CS599: Algorithm Design in Strategic Settings
Fall 2012
Lecture 10: Introduction to Bayesian Mechanism Design

Instructor: Shaddin Dughmi

Administrivia

- HW2 Due
- Projects
  - Email me topic choice, and paper list
  - Schedule additional meeting to discuss.

Outline

Outline

Recall: Mechanism Design Problem in Quasi-linear Settings

Public (common knowledge) inputs describes

- Set Ω of allocations.
- Typespace T_i for each player i.
  - $T = T_1 \times T_2 \times \ldots \times T_n$
- Valuation map $v_i : T_i \times \Omega \to \mathbb{R}$

Recall: Mechanism Design Problem in Quasi-linear Settings

Public (common knowledge) inputs describes

- Set Ω of allocations.
- Typespace T_i for each player i.
  - $T = T_1 \times T_2 \times \ldots \times T_n$
- Valuation map $v_i : T_i \times \Omega \to \mathbb{R}$

Bayesian Setting

Supplement with a prior distribution D on T.

Incentive-Compatibility

Incentive-compatibility (Dominant Strategy)

A mechanism (f, p) is dominant-strategy truthful if, for every player i, true type t_i, possible mis-report $\widetilde{t_i}$, and reported types t_{-i} of the others, we have

$$\mathbf{E}[v_i(t_i, f(t)) - p_i(t)] \geq \mathbf{E}[v_i(t_i, f(\widetilde{t_i}, t_{-i})) - p_i(\widetilde{t_i}, t_{-i})]$$

where the expectation is over random coins of the mechanism.

Incentive-Compatibility

Incentive-compatibility (Dominant Strategy)

A mechanism (f, p) is dominant-strategy truthful if, for every player i, true type t_i, possible mis-report $\widetilde{t_i}$, and reported types t_{-i} of the others, we have

$$\mathbf{E}[v_i(t_i, f(t)) - p_i(t)] \geq \mathbf{E}[v_i(t_i, f(\widetilde{t_i}, t_{-i})) - p_i(\widetilde{t_i}, t_{-i})]$$

where the expectation is over random coins of the mechanism.

Incentive-compatibility (Bayesian)

A mechanism (f, p) is Bayesian incentive compatible if, for every player i, true type t_i, possible mis-report $\widetilde{t_i}$, the following holds

where the xpectation is over random coins of the mechanism as well as $t_{-i} \sim D|t_i$

Examples

Vickrey Auction

- Allocation rule maps $b_1, \ldots, b_n$ to e_{i^*} for $i^* = \mathrm{argmax}_i\, b_i$
- Payment rule maps $b_1, \ldots, b_n$ to $p_1, \ldots, p_n$ where $p_{i^*} = b_{(2)}$, and $p_i = 0$ for $i \neq i^*$.

Dominant-strategy truthful.

First Price Auction

- Allocation rule maps $b_1, \ldots, b_n$ to e_{i^*} for $i^* = \mathrm{argmax}_i\, b_i$
- Payment rule maps $b_1, \ldots, b_n$ to $p_1, \ldots, p_n$ where $p_{i^*} = b_{(1)}$, and $p_i = 0$ for $i \neq i^*$.

For two players i.i.d $U[0, 1]$, players bidding half their value is a BNE.
Not Bayesian incentive compatible.

Examples

Modified First Price Auction

- Allocation rule maps $b_1, \ldots, b_n$ to e_{i^*} for $i^* = \operatorname{argmax}_i b_i$
- Payment rule maps $b_1, \ldots, b_n$ to $p_1, \ldots, p_n$ where $p_{i^*} = b_{(1)}/2$, and $p_i = 0$ for $i \neq i^*$.

For two players i.i.d $U[0, 1]$, Bayesian incentive compatible.

Bayesian vs Worst case

A priori, Bayesian AMD seems easier than prior-free

- Expand space of mechanisms: BIC weaker guarantee than IC
- Relax to average case guarantees: e.g. a mechanism that α-approximates welfare in expectation may be easier than worst-case
- Provides unambiguous notion of “the best algorithm/mechanism”, since inputs are weighted. Serves as a benchmark.

