_z^{\bar{v}_i} f(v)\, dv \right) dz \\ &= \int_0^{\bar{v}_i} x(z)(1 - F(z))\, dz.\end{aligned}$$

The integrand is nonnegative, and both terms of (21.4) are finite: $x \le 1$, and $\int_0^{\bar{v}_i}(1-F(z))\, dz = \mathbb{E}[V_i] < \infty$. Substituting back,

$$\begin{aligned}\mathbb{E}[p(V_i)] &= \int_0^{\bar{v}_i} x(v)\big(v f(v) - (1 - F(v))\big)\, dv \\ &= \int_0^{\bar{v}_i} x(v)\phi_i(v) f_i(v)\, dv \\ &= \mathbb{E}[\phi_i(V_i) x_i(V_i, v_{-i})].\end{aligned}$$

Finally average over V_{-i} and sum over i to obtain (21.3). ■

This is an equality of expectations, not a payment formula on every realized value profile. For example, a uniform buyer with value .9 who faces price .5 pays .5, while its virtual value is .8.

The normalization matters too. If $p_i(0, v_{-i}) = \kappa_i(v_{-i})$ instead, expected revenue acquires the additional term $\sum_i \mathbb{E}[\kappa_i(V_{-i})]$. Individual rationality of a zero-valued type requires $\kappa_i \le 0$; giving that type a positive subsidy reduces revenue.

21.5 Turn revenue maximization into allocation optimization

VCG in Lecture 18 implemented the welfare objective $\sum_i v_i x_i$. The revenue identity suggests replacing values by virtual values. To use Myerson's lemma, we must check that the resulting allocation is monotone.

Theorem 21.2 (Optimal mechanisms for regular values). *Let $\mathcal{X} \subseteq \{0,1\}^n$ be a fixed finite feasible set containing the no-service vector. Suppose each independent value distribution satisfies the assumptions above and is regular. For every reported profile, choose*

$$x^*(v) \in \arg\max_{x \in \mathcal{X}} \sum_i \phi_i(v_i) x_i, \tag{21.5}$$

using a fixed ordering of feasible vectors to break ties. Implement this allocation with normalized critical-value payments. The resulting DSIC, ex post IR mechanism maximizes expected revenue among normalized DSIC mechanisms.

Proof: Fix other reports. If agent i is served and increases its report, regularity makes its virtual value weakly increase. All allocations serving i improve by the same amount relative to allocations not serving i; if the increase is zero, the fixed tie rule makes the same choice. Thus the agent cannot lose service. The allocation is monotone and has the threshold implementation from Lecture 19.

At every profile, (21.5) achieves virtual surplus at least as large as any other feasible allocation, including any average of feasible randomized allocations. Taking expectations and applying (21.3) proves optimality. ■

Remark 21.3 (The larger Bayesian benchmark). *The same mechanism is optimal even among BIC, interim-IR mechanisms. Apply the interim payment identity from Lecture 20 and repeat the integral calculation with $q_i(v)$ in place of $x_i(v, v_{-i})$. If $u_i^{\text{int}}(0)$ is the zero type's interim utility, the result is*

$$\rho = \mathbb{E}\left[\sum_i \phi_i(V_i)x_i(V)\right] - \sum_i u_i^{\text{int}}(0).$$

Interim IR makes the last sum nonnegative. The pointwise maximum in (21.5) is therefore an upper bound for this larger class too, and the normalized DSIC construction attains it. Independence is needed for this interim argument as well.

21.6 One item and iid buyers: an optimal reserve auction

For one item, virtual-surplus maximization chooses the largest positive virtual value, with a choice to sell or withhold at zero virtual value. If buyers are iid, they have the same nondecreasing ϕ, so a highest-valued buyer also has a highest virtual value. Choose the highest actual value when virtual values tie; this tie convention remains monotone.

Define the monopoly reserve by the lower crossing

$$r^* = \inf\{v : \phi(v) \geq 0\}. \tag{21.6}$$

Use $r^* = 0$ if all types qualify and $r^* = \infty$ if none does. At an interior continuous crossing, $\phi(r^*) = 0$. A zero interval can give several equally good cutoffs. Endpoint ties can be specified consistently and have no effect on expected revenue for continuous values.

The optimal rule is now familiar:

1. If every bid is below r^*, withhold the item.
2. Otherwise give it to the highest bidder and charge

$$\max\{r^*, \text{highest competing bid}\}.$$

This is a second-price auction with a monopoly reserve. The critical-value theorem supplies its payments. The reserve is the optimal *one-buyer* monopoly price; under these iid regular assumptions it does not depend on the number of bidders.

