Algorithmic Game Theory
Introduction to Mechanism Design

Makis Arsenis

National Technical University of Athens

April 2016

Outline

1. Social Choice
   - Social Choice Theory
   - Voting Rules
   - Incentives
   - Impossibility Theorems

2. Mechanism Design
   - Single-item Auctions
   - The revelation principle
   - Single-parameter environment
   - Welfare maximization and VCG
   - Revenue maximization

Social Choice

Social Choice Theory

- Mathematical **theory** dealing with aggregation of **preferences**.
- Founded by Condorcet, Borda (1700's) and Dodgson (1800's).
- Axiomatic framework and impossibility result by Arrow (1951).
- Collective decision making, by **voting**, over **anything**:
  - Political representatives, award nominees, contest winners, allocation of tasks/resources, joint plans, meetings, food, ...
  - Web-page ranking, preferences in multi-agent systems.

Formal Setting

- Set $A, |A| = m$, of possible **alternatives** (candidates).
- Set $N = \{1, 2, \ldots, n\}$ of **agents** (voters).
- $\forall$ agent i has a (private) **linear order** $\succ_i \in L$ over alternatives A.

Social Choice

Social Choice Theory

- Mathematical **theory** dealing with aggregation of **preferences**.
- Founded by Condorcet, Borda (1700's) and Dodgson (1800's).
- Axiomatic framework and impossibility result by Arrow (1951).
- Collective decision making, by **voting**, over **anything**:
  - Political representatives, award nominees, contest winners, allocation of tasks/resources, joint plans, meetings, food, ...
  - Web-page ranking, preferences in multi-agent systems.

Formal Setting

- Set A, $|A| = m$, of possible **alternatives** (candidates).
- Set $N = \{1, 2, \ldots, n\}$ of **agents** (voters).
- $\forall$ agent i has a (private) **linear order** $\succ_i \in L$ over alternatives A.

Social Choice

Formal Setting

- **Social choice function** (or **mechanism**) $F : L^n \to A$ mapping the agent's preferences to an alternative.
- **Social welfare function** $W : L^n \to L$ mapping the agent's preferences to a total order on the alternatives.

Social Choice

Example (Colors of the local football club)

Preferences of the founders about the colors of the local club:

- 12 boys: Green $\succ$ Red $\succ$ Blue
- 10 boys: Red $\succ$ Green $\succ$ Blue
- 3 girls:Blue $\succ$ Red $\succ$ Green

Voting Rule allocating **(2, 1, 0)**.

Outcome: Red(35) $\succ$ Green(34) $\succ$ Blue(6).

With **plurality** voting **(1, 0, 0)**: Green(12) $\succ$ Red(10) $\succ$ Blue(3).

Which voting rule should we use?
Is there a notion of a “perfect” rule?

Social Choice

Example (Colors of the local football club)

Preferences of the founders about the colors of the local club:

- 12 boys: Green $\succ$ Red $\succ$ Blue
- 10 boys: Red $\succ$ Green $\succ$ Blue
- 3 girls:Blue $\succ$ Red $\succ$ Green

Voting Rule allocating **(2, 1, 0)**.

Outcome: Red(35) $\succ$ Green(34) $\succ$ Blue(6).

With **plurality** voting **(1, 0, 0)**: Green(12) $\succ$ Red(10) $\succ$ Blue(3).

Which voting rule should we use?
Is there a notion of a "perfect" rule?

Example (Colors of the local football club)

Preferences of the founders about the colors of the local club:

- 12 boys: Green $\succ$ Red $\succ$ Blue
- 10 boys: Red $\succ$ Green $\succ$ Blue
- 3 girls:Blue $\succ$ Red $\succ$ Green

Voting Rule allocating **(2, 1, 0)**.

Outcome: Red(35) $\succ$ Green(34) $\succ$ Blue(6).

With **plurality** voting **(1, 0, 0)**: Green(12) $\succ$ Red(10) $\succ$ Blue(3).

Which voting rule should we use?
Is there a notion of a "perfect" rule?

Social Choice

Example (Colors of the local football club)

Preferences of the founders about the colors of the local club:

- 12 boys: Green $\succ$ Red $\succ$ Blue
- 10 boys: Red $\succ$ Green $\succ$ Blue
- 3 girls:Blue $\succ$ Red $\succ$ Green

Voting Rule allocating **(2, 1, 0)**.

Outcome: Red(35) $\succ$ Green(34) $\succ$ Blue(6).

With **plurality** voting **(1, 0, 0)**: Green(12) $\succ$ Red(10) $\succ$ Blue(3).

Which voting rule should we use?
Is there a notion of a "perfect" rule?

Social Choice

Definition (Condorcet Winner)

Condorcet Winner is the alternative **beating every other** alternative in **pairwise election**.

Example (continued ...)