Bayesian vs Worst case

So What does it Buy us?

- Today: Non-trivial mechanisms for new objectives that were (arguably) hopeless in prior-free (like revenue).
- Tomorrow: Enables better polytime BIC approximate mechanisms for welfare (and other objectives)

Disadvantages of relaxing to BIC / average case guarantees

May be non-robust to discrepancies between the environment for which it was designed, and that in which it is deployed (overfitting)

- Bayesian Incentive Compatibility contingent on prior and common knowledge assumption.
- Average case approximation gurantee hinges on prior

Today

- We begin examining mechanism design in Bayesian settings, like we did in prior-free settings. We focus on additional design power afforded.
- First, we look at mechanisms that optimize revenue in single parameter settings.
  - Mechanisms with worst-case guarantees on revenue are not possible in prior-free settings (at least for uncontroversial benchmarks).
- Today: Myerson's revenue-optimal single item auction (2007 Nobel Prize)
- Later lectures: Revenue/Welfare in NP-hard single-parameter problems, multi-parameter problems.

Single-parameter Problems

Informally

- There is a single homogenous resource (items, bandwidth, clicks, spots in a knapsack, etc).
- There are constraints on how the resource may be divided up.
- Each player's private data is his "value (or cost) per unit resource."

Single-parameter Problems

Informally

- There is a single homogenous resource (items, bandwidth, clicks, spots in a knapsack, etc).
- There are constraints on how the resource may be divided up.
- Each player's private data is his "value (or cost) per unit resource."

Formally

- Set Ω of allocations is common knowledge.
- Each player i's type is a single real number t_i. Player i's type-space T_i is an interval in $\mathbb{R}$.
- Each allocation $x \in \Omega$ is a vector in $\mathbb{R}^n$.
- A player's utility for allocation x and payment p_i is $t_i x_i - p_i$.
- Bayesian assumption: Common prior D on T

Recall: Single-item Allocation

- Allocations: choice of player who wins the item
  - $\Omega = \{e_1, \ldots, e_n\}$
- Type: private value $v_i \in \mathbb{R}_+$ for the item. Typespace T_i is $\mathbb{R}_+$ or some closed interval in $\mathbb{R}_+$.
- For $x \in \Omega$ and $p \in \mathbb{R}^n_+$, utility is $u_i(x) = v_i x_i - p_i$

Why a Prior?

- For social welfare, input-by-input optimum achievable via a truthful mechanism (Vickrey)
  - Uncontroversial benchmark, matched in the worst case.
- For revenue, no longer the case.
  - Consider the analogous input-by-input optimum as a benchmark: give item to highest bidder and charge him his bid.
  - No incentive compatible mechanism achieves a constant factor approximation for every such input.
    - Easiest to see: deterministic. Must be posted price take-it-or-leave-it offer.
- With priors, can do better.
  - Single player, uniform $[0, 1]$
  - Posting a price of $1/2$ gets revenue $1/4$ in expectation, which is half the expected welfare.

The Prior

We make several assumptions on the prior distribution of player types to simplify/obtain results

- Player types drawn independently.
  - Let F_i denote the c.d.f of player i's value for the item.
  - Let f_i denote p.d.f, and $S_i = 1 - F_i$.
  - Let $F = F_1 \times \ldots \times F_n$ denote the distribution over type profiles.
- Assume $f_i(v) > 0$ for $v \in T_i$.

Outline

1. Bayesian Mechanism Design
2. Optimal Deterministic Single-Player Single-Item Auction
3. Reducing Revenue Maximization to Welfare Maximization
4. Myerson's Revenue-Optimal Auction

Optimal Single-player Deterministic Auction

- In order to build intuition, we examine the single player case
- For a single player, BIC = DSIC
- Recall: A mechanism is DSIC if its allocation rule is monotone
- For a deterministic mechanism, this is a posted price mechanism.

