

Mechanism Design Basics (continued)

Recap from last time: Each bidder i has a valuation (maximum willingness-to-pay) v_i for an item, earns utility (happiness) $u_i(\cdot)$ depending on the outcome.

Definition 1. For a deterministic mechanism with at most one winner, a bidder with *quasilinear utility* has utility

$$u_i(\cdot) = \begin{cases} v_i - p_i & \text{if } i \text{ wins and pays } p_i \\ 0 & \text{otherwise.} \end{cases}$$

Mechanisms (1) receive bids, (2) determine who gets the good, and (3) decide on payments. The question is, how should they do (2) and (3)?

Last time we began to discuss the benefits of a second-price auction (highest bidder wins and pays the second-highest bid) over options like a first-price auction or all-pay auction: because bidding in a second-price auction is straightforward. We'll formalize that now.

Claim 1 (Dominant-Strategy Incentive-Compatibility). In a second-price auction, every bidder has a *dominant strategy*: set its bid b_i equal to its private valuation v_i . That is, this strategy maximizes the utility of bidder i , no matter what the other bidders do.

Proof. [Hint: Consider two cases of outcomes.]

Claim 2 (Individual Rationality). In a second-price auction, every truth-telling bidder is guaranteed non-negative utility.

Proof.

Theorem 1 (Vickrey). *The Vickrey (second-price) auction satisfies the following three quite different and desirable properties:*

- (1) [**strong incentive guarantees**] *It is dominant-strategy incentive-compatible (DSIC) and individually rational (IR), i.e., Claims 1 and 2 hold.*
- (2) [**strong performance guarantees**] *If bidders report truthfully, then the auction maximizes the social surplus*

$$\sum_{i=1}^n v_i x_i,$$

where x_i is 1 if i wins and 0 if i loses, subject to the obvious feasibility constraint that $\sum_{i=1}^n x_i \leq 1$ (i.e., there is only one item).

- (3) [**computational efficiency**] *The auction can be implemented in polynomial time.*

In general, as we design mechanisms, we'll take the following design approach:

- Step 1: Assume, without justification, that bidders bid truthfully. Then, how should we assign bidders to slots so that properties (2) strong performance guarantees and (3) computational efficiency hold?
- Step 2: Given our answer to Step 1, how should we set selling prices so that property (1) strong incentive guarantees holds?

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 2. 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 $\mathbf{x} = (x_1, \dots, x_n)$ 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 3. 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 4. 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 5. 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 6 (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 7 (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.

Give an example of a monotone allocation rule:

Give an example of a non-monotone allocation rule:

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

Theorem 2 (Myerson's Lemma [1]). *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 (see (4), below).*

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 6) seems confusing, but is actually equivalent to and just as easy as checking if the allocation is monotone (Definition 7).

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!

Proof of Myerson's Lemma (Theorem 2). 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

$$\underbrace{\hspace{10em}}_{\text{utility of bidding } z \text{ given value } z} \geq \underbrace{\hspace{10em}}_{\text{utility of bidding } y \text{ given value } z} \quad (1)$$

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

$$\underbrace{\hspace{10em}}_{\text{utility of bidding } y \text{ given value } y} \geq \underbrace{\hspace{10em}}_{\text{utility of bidding } z \text{ given value } y} \quad (2)$$

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

$$z \cdot [x(y) - x(z)] \leq p(y) - p(z) \leq y \cdot [x(y) - x(z)] \quad (3)$$

From here, we can conclude:

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

Why?

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

$$p_i(b_i, \mathbf{b}_{-i}) =$$

or alternatively, after integration by parts,

$$p_i(b_i, \mathbf{b}_{-i}) = \quad (4)$$

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

$$u_i(b_i, \mathbf{b}_{-i}) =$$

which is maximized when $b_i =$

independent of \mathbf{b}_{-i} , as desired. □

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:
- **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 j th 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 j th 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.

References

- [1] Roger B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.