- 12 boys: Green $\succ$ Red $\succ$ Blue
- 10 boys: Red $\succ$ Green $\succ$ Blue
- 3 girls:Blue $\succ$ Red $\succ$ Green

(Green, Red) : (12, 13), (Green, Blue) : (22, 3), (Red, Blue) : (22, 3)
Therefore: Red is a Condorcet Winner!

Condorcet Paradox: Condorcet Winner may **not exist**:

- $a \succ b \succ c$
- $b \succ c \succ a$
- $c \succ a \succ b$

$(a, b) : (2, 1), (a, c) : (1, 2), (b, c) : (2, 1)$

Social Choice

Definition (Condorcet Winner)

Condorcet Winner is the alternative **beating every other** alternative in **pairwise election**.

Example (continued ...)

- 12 boys: Green $\succ$ Red $\succ$ Blue
- 10 boys: Red $\succ$ Green $\succ$ Blue
- 3 girls:Blue $\succ$ Red $\succ$ Green

(Green, Red) : $(12, 13)$, (Green, Blue) : $(22, 3)$, (Red, Blue) : $(22, 3)$
Therefore: Red is a Condorcet Winner!

Condorcet Paradox: Condorcet Winner may **not exist**:

- $a \succ b \succ c$
- $b \succ c \succ a$
- $c \succ a \succ b$

$(a, b) : (2, 1), (a, c) : (1, 2), (b, c) : (2, 1)$

Social Choice

Popular **Voting Rules**:

- **Plurality voting**: Each voter casts a single vote. The candidate with the most votes is selected.
- **Cumulative voting**: Each voter is given k votes, which can be cast arbitrarily.
- **Approval voting**: Each voter can cast a single vote for as many of the candidates as he/she wishes.
- **Plurality with elimination**: Each voter casts a single vote for their most-preferable candidate. The candidate with the fewer votes is eliminated etc.. until a single candidate remains.
- **Borda Count**: Positional Scoring Rule $(m-1, m-2, \dots, 0)$. (chooses a *Condorcet winner* if one exists).

Example (continued ...)

- 12 boys: Green $\succ$ Red $\succ$ Blue
- 10 boys: Red $\succ$ Green $\succ$ Blue
- 3 girls:Blue $\succ$ Red $\succ$ Green

Voting Rule allocating **(2, 1, 0)**.
Expected Outcome: Red(35) $\succ$ Green(34) $\succ$ Blue(6).

What really happens:

- 12 boys: Green $\succ$ Blue $\succ$ Red
- 10 boys: Red $\succ$ Blue $\succ$ Green
- 3 girls:Blue $\succ$ Red $\succ$ Green

Outcome: Blue(28) $\succ$ Green(24) $\succ$ Red(23).

Incentives

Example (continued ...)

- 12 boys: Green $\succ$ Red $\succ$ Blue
- 10 boys: Red $\succ$ Green $\succ$ Blue
- 3 girls:Blue $\succ$ Red $\succ$ Green

Voting Rule allocating **(2, 1, 0)**.
Expected Outcome: Red(35) $\succ$ Green(34) $\succ$ Blue(6).

What really happens:

- 12 boys: Green $\succ$ Blue $\succ$ Red
- 10 boys: Red $\succ$ Blue $\succ$ Green
- 3 girls:Blue $\succ$ Red $\succ$ Green

Outcome: Blue(28) $\succ$ Green(24) $\succ$ Red(23).

Arrow's Impossibility Theorem

Desirable Properties of Social Welfare Functions

- **Unanimity**: $\forall \succ \in L : W(\succ, \ldots, \succ) = \succ$.
- **Non dictatorial**: An agent $i \in N$ is a dictator if:

$$\forall \succ_1, \ldots, \succ_n \in L : W(\succ_1, \ldots, \succ_n) = \succ_i$$

- **Independence of irrelevant alternatives (IIA)**:
  $\forall a, b \in A$,
  $\forall \succ_1, \ldots, \succ_n, \succ'_1, \ldots, \succ'_n \in L$,
  if we denote $\succ = W(\succ_1, \ldots, \succ_n), \succ' = W(\succ'_1, \ldots, \succ'_n)$ then:

$$(\forall i \ a \succ_i b \Leftrightarrow a \succ'_i b) \Rightarrow (a \succ b \Leftrightarrow a \succ' b)$$

Theorem (Arrow, 1951)

If $|A| \geq 3$, any social welfare function W that satisfies unanimity and independence of irrelevant alternatives is dictatorial.

