DS 574 Algorithmic Mechanism Design **Lecture #3**
Fall 2026 Prof. Kira Goldner

Allocation and Payment Rules

Now, we formalize the concepts we've been using so far. A mechanism $M = (\mathbf{x}, \mathbf{p})$ is completely determined by its allocation rule $\mathbf{x}$ and payment rule $\mathbf{p}$.

Definition 1. An *allocation rule* x is a (potentially randomized) mapping from bidder actions (bids $\mathbf{b}$) to feasible outcomes in X.

In the single-item setting, what is the set of feasible outcomes X? We say $\mathbf{x} \in X$ where $\mathbf{x} = (x_1, \ldots, x_n)$ and x_i denotes how much of the item bidder i gets.

- At most 1 item is allocated: $\sum_{i=1}^{n} x_i \leq 1$.
- A bidder is either allocated or isn't: $x_i \in \{0, 1\}\, \forall i$.

What does this mean for a potentially randomized allocation rule $\mathbf{x}(\mathbf{b})$?

Definition 2. A *payment rule* $\mathbf{p}(\mathbf{b}) \in \mathbb{R}^n$ is a mapping from bidder actions (bids $\mathbf{b}$) to (non-negative) real numbers where $p_i(\mathbf{b})$ is the amount that bidder i pays in the outcome $\mathbf{x}(\mathbf{b})$.

Now we can formalize quasilinear utility in terms of general allocation and payment rules.

Definition 3. For a mechanism $M = (\mathbf{x}, \mathbf{p})$, a bidder with *quasilinear utility* has utility

$$u_i(\mathbf{b}) = v_i \cdot x_i(\mathbf{b}) - p_i(\mathbf{b}).$$

We'll narrow our attention to payment rules that satisfy

$$p_i(\mathbf{b}) \in [0, b_i \cdot x_i(\mathbf{b})]$$

for every i and $\mathbf{b}$. The constraint that $p_i(\mathbf{b}) \geq 0$ is equivalent to prohibiting the seller from paying the bidders. The constraint that $p_i(\mathbf{b}) \leq b_i \cdot x_i(\mathbf{b})$ ensures that a truth-telling bidder receives nonnegative utility (do you see why?).

Again, our goal is to design DSIC mechanisms:

Definition 4. A mechanism is *dominant-strategy incentive-compatible (DSIC)* if it is a bidder's dominant strategy to bid their true value, i.e. it maximizes their utility, *no matter what* the other bidders do. That is,

$$u_i(v_i, \mathbf{b}_{-i}) \geq u_i(z, \mathbf{b}_{-i}) \quad \forall z, \mathbf{b}_{-i}.$$

Myerson's Lemma

We now come to two important definitions. Both articulate a property of allocation rules.

Definition 5 (Implementable Allocation Rule)**.** An allocation rule $\mathbf{x}$ is *implementable* if there is a payment rule $\mathbf{p}$ such the sealed-bid auction $(\mathbf{x}, \mathbf{p})$ is DSIC.

Definition 6 (Monotone Allocation Rule)**.** An allocation rule x for a single-parameter environment is *monotone* if for every bidder i and bids $\mathbf{b}_{-i}$ by the other bidders, the allocation $x_i(z, \mathbf{b}_{-i})$ to i is nondecreasing in its bid z.

That is, in a monotone allocation rule, bidding higher can only get you more stuff.

For example, the single-item auction allocation rule that awards the good to the highest bidder is monotone: if you're the winner and you raise your bid (keeping other bids constant), you continue to win. By contrast, awarding the good to the second-highest bidder is a non-monotone allocation rule: if you're the winner and raise your bid high enough, you lose.

We state Myerson's Lemma in three parts; each is conceptually interesting and will be useful in later applications.

Theorem 1 (Myerson's Lemma **?**)**.** *Fix a single-parameter environment.*

(a) An allocation rule $\mathbf{x}$ *is implementable if and only if it is monotone.*

(b) If $\mathbf{x}$ *is monotone, then there is a unique payment rule such that the sealed-bid mechanism* $(\mathbf{x}, \mathbf{p})$ *is DSIC [assuming the normalization that* $b_i = 0$ *implies* $p_i(\mathbf{b}) = 0$*].*

(c) The payment rule in (b) is given by an explicit formula:

$$p_i(b_i, \mathbf{b}_{-i}) = b_i \cdot x_i(b_i, \mathbf{b}_{-i}) - \int_0^{b_i} x_i(z, \mathbf{b}_{-i}) dz.$$

Myerson's Lemma is the foundation on which we'll build most of our mechanism design theory. Let's review what it is saying.

Part (a): Finding an allocation rule that can be made DSIC (is implementable, Definition 5) seems confusing, but is actually equivalent to and just as easy as checking if the allocation is monotone (Definition 6).

