

Designing Incentive-Compatible Mechanisms

Algorithmic Game Theory

Winter 2022/23

Vickrey-Auction and Vickrey-Clarke-Groves Mechanisms

Characterization of Incentive Compatibility

Single-Parameter Mechanisms

Revelation Principle

Mechanisms and Approximation Algorithms

Revenue Maximization in Single-Parameter Domains

Mechanisms with Money

- ▶ Set A of possible **outcomes**.
- ▶ **Goal**: Choose a desired outcome $a \in A$.
- ▶ Players have quantifiable preferences over outcomes. Common currency enables utility transfer between players.
- ▶ Preference of player i is given by a **valuation function** $v_i : A \rightarrow \mathbb{R}$ from a commonly known set $V_i \subseteq \mathbb{R}^A$
- ▶ v_i is **private information** of player i .
- ▶ **Mechanism** to determine a good outcome $a \in A$:
 1. Ask every player i for a “bid” b_i , i.e., her valuation (direct revelation)
 2. Determine a desired outcome $a \in A$
 3. Determine **payments** m_i for every player i
- ▶ **Utility** of player i is $v_i(a) - m_i$, *quasi-linear* utilities.

Example: Sealed Bid Auction

A single item is sold to one customer.

Customer	1	2	3	4	5
Value	9	1	20	11	14

Bidders initially report values using a “sealed bid”.

Social Choice: Winner is bidder with highest bid.

Payments: Find payments to ensure incentive-compatibility

- ▶ No payments: Bidders try to bid unbounded high values.
- ▶ Payments = Bids: Bidders try to guess whether they are the highest bidder, estimate the second highest bid and bid a slightly higher value.

Vickrey Second Price Auction

Payment of the winner is the second largest bid.

Value	9	1	20	11	14
Payment	0	0	14	0	0
Utility	0	0	6	0	0

A mechanism is called **incentive compatible** if, for every bidder i and every set of bids of other players, truthful revelation of v_i is maximizing the utility for i .

Proposition

The Vickrey auction is incentive compatible.

Example

Value	?	?	20	?	?
Bid	5	11	x	2	14
Payment			14		
Utility			6		

Case 1: i wins with true value $x = 20$, then for all $x \geq 14$ utility 6, for $x < 14$ utility 0.

Value	?	?	20	?	?
Bid	5	11	x	2	24
Payment			0		
Utility			0		

Case 2: i loses with true value $x = 20$, then for all $x < 24$ utility 0, for $x \geq 24$ utility -4 .

Definitions

Direct Revelation Mechanism

- ▶ Notation: $V = V_1 \times \dots \times V_n$ and $v \in V$
- ▶ $v = (v_1, \dots, v_n)$, v_i is **type** of bidder i
- ▶ Bidder “bids”: Reports a type to the mechanism
- ▶ Social choice function $f : V \rightarrow A$, payment functions p_1, \dots, p_n
- ▶ $p_i : V \rightarrow \mathbb{R}$ specifies the amount player i pays.

Incentive Compatibility (IC)

- ▶ Consider every bidder i , every profile $v \in V$, and every alternative $v'_i \in V_i$.
- ▶ We denote outcomes by $a = f(v_i, v_{-i})$ and $b = f(v'_i, v_{-i})$
- ▶ Mechanism (f, p_1, \dots, p_n) is incentive compatible if the utility

$$v_i(a) - p_i(v_i, v_{-i}) \geq v_i(b) - p_i(v'_i, v_{-i})$$

Sealed-Bid Auction



Bidder	1	2	3	4	5
Value	9	1	20	11	14

- Outcomes $A = \{1, 2, 3, 4, 5\}$, where i means “ i wins”

Outcome	1	2	3	4	5
v_1	9	0	0	0	0
v_2	0	1	0	0	0
etc.					

- Social Choice: $f(v) = \arg \max_i \{v_i(i)\}$
- Payments: $p_i(v) = 0$ if $f(v) \neq i$,
otherwise $p_i(v) = \max_{j \neq i} v_j(j)$.

VCG Mechanism

Definition

A **Vickrey-Clarke-Groves (VCG) mechanism** is given by

- ▶ $f(v) \in \arg \max_{a \in A} \sum_i v_i(a)$, and
- ▶ for every $v \in V$ and every bidder i

$$p_i(v) = h_i(v_{-i}) - \sum_{j \neq i} v_j(f(v)) ,$$

with h_1, \dots, h_n being arbitrary functions $h_i : V_{-i} \rightarrow \mathbb{R}$.

Observations:

- ▶ VCG mechanism picks outcome a that maximizes **social welfare** $\sum_j v_j(a)$
- ▶ h_i does not depend on the own "bid" v_i
- ▶ Utility of player i when $f(v) = a$:

$$v_i(a) - p_i(v) = \sum_j v_j(a) - h_i(v_{-i})$$

VCG is IC

Theorem

Every VCG mechanism is incentive compatible.

Proof:

- ▶ Given types v , let $v'_i \neq v_i$ be a “lie” for bidder i
- ▶ Let $a = f(v)$ and $b = f(v'_i, v_{-i})$.
- ▶ Utility of i declaring v_i is $v_i(a) + \sum_{j \neq i} v_j(a) - h_i(v_{-i})$
- ▶ Utility of i declaring v'_i is $v_i(b) + \sum_{j \neq i} v_j(b) - h_i(v_{-i})$
- ▶ Utility is maximized when outcome maximizes social welfare $\sum_j v_j(x)$.
- ▶ VCG mechanism maximizes social welfare, $\sum_j v_j(a) \geq \sum_j v_j(b)$.
- ▶ By declaring v'_i bidder i , VCG picks b . However, b is optimal for i 's lie, but possibly suboptimal for her real utility.
- ▶ VCG aligns every bidder incentive with the social incentives. □

Desirable Properties of Payments

Definition

- ▶ A mechanism is (ex-post) **individually rational** if bidders always get non-negative utility, i.e. for all $v \in V$ we have $v_i(f(v)) - p_i(v) \geq 0$.
- ▶ A mechanism has **no positive transfers** if no bidder is ever paid money, i.e. for all $v \in V$ and all i we have $p_i(v) \geq 0$.

Definition (Clarke Rule)

The payment functions resulting from $h_i(v_{-i}) = \max_{b \in A} \sum_{j \neq i} v_j(b)$ are called Clarke pivot payment.

Clarke Rule

Using Clarke pivot payment the payments of bidder i become

$$p_i(v) = \max_{b \in A} \sum_{j \neq i} v_j(b) - \sum_{j \neq i} v_j(f(v))$$

Payment is the “total damage” that i causes to the other players by her presence in the system. Each player *internalizes externalities*.

Lemma

A VCG mechanism with Clarke pivot payments makes no positive transfers. If $v_i(a) \geq 0$ for all $v_i \in V_i$ and $a \in A$, then it is individually rational.

Clarke Rule

Proof:

- ▶ Let $a = f(v)$ and $b = \arg \max_{a' \in A} \sum_{j \neq i} v_j(a')$
- ▶ No positive transfers (by definition)

$$\sum_{j \neq i} v_j(b) - \sum_{j \neq i} v_j(a) \geq 0$$

- ▶ Individually rational

$$v_i(a) + \sum_{j \neq i} v_j(a) - \sum_{j \neq i} v_j(b) \geq \sum_j v_j(a) - \sum_j v_j(b) \geq 0$$



Example: Bilateral Trade



	trade	no-trade
Seller	$-v_s$	0
Buyer	v_b	0

- ▶ Trade occurs if $v_b > v_s$, no-trade if $v_s > v_b$
- ▶ Analyze VCG Mechanism, should not subsidize trade.