Arrow's Impossibility Theorem

Desirable Properties of Social Welfare Functions

- **Unanimity**: $\forall \succ \in L : W(\succ, \dots, \succ) = \succ$.
- **Non dictatorial**: An agent $i \in N$ is a dictator if:

$$\forall \succ_1, \dots, \succ_n \in L : W(\succ_1, \dots, \succ_n) = \succ_i$$

- **Independence of irrelevant alternatives (IIA)**:
  $\forall a, b \in A$,
  $\forall \succ_1, \dots, \succ_n, \succ'_1, \dots, \succ'_n \in L$,
  if we denote $\succ = W(\succ_1, \dots, \succ_n), \succ' = W(\succ'_1, \dots, \succ'_n)$ then:

$$(\forall i \; a \succ_i b \Leftrightarrow a \succ'_i b) \Rightarrow (a \succ b \Leftrightarrow a \succ' b)$$

Theorem (Arrow, 1951)

If $|A| \geq 3$, any social welfare function W that satisfies unanimity and independence of irrelevant alternatives is dictatorial.

Muller-Satterthwaite Impossibility Theorem

Desirable Properties of Social **Choice** Functions

- **Weak Pareto efficiency**: For all preference profiles:

$$(\forall i : a \succ_i b) \Leftrightarrow F(\succ_1, \ldots, \succ_n) \neq b$$

- **Non dictatorial**: For each agent $i, \exists \succ_1, \ldots, \succ_n \in L$:

$$F(\succ_1, \ldots, \succ_n) \neq \text{ agent's } i \text{ top alternative}$$

- **Monotonicity**:
  $\forall a, b \in A$,
  $\forall \succ_1, \ldots, \succ_n, \succ'_1, \ldots, \succ'_n \in L$ such that $F(\succ_1, \ldots, \succ_n) = a$,
  if $(\forall i : a \succ_i b \Leftrightarrow a \succ'_i b)$ then $F(\succ'_1, \ldots, \succ'_n) = a$.

Theorem (Muller-Satterthwaite, 1977)

If $|A| \geq 3$, any social choice function F that is weakly Pareto efficient and monotonic is dictatorial.

Muller-Satterthwaite Impossibility Theorem

Desirable Properties of Social **Choice** Functions

- **Weak Pareto efficiency**: For all preference profiles:

$$(\forall i : a \succ_i b) \Leftrightarrow F(\succ_1, \ldots, \succ_n) \neq b$$

- **Non dictatorial**: For each agent $i, \exists \succ_1, \ldots, \succ_n \in L$:

$$F(\succ_1, \ldots, \succ_n) \neq \text{ agent's } i \text{ top alternative}$$

- **Monotonicity**:
  $\forall a, b \in A$,
  $\forall \succ_1, \ldots, \succ_n, \succ'_1, \ldots, \succ'_n \in L$ such that $F(\succ_1, \ldots, \succ_n) = a$,
  if $(\forall i : a \succ_i b \Leftrightarrow a \succ'_i b)$ then $F(\succ'_1, \ldots, \succ'_n) = a$.

Theorem (Muller-Satterthwaite, 1977)

If $|A| \geq 3$, any social choice function F that is weakly Pareto efficient and monotonic is dictatorial.

Gibbard-Satterthwaite Theorem

Definition (Truthfulnes)

A social choice function F can be **strategically manipulated** by voter i if for some $\succ_1, \ldots, \succ_n, \in L$ and some $\succ'_i \in L$ we have:

$$F(\succ_1, \ldots, \succ'_i, \ldots, \succ_n) \succ_i F(\succ_1, \ldots, \succ_i, \ldots, \succ_n)$$

A social choice function that *cannot* be *strategically manipulated* is called **incentive compatible** or **truthful** or **strategyproof**.

Definition (Onto)

A social choice function F is said to be **onto** a set A if for every $a \in A$ there exist $\succ_1, \ldots, \succ_n \in L$ such that $F(\succ_1, \ldots, \succ_n) = a$.

Theorem (Gibbard 1973, Satterthwaite 1975)

Let F be a **truthful** *social choice function onto A, where $|A| \geq 3$, then F is a dictatorship.*

Gibbard-Satterthwaite Theorem

Definition (Truthfulnes)

A social choice function F can be **strategically manipulated** by voter i if for some $\succ_1, \ldots, \succ_n, \in L$ and some $\succ_i' \in L$ we have:

$$F(\succ_1, \ldots, \succ_i', \ldots, \succ_n) \succ_i F(\succ_1, \ldots, \succ_i, \ldots, \succ_n)$$

A social choice function that *cannot* be *strategically manipulated* is called **incentive compatible** or **truthful** or **strategyproof**.

Definition (Onto)

A social choice function F is said to be **onto** a set A if for every $a \in A$ there exist $\succ_1, \ldots, \succ_n \in L$ such that $F(\succ_1, \ldots, \succ_n) = a$.

