

Mechanism Design Basics

Definition 1. Each bidder i has a private *valuation* v_i that is its maximum willingness-to-pay for the item being sold.

Our default assumption is that a bidder's utility is modeled by quasilinear utility.

Definition 2. For a deterministic mechanism with at most one winner, a bidder with *quasilinear utility* has utility

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

Definition 3. A *dominant strategy* is a strategy (bid) that is guaranteed to maximize a bidder's utility *no matter what* the other bidders do.

Sealed-Bid Auctions:

- (1) Each bidder i privately communicates a bid b_i to the auctioneer—in a sealed envelope, if you like.
- (2) The auctioneer decides who gets the item (if anyone).
- (3) The auctioneer decides on a selling price.

How should we do (2) and (3)? For now, (2) will just be giving the item to the highest bidder. What about (3)?

Some potential auctions:

- First-price auction: the price is equal to the highest bid.
- Second-price auction: the price is equal to the second-highest bid.
- All-pay auction: *every bidder* (not just the winning bidder) pays their bid.*

*Note that we need to amend our definition of quasilinear utility already for the all-pay auction, since we only defined payments in terms of when the bidder wins. For now, we can modify it to

$$u_i(\cdot) = v_i \cdot \mathbb{1}[i \text{ wins}] - p_i$$

where p_i is i 's assigned payment. In the next class, we'll further modify it.

How should we bid in these auctions? It's not necessarily clear in first-price or all-pay, but it *is* clear in the second-price auction with a bit of reasoning: just bid your true value!

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.

[Hint: Consider two cases of outcomes.]

This claim implies that second-price auctions are particularly easy to participate in—you don't need to reason about the other bidders in any way (how many there are, what their valuations, whether or not they bid truthfully, etc.) to figure out how you should bid. Note this is completely different from a first-price auction. You should never bid your valuation in a first-price auction (that would guarantee zero utility), and the ideal amount to underbid depends on the bids of the other players

Proof. Fix an arbitrary player i , its valuation v_i , and the bids \mathbf{b}_{-i} of the other players. (Here \mathbf{b}_{-i} means the vector b of all bids, but with the i th component deleted. It's wonky notation but you need to get used to it.) We need to show that bidder i 's utility is maximized by setting $b_i = v_i$. (Recall v_i is i 's fixed valuation, while it can set its bid b_i to whatever it wants.)

Let $B = \max_{j \neq i} b_j$ denote the highest bid by some other bidder. What's special about a second-price auction is that, even though there are an infinite number of bids that i could make, only distinct outcomes can result. If $b_i < B$, then i loses and receives utility 0. If $b_i \geq B$, then i wins at price B and receives utility $v_i - B$.

We now consider two cases. First, if $v_i < B$, the highest utility that bidder i can get is $\max\{0, v_i - B\} = 0$, and it achieves this by bidding truthfully (and losing). Second, if $v_i \geq B$, the highest utility that bidder i can get is $\max\{0, v_i - B\} = v_i - B$, and it achieves this by bidding truthfully (and winning). \square

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

Proof. Losers all get utility 0. If bidder i is the winner, then its utility is $v_i - p$, where p is the second-highest bid. Since i is winner (and hence the highest bidder) and bid its true valuation, $p \leq v_i$ and hence $v_i - p \geq 0$. \square

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 (indeed, linear) 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 4. 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 5. 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 6. 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 7. 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 8 (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 9 (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 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:*

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

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 8) seems confusing, but is actually equivalent to and just as easy as checking if the allocation is monotone (Definition 9).

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

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

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

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

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

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

From here, we can conclude:

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

Why?

Acknowledgements

This lecture was developed in part using materials by Tim Roughgarden, and in particular, his book “Twenty Lectures on Algorithmic Game Theory” [2].

References

- [1] Roger B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.
- [2] Tim Roughgarden. *Twenty lectures on algorithmic game theory*. Cambridge University Press, 2016.