Example: Bilateral Trade

	trade	no-trade
Seller	$-v_s$	0
Buyer	v_b	0

- ▶ VCG payments for no-trade:

Seller payments: $h_s(v_b) - 0$, Buyer payments: $h_b(v_s) - 0$

No additional payments by the mechanism, so $h_s(v_b) = h_b(v_s) = 0$.

- ▶ VCG payments for trade:

Seller payments: $h_s(v_b) - v_b$, Buyer payments: $h_b(v_s) + v_s$

Seller receives v_b , but buyer pays only $v_s < v_b$.

- ▶ Not *budget-balanced*: VCG mechanism subsidizes trade!

Example: Procurement or Reverse Auction

- ▶ Auctioneer buys service
- ▶ Participants offer service, each one has (private) cost
- ▶ Auctioneer pays participants
- ▶ Negative utility, negative payments
- ▶ Vickrey reverse auction:
Pick participant with smallest bid, pay the second-smallest bid

Corollary

The Vickrey reverse auction is incentive compatible.

Vickrey Reverse Auction is IC

Case 1: If bidding his true value, player i wins.

Value	?	?	-7	?	?
Bid	-9	-11	x	-17	-14
Payment			-9		
Utility			2		

Case 2: If bidding his true value, player i loses.


Value	?	?	-12	?	?
Bid	-9	-11	x	-17	-24
Payment			0		
Utility			0		

Example: Buying a Path in a Network

Reverse auction:

Bidders are edges in a network. Each edge has private cost c_e for being used.

Mechanism wants to buy an s - t -path.





Example: Buying a Path in a Network

- ▶ Outcomes are s - t -paths in graph G
- ▶ VCG picks shortest path P^* for reported costs c_e (= maximizes wrt. values $v_e = -c_e$)
- ▶ Payments for edge $e \in P^*$ to the mechanism:

$$h_e(c_{-e}) - \sum_{e' \in P^*, e' \neq e} -c_{e'} = h_e(c_{-e}) + c(P^* - e)$$
- ▶ We use $h_e(c_{-e}) = \max_{P \in G_{-e}} \sum_{e \in P} -c_e = -c(P_{-e}^*)$, where P_{-e}^* is a shortest s - t -path in the graph that does not contain edge e .
 This choice for h_e is not exactly the Clarke Rule (Why?)
- ▶ Total payment for $e \in P^*$ is $c(P^* - e) - c(P_{-e}^*) \leq 0$, i.e., edge e receives money from the mechanism.
- ▶ Any edge $e \notin P^*$ has cost 0 and gets no payment.

Truth-telling is a dominant strategy!



Vickrey-Auction and Vickrey-Clarke-Groves Mechanisms

Characterization of Incentive Compatibility

Single-Parameter Mechanisms

Revelation Principle

Mechanisms and Approximation Algorithms

Revenue Maximization in Single-Parameter Domains

More incentive-compatible mechanisms?

The VCG mechanism is incentive compatible and maximizes social welfare f .

Are there other social choice functions f that can be *implemented*, i.e. augmented using suitable payments into incentive-compatible mechanisms?

Are there different types of incentive-compatible mechanisms besides VCG?

Direct Characterization

Proposition

A mechanism is incentive compatible if and only if it satisfies the following conditions hold for every bidder i and every v_{-i} :

- 1. The payment p_i does not depend on v_i , but only on the outcome, i.e., for every v_{-i} there exist prices $p_a \in \mathbb{R}$ such that for all v_i with $f(v_i, v_{-i}) = a$ we have $p_i(v_i, v_{-i}) = p_a$.*
- 2. The mechanism optimizes for each bidder, i.e., for every v_i it holds that $f(v_i, v_{-i}) \in \arg \max_{a \in A'} \{v_i(a) - p_a\}$, where A' is the set of alternatives in the range of $f(\cdot, v_{-i})$.*

Proof:

Conditions hold \Rightarrow IC: obvious.

Proof Direct Characterization

1. The payment $p_i = p_a$ does not depend on v_i , but only on the outcome $a = f(v_i, v_{-i})$.
2. The mechanism optimizes for each bidder.

IC \Rightarrow Conditions hold:

► Condition 1:

$v_i \neq v'_i$ yield same outcome for fixed v_{-i} . Payment $p_i(v_i, v_{-i}) > p_i(v'_i, v_{-i})$ then bidder i with v_i is motivated to lie v'_i .

► Condition 2:

If not, then there is a better outcome $a' \in \arg \max_a (v_i(a) - p_a)$ and some v'_i that gives $a' = f(v'_i, v_{-i})$. Hence, bidder i with v_i is motivated to lie v'_i . □

Affine Maximizer

Definition

A social choice function f is an **affine maximizer** if there is a subset $A' \subset A$, bidder weights $w_1, \dots, w_n \in \mathbb{R}$, and outcome weights $c_a \in \mathbb{R}$ for each $a \in A$, such that

$$f(v_1, \dots, v_n) \in \arg \max_{a \in A'} \left\{ c_a + \sum_i w_i v_i(a) \right\}.$$

Proposition

Suppose f is an affine maximizer, and h_i is an arbitrary function independent of v_i . Suppose bidder i with $w_i = 0$ pays $p_i(v) = 0$, and bidder i with $w_i > 0$ pays

$$p_i(v) = h_i(v_{-i}) - \frac{1}{w_i} \left(\sum_{j \neq i} w_j v_j(a) + c_a \right).$$

Then (f, p_1, \dots, p_n) is incentive compatible.

(Only) Affine Maximizers can be implemented

Proof:

- ▶ If $w_i = 0$, then i has no influence on the mechanism.
- ▶ With $p_i = 0$ same utility for every bid of i .
- ▶ If $w_i > 0$, then assume wlog $h_i = 0$. Utility of i if a is chosen:

$$v_i(a) + \frac{1}{w_i} \left(\sum_{j \neq i} w_j v_j(a) + c_a \right).$$

- ▶ Multiply by $w_i > 0$, expression is maximized when $c_a + \sum_j w_j v_j(a)$ is maximized.
- ▶ f affine maximizer, true type is a dominant strategy for i . □

Theorem (Roberts 1979)

Suppose $|A| \geq 3$, f is surjective, $V_i = \mathbb{R}^A$ for every i , and (f, p_1, \dots, p_n) is incentive compatible. Then f must be an affine maximizer.

Vickrey-Auction and Vickrey-Clarke-Groves Mechanisms

Characterization of Incentive Compatibility

Single-Parameter Mechanisms

Revelation Principle

Mechanisms and Approximation Algorithms

Revenue Maximization in Single-Parameter Domains

“Stuff times Value” Valuations

Single-parameter valuations have a simple structure:

- ▶ For every outcome $a \in A$, bidder i receives some amount of “stuff”
- ▶ Let $x_i(a) \in \mathbb{R}$ be the amount of “stuff” that bidder i gets in outcome a
- ▶ Valuation based on a single parameter:

Value per unit of stuff: $t_i \in \mathbb{R}$

Valuation function: $v_i(a) = t_i \cdot x_i(a)$

Definition

A **single-parameter domain** V_i is defined by (public) function $x_i : A \rightarrow \mathbb{R}$ and domain $[t_i^0, t_i^1]$. The set V_i contains all v_i such that there is $t_i^0 \leq t_i \leq t_i^1$ with

$$v_i(a) = t_i \cdot x_i(a) .$$

The single parameter t_i is private information.

Overload notation: v_i refers to both, valuation function and parameter t_i .