Theorem (Gibbard 1973, Satterthwaite 1975)

Let F be a **truthful** *social choice function onto A, where $|A| \geq 3$, then F is a dictatorship.*

Gibbard-Satterthwaite Theorem

Definition (Truthfulnes)

A social choice function F can be **strategically manipulated** by voter i if for some $\succ_1, \ldots, \succ_n, \in L$ and some $\succ'_i \in L$ we have:

$$F(\succ_1, \ldots, \succ'_i, \ldots, \succ_n) \succ_i F(\succ_1, \ldots, \succ_i, \ldots, \succ_n)$$

A social choice function that *cannot* be *strategically manipulated* is called **incentive compatible** or **truthful** or **strategyproof**.

Definition (Onto)

A social choice function F is said to be **onto** a set A if for every $a \in A$ there exist $\succ_1, \ldots, \succ_n \in L$ such that $F(\succ_1, \ldots, \succ_n) = a$.

Theorem (Gibbard 1973, Satterthwaite 1975)

Let F be a **truthful** *social choice function onto A, where $|A| \geq 3$, then F is a dictatorship.*

Gibbard-Satterthwaite Theorem

Escape Routes

- Randomization
- Monetary Payments
- Voting systems **Computationally Hard** to manipulate
- Restricted domain of preferences.
  - Approximation
  - Verification
  - ...

Outline

1. Social Choice
   - Social Choice Theory
   - Voting Rules
   - Incentives
   - Impossibility Theorems
2. Mechanism Design
   - Single-item Auctions
   - The revelation principle
   - Single-parameter environment
   - Welfare maximization and VCG
   - Revenue maximization

Example problem: Single-item Auctions

Sealed-bid Auction Format

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

Mechanism: Defines how we implement steps (2), and (3).

Mechanisms with Money

More formally:

Redefining our model

- Set Ω, $|\Omega| = m$, of possible **outcomes**.
- Set $N = \{1, 2, \ldots, n\}$ of **agents** (players).
- **Valuation vector** $\mathbf{v} = (v_1, \ldots, v_n) \in V$ where $v_i : \Omega \to \mathbb{R}$ is the (private) **valuation function** of each player.

Mechanism

- **Outcome function**: $f : V^n \to \Omega$
- **Payment vector**: $\mathbf{p} = (p_1, \ldots, p_n)$ where $p_i : V^n \to \mathbb{R}$.

Players have **quasilinear utilities**. For $\omega \in \Omega$, player i tries to maximize her utility $u_i(\omega) = v_i(\omega) - p$ where p is the monetary payment the player makes.

Mechanisms with Money

Possible objectives:

- Design **truthful** mechanisms that maximize the **Social Welfare**.
- Design **truthful** mechanisms that maximize the expected **revenue** of the seller.

Definition (Truthful)

A mechanism is **truthful** if for every agent i it is a *dominant strategy* to report her true valuation irrespective of the valuations of the other players.

Social Welfare: $\mathrm{SW}(\omega) = \sum_{i=1}^{n} v_i(\omega)$.

Revenue: $\mathrm{REV}(\mathbf{v}) = \sum_{i=1}^{n} p_i(\mathbf{v})$.

Single-item auctions

First price auction ?

- Give the item to the **highest bidder**.
- Charge him **its bid**.

Drawbacks

Hard to reason about:

- Hard to figure out (as a **participant**) how to bid.
- As a **seller** or auction designer, it's hard to predict what will happen.

First price auction ?

- Give the item to the **highest bidder**.
- Charge him **its bid**.

Drawbacks

Hard to reason about:

- Hard to figure out (as a **participant**) how to bid.
- As a **seller** or auction designer, it's hard to predict what will happen.

Single-item auctions

Second price auction

- Give the item to the **highest bidder**.
- Charge him the bid of the **second highest** bidder.

Theorem

The second price auction is **truthful**.

Proof.

Fix a player i, its valuation v_i and the bids $\mathbf{b}_{-i}$ of all the other players.
We need to show that u_i is maximized when $b_i = v_i$.
Let $B = \max_{j \neq i} b_j$

- if $b_i < B$: player i loses the item and $u_i = 0$.
- if $b_i > B$: player i wins the item at price B and $u_i = v_i - B$.
  - if $v_i < B$ then player i has negative utility.
  - if $v_i \geq B$ then he would also win the item even if she reported $b_i = v_i$ and she would have the same utility.

□

Single-item auctions

Second price auction

- Give the item to the **highest bidder**.
- Charge him the bid of the **second highest** bidder.

Theorem

The second price auction is **truthful**.

Proof.

Fix a player i, its valuation v_i and the bids $\mathbf{b}_{-i}$ of all the other players.
We need to show that u_i is maximized when $b_i = v_i$.
Let $B = \max_{j \neq i} b_j$

- if $b_i < B$: player i loses the item and $u_i = 0$.
- if $b_i > B$: player i wins the item at price B and $u_i = v_i - B$.
  - if $v_i < B$ then player i has negative utility.
  - if $v_i \geq B$ then he would also win the item even if she reported $b_i = v_i$ and she would have the same utility.

