

CSCI699: Topics in Learning and Game Theory
Lecture 4

Lecturer: Shaddin Dughmi

Scribes: Haifeng Xu, Reem Alfayez

1 Examples of Mechanism Design Problems

Example 1: Single Item Auctions.

There is a single item for sale and there are n players (i.e. bidders) competing for the item. Player i values the item at v_i , which is private data (a.k.a., private type) of player i . The outcome of the game is the choice of a winning player, and payment from each player. The utility of a player for an outcome is his value for the outcome minus his payment. The key thing here is that you as the principle get to design the rule of the game to yield desired equilibrium outcomes (e.g., maximizing welfare or the seller's revenue). This is why mechanism design is also sometimes called *reverse game theory*.

One example of mechanisms is the first price auction. That is, every bidder is asked to submit their bid for the item. The bidder with the highest bid wins the item and pays his bid. Another example is the second-price auction (a.k.a. Vickrey auction), in which every bidder is asked to submit their bid for the item, and the bidder with the highest bid wins the item and pays the second highest bid.

Example 2: Combinatorial Auctions.

This is a generalization of the previous single item auctions and considers the problem of selling a set T of m items. Player i 's private utility function is a mapping $v_i : 2^T \rightarrow \mathbb{R}$; $v_i(S)$ is player i 's value for the bundle $S \subseteq T$.

One goal for mechanism design in this setting is to maximize the welfare. Particularly, the goal is to partition ground set T into S_1, \dots, S_n to maximize $\sum_{i=1}^n v_i(S_i)$.

Example 3: Public Project

There are n players. Player i has private value v_i for building certain public project (e.g., a bridge). The outcome of the problem is the choice of whether or not to build the project and payment from each player covering the cost of the project if built. The

utility of a player for an outcome is his value for the project if built, minus his payment. The goal is, e.g., to build the project if sum of values exceeds its cost (i.e., maximizing welfare). Or the goal could also be to maximize revenue.

Example: Voting

Our final example considers voting. Here, there are n players and m candidates running for office. Each player has a total (or partial) preference order on the m candidates. The outcome of the problem is the choice of the winning candidate. The goal is usually to select the candidate that makes the most people “happy”, though the happiness here could be defined from various perspectives.

2 The General Mechanism Design Problem

The general mechanism design setting (prior free) is given by a tuple (N, \mathcal{X}, T, u) , where

- N is a finite set of players. Denote $N = \{1, \dots, n\}$ and $n = |N|$.
- \mathcal{X} is a set of outcomes.
- $T = T_1 \times \dots \times T_n$, where T_i is the set of types of player i . Each $\vec{t} = (t_1, \dots, t_n) \in T$ is called a type profile. The private type t_i encodes what player i knows but the rest of the world does not know.
- $u = (u_1, \dots, u_n)$, where $u_i : T_i \times \mathcal{X} \rightarrow \mathbb{R}$ is the utility function of player i .

The problem is in Bayesian setting if the above definition is supplement with a common prior belief \mathcal{D} over T .

As an example, the above definition instantiated in the single-item auction setting is as follows. The outcome is a choice $x \in \{e_1, \dots, e_n\}$ of the winning player and payment p_1, \dots, p_n from each player. Here, e_i is the i 'th basis vector (i.e. only the i 'th entry has value 1) denoting that player i wins the item. The private type of player i is her value $v_i \in \mathbb{R}_+$ and player i 's utility is $u_i(v_i, x) = v_i x_i - p_i$.

The above mechanism design setting can be viewed more abstractly as follows: a principal wants to communicate with players and aggregate their private data (types) into a choice of outcome. Such aggregation is captured by the *social choice function*.

Definition 1 (Social Choice Function). *A social choice function $f : T \rightarrow \mathcal{X}$ is a map from type profiles to outcomes.*

There are different rules for choosing a social choice function. The following two are typical:

- The principle may have a particular social choice function in mind (e.g. majority voting, utilitarian allocation of a single item, etc).
- The principle has an objective function $o : T \times X \rightarrow \mathbb{R}$, and want $f(T)$ to (approximately) maximize $o(T, f(T))$. To optimize the objective function in the worst case [in expectation] over T is called *prior free* [Bayesian] mechanism design.

For example, in single-item allocation problems, the objective could be the welfare $welfare(v, (x, p)) = \sum_i v_i x_i$ or the revenue $revenue(v, (x, p)) = \sum_i p_i$. Revenue is usually over expectation since point-wise revenue-optimal mechanism usually does not exist.