For one buyer, this rule is exactly the posted price derived at the start. For several buyers, competition raises the payment whenever the strongest competing bid exceeds the reserve.

21.7 Uniform values: a complete numerical benchmark

For $V \sim \text{Unif}[0,1]$, $\phi(v) = 2v - 1$ and $r^* = 1/2$.

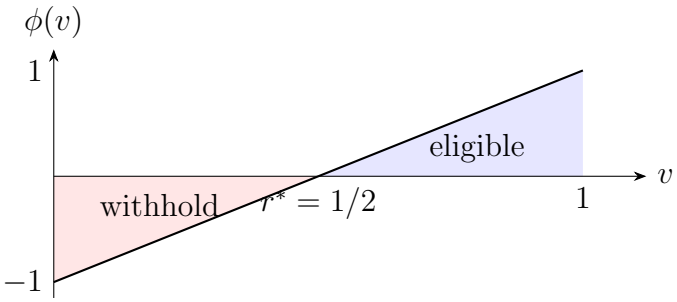

With two buyers, values $(.9, .7)$ lead to payment .7; values $(.9, .2)$ lead to payment .5; values $(.4, .2)$ lead to no sale. A reserve boosts payment when competition is weak but can sacrifice a sale.

The virtual-surplus identity gives the expected optimum directly. With n buyers, the maximum value has density nv^{n-1} on $[0,1]$. Therefore

$$\begin{aligned}\rho_n^* &= \int_{1/2}^{1} (2v-1) n v^{n-1}\, dv \\ &= \frac{2n}{n+1}\left(1 - 2^{-(n+1)}\right) - \left(1 - 2^{-n}\right) \\ &= \boxed{\frac{n - 1 + 2^{-n}}{n+1}.}\end{aligned} \tag{21.7}$$

In particular $\rho_1^* = 1/4$ and $\rho_2^* = 5/12$. With two buyers and no reserve, revenue is the smaller value, whose expectation is $\int_0^1 (1-z)^2\, dz = 1/3$. Thus the reserve improves expected revenue by $1/12$. Lecture 22 derives the general no-reserve benchmark $(n-1)/(n+1)$; compared with it, the uniform reserve gain is $2^{-n}/(n+1)$.

Optional *: independent does not mean identical

With independent but different distributions, virtual values can rank buyers differently from actual values. If $V_1 \sim \text{Unif}[0,1]$ and $V_2 \sim \text{Unif}[0,2]$, then $\phi_1(v) = 2v - 1$ and $\phi_2(v) = 2v - 2$.

At values $(.9, 1.2)$, virtual values are $(.8, .4)$, so buyer 1 wins despite having the lower actual value. Its critical report solves $2b_1 - 1 = .4$, giving payment .7. Simply applying separate reserves and then choosing the highest actual bid would miss this comparison.

21.8 Opportunity cost and the limits of regularity

There is a useful OR connection. If allocating the item costs the seller $c \geq 0$, expected profit equals expected virtual surplus with $\phi_i(v_i)$ replaced by $\phi_i(v_i) - c$. In the iid regular case the reserve crossing becomes $\phi(r) = c$, just as in Lecture 3's monopoly-pricing calculation. Opportunity cost changes the cutoff; incentive compatibility still determines how winners pay.

Optional *: what changes when regularity fails?

The expected virtual-surplus identity still holds, but pointwise maximization can be non-monotone. With one buyer, if $v < w$ yet $\phi(v) > 0 > \phi(w)$, the naive rule serves type v and rejects type w, contradicting Lecture 19's necessary condition. Full optimal-auction theory repairs this through *ironing*, related to concavifying the quantile revenue curve. We do not develop that construction here.

21.9 Takeaways

1. Truthfulness turns allocations into payments; averaging the payment identity turns revenue into expected virtual surplus.
2. Regularity makes pointwise virtual-surplus maximization implementable. For iid single-item buyers, this yields a second-price auction with a monopoly reserve.
3. Welfare and revenue can prefer different allocations. Withholding capacity can improve revenue even when the seller's physical cost is zero.

The construction needs distributional knowledge. Can additional competition substitute for choosing the statistically optimal reserve? That is Lecture 22's central question.

Recommended reading

Roughgarden, *Twenty Lectures on Algorithmic Game Theory*, Sections 5.1–5.2 and 5.4; Vohra and Krishnamurthi, Section 5.1.5 (reserves). The double-integral proof and reserve examples follow Banerjee's *Optimal Revenue Mechanisms*. For the original theory, see Myerson, "Optimal Auction Design", *Mathematics of Operations Research* 6(1), 58–73, 1981 (optional).