□

Single-item auctions

Second price auction

- Give the item to the **highest bidder**.
- Charge him the bid of the **second highest** bidder.

Theorem

The second price auction is **truthful**.

Proof.

Fix a player i, its valuation v_i and the bids $\mathbf{b}_{-i}$ of all the other players.
We need to show that u_i is maximized when $b_i = v_i$.
Let $B = \max_{j \neq i} b_j$

- if $b_i < B$: player i loses the item and $u_i = 0$.
- if $b_i > B$: player i wins the item at price B and $u_i = v_i - B$.
  - if $v_i < B$ then player i has negative utility.
  - if $v_i \geq B$ then he would also win the item even if she reported $b_i = v_i$ and she would have the same utility.

□

Single-item auctions

Some desirable characteristics of the second-price auction:

- **Strong incentive guarantees**: **truthful** and **individually rational** (every player has non-negative utility).
- **Strong performance guarantees**: the auction maximizes the **social welfare**.
- **Computational efficiency**: The auction can be implemented in **polynomial** (indeed linear) time.

Single-item auctions

Some desirable characteristics of the second-price auction:

- **Strong incentive guarantees**: **truthful** and **individually rational** (every player has non-negative utility).
- **Strong performance guarantees**: the auction maximizes the **social welfare**.
- **Computational efficiency**: The auction can be implemented in **polynomial** (indeed linear) time.

Single-item auctions

Some desirable characteristics of the second-price auction:

- **Strong incentive guarantees**: **truthful** and **individually rational** (every player has non-negative utility).
- **Strong performance guarantees**: the auction maximizes the **social welfare**.
- **Computational efficiency**: The auction can be implemented in **polynomial** (indeed linear) time.

Revelation Principle

Revisiting truthfulness:

$$\textbf{truthfulness} = \text{(every player has a dominant strategy)} + \text{(this strategy is to tell the truth)}$$

Are both conditions necessary?

Revelation Principle

Revisiting truthfulness:

truthfulness = (every player has a dominant strategy)
+ (this strategy is to tell the truth)

Are both conditions necessary?

Revelation Principle

Revelation Principle

For every mechanism M in which every participant has a **dominant strategy** (no matter what its private information), there is an equivalent **truthful** *direct-revelation* mechanism M'

Proof.

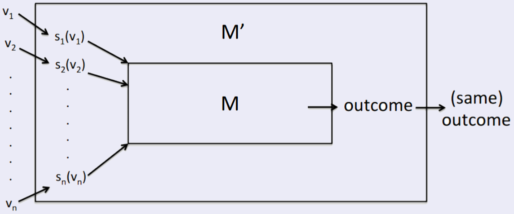

□

Single-parameter environment

Single-parameter environment

A special case of the general mechanism design setting able to model simple auction formats:

- n bidders
- Each bidder i has a **valuation** $v_i \in \mathbb{R}$ which is her value "per unit of stuff" she gets.
- A **feasible set** $\mathcal{X}$. Each element of $\mathcal{X}$ is an n-vector $(x_1, \dots, x_n)$, where x_i denotes the "amount of stuff" that player i gets.

For example:

- In a single-item auction, $\mathcal{X}$ is the set of 0-1 vectors that have at most one 1 (i.e. $\sum_{i=1}^{n} x_i \leq 1$).
- With k identical goods and the constraint the each customer gets at most one, the feasible set is the 0-1 vectors satisfying $\sum_{i=1}^{n} x_i \leq k$.

Single-parameter environment

Sealed-bid auctions in the single-parameter environment

1. Collect bids $\mathbf{b} = (b_1, \dots, b_n)$.
2. **Allocation rule**: Choose a feasible allocation $\mathbf{x}(\mathbf{b}) \in \mathcal{X} \subset \mathbb{R}^n$.
3. **Payment rule**: Choose payments $\mathbf{p}(\mathbf{b}) \in \mathbb{R}^n$.

The **utility** of bidder i is: $u_i(\mathbf{b}) = v_i \cdot x_i(\mathbf{b}) - p_i(\mathbf{b})$.

Definition (Implementable Allocation Rule)

An allocation rule x for a single-parameter environment is **implementable** if there is a payment rule p such the sealed-bid auction (x, p) is **truthful** and **individually rational**.

Definition (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.

Myerson's Lemma

Meyrson's Lemma

Fix a single-parameter environment.

