| | |
|---|---|
| **DS 574 Algorithmic Mechanism Design** | **Lecture #3 Worksheet** |
| **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). Give an example:

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.*

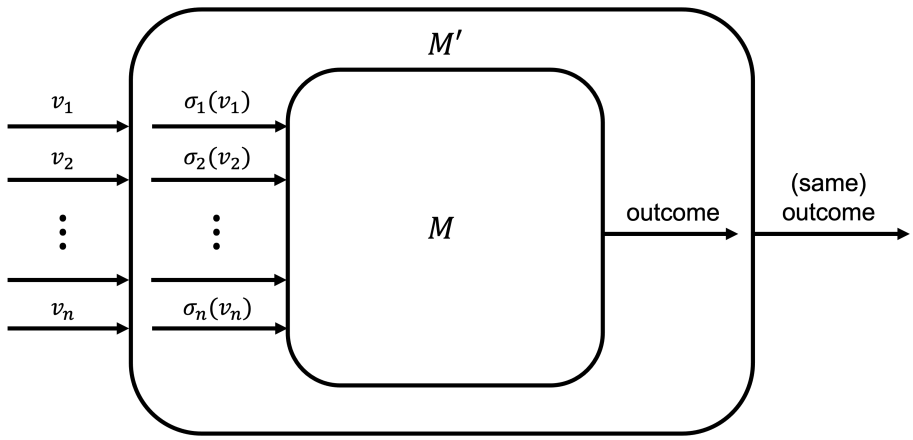


Figure 1: Proof of the Revelation Principle. Construction of the direct-revelation mechanism $M'$, given a mechanism $M$ with dominant strategies.

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 $F_i$.

- CDF $F_i(x) = \Pr_{v_i \sim F_i}[v_i \leq x]$.
- PDF $f_i(x) = \frac{d}{dx} F_i(x)$.
- Joint distribution $\mathbf{F}$ or $\vec{F}$.

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 $\sigma = (\sigma_1, \ldots, \sigma_n)$ 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 $\mathbf{v}_{-i} \sim \mathbf{F}_{-i} \mid_{v_i}$.

**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.*

**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.*

## 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

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

**Exercises** (optional):

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