Optimal Single-player Deterministic Auction

- In order to build intuition, we examine the single player case
- For a single player, BIC = DSIC
- Recall: A mechanism is DSIC if its allocation rule is monotone
- For a deterministic mechanism, this is a posted price mechanism.

Question

Find the revenue maximizing posted price for a player with value drawn from $U([0,1])$. How about $U([1,2])$? How about $Exp(1)$?

Optimal Single-player Deterministic Auction

- In order to build intuition, we examine the single player case
- For a single player, BIC = DSIC
- Recall: A mechanism is DSIC if its allocation rule is monotone
- For a deterministic mechanism, this is a posted price mechanism.

Question

Find the revenue maximizing posted price for a player with value drawn from $U([0,1])$. How about $U([1,2])$? How about $Exp(1)$?

More generally, for a distribution F, Find price v maximizing $vS(v)$.

Quantiles

We will perform a convenient change of variables.

Definition

Fix a c.d.f F with $S = 1 - F$. We define the quantile of v in the support of F as

$$q(v) = S(v).$$

Quantiles

We will perform a convenient change of variables.

Definition

Fix a c.d.f F with $S = 1 - F$. We define the quantile of v in the support of F as

$$q(v) = S(v).$$

Observations

- Examples: $U([0,1])$, $Exp(1)$
- The quantile of v is the probability of sale when we post price v.
- The quantile of v, for $v \sim F$, is always uniformly distributed in $[0,1]$.

Quantiles

We will perform a convenient change of variables.

Definition

Fix a c.d.f F with $S = 1 - F$. We define the quantile of v in the support of F as

$$q(v) = S(v).$$

Observations

- Examples: $U([0,1])$, $Exp(1)$
- The quantile of v is the probability of sale when we post price v.
- The quantile of v, for $v \sim F$, is always uniformly distributed in $[0,1]$.

- For mathematical convenience, we will parametrize valuations by their quantiles, as we will see next.
- For notational convenience, we also use $v(q)$ to denote the value v with quantile q. Note that $v(q) = S^{-1}(q)$.

Revenue Curves

Definition

Fix a c.d.f F. The revenue curve $R(.)$ specifices the posted-price revenue as a function of probability of sale (i.e. quantile). Specifically, $R(q) = v(q) \cdot q$.

- For $U[0,1]$ it is $q(1-q)$
- For $Exp(1)$ it is $-q \ln q$.

Revenue Curves

Definition

Fix a c.d.f F. The revenue curve $R(.)$ specifices the posted-price revenue as a function of probability of sale (i.e. quantile). Specifically, $R(q) = v(q) \cdot q$.

- For $U[0,1]$ it is $q(1-q)$
- For $Exp(1)$ it is $-q \ln q$.

- We can find the optimal sale price / sale probability by finding the maximum of R.
- In the above examples, since the curves are concave it suffices to sell at the price corresponding to the point where R has zero derivative.

Marginal Revenue and Virtual Value

Definition

The Marginal Revenue at q is $R'(q)$. Specifically, this is the rate of increase of revenue as a function of probability of sale.

$$R'(q) = \frac{d}{dq}(v(q) \cdot q) = v(q) - \frac{q}{f(v(q))}$$

In other words: $R'(q)dq$ is the additional revenue generated by lowering the price so as to sell to dq additional customers in expectation.

Marginal Revenue and Virtual Value

Definition

The Marginal Revenue at q is $R'(q)$. Specifically, this is the rate of increase of revenue as a function of probability of sale.

$$R'(q) = \frac{d}{dq}(v(q) \cdot q) = v(q) - \frac{q}{f(v(q))}$$

In other words: $R'(q)dq$ is the additional revenue generated by lowering the price so as to sell to dq additional customers in expectation.

Definition

The virtual value $\phi(v)$ of a player with value v at quantile q is $R'(q)$, or equivalently:

$$\phi(v) = v - \frac{S(v)}{f(v)}$$

Interpretation when Revenue is Concave

Observe

- When revenue curve is concave, optimal auction lowers the posted price so long as marginal revenue at the price is nonnegative.
- Equivalently: Allocation rule awards item to player so long as his virtual value is positive, and then uses the threshold payment rule suggested by myseron's lemma!