To perform the aggregation, the principal needs to run a protocol, which is called a *mechanism*.

Definition 2. A *mechanism* is a pair (A, g) , where

- $A = A_1 \times \dots \times A_n$, where A_i is the set of possible actions (e.g., bids in auction settings) of player i in the protocol. A is the set of action profiles.
- $g : A \rightarrow \mathcal{X}$ is an outcome function.

The resulting game of mechanism design is a game of incomplete information where when players play $a \in A$, player i 's utility is $u_i(t_i, g(a))$ when his private type is t_i .

As an example, in the first price auction, $A_i = \mathbb{R}$ contains all the possible bids that a player can submit and $g(b_1, \dots, b_n) = (x, p)$ where $x_{i^*} = 1, p_{i^*} = b_{i^*}$ for $i^* = \arg \max_i b_i$ and $x_i = 0, p_i = 0$ for all $i \neq i^*$.

We say a mechanism (A, g) implements the social choice function $f : T \rightarrow \mathcal{X}$ in dominant-strategy/Bayes-Nash equilibrium if: (1) there is a strategy profile $s = (s_1, \dots, s_n)$ with $s_i : T_i \rightarrow A_i$ such that $s_i : T_i \rightarrow A_i$ is a dominant-strategy/Bayes-Nash equilibrium in the resulting incomplete information game; (2) $g(s_1(t_1), s_2(t_2), \dots, s_n(t_n)) = f(t_1, t_2, \dots, t_n)$ for all $t \in T$

For example, in a single-item auction with two players whose valuation is drawn i.i.d. from $U[0, 1]$, the first price auction implements in BNE the following social choice function: give the item to the player with the highest value and charges him half his value (see the proof as a simple exercise in Lecture 3).

On the other hand, the Vickrey auction implements in DSE the following social choice function: give the item to the player with the highest value and charges him

the second highest value. This is because in Vickrey auction, the dominant bidding strategy for each player is to bid her true valuation.

The task of mechanism design is the follows (defined abstractly here, a more detailed version will be provided later): Given a notion of a “good” social choice function from T to \mathcal{X} , find: (a) a mechanism which consists of an action space $A = (A_1, \dots, A_n)$, an outcome function $g : A \rightarrow \mathcal{X}$; (b) an equilibrium (s_1, \dots, s_n) of the resulting game of the mechanism, with the goal of implementing a “good” social choice function $f(t_1, \dots, t_n) = g(s_1(t_1), \dots, s_n(t_n))$.

This task seems like a complicated, multivariate search problem. Luckily, the *revelation principle* reduces the search space to just $g : T \rightarrow \mathcal{X}$.

3 The Revelation Principle and Incentive Compatibility

Definition 3 (Direct Revelation). *A mechanism (A, g) is a direct revelation mechanism if $A_i = T_i$ for all i .*

That is, in a direct revelation mechanism, players simultaneously report types (not necessarily truthfully) to the mechanism. The outcome function in such mechanisms can simply be described via the function $g : T \rightarrow \mathcal{X}$.

Definition 4 (Incentive-Compatibility). *A direct-revelation mechanism is dominant-strategy/Bayesian incentive-compatible (a.k.a truthful) if the truth-telling, i.e., reporting the true type, is a dominant-strategy/Bayes-Nash equilibrium in the resulting incomplete-information game.*

Note that a direct revelation incentive-compatible mechanism implements its outcome function $g : T \rightarrow \mathcal{X}$, by definition. Therefore, the social choice function is the mechanism!!

The following are a few illustrative examples.

- **Vickrey Auction** is a direct revelation dominant-strategy incentive-compatible mechanism.
- **First Price Auction** is a direct revelation mechanism but is not Bayesian incentive compatible (since player will bid half of their true valuation at equilibrium).

- **Posted price auction** is the auction that simply posts a fixed price to players in sequence until one accepts the price. This auction format is not direct revelation, thus incentive compatibility is also not well-defined in this case.

[Revelation Principle] *If there is a mechanism implementing social choice function f in dominant-strategy/Bayes-Nash equilibrium, then there is a direct revelation, dominant-strategy/Bayesian incentive-compatible mechanism implementing f .*

This simplifies the task of mechanism design.