1. An allocation rule x is **implementable** iff it's **monotone**.
2. If x is **monotone**, then there is a *unique* payment rule such that the sealed-bid mechanism (x, p) is **truthful** (assuming the normalization that $b_i = 0$ implies $p_i(b) = 0$).
3. The payment rule in (2) is given by an explicit formula:

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

Myerson's Lemma

Proof:

- **implementable** $\Rightarrow$ **monotone**, payments derived from **(3)**.
  Fix a bidder i and everybody else's valuations $\mathbf{b}_{-i}$.
  Notation: $x(z), p(z)$ instead of $x_i(z, \mathbf{b}_{-i}), p_i(z, \mathbf{b}_{-i})$.
  Suppose $(\mathbf{x}, \mathbf{p})$ is a truthful mechanism and consider $0 \leq y \leq z$.
  - Bidder i has real valuation y but instead bids z. Truthfulness implies:

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

  - Bidder i has real valuation z but instead bids y. Truthfulness implies:

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

Myerson's Lemma

Proof (cont.):

Combining (1), (2):

$$y \cdot [x(z) - x(y)] \leq p(z) - p(y) \leq z \cdot [x(z) - x(y)] \tag{3}$$

$$(3) \Rightarrow (z - y) \cdot [x(z) - x(y)] \geq 0 \Rightarrow x_i(\cdot, b_{-i}) \uparrow$$

Thus the allocation rule is **monotone**.

$$(3) \Rightarrow y \cdot \frac{x(z) - x(y)}{z - y} \leq \frac{p(z) - p(y)}{z - y} \leq z \cdot \frac{x(z) - x(y)}{z - y}$$

Myerson's Lemma

Proof (cont.):

Taking the limit as $y \to z$:

$$\begin{aligned} z \cdot x'(z) \leq p'(z) \leq z \cdot x'(z) &\Rightarrow p'(z) = z \cdot x'(z) \\ &\Rightarrow \int_0^{b_i} p'(z)\, dz = \int_0^{b_i} z \cdot x'(z)\, dz \\ &\Rightarrow p(z) = p(0) + \int_0^{b_i} z \cdot x'(z)\, dz \end{aligned}$$

Assuming normalization $p(0) = 0$ and reverting back to the formal notation:

$$p_i(b_i, \mathbf{b}_{-i}) = \int_0^{b_i} z \frac{d}{dz} x(z)\, dz$$

Myerson's Lemma

Proof (cont.):

- **monotone** $\Rightarrow$ **implementable** with payments from **(3)**.

Proof by pictures (and whiteboard):

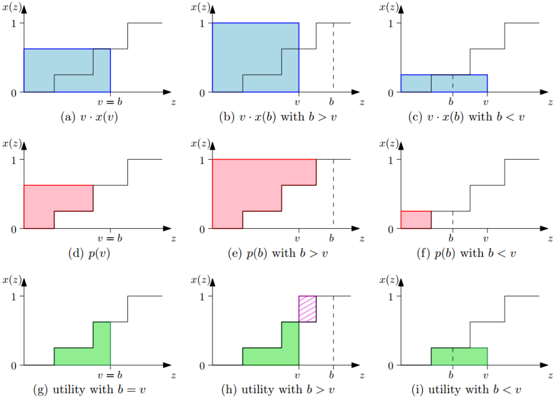

(a) $v \cdot x(v)$ (b) $v \cdot x(b)$ with $b > v$ (c) $v \cdot x(b)$ with $b < v$

(d) $p(v)$ (e) $p(b)$ with $b > v$ (f) $p(b)$ with $b < v$

(g) utility with $b = v$ (h) utility with $b > v$ (i) utility with $b < v$

Welfare maximization in multi-parameter environment

The model

- Set Ω, $|\Omega| = m$, of possible **outcomes**.
- Set $N = \{1, 2, \ldots, n\}$ of **agents** (players).
- **Valuation vector** $\mathbf{v} = (v_1, \ldots, v_n) \in V$ where $v_i : \Omega \to \mathbb{R}$ is the (private) **valuation function** of each player.

Mechanism

- **Allocation Rule**: $x : V^n \to \Omega$.
- **Payment vector**: $\mathbf{p} = (p_1, \ldots, p_n)$ where $p_i : V^n \to \mathbb{R}$.

We are interested in the following **welfare maximizing** allocation rule:

$$x(\mathbf{b}) = \underset{\omega \in \Omega}{\operatorname{argmax}} \sum_{i=1}^{n} b_i(\omega)$$

VCG

Idea: Each player tries to maximize $u_i(\mathbf{b}) = v_i(\omega^*) - p(\mathbf{b})$ where $\omega^* = x(\mathbf{b})$. If we could design the payments in a way that maximizing one's utility is equivalent to trying to maximize the social welfare then we are done!

Notice that

$$\text{SW}(\omega^*) = b_i(\omega^*) + \sum_{j \neq i} b_j(\omega^*) = b_i(\omega^*) - \underbrace{\left[-\sum_{j \neq i} b_j(\omega^*) \right]}_{p(\mathbf{b})} = u_i(\omega^*)$$

VCG

Idea: Each player tries to maximize $u_i(\mathbf{b}) = v_i(\omega^*) - p(\mathbf{b})$ where $\omega^* = x(\mathbf{b})$. If we could design the payments in a way that maximizing one's utility is equivalent to trying to maximize the social welfare then we are done!