Interpretation when Revenue is Concave

Observe

- When revenue curve is concave, optimal auction lowers the posted price so long as marginal revenue at the price is nonnegative.
- Equivalently: Allocation rule awards item to player so long as his virtual value is positive, and then uses the threshold payment rule suggested by myseron's lemma!

- Because of truthfulness, the posted price is uniquely determined by the allocation rule.
- The allocation rule inducing the optimal mechanism is the one that sells to the player if and only if his virtual value is nonnegative.

Interpretation when Revenue is Concave

Observe

- When revenue curve is concave, optimal auction lowers the posted price so long as marginal revenue at the price is nonnegative.
- Equivalently: Allocation rule awards item to player so long as his virtual value is positive, and then uses the threshold payment rule suggested by myseron's lemma!

- Because of truthfulness, the posted price is uniquely determined by the allocation rule.
- The allocation rule inducing the optimal mechanism is the one that sells to the player if and only if his virtual value is nonnegative.

Upshot

The allocation rule of the revenue maximizing single-player, single item auction is the one that maximizes virtual welfare!

Regular Distributions

Definition

- A distribution is regular if the corresponding revenue curve $R(q)$ is concave.
- Equivalently, if $R'(q)$ is monotone non-increasing.
- Equivalently, if $\phi(v)$ is monotone non-decreasing.

Regular Distributions

Definition

- A distribution is regular if the corresponding revenue curve $R(q)$ is concave.
- Equivalently, if $R'(q)$ is monotone non-increasing.
- Equivalently, if $\phi(v)$ is monotone non-decreasing.

We restrict our attention to regular distributions in this lecture, as they guarantee that virtual welfare maximization is monotone. Moreover, they include most natural distributions: uniform, normal, exponential, and more...

Outline

Coming Up

We generalize the intuition from the previous section. We consider a single-item allocation setting where players' values are drawn from independent regular distributions.

Lemma (Myerson's Virtual Surplus Lemma)

Let $M = (\mathcal{A}, p)$ be a BIC mechanism where a player bidding zero pays nothing in expectation. The expected revenue of M is equal to the expected virtual welfare served by $\mathcal{A}$.

Theorem

The revenue optimal BIC mechanism for selling a single item is that which, on each valuation profile, awards the item to the player with the highest nonnegative virtual value, and discards the item if all virtual values are negative.

Stages of a Bayesian Game

For terminology, it will be helpful to formalize the “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.

Of particular interest to us is the interim stage, because it is the stage when players make decisions.

- The interim allocation rule for player i is a function $x_i(v_i)$ of player i’s type, evaluating to the probability (in equilibrium) of player i receiving the item in expectation over draws of other players’ types and the randomness of the mechanism.
- Similarly, the interim payment rule.

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

Vickrey Auction

- $x_i(v_i) = v_i$
- $p_i(v_i) = v_i/2$.

First Price Auction

- $x_i(v_i) = v_i$
- $p_i(v_i) = v_i/2$

Recall: Myerson's Monotonicity Lemma (Dominant Strategy)

A mechanism (x, p) for a single-parameter problem is dominant-strategy truthful if and only if for every player i and fixed reports b_{-i} of other players,

- $x_i(b_i)$ is a monotone non-decreasing function of b_i
- $p_i(b_i)$ is an integral of $b_i \, dx_i$.

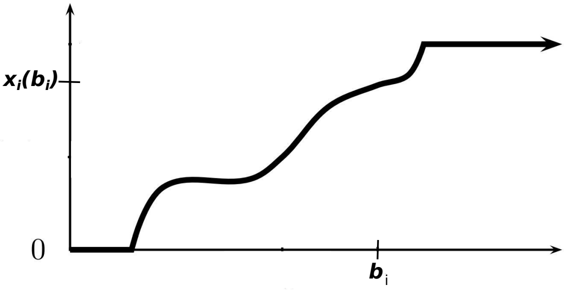

Recall: Myerson's Monotonicity Lemma (Dominant Strategy)

A mechanism (x, p) for a single-parameter problem is dominant-strategy truthful if and only if for every player i and fixed reports b_{-i} of other players,

- $x_i(b_i)$ is a monotone non-decreasing function of b_i
- $p_i(b_i)$ is an integral of $b_i \, dx_i$.