The (simplified) task of mechanism design is the follows: Given a notion of a “good” social choice function from T to \mathcal{X} , find such a social choice function $f : T \rightarrow \mathcal{X}$ such that truth-telling is an equilibrium in the following mechanism: solicit reports $\tilde{t}_i \in T_i$ from each player i (simultaneous, sealed bid) and then choose outcome $f(\tilde{t}_1, \dots, \tilde{t}_n)$.

Let us consider again the single-item auction with 2 players whose valuation is drawn i.i.d from $U[0, 1]$. In the first-price auction, a principle solicits bids b_1, b_2 , and then gives the item to the highest bidder and charges him his bid. Recall that the strategies where each player reports half their value forms a Bayesian Nash equilibrium (BNE). In other words, when player 1 knows his value v_1 , and faces player 2 who is bidding uniformly from $[0, 1/2]$, he maximizes his expected utility $(v_1 - b_1) \cdot 2b_1$ by bidding $b_1 = v_1/2$. And vice versa. Therefore, the first price auction implements in BNE the social choice function which gives the item to the highest bidder, and charges him half his bid.

This shows that first price auction is not an incentive-compatible mechanism. Nevertheless, we can modify it to become an incentive-compatible mechanism. Intuitively, we will integrate the player’s strategic behavior into the mechanism and lie “on behalf of the player”, so that the player does not need to lie any more.

Definition 5 (modified first-price Auction for two bidders, i.i.d. $U(0,1)$ valuations).

1. *Solicit bids b_1, b_2 .*
2. *Give the item to the highest bidder, charging him half his bid (equivalently, simulate a first price auction where bidders bid $b_1/2, b_2/2$).*

Claim 6. *Truth-telling is a BNE in the modified first-price auction.*

Proof. Assume player 2 bids truthfully. Player 1 faces a (simulated) first price auction where his own bid is halved before participating, and player 2 bids uniformly from $[0, 1/2]$. To respond optimally in the simulation, he bids $b_1 = v_1$ and lets the mechanism halve his bid on his behalf. \square

Therefore, the modified first price auction implements the same social-choice function in equilibrium, but is truthful.

Proof of the Revelation Principle in Bayesian Settings

Consider any mechanism (A, g) , with BNE strategies $s_i : T_i \rightarrow A_i$, that implements social choice function $f(t_1, \dots, t_n) = g(s_1(t_1), \dots, s_n(t_n))$ in BNE. For all i and t_i , action $s_i(t_i)$ maximizes player i 's expected utility when other players are playing $s_{-i}(t_{-i})$ for $t_{-i} \sim \mathcal{D}|t_i$.

Now we consider the following modified mechanism: (1) Solicit reported types $\tilde{t}_1, \dots, \tilde{t}_n$; (2) Choose outcome $f(\tilde{t}_1, \dots, \tilde{t}_n) = g(s_1(\tilde{t}_1), \dots, s_n(\tilde{t}_n))$. The second step of the mechanism is equivalently to simulating (A, g) when players play $s_i(t_i)$. Assume all players other than i report truthfully. Then when i 's type is t_i , other players play $s_{-i}(t_{-i})$ for $t_{-i} \sim \mathcal{D}|t_i$ in the simulated mechanism. As stated above, his best response in simulation is $s_i(t_i)$. The mechanism transforms his bid by applying s_i , so player i 's best response is to bid his true type t_i .

Note, the proof for Dominant strategy equilibrium case is similar.

4 Mechanisms with Money: The Quasilinear Utility Model

To make much of modern mechanism design possible, we assume that

1. The set of outcomes has a particular structure: every outcome includes a payment to or from each player
2. Player utilities vary linearly with their payment.

Examples include single-item allocation, public project. Notice that single-item allocation without money and voting do *not* fall in this setting.

Definition 7 (The Quasi-linear Setting). *Formally, $\mathcal{X} = \Omega \times \mathbb{R}^n$.*

- Ω is the set of allocations
- For $(\omega, p_1, \dots, p_n) \in \mathcal{X}$, p_i is the payment from (or to if $p_i < 0$) player i .

and player i 's utility function $u_i : T_i \times \mathcal{X} \rightarrow \mathbb{R}$ takes the form $u_i(t_i, (\omega, p_1, \dots, p_n)) = v_i(t_i, \omega) - p_i$ for some valuation function v_i . In this case, we say players have quasilinear utilities.

