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

Lecture 18 — Mechanism Design: From Optimization to Truthful Implementation

Lecturer: Sid Banerjee *Date: October 27, 2026 (draft)*

18.1 We can optimize an allocation—but whose values do we use?

In Lecture 1 we computed efficient allocations from known values, including assignments of heterogeneous items to buyers. In practice, those values often belong to the participants' private information. Asking people to report their values does not by itself make the reports reliable. A payment rule must make truthful reporting attractive.

Our central question is: *How can an optimization procedure obtain the truthful inputs it needs?* We begin with a single-item auction, then show how the Vickrey–Clarke–Groves (VCG) mechanism implements welfare maximization for general valuation functions. No distribution of values is needed for the incentive guarantee.

18.2 Outcomes, private valuations, and reports

Let Ω be a public, finite set of feasible outcomes. Agent i has a private valuation function $v_i(\omega)$ for $\omega \in \Omega$ and submits a reported function $b_i(\omega)$. A direct-revelation mechanism chooses an outcome $\omega(b)$ and a payment $p_i(b)$ from each agent to the seller. With quasilinear utility,

$$u_i(v_i; b) = v_i(\omega(b)) - p_i(b). \tag{18.1}$$

An outcome can specify an assignment, a set of agents served, or a public decision. The private information may contain several numbers: a unit-demand buyer can value items A and B differently.

Definition 1 (DSIC). *A direct mechanism is dominant-strategy incentive compatible if, for every agent i, true valuation function v_i, permissible false report b_i, and reports b_{-i} of all other agents,*

$$v_i(\omega(v_i, b_{-i})) - p_i(v_i, b_{-i}) \geq v_i(\omega(b_i, b_{-i})) - p_i(b_i, b_{-i}). \tag{18.2}$$

Truthful reporting is then a *weakly dominant strategy*: it maximizes utility for every possible report profile of the others. The others need not report truthfully. This is a pointwise incentive statement, rather than a prediction based on beliefs about their behavior.

Individual rationality (IR) is a separate requirement: a truthful participant obtains at least its outside-option utility, taken to be zero. The payment signs and feasibility assumptions will determine when VCG also has this property.

18.3 The first example: a second-price auction

There is one item and a seller with zero cost. Agent i values receiving it at $v_i \geq 0$ and values not receiving it at zero. Its report reduces to a single number b_i. The efficient allocation gives the item to the highest report, with a fixed priority order for ties. A *second-price*, or Vickrey, auction charges the winner the highest competing report and charges losers zero.

Theorem 18.1 (Second-price truthfulness). *Truthful bidding is a weakly dominant strategy in the second-price auction. A truthful bidder obtains nonnegative utility, and truthful reports produce a welfare-maximizing allocation.*

Proof: Fix agent i and all other bids, and write $m_i = \max_{j \neq i} b_j$, using $m_i = 0$ if there are no rivals. Any winning bid makes i pay m_i and gives utility $v_i - m_i$; any losing bid gives zero. If $v_i > m_i$, truthful bidding wins and obtains positive utility; losing instead is worse. If $v_i < m_i$, truthful bidding loses, while winning would give negative utility. At $v_i = m_i$, both outcomes give zero, so the tie rule does not change the conclusion. Thus no deviation improves on truthfulness, and truthful utility is nonnegative. With all bids truthful, the highest-valued agent wins, maximizing welfare. ■

The price has an externality interpretation. Without the winner, the other agents could obtain welfare m_i. With the winner receiving the item, they obtain zero. The winner pays the welfare loss its presence causes to the others. A losing agent causes no such loss and pays zero.

18.4 General VCG: charge the externality

Given reports b, choose an exact reported-welfare maximizer:

$$\omega^*(b) \in \arg\max_{\omega \in \Omega} \sum_{j=1}^{n} b_j(\omega). \tag{18.3}$$

Use a fixed ordering of outcomes to select a maximizer when there is a tie.

For agent i, compute the benchmark

$$H_i(b_{-i}) = \max_{\omega \in \Omega} \sum_{j \neq i} b_j(\omega). \tag{18.4}$$

The *Clarke-pivot payment* is

$$\boxed{p_i(b) = H_i(b_{-i}) - \sum_{j \neq i} b_j(\omega^*(b)).} \tag{18.5}$$

It is the best welfare the other agents could obtain minus their welfare in the chosen outcome. All quantities in this formula use reports; when reports are truthful, they are actual welfare quantities.