Examples

Simple Examples:

- ▶ Single-item auction: $x_i(a) \in \{0, 1\}$ and $\sum_i x_i(a) \leq 1$.
- ▶ k identical items: $x_i(a) \in \{0, 1, \dots, k\}$ and $\sum_i x_i(a) \leq k$.
- ▶ s - t -path: $x_e(a) \in \{0, 1\}$ and $P(x) = \{e \mid x_e(a) = 1\}$ is an s - t -path in G .

Sponsored-Search Auction:

- ▶ A search result page has several advertisement slots
- ▶ Search engine auctions off the slots to advertisers
- ▶ Slot k has a known **click-through rate (CTR)** $\alpha_k \geq 0$
- ▶ Firm i has private **value v_i per click** for its ad
- ▶ An outcome $a \in A$ is a matching of adslots to firms
- ▶ $x_i(a) = \alpha_k$ if firm i gets a slot k , and $x_i(a) = 0$ otherwise
- ▶ Valuation of firm i is $v_i(a) = v_i \cdot x_i(a)$

Are there IC mechanisms for single-parameter domains that are not affine maximizers?

Sponsored-Search Auctions

+You Search Images Maps Play YouTube Gmail Documents Calendar Translate More +

Google

Search About 257,000,000 results (0.26 seconds)

Web

Images

Maps

Videos

News

Shopping

More

Aachen

Change location

Show search tools

Ads related to rental car Why these ads?

Best Car Rental Prices - Guaranteed! Book Online Today
www.rentalcars.com/ - ★★★★★ 3,660 seller reviews
 Low Cost **Rental** from Rentalcars.com

Car Rental Low Cost	No Credit Card Fees
Best Price Guarantee	Leading Brands
Free Amendments	2 Million+ Customers

Car Rentals - No Bidding. No Guessing | CarRentals.com
www.carrentals.com/
 Discount Economy **Rental Cars** from \$10/Day.
 ↳ Los Angeles - Las Vegas - Orlando - Phoenix


SIXT Online Car Rental | sixt.de
www.sixt.de/Car_Rental
 Aktuelle Modelle aller Klassen & keine KM-Begrenzung ab 23 €/Tag!
 ↳ Mercedes Benz Roadster ab €45/Tag - Porsche Special: jetzt ab € 249/Tag

Enterprise Rent-A-Car - Rental Cars at Low Rates
www.enterprise.com/
 Reserve a car rental from Enterprise **Rent-A-Car** at low rates. Choose from more than 6000 **rental car** locations at major airports and neighborhood locations.
 ↳ Car Rental Locations in - Contact Us - Rental Cars at Everyday - Canada

Avis Car Rental - Rent A Car with Avis
www.avis.com/ - United States
 Make a car reservation from Avis **Rent A Car** car rental. Book low rates online and reserve a **rental car**. Worried about the environment? Choose one of our ...

Hertz Rent-a-Car - Rental Car Discounts, Coupons and Great Rates
www.hertz.com/

Map for rental car



©2012 Google. Map data ©2012 GeoBasis-DE/BKG (©2009), Google

Ads - Why these ads?

Find 100s Of Car Rentals
www.kayak.com/
 Don't Overpay For Your **Car Rental**. Compare **Rental** Options On One Site. 124,958 people +1'd or follow KAYAK

Rental Car International
www.autoeurope.de/Rental+Car
 Discover the greatest spots with Autoeurope **car** rentals. book here!

Cheap Car Rental Germany
www.germanycarrentalsite.com/
 Find Best **Car Rental** Deals Online. Low Prices. No Taxes. Book Now.

Car Hire in Italy Offers
www.ciaocars.com/
 Great Rates! No Book or Change Fees Book online. Pay on Arrival. Easy.

Example: Second-Highest Bid Wins

We auction a single good and assign it to the **second-highest bidder**. Are there payments such that the resulting mechanism is IC?

Consider some bidder i and fix the other bids v_{-i} .

It holds $x_i(a) \in \{0, 1\}$. Direct characterization shows: i pays one of two prices, p_i^1 or p_i^0 , depending on whether she is second-highest bidder or not.

Suppose y is a bid that makes i the second-highest bidder, and z is one that makes her the highest bidder, with $y < z$.

If $v_i = y$, then i shall not want to lie z . Hence: $y \cdot 1 - p_i^1 \geq y \cdot 0 - p_i^0$.

If $v_i = z$, then i shall not want to lie y . Hence: $z \cdot 0 - p_i^0 \geq z \cdot 1 - p_i^1$.

This implies $y \geq z$, a contradiction.

There are no payments that yield an IC mechanism. The social choice function is **not monotone** – a higher bid cannot reduce the received amount of stuff.

Monotonicity

Definition

A social choice function f on a single parameter domain f is called **monotone in v_i** if for every v_{-i} and every $v'_i \in V_i$ with $v'_i \geq v_i$

$$x_i(f(v'_i, v_{-i})) \geq x_i(f(v_i, v_{-i})) .$$

Normalized mechanism: Using the smallest bid t_i^0 , bidder i never gets stuff and always pays nothing, i.e., $x_i(t_i^0, v_{-i}) = 0$ and $p_i(t_i^0, v_{-i}) = 0$ for every v_{-i} .

Characterization

Theorem (Myersons Lemma)

A normalized mechanism (f, p_1, \dots, p_n) on a single parameter domain is incentive compatible if and only if the following conditions hold:

- ▶ f is monotone in every v_i , and
- ▶ the payments are given by

$$p_i(v_i, v_{-i}) = v_i \cdot x_i(f(v)) - \int_{t_i^0}^{v_i} x_i(f(t, v_{-i})) dt.$$

Proof:

Fix v_{-i} . Let $y < z$ be two possible private values of i .

We write $a_y = f(y, v_{-i})$ and $a_z = f(z, v_{-i})$.

Proof Myerson's Lemma

IC implies:

$$y \cdot x_i(a_y) - p_i(a_y) \geq y \cdot x_i(a_z) - p_i(a_z) \quad (1)$$

and

$$z \cdot x_i(a_z) - p_i(a_z) \geq z \cdot x_i(a_y) - p_i(a_y) \quad (2)$$

Sum (1) and (2) and rearrange:

$$z \cdot (x_i(a_z) - x_i(a_y)) \geq y \cdot (x_i(a_z) - x_i(a_y))$$


Since $z > y$, we know $x_i(a_z) \geq x_i(a_y)$. Hence, IC \Rightarrow f monotone.

We next show that (IC \wedge f monotone) \Rightarrow payments as given in the Lemma.

We show this only for the special case with $x_i(a) \in \mathbb{N}$.

Proof Myerson's Lemma

Suppose x_i monotone and $x_i(a) \in \{0, 1, 2, \dots, k\}$, a step function. x_i jumps at $z_1 \leq z_2 \leq \dots \leq z_\ell$ by k_1, k_2, \dots, k_ℓ , where $\sum_{j=1}^{\ell} k_j \leq k$.