Notice that

$$\mathrm{SW}(\omega^*) - h(\mathbf{b}_{-i}) = b_i(\omega^*) + \sum_{j \neq i} b_j(\omega^*) - h(\mathbf{b}_{-i})$$

$$= b_i(\omega^*) - \underbrace{\left[h(\mathbf{b}_{-i}) - \sum_{j \neq i} b_j(\omega^*) \right]}_{p(\mathbf{b})} = u_i(\omega^*)$$

Groves Mechanisms

Every mechanism of the following form is **truthful**:

$$x(\mathbf{b}) = \operatorname*{argmax}_{\omega \in \Omega} \sum_{i=1}^{n} b_i(\mathbf{b})$$

$$p(\mathbf{b}) = h(\mathbf{b}_{-i}) - \sum_{j \neq i} b_j(x(\mathbf{b}))$$

Clarke tax:

$$h(\mathbf{b}_{-i}) = \max_{\omega \in \Omega} \sum_{j \neq i} b_j(\omega)$$

VCG

The VCG mechanism

The Vickrey-Clarke-Grooves mechanism is **truthful**, **individually rational** and exhibits **no positive transfers** ($\forall i : p_i(\mathbf{b}) \geq 0$):

$$x(\mathbf{b}) = \operatorname*{argmax}_{\omega \in \Omega} \sum_{i=1}^{n} b_i(\mathbf{b})$$

$$p(\mathbf{b}) = \max_{\omega \in \Omega} \sum_{j \neq i} b_j(\omega) - \sum_{j \neq i} b_j(x(\mathbf{b}))$$

Proof.

- **Truthfulness**: Follows from the general Groove mechanism.
- **Individual rationality**:

$$u_i(\mathbf{b}) = \ldots = \mathrm{SW}(\omega^*) - \max_{\omega \in \Omega} \sum_{j \neq i} b_j(\omega) \geq \mathrm{SW}(\omega^*) - \max_{\omega \in \Omega} \sum_{j=1}^{n} b_j(\omega) = 0$$

- **No positive transfers**: $\max_{\omega \in \Omega} \sum_{j \neq i} b_j(\omega) \geq \sum_{j \neq i} b_j(x(\mathbf{b}))$.

□

Revenue maximization

As opposed to welfare maximization, maximizing revenue is impossible to achieve **ex-post** (without knowing v_i's beforehand). For example: One item and one bidder with valuation v_i.

Bayesian Model

- A **single-parameter environment**.
- The private valuation v_i of participant i is assumed to be drawn from a **distribution** F_i with density function f_i with support contained in $[0, v_{\max}]$. We also assume the F_i's are **independent**.
- The distributions $F_1, \ldots, F_n$ are **known in advance** to the mechanism designer.

Note: The realizations $v_1, \ldots, v_n$ of bidders' valuations are private, as usual.

We are interested in designing **truthful** mechanisms that maximize the **expected revenue** of the seller.

Revenue maximization

Single-bidder, single-item auction

- The space of direct-revelation truthful mechanisms is small: they are precisely the "**posted prices**", or take-it-or-leave-it offers (because it has to be monotone!)
- Suppose we sell at price r. Then:

$$\mathbb{E}[\text{ Revenue }] = \underbrace{r}_{\text{revenue of a sale}} \cdot \underbrace{(1 - F(r))}_{\text{probability of a sale}}$$

- We chose the price r that maximizes the above quantity.

Example

If F is the **uniform** distribution on $[0, 1]$ then $F(x) = x$ and so:

$$\mathbb{E}[\text{ Revenue }] = r \cdot (1 - r)$$

which is maximized by setting $\mathbf{r = 1/2}$, achieving an expected revenue of $\mathbf{1/4}$.

Revenue maximization

Single-bidder, single-item auction

- The space of direct-revelation truthful mechanisms is small: they are precisely the "**posted prices**", or take-it-or-leave-it offers (because it has to be monotone!)
- Suppose we sell at price r. Then:

$$\mathbb{E}[\text{ Revenue }] = \underbrace{r}_{\text{revenue of a sale}} \cdot \underbrace{(1 - F(r))}_{\text{probability of a sale}}$$

- We chose the price r that maximizes the above quantity.

Example

If F is the **uniform** distribution on $[0, 1]$ then $F(x) = x$ and so:

$$\mathbb{E}[\text{ Revenue }] = r \cdot (1 - r)$$

which is maximized by setting $\mathbf{r = 1/2}$, achieving an expected revenue of $\mathbf{1/4}$.