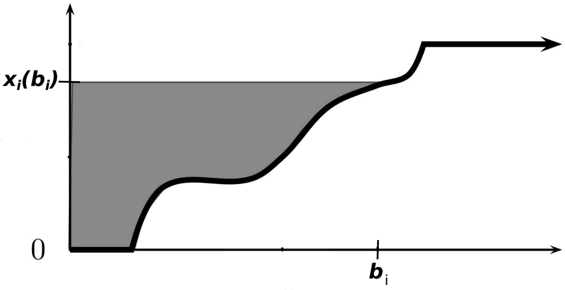

Recall: Myerson's Monotonicity Lemma (Dominant Strategy)

A mechanism (x, p) for a single-parameter problem is dominant-strategy truthful if and only if for every player i and fixed reports b_{-i} of other players,

- $x_i(b_i)$ is a monotone non-decreasing function of b_i
- $p_i(b_i)$ is an integral of $b_i \, dx_i$.

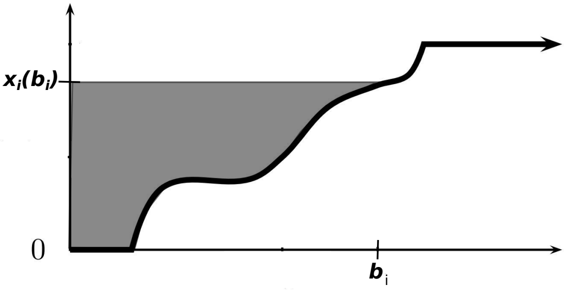

The mention of many players, and a dominant strategy, is a red herring.

Myerson's Monotonicity Lemma (Single Player)

Consider a 1 player game (i.e. decision problem) of incomplete information. The player has type $v \in \mathbb{R}$, action set $b \in \mathbb{R}$, and utility function $vx(b) - p(b)$ for some allocation rule x and payment rule p. Truth-telling is a best response (i.e. best decision) iff

- $x(b)$ is a monotone non-decreasing function of b
- $p(b)$ is an integral of $b\, dx$.

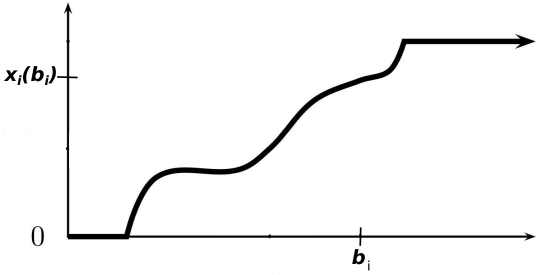

Myerson's Monotonicity Lemma (Single Player)

Consider a 1 player game (i.e. decision problem) of incomplete information. The player has type $v \in \mathbb{R}$, action set $b \in \mathbb{R}$, and utility function $vx(b) - p(b)$ for some allocation rule x and payment rule p. Truth-telling is a best response (i.e. best decision) iff

- $x(b)$ is a monotone non-decreasing function of b
- $p(b)$ is an integral of $b\, dx$.

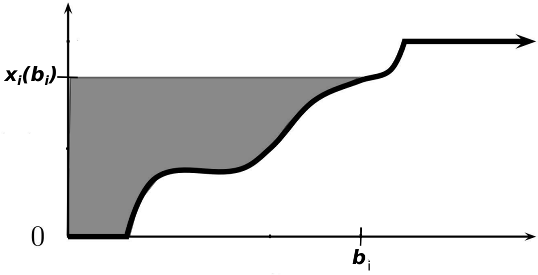

Myerson's Monotonicity Lemma (Single Player)

Consider a 1 player game (i.e. decision problem) of incomplete information. The player has type $v \in \mathbb{R}$, action set $b \in \mathbb{R}$, and utility function $vx(b) - p(b)$ for some allocation rule x and payment rule p. Truth-telling is a best response (i.e. best decision) iff

- $x(b)$ is a monotone non-decreasing function of b
- $p(b)$ is an integral of $b\, dx$.