Proof Myerson's Lemma

(1) and (2) yield

$$z \cdot (x_i(a_z) - x_i(a_y)) \geq p_i(a_z) - p_i(a_y) \geq y \cdot (x_i(a_z) - x_i(a_y))$$

In addition, $p_i(a_z) = p_i(a_y)$ if $x_i(a_z) = x_i(a_y)$. Set $z = z_i$ and $y = z_i - \varepsilon$, then with $\varepsilon \rightarrow 0$ we see that p_i jumps at z_i by $z_i k_i$. Thus

$$p_i(a_z) = \sum_{j: z_j \leq z} z_j k_j = z \cdot x_i(a_z) - \int_{t_i^0}^z x_i(a_t) dt .$$




Proof Myerson's Lemma

(1) and (2) yield

$$z \cdot (x_i(a_z) - x_i(a_y)) \geq p_i(a_z) - p_i(a_y) \geq y \cdot (x_i(a_z) - x_i(a_y))$$

In addition, $p_i(a_z) = p_i(a_y)$ if $x_i(a_z) = x_i(a_y)$. Set $z = z_i$ and $y = z_i - \varepsilon$, then with $\varepsilon \rightarrow 0$ we see that p_i jumps at z_i by $z_i k_i$. Thus

$$p_i(a_z) = \sum_{j: z_j \leq z} z_j k_j = z \cdot x_i(a_z) - \int_{t_i^0}^z x_i(a_t) dt .$$




Proof Myerson's Lemma

(1) and (2) yield

$$z \cdot (x_i(a_z) - x_i(a_y)) \geq p_i(a_z) - p_i(a_y) \geq y \cdot (x_i(a_z) - x_i(a_y))$$

In addition, $p_i(a_z) = p_i(a_y)$ if $x_i(a_z) = x_i(a_y)$. Set $z = z_i$ and $y = z_i - \varepsilon$, then with $\varepsilon \rightarrow 0$ we see that p_i jumps at z_i by $z_i k_i$. Thus

$$p_i(a_z) = \sum_{j: z_j \leq z} z_j k_j = z \cdot x_i(a_z) - \int_{t_i^0}^z x_i(a_t) dt .$$




Proof Myerson's Lemma

(1) and (2) yield

$$z \cdot (x_i(a_z) - x_i(a_y)) \geq p_i(a_z) - p_i(a_y) \geq y \cdot (x_i(a_z) - x_i(a_y))$$

In addition, $p_i(a_z) = p_i(a_y)$ if $x_i(a_z) = x_i(a_y)$. Set $z = z_i$ and $y = z_i - \varepsilon$, then with $\varepsilon \rightarrow 0$ we see that p_i jumps at z_i by $z_i k_i$. Thus

$$p_i(a_z) = \sum_{j: z_j \leq z} z_j k_j = z \cdot x_i(a_z) - \int_{t_i^0}^z x_i(a_t) dt .$$




Proof Myerson's Lemma

(1) and (2) yield

$$z \cdot (x_i(a_z) - x_i(a_y)) \geq p_i(a_z) - p_i(a_y) \geq y \cdot (x_i(a_z) - x_i(a_y))$$

In addition, $p_i(a_z) = p_i(a_y)$ if $x_i(a_z) = x_i(a_y)$. Set $z = z_i$ and $y = z_i - \varepsilon$, then with $\varepsilon \rightarrow 0$ we see that p_i jumps at z_i by $z_i k_i$. Thus

$$p_i(a_z) = \sum_{j: z_j \leq z} z_j k_j = z \cdot x_i(a_z) - \int_{t_i^0}^z x_i(a_t) dt .$$




Proof Myerson's Lemma

For every IC mechanism: (1) monotone f and (2) payments as in the Lemma.

Finally, are these two conditions also sufficient, i.e., is every mechanism with these conditions also IC?

Valuation with truthful bid:




Proof Myerson's Lemma

For every IC mechanism: (1) monotone f and (2) payments as in the Lemma.

Finally, are these two conditions also sufficient, i.e., is every mechanism with these conditions also IC?

Payments with truthful bid:




Proof Myerson's Lemma

For every IC mechanism: (1) monotone f and (2) payments as in the Lemma.

Finally, are these two conditions also sufficient, i.e., is every mechanism with these conditions also IC?

Utility with truthful bid:




Proof Myerson's Lemma

For every IC mechanism: (1) monotone f and (2) payments as in the Lemma.

Finally, are these two conditions also sufficient, i.e., is every mechanism with these conditions also IC?

Valuation with bid $z' > v_i$:




Proof Myerson's Lemma

For every IC mechanism: (1) monotone f and (2) payments as in the Lemma.

Finally, are these two conditions also sufficient, i.e., is every mechanism with these conditions also IC?

Payments with bid $z' > v_i$:




Proof Myerson's Lemma

For every IC mechanism: (1) monotone f and (2) payments as in the Lemma.

Finally, are these two conditions also sufficient, i.e., is every mechanism with these conditions also IC?

Utility with bid $z' > v_i$ has not improved!




Proof Myerson's Lemma

For every IC mechanism: (1) monotone f and (2) payments as in the Lemma.

Finally, are these two conditions also sufficient, i.e., is every mechanism with these conditions also IC?

Valuation with bid $z' < v_i$:




Proof Myerson's Lemma

For every IC mechanism: (1) monotone f and (2) payments as in the Lemma.

Finally, are these two conditions also sufficient, i.e., is every mechanism with these conditions also IC?

Payment with bid $z' < v_i$:




Proof Myerson's Lemma

For every IC mechanism: (1) monotone f and (2) payments as in the Lemma.

Finally, are these two conditions also sufficient, i.e., is every mechanism with these conditions also IC?

Utility with bid $z' < v_i$ has not improved!




Example: Buying a Path in a Network (Part 2)

Reverse Auction and Min-Max-Paths:

Bidders are edges in a network. Each edge e has private cost c_e for being used.

Mechanism wants to buy an s - t -path.





Choose a path P^* that **minimizes the maximum cost of any edge in the path.**

Min-Max is monotone!

If e reduces her bid, she can only join or stay in P^* . Monotone $x_i(f(v_i, v_{-i})) \in \{0, 1\}$, at most one step. IC:

$e \notin P^*$ gets no payment.

$e \in P^*$ gets maximum edge cost on min-max s - t -path in $G - \{e\}$



Vickrey-Auction and Vickrey-Clarke-Groves Mechanisms

Characterization of Incentive Compatibility

Single-Parameter Mechanisms

Revelation Principle

Mechanisms and Approximation Algorithms

Revenue Maximization in Single-Parameter Domains

Mechanisms with Communication

All results so far apply to mechanisms with direct revelation.

Are there fundamentally different mechanisms with more complex communication?

For example, a mechanism could ask in k rounds sequentially some yes/no questions, and the bidders must react to that. Or a mechanism would present in every round two outcomes and ask each bidder which outcome she likes better. Or some other interaction rule, or...

For general communication between mechanism and bidder i we assume that, for every bidder i , there is a set X_i of **possible actions**. Each $x_i \in X_i$ represents a **collection of answers** bidder i can use to reply to the questions of the mechanism.

General Mechanisms with Action Spaces

General Mechanism with Action Space:

- ▶ **Action space** X_i for bidder i , we set $X = X_1 \times \dots \times X_n$.
- ▶ **Strategy** $s_i : V_i \rightarrow X_i$ maps every possible valuation $v_i \in V_i$ to an action.
- ▶ Every bidder i picks a strategy s_i and, hence, the action $x_i = s_i(v_i)$.
- ▶ **Social choice function** $g : X \rightarrow A$ maps chosen actions to an outcome
- ▶ **Payment** $p_i : X \rightarrow \mathbb{R}$ depends on chosen actions
- ▶ **Quasi-linear Utility**: $u_i(x) = v_i(g(x)) - p_i(x)$

Direct revelation is the case $X_i = V_i$. With her strategy a bidder directly reports her (possibly incorrect) private valuation. More generally, the set of actions X_i is not necessarily identical to the set of valuations V_i . Using strategy s_i a bidder determines for every possible private valuation a choice of action (i.e., the collection of answers it gives to the mechanism).