Revenue maximization

General setting of multi-player single-parameter environment:

Theorem (Myerson, 1981)

$$\mathbb{E}_{\mathbf{v}\sim\mathbf{F}}\left[\sum_{i=1}^{n} p_i(\mathbf{v})\right] = \mathbb{E}_{\mathbf{v}\sim\mathbf{F}}\left[\sum_{i=1}^{n} \phi_i(v_i)\cdot x_i(v_i)\right]$$

where:

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

is called **virtual welfare**.

Revenue maximization

Proof:

Step 1: Fix i, $\mathbf{v}_{-i}$. By Myerson's payment formula:

$$\mathbb{E}_{v_i \sim F_i}\left[p_i(\mathbf{v})\right] = \int_0^{v_{\max}} p_i(\mathbf{v}) f_i(v_i)\, dv_i = \int_0^{v_{\max}} \left[\int_0^{v_i} z \cdot x_i'(z, \mathbf{v}_{-i})\, dz\right] f_i(v_i)\, dv_i$$

Step 2: Reverse integration order:

$$\int_0^{v_{\max}} \left[\int_0^{v_i} z \cdot x_i'(z, \mathbf{v}_{-i})\, dz\right] f_i(v_i)\, dv_i = \int_0^{v_{\max}} \left[\int_z^{v_{\max}} f_i(v_i)\, dv_i\right] z \cdot x_i'(z, \mathbf{v}_{-i})\, dz$$

$$= \int_0^{v_{\max}} (1 - F_i(z)) \cdot z \cdot x_i'(z, \mathbf{v}_{-i})\, dz$$

Proof (cont.):

Step 3: Integration by parts:

$$\int_0^{v_{\max}} \underbrace{(1 - F_i(z)) \cdot z}_{f} \cdot \underbrace{x_i'(z, \mathbf{v}_{-i})}_{g'} \, dz$$

$$= \underbrace{(1 - F_i(z)) \cdot z \cdot x_i(z, \mathbf{v}_{-i})\big|_0^{v_{\max}}}_{=0-0} - \int_0^{v_{\max}} x_i(z, \mathbf{v}_{-i}) \cdot (1 - F_i(z) - z f_i(z)) \, dz$$

$$= \int_0^{v_{\max}} \underbrace{\left(z - \frac{1 - F_i(z)}{f_i(z)}\right)}_{:=\varphi_i(z)} x_i(z, \mathbf{v}_{-i}) f_i(z) \, dz$$

Revenue Maximization

Proof (cont.):

Step 4: To *simplify* and help *interpret* the expression we introduce the **virtual valuation** $\varphi_i(v_i)$:

$$\varphi(v_i) = \underbrace{v_i}_{\text{what you'd like to charge } i} - \underbrace{\frac{1 - F_i(v_i)}{f_i(v_i)}}_{\text{``information rent'' earned by bidder } i}$$

Summary:

$$\mathbb{E}_{v_i \sim F_i}[p_i(\mathbf{v})] = \mathbb{E}_{v_i \sim F_i}[\varphi(v_i) \cdot x_i(\mathbf{v})] \tag{4}$$

Revenue Maximization

Proof (cont.):

Step 5: Take the expectation, with respect to $\mathbf{v}_{-i}$ of both sides of (4):

$$\mathbb{E}_{\mathbf{v}}[p_i(\mathbf{v})] = \mathbb{E}_{\mathbf{v}}[\varphi_i(v_i) \cdot x_i(\mathbf{v})]$$

Step 6: Apply linearity of expectation twice:

$$\mathbb{E}_{\mathbf{v}}\left[\sum_{i=1}^{n} p_i(\mathbf{v})\right] = \sum_{i=1}^{n} \mathbb{E}_{\mathbf{v}}[p_i(\mathbf{v})] = \sum_{i=1}^{n} \mathbb{E}_{\mathbf{v}}[\varphi_i(v_i) \cdot x_i(\mathbf{v})] = \mathbb{E}_{\mathbf{v}}\left[\sum_{i=1}^{n} \varphi_i(v_i) \cdot x_i(\mathbf{v})\right]$$

□

Revenue Maximization

Conclusion

MAXIMIZING **REVENUE** $\Leftrightarrow$ MAXIMIZING **VIRTUAL WELFARE**

Example: Single-item auction with i.i.d. bidders

Assuming that the distributions F_i are such that $\phi_i(v_i)$ is monotone (such distributions are called **regular**) then a **second-price** auction on *virtual valuations* with reserve price $\phi^{-1}(0)$ maximizes the revenue.

Revenue Maximization

Conclusion

MAXIMIZING **REVENUE** $\Leftrightarrow$ MAXIMIZING **VIRTUAL WELFARE**

Example: Single-item auction with i.i.d. bidders

Assuming that the distributions F_i are such that $\phi_i(v_i)$ is monotone (such distributions are called **regular**) then a **second-price** auction on *virtual valuations* with reserve price $\phi^{-1}(0)$ maximizes the revenue.