Lecture 19 — Single-Parameter Mechanisms: Monotonicity and Threshold Payments

Lecturer: Sid Banerjee *Date: October 29, 2026 (draft)*

19.1 Which other allocations can be made truthful?

Lecture 18 showed that VCG implements an efficient allocation using externality payments. In a single-item auction this is the highest-bid allocation with second-price payments. But welfare maximization is only one possible objective. Which other allocation rules, including inefficient ones, admit truthful payments?

Today we restrict the private information to one value per agent and focus on deterministic *win/lose* service. This restriction lets us characterize all implementable allocations, rather than only the efficient VCG choice. The next lecture allows intermediate allocation levels, including probabilities of service.

19.2 A single-parameter service environment

There are n agents, each with a private nonnegative value v_i for service. Public feasibility constraints specify a set

$$\mathcal{X} \subseteq \{0,1\}^n$$

of permissible service vectors. A *direct-revelation mechanism* asks each agent for a report $b_i \geq 0$ and chooses

$$x(b) \in \mathcal{X}, \qquad p(b) = (p_1(b), \ldots, p_n(b)).$$

Agent i's utility is

$$u_i(v_i; b) = v_i x_i(b) - p_i(b). \tag{19.1}$$

Only v_i is private: resource requirements, feasible sets, and service quality are public and do not depend on the reports. This is why the environment is *single parameter*.

Examples include a single item ($\sum_i x_i \leq 1$), k identical items with unit-demand buyers ($\sum_i x_i \leq k$), and a capacity constraint $\sum_i a_i x_i \leq c$ with publicly known resource requirements a_i. In the last example v_i is the value of completing agent i's entire service, rather than a value per unit of resource.

Definition 1 (Dominant-strategy incentive compatibility). *The mechanism is DSIC if, for every i, true value v_i, alternative report b_i, and vector b_{-i} of other reports,*

$$v_i x_i(v_i, b_{-i}) - p_i(v_i, b_{-i}) \geq v_i x_i(b_i, b_{-i}) - p_i(b_i, b_{-i}). \tag{19.2}$$

The other agents need not be truthful. DSIC is a pointwise guarantee against every possible b_{-i}, with no value distribution required. *Ex post IR* is the separate requirement that truthful utility be nonnegative for every such profile.

For the second-price auction, fixing b_{-i} leaves agent i a simple offer: receive the item at price $\max_{j \neq i} b_j$, or do not receive it and pay zero. The seller can compute that price without using i's own report.

19.3 Truthfulness forces monotonicity

Fix an agent and the other reports. For this section abbreviate

$$x(z) = x_i(z, b_{-i}), \qquad p(z) = p_i(z, b_{-i}).$$

A higher report must not make service less likely. The next proof derives that claim from incentives alone.

Lemma 19.1 (Necessary monotonicity). *If a single-parameter mechanism is DSIC, then $x_i(z, b_{-i})$ is nondecreasing in z, for every fixed b_{-i}.*

Proof: Take two values $v > w$. DSIC for type v against report w, and type w against report v, gives

$$\begin{aligned} vx(v) - p(v) &\geq vx(w) - p(w), \\ wx(w) - p(w) &\geq wx(v) - p(v). \end{aligned}$$

Adding cancels the payments:

$$(v - w)\big(x(v) - x(w)\big) \geq 0.$$

Because $v - w > 0$, we must have $x(v) \geq x(w)$. ■

For example, allocating a single item to the *second*-highest bidder is not monotone: a winning agent can raise its report above all rivals and lose. No choice of payments can make that allocation DSIC. This is a restriction on the allocation rule itself, rather than a failure to find a clever price.

19.4 A monotone binary rule is a threshold rule

For a monotone $x_i(\cdot, b_{-i}) \in \{0, 1\}$, define the critical value

$$\theta_i(b_{-i}) = \inf\{z \geq 0 : x_i(z, b_{-i}) = 1\}. \tag{19.3}$$

Use $\theta_i = \infty$ if the agent never receives service. A report below a finite threshold loses; one above it wins. At the threshold itself either decision is allowed, according to the specified allocation rule. A zero threshold covers service available at arbitrarily small positive reports, including an always-served agent.

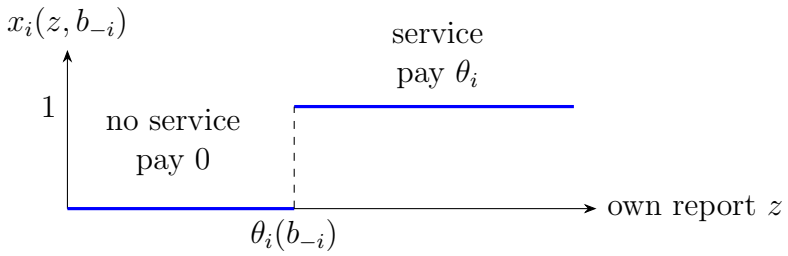

Raising a winning report changes neither the service nor its critical price.

Theorem 19.2 (Threshold implementation). *Let $x(b) \in \mathcal{X}$ be a feasible, deterministic, binary allocation rule, monotone in each agent's own report. Charge each winner its critical value and each loser zero:*

$$p_i(b) = \begin{cases} \theta_i(b_{-i}), & x_i(b) = 1, \\ 0, & x_i(b) = 0. \end{cases} \tag{19.4}$$

Then the mechanism is DSIC, ex post IR, and makes no payments to agents. It satisfies $p_i(0, b_{-i}) = 0$.