For example, in single item allocation problems, $\Omega = \{e_1, \dots, e_n\}$ and $u_i(t_i, (\omega, p_1, \dots, p_n)) = t_i w_i - p_i$.

Using the revelation principle, we can further simplify the task of mechanism design in quasilinear settings, as follows.

Definition 8 (Mechanism Design in Quasilinear Settings). *Find a “good” allocation rule $f : T \rightarrow \Omega$ and payment rule $p : T \rightarrow \mathbb{R}^n$ such that the following mechanism is incentive-compatible:*

1. Solicit reports $\tilde{t}_i \in T_i$ from each player i (simultaneous, sealed bid);
2. Choose allocation $f(\tilde{t})$;
3. Charge player i payment $p_i(\tilde{t})$.

We will think of the mechanism as the pair (f, p) . Sometimes, we abuse notation and think of type t_i directly as the valuation $v_i : \Omega \rightarrow \mathbb{R}$. In this case incentive compatibility can be simplified.

Definition 9. *[Incentive-compatibility (dominant strategy, quasilinear settings)] A mechanism (f, p) is dominant-strategy truthful if, for every player i , true type t_i , possible mis-report \tilde{t}_i , and reported types t_{-i} of the others, we have*

$$v_i(t_i, f(t)) - p_i(t) \geq v_i(t_i, f(\tilde{t}_i, t_{-i})) - p_i(\tilde{t}_i, t_{-i}).$$

If (f, p) has randomness, add expectation sign to both sides of the above inequality.

The notion for Bayesian Nash equilibrium is defined similarly.

Definition 10. *[Incentive-compatibility (Bayesian, quasilinear settings)] A mechanism (f, p) is dominant-strategy truthful if, for every player i , true type t_i , possible mis-report \tilde{t}_i , the following holds in expectation over $t_{-i} \sim \mathcal{D}|t_i$:*

$$\mathbb{E} \left[v_i(t_i, f(t)) - p_i(t) \right] \geq \mathbb{E} \left[v_i(t_i, f(\tilde{t}_i, t_{-i})) - p_i(\tilde{t}_i, t_{-i}) \right].$$

The following are some examples.

1. **Vickrey auction:** the allocation rule maps b_1, \dots, b_n to e_{i^*} for $i^* = \arg \max_i b_i$ and the payment rule maps b_1, \dots, b_n to p_1, \dots, p_n where $p_{i^*} = b_{(2)}$, and $p_i = 0$ for $i \neq i^*$. This auction is dominant-strategy truthful.

2. **First Price Auction:** the allocation rule maps b_1, \dots, b_n to e_{i^*} for $i^* = \arg \max_i b_i$ and the payment rule maps b_1, \dots, b_n to p_1, \dots, p_n where $p_{i^*} = b_{(1)}$, and $p_i = 0$ for $i \neq i^*$. For two players with valuations drawn i.i.d from $U[0, 1]$, players bidding half their value is a BNE. This auction is not Bayesian incentive compatible.

Modified First Price Auction: the allocation rule maps b_1, \dots, b_n to e_{i^*} for $i^* = \arg \max_i b_i$ and the payment rule maps b_1, \dots, b_n to p_1, \dots, p_n where $p_{i^*} = b_{(1)}/2$, and $p_i = 0$ for $i \neq i^*$. For two players with valuations drawn i.i.d from $U[0, 1]$, this auction is Bayesian incentive compatible.

5 Maximizing Welfare: The VCG Mechanism

In quasilinear setting, a simple mechanism is DSE and maximizes the social welfare P

$$\sum_i v_i(\omega) \quad (1)$$

Vickrey Clarke Groves (VCG) Mechanism:

- Solicit type v_i from each player i
- Choose allocation

$$\omega^* \in \operatorname{argmax}_{\omega \in \Omega} \sum_i v_i(\omega) \quad (2)$$

- Charge each player i payment

$$p_i(v) = h_i(v_{-i}) - \sum_{j \neq i} v_j(\omega^*) \quad (3)$$

Allocation rule maximizes welfare exactly over Ω . Player i is paid the reported value of others for the chosen allocation, less a pivot term $h_i(v_{-i})$ independent of his own bid, and in most cases the "right" pivot term is $\max_{\omega \in \Omega} \sum_{j \neq i} v_j(\omega)$. Payment $p_i(v)$ is player i 's is externality: $0 \leq p_i(v) \leq v_i(\omega^*)$

Theorem: VCG is dominant-strategy truthful.