Part (b): If an allocation rule *is* implementable (can be made to be DSIC), then there's no ambiguity in what the payment rule should be.

Part (c): There's a simple and explicit formula for this!

Myerson's Lemma is what makes Steps 4 and 5 of our design approach legal: (a) says the monotonicity check is the *right* check, and (b)–(c) say the payments are not ours to choose.

Proof of Myerson's Lemma (Theorem 1). As shorthand, write $x(z)$ and $p(z)$ for the allocation $x_i(z, \mathbf{b}_{-i})$ and payment $p_i(z, \mathbf{b}_{-i})$ of i when it bids z, respectively.

Suppose $(\mathbf{x}, \mathbf{p})$ is DSIC, and consider any $0 \leq y < z$. Because bidder i might well have private valuation z and can submit the false bid y if it wants, DSIC demands that

$$\underbrace{z \cdot x(z) - p(z)}_{\text{utility of bidding } z \text{ given value } z} \geq \underbrace{z \cdot x(y) - p(y)}_{\text{utility of bidding } y \text{ given value } z} \tag{1}$$

Similarly, since bidder i might well have the private valuation y and could submit the false bid z, $(\mathbf{x}, \mathbf{p})$ must satisfy

$$\underbrace{y \cdot x(y) - p(y)}_{\text{utility of bidding } y \text{ given value } y} \geq \underbrace{y \cdot x(z) - p(z)}_{\text{utility of bidding } z \text{ given value } y} \tag{2}$$

Rearranging inequalities (1) and (2) yields the following sandwich, bounding $p(z) - p(y)$ from below and above:

$$y \cdot [x(z) - x(y)] \leq p(z) - p(y) \leq z \cdot [x(z) - x(y)] \tag{3}$$

Translation: if raising your bid from y to z buys you $[x(z) - x(y)]$ more stuff, the extra you're charged for it must be at least the low value y per unit and at most the high value z per unit. Everything in Myerson's Lemma falls out of squeezing this sandwich.

From here, we can conclude:

- $\mathbf{x}$ must be monotone.
- $p'(z) = z \cdot x'(z)$.

Why? First, if x is not monotone, the inequalities in (3) would be violated. Second, assuming x is differentiable, by dividing (3) by $z - y$ and taking the limit as $y \to z$, we obtain $p'(z) = z \cdot x'(z)$. Even for non-differentiable x, we obtain a similar equation in terms of the change in the allocation at z.

Assuming that $p(0) = 0$ then gives the payment identity

$$p_i(b_i, \mathbf{b}_{-i}) = \int_0^{b_i} z \cdot \frac{d}{dz} x_i(z, \mathbf{b}_{-i}) dz$$

or alternatively, after integration by parts,

$$p_i(b_i, \mathbf{b}_{-i}) = b_i \cdot x_i(b_i, \mathbf{b}_{-i}) - \int_0^{b_i} x_i(z, \mathbf{b}_{-i}) dz \tag{4}$$

for every bidder i, bid b_i, and bids $\mathbf{b}_{-i}$ by the others.

Equation (3) tells us that this is the only payment rule that could possibly be DSIC. But does it in fact satisfy DSIC when x is monotone?

Bidder i's utility will then be

$$u_i(b_i, \mathbf{b}_{-i}) = v_i \cdot x_i(b_i, \mathbf{b}_{-i}) - p_i(b_i, \mathbf{b}_{-i}),$$

or with the payment identity,

$$u_i(b_i, \mathbf{b}_{-i}) = (v_i - b_i) \cdot x_i(b_i, \mathbf{b}_{-i}) + \int_0^{b_i} x_i(z, \mathbf{b}_{-i}) dz$$

which for monotone $\mathbf{x}$ is maximized when $b_i = v_i$, independent of $\mathbf{b}_{-i}$, as desired. $\square$
Sanity check: plug the single-item allocation rule into the payment identity. Here $x_i(z, \mathbf{b}_{-i})$ is 0 until z hits $B = \max_{j \neq i} b_j$ and 1 after, so the integral is $b_i - B$ and the formula gives $p_i = b_i - (b_i - B) = B$. Second price! Our general formula had better reproduce the one auction we already trust.

Single-Parameter Environments

All of our definitions and Myerson's Lemma actually apply to a more general setting which we call *single-parameter environments*. The main idea here is that each bidder i only has a single piece of private information, like their value v_i, that needs to be elicited in order to run the mechanism. Here are some other examples of non-single-item yet single-parameter environments.