Revelation Principle

Consider a strategy profile $s(v) = (s_1, \dots, s_n)$, and suppose s is a dominant-strategy equilibrium for the general mechanism. Let $f(v) = g(s(v))$. We say the mechanism **implements the social choice function f in dominant strategies**.

For an IC mechanism with direct revelation, truth-telling is a dominant strategy for every bidder. Formally, for such a mechanism there is a dominant-strategy equilibrium s with $s_i(v_i) = v_i$ for all $v_i \in V_i$ and every bidder i .

The revelation principle says that complex communication cannot entail fundamentally different mechanisms with dominant-strategy equilibria. Thus, we can continue to restrict attention to mechanisms with direct revelation.

Revelation Principle


Proposition (Revelation Principle)

There is general mechanism M that implements f in dominant strategies.

\iff

There is IC mechanism M' with direct revelation and social choice function f .

Proof:



□

Revelation Principle


Proposition (Revelation Principle)

There is general mechanism M that implements f in dominant strategies.

\iff

There is IC mechanism M' with direct revelation and social choice function f .

Proof:



Vickrey-Auction and Vickrey-Clarke-Groves Mechanisms

Characterization of Incentive Compatibility

Single-Parameter Mechanisms

Revelation Principle

Mechanisms and Approximation Algorithms

Revenue Maximization in Single-Parameter Domains

Knapsack Auction

Myerson's Lemma implies that **designing incentive-compatible mechanisms** reduces to **designing monotone social choice functions**. This raises issues with computational complexity.

As an example consider a **Knapsack Auction**:

A TV station wants to fill a **commercial break of G seconds** with spots. There is a set I of n firms that would like to broadcast their spot. Every firm $i \in I$

- ▶ delivers **spot of length** $g_i \leq G$ seconds (g_i public knowledge),
- ▶ has **valuation** $v_i \geq 0$ if her spot is included (v_i private information), and valuation 0 otherwise.

The knapsack auction is obviously a single-parameter domain. Let us first consider a VCG mechanism.

VCG Mechanism for the Knapsack Auction

VCG Mechanism

- ▶ Query valuations v_i from every firm $i \in I$
- ▶ Choose subset $S \subseteq I$ of spots that maximizes social welfare:

$$f(v) = \arg \max_{S \subseteq I} \left\{ \sum_{i \in S} v_i \mid \sum_{i \in S} g_i \leq G \right\}$$

- ▶ Payments $p_i(v)$ as given by Myerson's Lemma

VCG must **compute optimal solutions for the knapsack problem**, but this problem is **NP-hard**. Thus, there is a tension between three desirable properties of the mechanism:

- (1) incentive compatible
- (2) maximizes social welfare
- (3) polynomial-time computation

Complexity of Incentive-Compatible Mechanisms

The conflict arises between properties (2) and (3). For decades, this conflict has been studied in the area of **approximation algorithms**. By using these algorithms, we soften property (2) into

- (1) incentive compatible
- (2') approximates social welfare as good as possible
- (3) polynomial-time computation

However, we cannot use **arbitrary** approximation algorithms. Since we need to fulfill (1), there must **exist payments** that yield an incentive-compatible mechanism. In single-parameter domains we must design **monotone approximation algorithms** with good performance.

Central issue in algorithmic mechanism design: How much social welfare is lost due to the additional requirement of incentive compatibility?

Approximation Algorithms and Mechanism Design

How well do monotone approximation algorithms perform compared to arbitrary approximation algorithms?

Approximation Ratio

- ▶ We denote by S^* an optimal subset of spots.
- ▶ **c -approximation algorithm**: Returns subset $T \subseteq I$ with

$$\sum_{i \in T} v_i \geq \frac{1}{c} \cdot \sum_{i \in S^*} v_i$$

- ▶ A trivial n -approximation:
Choose a single spot with maximum value. IC is trivial – we treat the commercial break as a single item and give it to the highest bidder (and use the second-highest valuation as payment for an IC mechanism)

Too easy – in “Theoretische Informatik 1” we proved:

Theorem

The knapsack problem has a fully-polynomial-time approximation scheme (FPTAS), i.e., for every $\varepsilon > 0$ we can compute a $(1 + \varepsilon)$ -approximate solution in time $O(n^3/\varepsilon)$.

Unfortunately, this algorithm is **not monotone** (Exercise)

Greedy-Algorithm for the Knapsack Auction

INPUT: (g_i, v_i) for every firm $i \in I$

OUTPUT: Set S of chosen spots.

1. Sort firms:

$$\frac{v_1}{g_1} \geq \dots \geq \frac{v_n}{g_n}$$

2. Set $S' \leftarrow \emptyset$ and $j \leftarrow 1$, denote $v_{\max} = \max_j v_j$

3. While $(g_j + \sum_{k \in S'} g_k) \leq G$ do:

4. $S' \leftarrow S' \cup \{j\}$ and $j \leftarrow j + 1$

5. If $v_{\max} > \sum_{k \in S'} v_k$ then $S \leftarrow \arg \max_j v_j$; else $S \leftarrow S'$

Theorem

Greedy is 2-approximate and monotone. There is an IC mechanism for the knapsack auction that guarantees at least half of the optimal social welfare.

Examples

The total length is $G = 100$ seconds.

Firm	1	2	3	4	5
v_i	45	20	45	40	50
g_i	15	25	60	50	90

After sorting in step 1 we obtain the order of firms (1,4,2,3,5):

$$45/15 \geq 40/50 = 20/25 \geq 45/60 \geq 50/90.$$

The loop in steps 2-4 computes $S' = \{1, 4, 2\}$.

In step 5

$$50 = v_{\max} < \sum_{j \in S'} v_j = 105.$$

The result is, thus, $S = \{1, 4, 2\}$ with welfare 105.

Optimum: $S^* = \{1, 2, 3\}$ with welfare 110.

Examples

The total length is $G = 100$ seconds.

Firm	1	2	3	4	5
v_i	45	20	45	40	260
g_i	15	25	60	50	90

After sorting in step 1 we obtain the order of firms (1,5,4,2,3):

$$45/15 \geq 260/90 \geq 40/50 = 20/25 \geq 45/60.$$

The loop in steps 2-4 computes $S' = \{1\}$.

In step 5

$$260 = v_{\max} > \sum_{j \in S'} v_j = 45.$$

The result is, thus, $S = \{5\}$ with welfare 260.

Optimum: $S^* = \{5\}$ with welfare 260.

2-Approximation

Proof:

We can directly observe that Greedy is monotone (Exercise).

To bound the approximation ratio we resort to the **fractional relaxation**, in which every spot i can be broken into arbitrary pieces, and we can send any **fraction** $x_i \in [0, 1]$.

For the fractional relaxation we optimize:

$$f_{\text{frac}}(v) = \arg \max_{x \in [0, 1]^n} \left\{ \sum_i x_i v_i \mid \sum_i x_i g_i \leq G \right\}$$

The fractional relaxation allows more solutions. Hence, the **optimal fractional solution** x^* can only be better than the optimal (binary) solution S^* to the knapsack problem:

$$\sum_{i \in S^*} v_i \leq \sum_{i \in I} x_i^* v_i.$$

2-Approximation

x^* yields as much value per second as possible for the ad break. Suppose the spots are numbered w.r.t. value per second $v_1/g_1 \geq \dots \geq v_n/g_n$. We choose as many seconds as possible from spot 1, then as many as possible from spot 2, then... until G seconds are chosen.