18.4.1 When does the benchmark mean removing an agent?

In the allocation problems considered here, an agent's value depends on its own service or item, and is zero if it receives nothing. Assume:

1. Deleting i's service from a feasible allocation leaves a feasible allocation for the other agents, with their values unchanged.
2. Every allocation feasible without i can be extended to a feasible full outcome by leaving i unserved, again preserving all other values.

These conditions hold for our unit-demand assignment problem and for ordinary capacity allocation with optional service. They identify (18.4) with optimizing over the smaller problem in which i is absent. They also give a zero-utility outside option. They should be checked anew in settings where an agent brings resources or its presence changes what others can do.

The DSIC proof below only requires a public fixed outcome set, exact welfare maximization, and H_i independent of i's report. The removal conditions explain the additional payment and participation guarantees for the pivot choice.

18.5 Why truthful reporting maximizes utility

Theorem 18.2 (VCG truthfulness and efficiency). *For a finite public outcome set and quasilinear utilities, the allocation rule (18.3) with payments (18.5) is DSIC and welfare maximizing under truthful reports. Under the service-removal assumptions above, it also has nonnegative payments and guarantees a truthful agent nonnegative utility against every profile of other reports.*

Proof: Fix agent i, its true valuation v_i, and the other reports b_{-i}. Let ω be the outcome induced by any report from i. Substituting its payment into its true utility gives

$$u_i = v_i(\omega) + \sum_{j \neq i} b_j(\omega) - H_i(b_{-i}). \tag{18.6}$$

The last term is independent of i's report. When i reports v_i, the mechanism chooses an outcome maximizing exactly the first two terms over all feasible ω. Any false report can only select another feasible outcome, whose value of those terms is no larger. Therefore the truthful report maximizes utility. With everyone truthful, (18.3) maximizes actual total welfare.

For participation, extend a benchmark-optimal outcome without i by leaving i unserved. Its value of the first two terms of (18.6) is H_i, since i then has value zero. The truthful maximizer does at least as well, so truthful utility is nonnegative. Finally, restricting the chosen outcome to the other agents gives a feasible outcome without i; its reported welfare cannot exceed H_i. Thus (18.5) is nonnegative. ■

A fixed tie order makes the allocation well defined and will be useful for later monotonicity arguments. In the VCG proof itself, the essential property is exact maximization: every selected maximizer gives the same value of the objective in (18.6) under truthful reporting. The benchmark H_i must not be changed using i's own report.

This is a design principle: the payment makes each agent evaluate the chosen outcome using its own true value plus everyone else's reported value. Its incentive is therefore aligned with the optimization performed by the mechanism.

18.6 A two-item assignment, with all payments computed

There are two items, A and B, and three unit-demand buyers. Each buyer receives at most one item; each item is assigned at most once. Values for receiving no item are zero:

	A	B
1	10	3
2	9	8
3	0	7

The efficient assignment is $1 \to A$, $2 \to B$, with welfare 18. The alternatives $1 \to A$, $3 \to B$ and $2 \to A$, $3 \to B$ give 17 and 16, respectively. The other assignments do no better.

To compute buyer 1's payment, remove buyer 1 and solve the remaining assignment problem. Assigning A to 2 and B to 3 gives $H_1 = 9 + 7 = 16$. In the full chosen assignment, the other buyers obtain only 8. Therefore

$$p_1 = 16 - 8 = 8, \qquad u_1 = 10 - 8 = 2.$$

Without buyer 2, buyers 1 and 3 obtain $H_2 = 10 + 7 = 17$. In the full assignment, the others obtain 10, so

$$p_2 = 17 - 10 = 7, \qquad u_2 = 8 - 7 = 1.$$

Without buyer 3, the original assignment remains feasible and optimal, so $H_3 = 18$ and $p_3 = 18 - 18 = 0$. Buyer 3 is unserved and has utility zero. We can check the accounting:

$$\underbrace{18}_{\text{welfare}} = \underbrace{(8+7)}_{\text{seller revenue}} + \underbrace{(2+1+0)}_{\text{buyer utilities}}.$$

Computing the allocation is only the first optimization. Computing each winner's payment requires solving a counterfactual optimization with that buyer removed. This is the assignment problem from Lecture 1, now paired with an incentive rule for acquiring its inputs.