- **Single-item:** A seller has a single item to sell. The set of feasible outcomes X satisfy $\sum_{i=1}^{n} x_i \leq 1$ and $x_i \in \{0, 1\}$.
- k **identical items:** A seller has k identical items to sell and each buyer gets at most one. The set of feasible outcomes X satisfy $\sum_{i=1}^{n} x_i \leq k$ and $x_i \in \{0, 1\}$.
- **Sponsored search:** There are k advertising slots, each with click-through-rate α_j. A buyer i gets value $v_i \cdot \alpha_j$ from winning the jth slot. The set of feasible outcomes X satisfy $\sum_{i=1}^{n} x_i \leq \sum_{j=1}^{k} \alpha_j$ and $x_i \in \{\alpha_j\}_{j=1}^{k} \cup \{0\}$ where $x_i = \alpha_j$ if bidder i is assigned the jth slot.

Exercise (optional): Graph an allocation rule as a function of a single-bidder (hold $\mathbf{b}_{-i}$ fixed) with value on the x-axis and allocation on the y-axis. Show that for a DSIC auction, Myerson's Lemma implies that the payment is the area to the left of the allocation curve, and a bidder's utility is the area under the allocation curve.

The Revelation Principle

So far, we've been investigating *Dominant-Strategy Incentive-Compatible (DSIC)* mechanisms. To be DSIC, this means that

(1) Every participant in the mechanism has a dominant strategy, no matter what their private valuation is.

(2) This dominant strategy is *direct revelation*, where the participant truthfully reports all of their private information to the mechanism.

There are mechanisms that satisfy (1) but not (2). To give a silly example, imagine a single-item auction in which the seller, given bids $\mathbf{b}$, runs a Vickrey auction on the bids $3\mathbf{b}$. Every bidder's dominant strategy is then to bid $v_i/3$.

For a formal definition of a direct revelation mechanism:

Definition 7. A mechanism is *direct revelation* if it is single-round, sealed-bid, and has action space equal to the type (value) space. That is, an agent can bid any type they might have, and an agent's action *is* bidding a type.

The Revelation Principle and the Irrelevance of Truthfulness

The Revelation Principle states that, given requirement (1), there is no need to relax requirement (2): it comes "for free."

Theorem 2 (Revelation Principle for DSIC Mechanisms)**.** *For every mechanism M in which every participant has a dominant strategy (no matter what their private information), there is an equivalent direct-revelation DSIC mechanism M'.*

Equivalent here means that as a function of the *valuation profile* (not bids), the allocation and payment $(x(\mathbf{v}), p(\mathbf{v}))$ are equivalent in both M and M'.

Proof. The proof uses a simulation argument; see Figure 1. By assumption, in mechanism M, every bidder i has a dominant strategy $\sigma_i(v_i)$ whatever their v_i.

We construct the following mechanism M', the mechanism takes over the responsibility of applying the dominant strategy. Precisely, (direct-revelation) mechanism M' accepts sealed bids $b_1, \ldots, b_n$ from the players. It submits the bids $\sigma_1(b_1), \ldots, \sigma_n(b_n)$ to the mechanism M, and chooses the same outcome (e.g., winners of an auction and selling prices) that M does.

Mechanism M' is DSIC: If a participant i has private information v_i, then submitting a bid other than v_i can only result in M' playing a strategy other than $\sigma_i(v_i)$ in M, which can only decrease i's utility. □

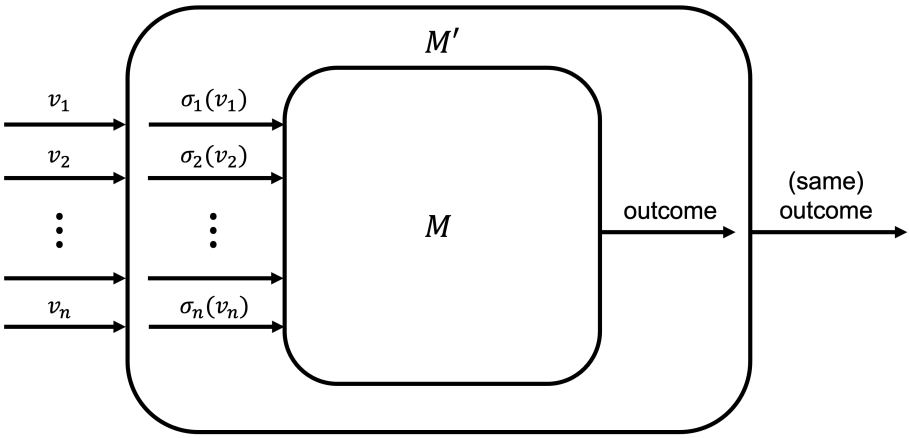

Figure 1: Proof of the Revelation Principle. Construction of the direct-revelation mechanism M', given a mechanism M with dominant strategies.

The takeaway from the Revelation Principle (Theorem 2) is that **it is without loss to design direct revelation mechanisms**. That is, you might as well require your mechanism to be **incentive-compatible**.

Acknowledgements

This lecture was developed in part using materials by Tim Roughgarden, and in particular, his book "Twenty Lectures on Algorithmic Game Theory" [**?**].