Proof: Fix b_{-i} and put $\theta = \theta_i(b_{-i})$. If $v_i > \theta$, truthful reporting wins and gives $v_i - \theta > 0$. Any other winning report gives the same utility, and losing gives zero. If $v_i < \theta$, truthful reporting loses and gives zero; winning instead would give $v_i - \theta < 0$. At $v_i = \theta$, both outcomes give zero. An agent with threshold ∞ always loses. These cases prove DSIC and IR. Payments are nonnegative; a zero report can win only at threshold zero, so the normalization follows. ■

Feasibility is not supplied by this theorem: we must first have a joint allocation rule $x(b) \in \mathcal{X}$. Independently inventing a threshold for every agent might promise service to too many agents.

19.4.1 Why the critical price is the right price

DSIC makes payments constant among reports producing the same allocation. Indeed, compare two winning reports in both directions in (19.2); their allocations coincide, so their payments must coincide. The same argument applies to two losing reports.

Suppose $0 < \theta < \infty$ and normalize the losing payment to zero. Let c be the common winning payment. A losing type $w < \theta$ must prefer not to buy, so $c \geq w$. A winning type $v > \theta$ must

prefer buying, so $c \le v$. Taking w up to θ and v down to θ forces $c = \theta$. At threshold zero, the same limiting argument forces a normalized winning payment of zero; a never-served agent pays zero. Lecture 20 makes the general normalization and uniqueness statements precise.

19.5 Examples: the same principle under different constraints

19.5.1 One item, with or without a reserve

Allocating to the highest bid is monotone. Agent i's critical value is the largest competing bid, so threshold payments reproduce the second-price auction—the efficient single-item VCG mechanism from Lecture 18. With a publicly announced reserve $r \ge 0$, the agent must beat both that reserve and its rivals:

$$\theta_i(b_{-i}) = \max\{r, \max_{j \ne i} b_j\}.$$

If every bid is below r, there is no sale. The winner pays the maximum above, even if it was the only bidder to clear the reserve. This rule can be inefficient: at values $(.4, .3)$ and reserve .5, the item is withheld despite a positive gain from trade. It is nevertheless monotone and DSIC with threshold payments. Truthful implementability does not require welfare maximization.

19.5.2 Several identical items

With k units and unit-demand buyers, serve the k highest bids (or everyone if $k \ge n$). For $k < n$, each winner's threshold is the kth highest bid among its rivals, which is the $(k+1)$st highest bid overall. Each winner pays that threshold; losers pay zero. If $k \ge n$, every agent can be served at threshold zero. This argument relies on unit demand: multiunit demand introduces additional private information.

19.5.3 A small capacity allocation problem

There are two units of capacity. Agent 1 requires both units; agents 2 and 3 require one each. Public sizes are $(a_1, a_2, a_3) = (2, 1, 1)$. With nonnegative reports, a welfare-maximizing rule compares

$$b_1 \quad \text{with} \quad b_2 + b_3.$$

Use a fixed tie rule. If b_1 is larger, serve agent 1; otherwise serve agents 2 and 3. This is monotone in each report. Its thresholds are

$$\theta_1 = b_2 + b_3, \qquad \theta_2 = \max\{0, b_1 - b_3\}, \qquad \theta_3 = \max\{0, b_1 - b_2\}.$$

For $b = (9, 6, 5)$, agents 2 and 3 win, paying 4 and 3. Their combined value is 11, while seller revenue is 7. Agent 2 need only contribute enough to make the pair beat agent 1. The critical price is therefore not simply the largest losing bid.

These payments also agree with the VCG externalities. Without agent 2, the best allocation earns reported welfare 9, while the other served agent contributes 5 in the chosen allocation, giving $p_2 = 9 - 5 = 4$. Similarly $p_3 = 9 - 6 = 3$. For this efficient rule, VCG and the threshold characterization coincide; the characterization also covers the inefficient reserve rule above.

More generally, exact maximization of $\sum_i b_i x_i$ over a fixed finite feasible set, with a fixed ordering to break ties between allocations, is monotone in each own bid. If a chosen allocation serves i, raising only b_i increases the score of every allocation serving i by the same amount and improves those scores relative to all allocations not serving i. The capacity example is a small instance of this reasoning. Computing an optimum and finding critical values can nevertheless be computationally demanding.

19.6 Takeaways and next question

1. Truthfulness rules out nonmonotone allocations: the two incentive constraints already imply this restriction.
2. A feasible monotone binary allocation is implemented by charging a winner its critical value. The resulting mechanism is DSIC and voluntarily acceptable.
3. The allocation problem and incentive implementation are different tasks. Threshold payments apply to monotone rules whether or not those rules are efficient.

What replaces a single critical value when an agent's allocation rises gradually, or is a probability? The next lecture derives the payment identity that answers this question.

Recommended reading

Core.

Roughgarden, *Twenty Lectures on Algorithmic Game Theory*, Sections 3.1–3.3. Section 4.1 gives a larger knapsack-auction example.

Course source.

Banerjee, Spring 2017 ORIE 4154, *Single-Parameter Settings and Myerson's Lemma*, especially the model and two-report monotonicity argument. We separate its binary threshold case here from the general payment identity in Lecture 20.