18.7 What VCG settles, and what remains open

VCG maximizes welfare, not seller revenue. With one buyer and one item, it allocates whenever the value is positive and charges zero: no other buyer suffers an externality. For a Uniform$[0, 1]$ buyer, a posted price of $1/2$ instead earns expected revenue $1/4$. That price deliberately forgoes some efficient trades. We will study truthful revenue maximization after characterizing the available allocation rules.

VCG also assumes quasilinear utilities and exact welfare maximization. A hard allocation problem remains hard, and reporting a value for many possible outcomes can itself be demanding. Simply replacing the optimizer by a heuristic and retaining the payment formula does not establish DSIC: the key maximization step in the proof may fail. Lecture 26 will revisit this computational question.

18.8 Takeaways and next question

1. A mechanism specifies both an allocation and how payments respond to reports. DSIC requires truthful reporting to be optimal against every report profile of the others.
2. VCG implements exact welfare maximization by charging each agent its externality. The service-removal assumptions make the pivot payments nonnegative and participation individually rational.
3. Welfare maximization is one truthful design, but it need not maximize revenue or be computationally easy.

Must a truthful mechanism maximize welfare? Which other allocation rules admit truthful payments? Lecture 19 answers these questions for single-parameter service environments.

Recommended reading. Roughgarden, *Twenty Lectures on Algorithmic Game Theory*, Sections 2.1–2.4 (Vickrey), 7.1–7.2 (VCG), and 7.3 (limitations). The quasilinear model and second-price proof build on Banerjee's Spring 2017 *Strategic Models and Mechanism Design*; the VCG development follows Roughgarden.

Optional * appendix: posted prices and first-price bidding

This optional background develops posted-price calculations and equilibrium bidding in a first-price auction. It goes beyond the main DSIC-based unit, whose results do not require this equilibrium analysis. Here there is one item, zero seller cost, and scalar private values V_i rather than general valuation functions.

A posted price facing several buyers

Suppose $V_1, \dots, V_n$ are independent and identically distributed (iid), with a continuous CDF F. Offer the item at the same price p to all buyers, using a fixed priority order to serve the first willing buyer. There is only *one* item: the seller earns p once if at least one value exceeds p. Thus random revenue and expected revenue are

$$R^{\text{post}}(p) = p\mathbf{1}\{\max_i V_i \geq p\}, \qquad \rho_n^{\text{post}}(p) = p\big(1 - F(p)^n\big). \tag{18.7}$$

Independence gives the product $F(p)^n$. With independent, nonidentical distributions it becomes $\prod_i F_i(p)$. A sale need not go to the highest-valued buyer: several buyers may accept, and priority then determines the allocation.

For iid Uniform$[0,1]$ values, $\rho_n^{\text{post}}(p) = p(1-p^n)$ on $[0,1]$. Differentiating gives

$$1 - (n+1)p^n = 0, \qquad p_n = (n+1)^{-1/n}, \qquad \rho_n^{\text{post}}(p_n) = \frac{n}{n+1}(n+1)^{-1/n}.$$

The price balances a larger payment against the chance of no sale. Notice that this is a different problem from offering one unit to *each* of n buyers, whose expected revenue would be $np(1-F(p))$.

Order statistics: value and competitive pressure

Write values in decreasing order:

$$V_{(1)} \geq V_{(2)} \geq \cdots \geq V_{(n)}.$$

These are random *order statistics*; lowercase $v_{(j)}$ denotes a realized order statistic. The largest value determines the welfare of an efficient allocation. The second-largest value will determine revenue in a truthful second-price auction without a reserve.

For iid continuous values,

$$\mathbb{P}(V_{(1)} \leq z) = F(z)^n, \tag{18.8}$$
$$\mathbb{P}(V_{(2)} \leq z) = F(z)^n + n(1-F(z))F(z)^{n-1}. \tag{18.9}$$

The second event means that at most one value exceeds z: either none does, or exactly one does.

The winner's value and the competitive price are distinct quantities.