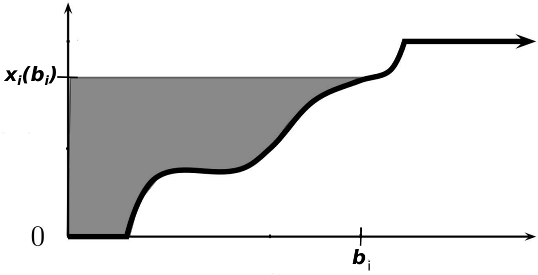

Need x to be independent of v for this to hold

Myerson's Monotonicity Lemma (BIC)

Consider a mechanism for a single-parameter problem in a Bayesian setting where player values are independent. Let $x_i(b_i)$ and $p_i(b_i)$ be the interim allocatin/payment rules faced by player i when other players play the truth-telling strategy. The mechanism is BIC if and only if:

- $x_i(b_i)$ is a monotone non-decreasing function of b_i
- $p_i(b_i)$ is an integral of $b_i \, dx_i$.

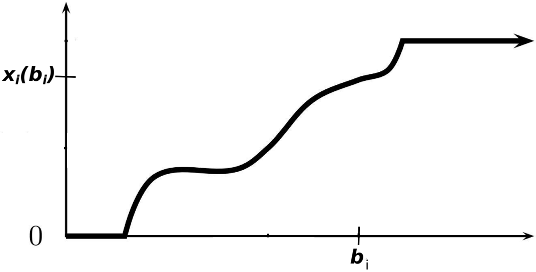

Myerson's Monotonicity Lemma (BIC)

Consider a mechanism for a single-parameter problem in a Bayesian setting where player values are independent. Let $x_i(b_i)$ and $p_i(b_i)$ be the interim allocatin/payment rules faced by player i when other players play the truth-telling strategy. The mechanism is BIC if and only if:

- $x_i(b_i)$ is a monotone non-decreasing function of b_i
- $p_i(b_i)$ is an integral of $b_i \; dx_i$.

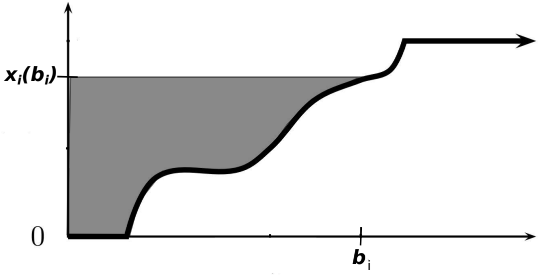

Myerson's Monotonicity Lemma (BIC)

Consider a mechanism for a single-parameter problem in a Bayesian setting where player values are independent. Let $x_i(b_i)$ and $p_i(b_i)$ be the interim allocatin/payment rules faced by player i when other players play the truth-telling strategy. The mechanism is BIC if and only if:

- $x_i(b_i)$ is a monotone non-decreasing function of b_i
- $p_i(b_i)$ is an integral of $b_i\ dx_i$.

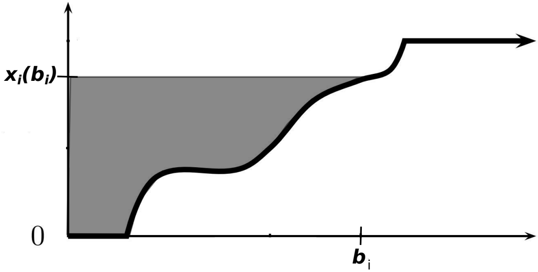

Needed independence of types so $x_i(b_i)$ does not depend on the player i's type.