Proof:

- Fix reports v_{-i} of players other than i .
- Assume player i 's true valuation is v_i

- Player i 's utility when reporting \hat{v}_i is given by:

$$u_i(\hat{v}_i) = v_i(\omega^*) + \sum_{j \neq i} v_j(\omega^*) - h_i(v_{-i})$$

$$\text{where } \omega^* \in \operatorname{argmax}_{\omega \in \Omega} (\hat{v}_i(\omega) + \sum_{j \neq i} v_j(\omega))$$

- Since the pivot term is independent of player i 's bid, maximizing $u_i(\hat{v}_i)$ is equivalent to maximizing: $v_i(\omega^*) + \sum_{j \neq i} v_j(\omega^*)$
- Setting $\hat{v}_i = v_i$ then maximizes the above expression.

Interpretation: allow the mechanism to optimize player i 's utility on his behalf.

Example: Single-item Allocation:

- Welfare maximizing outcome: Allocate to player with highest value
- Externality of i : second-highest value if i wins, 0 otherwise.

VCG is the second-price (Vickrey) auction in the special case of single-item allocation.

6 Maximizing Revenue

Maximizing revenue is well understood in the case of single-parameter problems:

Single-parameter problem (informally):

- There is a single homogenous resource.
- Constraints on how much of the resource each player can get.
- Each player's type is his "value (or cost) per unit resource."

Canonical example: single-item allocation:

- Resource: one unit of item.
- Outcomes Ω : vectors (x_1, \dots, x_n) with $x_i \geq 0$ and $\sum_i x_i \leq 1$ where x_i is probability player i gets item
- Player i 's type is $v_i \geq 0$ (value for item)

$$u_i(x, p) = v_i x_i - p_i$$

Makes most sense in Bayesian setting with independent types (prior $\mathbb{F} = \mathbb{F}_1 \times \dots \times \mathbb{F}_n$ on (v_1, \dots, v_n))

Bayesian Revenue Maximization (Single Parameter):

Given prior \mathbb{F} on type profiles $\mathbb{T} \subseteq \mathbb{R}^n$, find allocation rule $x : T \rightarrow \Omega$ (recall $\Omega \subseteq R^n$) and payment rules $p : T \rightarrow \subseteq R^n$ such that

- (x, p) is a BIC direct revelation mechanism Bidding $b_i = v_i$ maximizes $\mathbb{E}_{v_i \sim \mathbb{F}_{-i}} [v_i x_i(b_i, v_{-i}) - p_i(b_i, v_{-i})]$
- $Rev(x, p) = \mathbb{E}_{v \sim \mathbb{F}} \sum_i p_i(v)$ is as large as possible.

Myerson characterized the optimal solution for single-item auctions, and it generalizes easily to single-parameter environments¹.

Stages of a Bayesian Game

Stages of a Bayesian game of mechanism design:

- Ex-ante: Before players learn their types
- Interim: A player learns his type, but not the types of others.
- Ex-post All player types are revealed.

Interim stage is when players make decisions.

- The interim allocation rule for player i tells us what the probability of winning (expected amount of resource) is as a function of player i 's bid, in expectation over other player's truthful reports. $\bar{x}_i(b_i) = \mathbb{E}_{v_{-i} \sim \mathbb{F}_{-i}} [x_i(b_i, v_{-i})]$
- Similarly, the interim payment rule. $\bar{p}_i(b_i) = \mathbb{E}_{v_{-i} \sim \mathbb{F}_{-i}} [p_i(b_i, v_{-i})]$
- BIC: Bidding $b_i = v_i$ maximizes $v_i \bar{x}_i(b_i) - \bar{p}_i(b_i)$

If BIC, then $Rev(x, p) = \sum_i \mathbb{E}_{v_i \sim \mathbb{F}_i} \bar{p}(v_i)$

Assume two players drawn independently from U [0, 1].