This is exactly the approach of Greedy in steps 2-4! At termination, however, the fractional solution could include an additional fraction of the next spot j' in the order:

$$\sum_{i \in I} x_i^* v_i = \sum_{k \in S'} 1 \cdot v_k + x_{j'}^* v_{j'}$$


Hence, we obtain an approximation ratio of

$$\begin{aligned} \frac{\sum_{k \in S^*} v_k}{\sum_{k \in S} v_k} &= \frac{\sum_{k \in S^*} v_k}{\max \{v_{\max}, \sum_{k \in S'} v_k\}} \leq \frac{\sum_{k \in S'} v_k + x_{j'}^* v_{j'}}{\max \{v_{\max}, \sum_{k \in S'} v_k\}} \\ &\leq 2 \cdot \frac{\sum_{k \in S'} v_k + x_{j'}^* v_{j'}}{\sum_{k \in S'} v_k + v_{\max}} \leq 2. \end{aligned}$$



Payments


In this problem, every firm gets a **binary** amount of stuff – for outcome $a \in A$, the spot i is either included ($x_i(a) = 1$) or not ($x_i(a) = 0$). Every incentive-compatible mechanism yields a **monotone, binary step function** x_i . The value where x_i jumps from 0 to 1 is called **critical value** $c_i(v_{-i})$. Obviously, it depends on the bids v_{-i} of other firms.



A normalized mechanism sets $p_i(v) = 0$ for spots i that are not included. If spot i is included, Myerson's Lemma implies $p_i(v) = c_i(v_{-i}) \cdot 1$, i.e., firm i pays (given fixed bids of other firms) her **smallest bid that guarantees inclusion of her spot**.

Payments

In this problem, every firm gets a **binary** amount of stuff – for outcome $a \in A$, the spot i is either included ($x_i(a) = 1$) or not ($x_i(a) = 0$). Every incentive-compatible mechanism yields a **monotone, binary step function** x_i . The value where x_i jumps from 0 to 1 is called **critical value** $c_i(v_{-i})$. Obviously, it depends on the bids v_{-i} of other firms.



A normalized mechanism sets $p_i(v) = 0$ for spots i that are not included. If spot i is included, Myerson's Lemma implies $p_i(v) = c_i(v_{-i}) \cdot 1$, i.e., firm i pays (given fixed bids of other firms) her **smallest bid that guarantees inclusion of her spot**.

FPTAS

Consider the fully-polynomial-time approximation scheme for the knapsack problem.

INPUT: (g_i, v_i) for every firm $i \in I$ and $\varepsilon > 0$

OUTPUT: Set S of chosen spots.

1. Let $v_{\max} = \max_i v_i$ and $s = \varepsilon \cdot v_{\max}/n$
2. Round all valuations to integers: $v'_i = \lfloor v_i/s \rfloor$
3. Solve the knapsack problem with rounded valuations v'_i optimally using dynamic programming
4. Let S' be the optimal solution for valuations v'_i
5. $S \leftarrow S'$.

Dynamic programming in step 3 takes time $O(n^2 \cdot \max_i v'_i)$. By rounding we know $v'_i \in \{0, 1, \dots, \lfloor n/\varepsilon \rfloor\}$. Thus, for constant $\varepsilon > 0$ the algorithm runs in polynomial time $O(n^3/\varepsilon)$.

Monotone FPTAS

The scheme is not monotone, because s depends on v_{\max} . If we could set the granularity in step 1 to a constant $s = \delta > 0$ independent of v_1, \dots, v_n , then the scheme would become monotone (Exercise).

Is there a single constant value δ using which we can always guarantee (without knowledge about the valuations) to obtain a $(1 + \varepsilon)$ -approximation? No!

Instead, we run the algorithm repeatedly, for **infinitely many** constant values δ . Then we choose the **best solution among all these infinitely many runs**.

The scheme is monotone in v_i for every single run. Social welfare is monotone in v_i . Therefore, choosing the best solution among all runs yields an algorithm that is **monotone in v_i** .

Infinite “FPTAS”

INPUT: (g_i, v_i) for every firm $i \in I$ and $\varepsilon > 0$

OUTPUT: Set S of chosen spots.

1. For all $k = \dots, -2, -1, 0, 1, 2, \dots$ do:
2. Set $s(k) = \varepsilon \cdot 2^k / n$
3. Round valuations: $v_i(k) = \min\{s(k) \cdot \lfloor v_i / s(k) \rfloor, 2^k\}$
4. Solve problem with rounded valuations (dyn. prog.)
5. Let $S(k)$ be the optimum solution for rounded valuations
6. Set $S \leftarrow \arg \max_{S(k)} \sum_{i \in S(k)} v_i(k)$
(tie breaking w.r.t. smaller k)

For the value $k^* = \lceil \log_2(v_{\max}) \rceil$ we see

$$\varepsilon \cdot v_{\max} / n \leq s(k^*) \leq \varepsilon \cdot 2 \cdot v_{\max} / n.$$

Hence, S'_{k^*} (and thus S) guarantees approximation ratio at most $(1 + 2\varepsilon)$.

True FPTAS

It is possible to show that the infinite scheme has to be called only for relatively few values $k \in \{k^* - \lceil \log_2 n \rceil - 2, \dots, k^*\}$. For other values of k **no better solutions** are obtained.

Hence, we **do not need infinitely** many runs. At most $\log_2(n) + 4$ **many** runs for the correct range of k suffice. The correct range of k depends on k^* and, hence, depends on v_1, \dots, v_n . But this does **not mean that we restrict k to this range** – it just means that the optimal solutions over all infinitely many constant values of k **must be located in this range**. Thus, the monotonicity arguments for constant values of k continue to hold.

For every run, dynamic programming takes time $O(n^2 \cdot \max_i v_i(k)/s(k))$. The smallest considered value $k^* - \lceil \log_2 n \rceil - 2$ yields the finest granularity and the largest bound on the running time.

True FPTAS

We see that

$$\begin{aligned} & \max_i \frac{v_i(k^* - \lceil \log_2 n \rceil - 2)}{s(k^* - \lceil \log_2 n \rceil - 2)} \\ & \leq \left\lfloor \frac{v_{\max}}{\varepsilon \cdot 2^{k^* - \lceil \log_2 n \rceil - 2} / n} \right\rfloor \leq \left\lfloor \frac{n \cdot 2^{k^*}}{\varepsilon \cdot 2^{k^* - \log_2(n) - 3}} \right\rfloor \\ & \leq \lfloor 8n^2 / \varepsilon \rfloor . \end{aligned}$$

Hence, dynamic programming needs time at most $O(n^4/\varepsilon)$ for every one of the $O(\log n)$ many runs.

Theorem

There is a monotone FPTAS for the knapsack problem with running time $O(n^4 \log n / \varepsilon)$. There are incentive-compatible mechanisms for the knapsack auction with polynomial running time, which guarantee a $1/(1 + \varepsilon)$ -fraction of the optimal social welfare, for every constant $\varepsilon > 0$.

Vickrey-Auction and Vickrey-Clarke-Groves Mechanisms

Characterization of Incentive Compatibility

Single-Parameter Mechanisms

Revelation Principle

Mechanisms and Approximation Algorithms

Revenue Maximization in Single-Parameter Domains

Revenue Maximization

We have used money only as a **means to enable incentive compatibility**. Now let's consider **money as objective** of the mechanism.

Single-Item-Auction with Single Bidder

IC mechanisms are *fixed-price mechanisms*:

- ▶ Choose price $p \geq 0$ (possibly at random) **independent of bid**.
- ▶ Sell item iff $v_i \geq p$.