The uniform example

For a nonnegative random variable $Z \le 1$, $\mathbb{E}[Z] = \int_0^1 \mathbb{P}(Z > z)\, dz$. Applying this to (18.8) gives

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

For $n \ge 2$, equation (18.9) gives

$$\begin{aligned}\mathbb{E}[V_{(2)}] &= \int_0^1 \left[1 - z^n - n(1-z)z^{n-1}\right] dz \\ &= 1 - \frac{1}{n+1} - n\left(\frac{1}{n} - \frac{1}{n+1}\right) = \frac{n-1}{n+1}.\end{aligned}$$

With only one bidder we define the missing competing bid to be 0, so a second-price auction without a reserve earns zero.

First-price auctions: why shade a bid?

In a first-price auction, the highest bidder wins and pays its own bid; losers pay zero. Winning at bid b_i gives utility $v_i - b_i$. Bidding one's value leaves no surplus even when one wins. A lower bid offers more surplus conditional on winning, but reduces the chance of winning. There is generally no best bid that works against every possible rival bid.

A *bidding strategy* $\beta_i(v_i)$ specifies a bid for each value. A strategy profile is a *Bayesian Nash equilibrium* (BNE) if, for every buyer and value, the prescribed bid maximizes expected utility given the other buyers' strategies and the conditional distribution of their values. This is a statement about beliefs and best responses, rather than a dominant-strategy guarantee.

A symmetric uniform equilibrium (optional)

Assume $n \ge 2$, iid Uniform$[0, 1]$ private values, risk-neutral quasilinear buyers, no reserve, and unrestricted nonnegative bids. We claim that

$$\beta(v) = \frac{n-1}{n} v \tag{18.10}$$

is a symmetric BNE. Symmetric means that every buyer uses the same strategy.

To verify all deviations, put $\alpha = (n-1)/n$ and suppose the other buyers use $\beta(v) = \alpha v$. A bid in $[0, \alpha]$ can be written as αz for $z \in [0, 1]$. It wins when all $n-1$ rival values are below z, so the expected utility of a buyer with value v is

$$g_v(z) = (v - \alpha z) z^{n-1}.$$

For $z > 0$,

$$g_v'(z) = z^{n-2}\big((n-1)v - n\alpha z\big) = (n-1)z^{n-2}(v - z).$$

The utility rises up to $z = v$ and falls thereafter; the boundary cases $v = 0, 1$ give the same conclusion. A bid above α wins with probability one but pays more than bidding α, so it cannot improve on the best bid in the stated range. Thus the prescribed bid is a global best response, proving the claim.

This argument uses risk neutrality, the uniform distribution, independence, and symmetry. Truthfulness of the second-price auction needed none of these distributional calculations.

Comparing revenues while accounting for behavior

Under truthful bidding, second-price revenue is $R_n^{\mathrm{SP}} = V_{(2)}$. In the first-price equilibrium just derived, revenue is $R_n^{\mathrm{FP}} = \alpha V_{(1)}$. Hence

$$\rho_n^{\mathrm{FP}} := \mathbb{E}[R_n^{\mathrm{FP}}] = \frac{n-1}{n}\frac{n}{n+1} = \frac{n-1}{n+1} = \mathbb{E}[R_n^{\mathrm{SP}}] =: \rho_n^{\mathrm{SP}}. \tag{18.11}$$

The expected revenues match, although payments on a particular value profile can differ. Comparing the same *bids* under different payment rules would miss the change in bidding behavior.

An optimized posted price and a no-reserve auction need not have the same revenue. With two uniform buyers, the posted-price optimum is $2/(3\sqrt{3}) \approx 0.385$, above the second-price revenue $1/3$. With three buyers, the second-price revenue is $1/2$, above the posted-price optimum $\frac{3}{4}\,4^{-1/3} \approx 0.472$. There is no general ranking established by these examples.

As n grows in this uniform model, expected efficient welfare $n/(n+1)$ and second-price revenue $(n-1)/(n+1)$ both approach 1. A thicker market creates more competitive pressure. This calculation holds the distribution of each additional buyer fixed; it says nothing yet about the cost of recruiting that buyer.

Optional appendix reading. Vohra and Krishnamurthi, *Principles of Pricing*, Sections 5.1.1–5.1.6 and 5.4.1; Karlin and Peres, *Game Theory, Alive*, Section 14.4. These calculations are background rather than prerequisites for our DSIC-based unit.