Vickrey Auction

$$\bar{x}_i(v_i) = v_i$$

$$\bar{p}_i(v_i) = v_i^2 / 2$$

First Price Auction

$$\bar{x}_i(v_i) = v_i$$

$$\bar{p}_i(v_i) = v_i^2 / 2$$


From now on we will write $x_i(b_i) = \bar{x}_i b_i$ to avoid cumbersome notation

¹Think of single-item auctions in upcoming discussion

Myerson's Monotonicity Lemma:

Consider a mechanism for a single-parameter problem in a Bayesian setting where player values are independent. A direct-revelation mechanism with interim allocation rule x and payment rule p is BIC if and only if for each player i :

- $x_i(b_i)$ is a monotone non-decreasing function of b_i
- $p_i(b_i)$ is an integral of $b_i dx_i$. Specifically, when $p_i(0) = 0$ then $p_i(b_i) = \int_{b=0}^{b_i} x_i(b) db$.

**Interpretation of Myerson's Monotonicity Lemma:**

The higher a player bids, the higher the probability of winning. For each additional sliver ϵ of winning probability, pays at a rate equal to the minimum bid needed to acquire that sliver. Recall: second price auction².

Corollaries of Myerson's Monotonicity Lemma:

Corollaries:

- Interim allocation rule uniquely determines interim payment rule.
- Expected revenue depends only on the allocation rule

Theorem (Revenue Equivalence):

Any two auctions with the same interim allocation rule in BNE have the same expected revenue in the same BNE.

Revenue as Virtual Welfare:

Define the virtual value of player i as a function of his value v_i

²See readings for proof of Myerson's monotonicity Lemma

$$\phi_i(v_i) = v_i - \frac{1 - F_i(v_i)}{f_i(v_i)}$$


Myerson's Virtual Welfare Lemma:

Consider a BIC mechanism M with interim allocation rule x and payment rule p , and assume that $p_i(0) = 0$ for all i . The expected revenue of M is equal to the expected virtual welfare served.

$$\sum_i \mathbb{E}_{v_i \sim F_i} [\phi(v_i)x(v_i)]$$

In single-item auction, this is the expected virtual value of the winning bidder.

Proof



$$\begin{aligned} \mathbb{E}_{v \sim F_i} [p_i(v)] &= \int_v \left[vx_i(v) - \int_{b=0}^v x_i(b) db \right] f_i(v) dv \\ &= \int_v vx_i(v) f_i(v) dv - \int_v \int_{b \leq v} x_i(b) f_i(v) db dv \\ &= \int_v vx_i(v) f_i(v) dv - \int_b x_i(b) \int_{v \geq b} f_i(v) dv db \\ &= \int_v vx_i(v) f_i(v) dv - \int_b x_i(b) (1 - F_i(b)) db \\ &= \int_v [vx_i(v) f_i(v) - x_i(v) (1 - F_i(v))] dv \\ &= \int_v f_i(v) x_i(v) \left[v - \frac{1 - F_i(v)}{f_i(v)} \right] dv = \int_v f_i(v) x_i(v) \phi_i(v) dv \end{aligned}$$

Myerson's Revenue-Optimal Auction:

- Solicit player values.
- If at least one player has nonnegative virtual value, then give the item to the player i with the highest virtual value $\phi_i(v_i) \geq 0$. Otherwise, nobody gets the item.
- Charge the minimum bid needed to win $\phi^{-1}(\max(0, (\max_{j \neq i} \phi_j(v_j))))$.
Check: satisfies Myerson's condition on interim payment.

Observations

- The allocation rule maximizes virtual welfare point-wise.
- Therefore, it maximizes expected virtual welfare over all allocation rules.
- By Myerson's virtual welfare Lemma, revenue is at least that of any BIC mechanism (since any BIC mechanism's revenue is equal to expected virtual welfare).

A Wrinkle:

If the allocation rule of the mechanism we just defined is non-monotone, it would still have revenue at least that of the optimal BIC mechanism if players happened to report truthfully, but it wouldn't be truthful itself. Fortunately, the virtual welfare maximization is monotone when the distributions are regular: $\phi_i(v) = v - \frac{1-F_i(v)}{f_i(v)}$ is nondecreasing in v . We can conclude that when distributions are regular, the VV maximizing auction (aka Myerson's optimal auction) is the revenue-optimal BIC mechanism. Most natural distributions are regular (Gaussian, uniform, exp, etc). Additionally, it can be extended to non-regular distributions via ironing.

Thoughts:

Myerson's optimal auction is noteworthy for many reasons:

- Matches practical experience: when players i.i.d regular, optimal auction is Vickrey with reserve price $\phi^{-1}(0)$.
- Applies to single parameter problems more generally
- Revenue maximization reduces to welfare maximization for these problems
- The optimal BIC mechanism just so happens to be DSIC and deterministic!!