Maximize social welfare: $p = 0$.

Maximize revenue: ??

For meaningful revenue maximization we need (partial) information about possible valuations of the bidders. Otherwise, the achieved revenue can be arbitrarily smaller than the optimal revenue.

Average-Case and Distributions

- ▶ Single-parameter domain for every bidder i
- ▶ **Distribution** \mathcal{V}_i for private parameter, $v_i \sim \mathcal{V}_i$
- ▶ **Vector of distributions** $\mathcal{V} = (\mathcal{V}_1, \dots, \mathcal{V}_n)$
- ▶ Private value of bidder i drawn **independently** from \mathcal{V}_i : Bidder i has the same distribution over v_i , **no matter what** values have been drawn for other bidders.
- ▶ Mechanism based on distributions, but **pointwise IC**: Truth-telling is dominant strategy for every bidder i , for every possible value v_i , and for all possible v_{-i}
- ▶ Bidder does not know distributions (i.e., any knowledge about distributions does not change incentive to tell the truth)
- ▶ Distributions matter **only in design and analysis of the mechanism**, but shall have no effect for the strategic behavior of bidders.

Distributions

The **cumulative distribution function (CDF)** $F_i(x)$ for distribution \mathcal{V}_i is

$$F_i(x) = \Pr_{v_i \sim \mathcal{V}_i} [v_i \leq x].$$

It has the **density function** $f_i(x)$, and it holds $F_i(x) = \int_{-\infty}^x f_i(x) dx$.

Example single-item auction with single bidder:

Using price p we obtain **revenue** $p \cdot (1 - F_i(p))$. Suppose \mathcal{V}_i uniform over $[0, 1]$, then $F_i(x) = x$ for $x \in [0, 1]$. Optimal revenue $1/4$ with $p = 1/2$.

Definition

An **optimal mechanism** is an incentive-compatible mechanism (f, p_1, \dots, p_n) that maximizes **expected revenue**

$$\mathbb{E}_{v \sim \mathcal{V}} \left[\sum_i p_i(v) \right].$$

Instead of analyzing payments directly, we consider a slightly different quantity.

Virtual Values

Definition

For bidder i , let v_i be the value, F_i the CDF, and f_i the density function. The **virtual value** of bidder i is

$$\varphi_i(v_i) = v_i - \frac{1 - F(v_i)}{f_i(v_i)} .$$

We have $v_i \geq \varphi(v_i)$ always. It is possible that $v_i \geq 0$ and $\varphi_i(v_i) \leq 0$.

Intuition: We would like to set v_i as price, but we have to “sacrifice” an amount of $(1 - F(v_i))/f_i(v_i)$ for truthful information.

Example with uniform distribution over $[0,1]$:

- ▶ $F(x) = x$ and $f(x) = 1$ for $x \in [0, 1]$.
- ▶ Hence: $\varphi(v_i) = v_i - (1 - v_i)/1 = 2v_i - 1$

Virtual Values and Payments

For every bidder the *expected* payments equal the *expected* virtual value.

Lemma

If (f, p_1, \dots, p_n) is an incentive-compatible mechanism in a single-parameter domain, and \mathcal{V}_i is the CDF of bidder i , then for every bidder i and every v_{-i}

$$\mathbb{E}_{v_i \sim \mathcal{V}_i}[p_i(v_i, v_{-i})] = \mathbb{E}_{v_i \sim \mathcal{V}_i}[\varphi_i(v_i) \cdot x_i(f(v_i, v_{-i}))] .$$

We will prove this lemma in the end of the section.

Instead of total payment we consider **virtual welfare** $\sum_i \varphi_i(v_i) \cdot x_i(f(v))$.

Expected Payments and Virtual Welfare

The lemma implies the main result: The *expected payments* equal the *expected virtual welfare*.

Theorem

If (f, p_1, \dots, p_n) is an incentive-compatible mechanism in a single-parameter domain, and \mathcal{V} is the vector of CDFs, then

$$\mathbb{E}_{v \sim \mathcal{V}} \left[\sum_i p_i(v) \right] = \mathbb{E}_{v \sim \mathcal{V}} \left[\sum_i \varphi_i(v_i) \cdot x_i(f(v)) \right].$$

Therefore, in order to maximize revenue we can concentrate on maximizing virtual welfare. This has a lot of similarities with maximizing social welfare.

Proof

Proof (Theorem):

We use the statement of the lemma and consider the expectation over v_{-i} :

$$\begin{aligned}\mathbb{E}_{v \sim \mathcal{V}}[p_i(v)] &= \mathbb{E}_{v_{-i} \sim \mathcal{V}_{-i}} \mathbb{E}_{v_i \sim \mathcal{V}_i} [p_i(v_i, v_{-i})] \\ &= \mathbb{E}_{v_{-i} \sim \mathcal{V}_{-i}} \mathbb{E}_{v_i \sim \mathcal{V}_i} [\varphi_i(v_i) \cdot x_i(f(v_i, v_{-i}))] \\ &= \mathbb{E}_{v \sim \mathcal{V}} [\varphi_i(v_i) \cdot x_i(f(v))] .\end{aligned}$$

Using linearity of expectation:

$$\begin{aligned}\mathbb{E}_{v \sim \mathcal{V}} \left[\sum_i p_i(v) \right] &= \sum_i \mathbb{E}_{v \sim \mathcal{V}} [p_i(v)] \\ &= \sum_i \mathbb{E}_{v \sim \mathcal{V}} [\varphi_i(v_i) \cdot x_i(f(v))] \\ &= \mathbb{E}_{v \sim \mathcal{V}} \left[\sum_i \varphi_i(v_i) \cdot x_i(f(v)) \right] . \quad \square\end{aligned}$$

Optimal Auctions

An optimal IC mechanism (maximizes expected payments, and hence) **maximizes expected virtual welfare!**

Other direction: Is a mechanism that maximizes expected virtual welfare also an optimal IC mechanism?

Yes, but only if the virtual welfare is *monotone in every v_i* , since this is necessary for the mechanism to be IC. A sufficient condition for monotone virtual welfare are regular distributions:

Definition

For a **regular** distribution \mathcal{V}_i the virtual value $\varphi_i(v_i) = v_i - \frac{1 - F_i(v)}{f_i(v)}$ is non-decreasing in v_i .

Corollary

An optimal mechanism with maximal expected revenue in a single-parameter domain with regular distributions $\mathcal{V}_1, \dots, \mathcal{V}_n$ optimizes the expected virtual welfare of the bidders.

Optimal Mechanisms for Regular Distributions

Two generalizations:

- We assume the bidders know all distributions and apply bidding strategies. They issue bids depending on (the realization of) their private value and the bidding strategies of other bidders and their (random) private values. A mechanism is *Bayes-IC* if truth-telling is a Nash equilibrium in this game (so-called Bayes-Nash equilibrium). Again, maximizing expected virtual welfare yields optimal expected revenue. For regular distributions this even yields an optimal Bayes-IC mechanism.
- For non-regular distributions there is a technique to make virtual welfare monotone (so-called *ironing*). Hence, the optimal expected revenue for non-regular distributions can be obtained by optimizing the (ironed) expected virtual welfare.

Optimal mechanisms are surprisingly simple!