Monotonicity Lemma for Quantiles

Let x_i and p_i be a function of the quantile of the player's report rather than the report itself.

Myerson's Monotonicity Lemma (BIC)

Consider a mechanism for a single-parameter problem in a Bayesian setting where player values are independent. Let $x_i(q_i)$ and $p_i(q_i)$ be the interim allocation/payment rules faced by player i when other players play the truth-telling strategy. The mechanism is BIC if and only if:

- $x_i(q_i)$ is a monotone non-increasing function of q_i
- $p_i(q_i)$ is an integral of $v_i(q_i)dx_i = v_i(q_i)x_i'(q_i)dq_i$. Doing the integration:

$$p_i(q_i) = p_i(1) - \int_{r=q_i}^{1} v_i(r)x_i'(r)dr$$

Corollaries of Myerson's Monotonicity Lemma

Corollaries

- The Interim allocation rule uniquely determins the 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: Myerson's Virtual Surplus Lemma

Lemma (Myerson's Virtual Surplus Lemma)

Let $M = (\mathcal{A}, p)$ be a BIC mechanism where a player bidding zero pays nothing in expectation. The expected revenue of M is equal to the expected virtual welfare served by $\mathcal{A}$.

Proof

We take the expected payment of player i.

$$\mathop{\mathbf{E}}_{q_i}[p_i(q_i)] = -\int_{q_i=0}^{1}\int_{r=q_i}^{1} v_i(r)x_i'(r)drdq_i$$

$$\cdots$$

$$= \int_{q_i=0}^{1} R_i'(q_i)x_i(q_i)dq_i$$

$$= \int_{v_i} \phi_i(v_i)x_i(v_i)f_i(v_i)dv_i$$

Outline

Myerson's Optimal Auction

1. Solicit player values
2. Give the item to the player i with the highest non-negative virtual value $\phi_i(v_i)$
3. Charge the corresponding critical payment: $\phi_i^{-1}(\max(0, (\max_{j \neq i} \phi_j(v_j))))$

Myerson's Optimal Auction

1. Solicit player values
2. Give the item to the player i with the highest non-negative virtual value $\phi_i(v_i)$
3. Charge the corresponding critical payment: $\phi_i^{-1}(\max(0, (\max_{j \neq i} \phi_j(v_j))))$

Observations

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

Myerson's Optimal Auction

1. Solicit player values
2. Give the item to the player i with the highest non-negative virtual value $\phi_i(v_i)$
3. Charge the corresponding critical payment: $\phi_i^{-1}(\max(0, (\max_{j \neq i} \phi_j(v_j))))$

Observations

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

Are we done?

A Wrinkle

Not really... What 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

A Wrinkle

Not really... What 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

Virtual welfare maximization is monotone when the distributions are regular!!

Regularity

- We know that welfare maximization is monotone in value
- Similarly, virtual welfare maximization is monotone in virtual value, which in turn is monotone in value when the distributions are regular!

Regularity

- We know that welfare maximization is monotone in value
- Similarly, virtual welfare maximization is monotone in virtual value, which in turn is monotone in value when the distributions are regular!

Conclude

When distributions are regular, the VV maximizing auction (aka Myerson's optimal auction) is the revenue-optimal BIC mechanism!

Regularity

- We know that welfare maximization is monotone in value
- Similarly, virtual welfare maximization is monotone in virtual value, which in turn is monotone in value when the distributions are regular!

Conclude

When distributions are regular, the VV maximizing auction (aka Myerson's optimal auction) is the revenue-optimal BIC mechanism!

Regularity is a mild assumption: Includes uniform, gaussian, exponential, ...

Thoughts

Myerson's optimal auction is noteworth 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 (next lecture)
- Revenue maximization reduces to welfare maximization for these problems
- The optimal BIC mechanism just so happens to be DSIC and deterministic!!

Next time

- Beyond regularity (Ironing)
- Beyond single item
- Approximation of revenue and welfare when welfare maximization (eq revenue maximization) is NP-hard