DS 574 Algorithmic Mechanism Design	**Lecture #3**
Fall 2023	Prof. Kira Goldner

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 1. 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 1 (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 $b_1, \dots, b_n$ from the players. It submits the bids $\sigma_1(b_1), \dots, \sigma_n(b_n)$ 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. $\square$

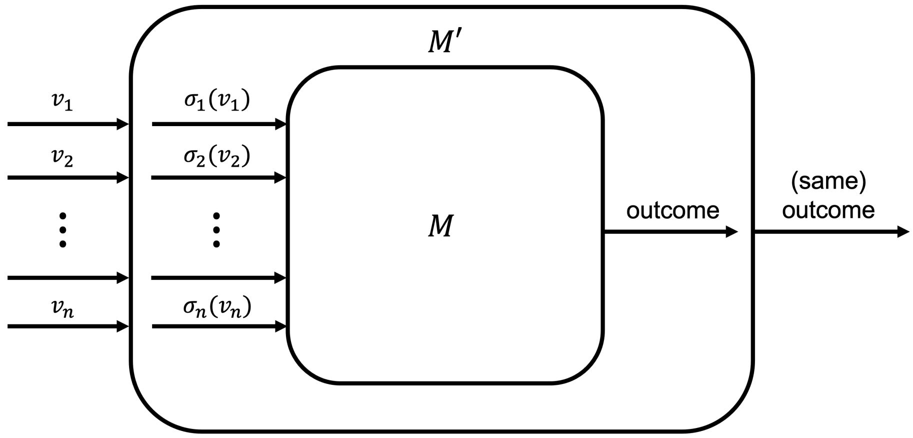

Figure 1: Proof of the Revelation Principle. Construction of the direct-revelation mechanism M', given a mechanism M with dominant strategies.

The takeaway from the Revelation Principle (Theorem 1) is that **it is without loss to design direct revelation mechanisms**. That is, you might as well require your mechanism to be **incentive-compatible**.

Beyond Dominant-Strategy: Bayesian Settings

There are many reasons why we can't always require dominant strategies when design mechanisms.

(1) Requiring such a strong concept might not be tractable.

(2) Agents do not always have dominant strategies! What then?

We'll now introduce the Bayesian setting.

Suppose the valuation v_i of bidder i is drawn from a prior distribution F_i.

- We abuse notation and let F_i denote the cumulative distribution function (CDF) of the distribution; that is, $F_i(x) = \Pr_{v_i \sim F_i}[v_i \leq x]$.
- We use $f_i(x)$ to denote the probability density function (pdf) of the distribution; that is, $f_i(x) = \frac{d}{dx} F_i(x)$.
- We use $\mathbf{F}$ or $\vec{F}$ to denote the *joint distribution* of the marginal buyer distributions F_i. That is, if the buyers are independently distributed, then $\mathbf{F}$ is the product distribution $\mathbf{F} = \times_i F_i$. (Note however that buyers are not necessarily independently distributed in all settings.)

Unless otherwise noted, we assume that the prior distribution $\mathbf{F}$ is *common knowledge* to all bidders and the mechanism designer (the seller).

Definition 2. A *Bayes-Nash equilibrium (BNE)* for a joint distribution $\mathbf{F}$ is a strategy profile $\sigma = (\sigma_1, \dots, \sigma_n)$ such that for all i and v, $\sigma_i(v_i)$ is a best-response when other agents play $\sigma_{-i}(\mathbf{v}_{-i})$ when $\mathbf{v}_{-i} \sim \mathbf{F}_{-i}\,|_{v_i}$.

Claim 1. Consider two identically and independently drawn bidders from $F = U[0,1]$. It is a (symmetric) BNE for each bidder to bid $\sigma_i(v_i) = v_i/2$ in the first-price auction.

Proof. Suppose bidder 2 is playing the strategy $\sigma_2(v_2) = v_2/2$. We prove that bidder 1's best-response is $\sigma_1(v_1) = v_1/2$. For now, call the bid given by $\sigma_1(v_1) = b_1$.

Given bidder 2's strategy, bidder 1's expected utility is

$$\begin{aligned}
\mathbb{E}_{v_2}[u_i(\sigma_1(v_1), \sigma_2(v_2))] &= v_1 \cdot \mathbb{E}_{v_2}[x_1(\sigma_1(v_1), \sigma_2(v_2))] - \mathbb{E}_{v_2}[p_1(\sigma_1(v_1), \sigma_2(v_2))] \\
&= v_1 \cdot \Pr_{v_2}[b_1 > v_2/2] - \mathbb{E}_{v_2}[b_1 \mid b_1 > v_2/2] && x, p \text{ in FPA} \\
&= \Pr_{v_2}[b_1 > v_2/2] \cdot [v_1 - b_1] && \text{def } \mathbb{E} \\
&= F(2b_1) \cdot [v_1 - b_1] && \text{def } F(\cdot) \\
&= (2b_1) \cdot [v_1 - b_1] && F = U[0,1] \\
&= 2b_1 v_1 - 2b_1^2 \\
\frac{d}{db_1}\mathbb{E}_{v_2}[u_i(b_1, v_2/2)] &= 2v_1 - 4b_1 && \text{differentiate to max} \\
\implies b_1 &= v_1/2
\end{aligned}$$

Hence bidder 1's best-response strategy is to bid $\sigma_1(v_1) = v_1/2$ in response to $\sigma_2(v_2) = v_2/2$, and thus these strategies are a BNE. □

Theorem 2 (Revenue Equivalence)**.** *The payment rule and revenue of a mechanism is uniquely determined by its allocation. Hence, any two mechanisms with the same allocation must earn the same revenue.*

What is this theorem a corollary of? Prove this for the first-price auction and the Vickrey (second-price) auction in the above setting!

Proof. This is just a corollary of Myerson's Lemma! As we pointed out, the *only* variables in the payment identity are the allocation rule! Payment is 100% determined by the allocation rule! Then two mechanisms with the same allocation *must* have the same payments.

Consider the first-price auction and the second-price auction each with two bidders i.i.d. from $U[0,1]$. Let V^1 and V^2 denote the random variables that are the highest and second-highest draws from $U[0,1]$, respectively. Note that two draws from the uniform distribution evenly divide the interval in expectation: $\mathbb{E}_{V^1,V^2 \sim U[0,1]}[V^1] = 2/3$ and $\mathbb{E}_{V^1,V^2 \sim U[0,1]}[V^2] = 1/3$.

In the first-price auction, the item is allocated to V^1 at a payment of its BNE bid of $V^1/2$. Then the expected winner's payment (and thus revenue) is $\frac{1}{2}\mathbb{E}[V^1] = 1/3$.

In the second-price auction, the item is allocated to V^1 at a payment of the second-highest bid $b^2 = V^2$, since Vickrey is DSIC. Then the expected winner's payment (and thus revenue) is $\mathbb{E}[V^2] = 1/3$. □

Bayesian Settings

Using notions from the Bayesian setting and how bidders Bayesian update as they learn information, we define three stages of the auction:

1. *ex ante*: Before any information has been drawn; i only knows $\mathbf{F}$.
2. *interim*: Values v_i have been drawn; i only knows their own valuation, and thus the updated prior $\mathbf{F} \mid_{v_i}$.
3. *ex post*: The auction has run and concluded. All bidders know all $v_1, \ldots, v_n$.

Typically we discuss the *ex post* allocation and payment rules as a function of all of the values. However, in the Bayesian setting, to reason about BIC, it often makes sense to take in terms of *interim* allocation and payment rules which have the same information as bidder i before the auction is run.

Definition 3. We define the *interim* allocation and payment rules in expectation over the updated Bayesian prior given i's valuation:

$$x_i(v_i) = \Pr_{\mathbf{F}}[x_i(\mathbf{v}) = 1 \mid v_i] = \mathbb{E}_{\mathbf{F}}[x_i(\mathbf{v}) \mid v_i]$$

and

$$p_i(v_i) = \mathbb{E}_{\mathbf{F}}[p_i(\mathbf{v}) \mid v_i].$$

Our definition of Bayesian Incentive-Compatibility then follows:

Definition 4. A mechanism with *interim* allocation rule x and *interim* payment rule p is Bayesian Incentive-Compatible (BIC) if

$$v_i x_i(v_i) - p_i(v_i) \geq v_i x_i(z) - p_i(z) \quad \forall i, v_i, z.$$

Exercise (optional):

- Extend Myerson's Lemma and the payment identity for Bayesian Incentive-Compatible (BIC) mechanisms.
- Extend the Revelation Principle for BIC mechanisms.

Acknowledgements

This lecture was developed in part using materials by Tim Roughgarden and Jason Hartline, and in particular, the books "Twenty Lectures on Algorithmic Game Theory" and "Mechanism Design and Approximation" [1, 2].

References

[1] Jason D. Hartline. Mechanism design and approximation. *Book draft. October*, 122, 2013.

[2] Tim Roughgarden. *Twenty lectures on algorithmic game theory*. Cambridge University Press, 2016.