Single-item auction with n bidders and possibly different regular distributions:

- ▶ Item assigned to bidder with **maximal virtual value** $\max_i \varphi_i(v_i)$. What if $\max_i \varphi_i(v_i)$ is **negative**? Then the item is **not assigned at all**.
- ▶ The value $\varphi_i^{-1}(0)$ is a **reservation price for bidder i** : v_i must be high enough to yield $\varphi_i(v_i) \geq 0$, otherwise she has no chance to get the item.
- ▶ If i gets the item, she pays the maximum of reservation price and second-highest bid – where “second-highest bid” stems from the bidder with **second-highest virtual value**. This second-highest virtual value must be translated into a **second-highest bid from i 's perspective**:
 $\max(\varphi_i^{-1}(0), \varphi_i^{-1}(\max_{j \neq i} \varphi_j(v_j)))$.
- ▶ Example with all \mathcal{V}_i identical and uniform on $[0, 1]$:
 All functions $\varphi_i(x) = 2x - 1$, all reservation prices $\varphi^{-1}(0) = 1/2$. It holds $\varphi_i^{-1}(\varphi_j(x)) = x$. The item is assigned to the highest bidder i if her bid $v_i \geq \varphi^{-1}(0) = 1/2$. Then she pays $\max(1/2, \max_{j \neq i} v_j)$. Optimal auction is a **Vickrey-Second-Price Auction with Reservation Prices!**

Proof of Lemma

Proof Sketch (Lemma):

Suppose $a(t) = f(t, v_{-i})$ for fixed bids v_{-i} . The goal is to show:

$$\mathbb{E}_{v_i \sim \mathcal{V}_i} [p_i(v_i, v_{-i})] = \mathbb{E}_{v_i \sim \mathcal{V}_i} [\varphi_i(v_i) \cdot x_i(a(v_i))] .$$

We use Myerson's Lemma. Wlog $t_i^0 = 0$, then the payments satisfy

$$\begin{aligned} p_i(v_i, v_{-i}) &= v_i \cdot x_i(a(v_i)) - \int_0^{v_i} x_i(a(t)) dt \\ &= \int_0^{v_i} t \cdot x'_i(a(t)) dt \end{aligned}$$

using integration by parts. We assume x to be differentiable. If x_i is monotone and bounded, then the proof follows with some more arguments and a suitable interpretation of the derivative x'_i .

Proof of Lemma

Step 1:

The expected revenue from bidder i given fixed bids v_{-i} is

$$\begin{aligned} \mathbb{E}_{v_i \sim \mathcal{V}_i} [p_i(v_i, v_{-i})] &= \int_{z=0}^{t_i^1} p_i(z, v_{-i}) f_i(z) dz \\ &= \int_{z=0}^{t_i^1} \left[\int_{t=0}^z t \cdot x'_i(a(t)) dt \right] f_i(z) dz \end{aligned}$$

The first equation uses independence of distributions – this implies that the fixed v_{-i} have no influence on \mathcal{V}_i .

Step 2:

We have to simplify the formula and exchange integrations:

$$\begin{aligned} \int_{z=0}^{t_i^1} \left[\int_{t=0}^z t \cdot x'_i(a(t)) dt \right] f_i(z) dz &= \int_{t=0}^{t_i^1} \left[\int_{z=t}^{t_i^1} f_i(z) dz \right] t \cdot x'_i(a(t)) dt \\ &= \int_{t=0}^{t_i^1} (1 - F_i(t)) \cdot t \cdot x'_i(a(t)) dt \end{aligned}$$

which makes the expression clearer.

Proof of Lemma

Step 3:

We again try to apply integration by parts and use

$$g(t) = (1 - F_i(t)) \cdot t \quad \text{and} \quad h'(t) = x'_i(a(t))$$

Integration by parts yields

$$\begin{aligned} \mathbb{E}_{v_i \sim \mathcal{V}_i} [p_i(v_i, v_{-i})] &= (1 - F_i(t)) \cdot t \cdot x_i(a(t)) \Big|_0^{t_i^1} \\ &\quad - \int_{t=0}^{t_i^1} x_i(a(t)) \cdot (1 - F_i(t) - t \cdot f_i(t)) dt \\ &= \int_{t=0}^{t_i^1} \left(t - \frac{1 - F_i(t)}{f_i(t)} \right) \cdot x_i(a(t)) \cdot f_i(t) dt \\ &= \int_{t=0}^{t_i^1} \varphi_i(t) \cdot x_i(a(t)) \cdot f_i(t) dt \\ &= \mathbb{E}_{v_i \sim \mathcal{V}_i} [\varphi_i(t) \cdot x_i(a(v_i))] \end{aligned}$$

as desired. □

An Alternative

Although the optimal auction is conceptually simple, it can be difficult to implement in practice. Even for selling a single item we might need up to n different reserve prices and virtual values, and, hence, exact knowledge about every CDF F_i and every density f_i .

In contrast, in the context of single-item auctions there is a simple alternative for more revenue – more competition!

The following result considers the revenue of single-item auctions with **identical regular** distributions for all bidders. We need just one extra bidder to make the revenue of the simple Vickrey auction better than the revenue of the optimal auction.

Extra Competition

Theorem (Bulow, Klemperer 1996)

Suppose \mathcal{V} is a regular distribution and $n \in \mathbb{N}$. Let p be the payments of the Vickrey second-price auction with $n + 1$ bidders and p^* the payments for the optimal (for \mathcal{V}) auction with n bidders. Then

$$\mathbb{E}_{v \sim \mathcal{V}^{n+1}} \left[\sum_{i=1}^{n+1} p_i(v) \right] \geq \mathbb{E}_{v \sim \mathcal{V}^n} \left[\sum_{i=1}^n p_i^*(v) \right].$$

Proof:

For the analysis, we rely on a **fictitious auction**:

1. Simulate the optimal n -bidder auction for \mathcal{V} on bidders $1, \dots, n$
2. If the item does not get assigned, give it to bidder $n + 1$ for free.

Obvious properties:

- ▶ The expected revenue of the fictitious auction for $n + 1$ bidders is exactly the expected revenue of the optimal auction for n bidders.
- ▶ The fictitious auction always assigns the item to exactly one bidder.

Proof

Now consider the **optimal auction for $n + 1$ bidders that must always assign the item**. This auction maximizes the expected virtual welfare (subject to the constraint that it must always assign the item). Also, the auction **always assigns the item to the bidder with highest virtual value**, even if the **best virtual value is negative**.

The Vickrey auction always assigns the item to the highest bidder. Since \mathcal{V} is regular, the bidder with highest value is also the bidder with highest virtual value. Therefore, the Vickrey auction is **precisely the optimal auction that always assigns the item**.

The fictitious auction for $n + 1$ bidders must always assign the item and obtains the revenue of the optimal auction for n bidders with distribution \mathcal{V} .

The Vickrey auction for $n + 1$ bidders has the best revenue (wrt. \mathcal{V}) of all auctions that must always assign the item. □

Literature

- ▶ Myerson. Optimal Auction Design. *Mathematics of Operations Research* 6(1):58–83, 1981.
- ▶ Mu’Alem, Nisan. Truthful Approximation Mechanisms for Restricted Combinatorial Auctions. *Games and Economic Behavior* 64(2):612–631, 2008.
- ▶ Briest, Krysta, Vöcking. Approximation Mechanisms for Utilitarian Mechanism Design. *STOC 2005*.
- ▶ Bulow, Klemperer. Auctions versus Negotiations. *American Economic Review* 86(1):180–194, 1996.
- ▶ Kirkegaard. A Short Proof of the Bulow-Klemperer Auctions vs. Negotiations Result. *Economic Theory* 28(2):449–452, 2006.

Literature

- ▶ Nisan, Roughgarden, Tardos, Vazirani. Algorithmic Game Theory, 2007. (Chapters 9 and 11)
- ▶ Roughgarden. Twenty Lectures on Algorithmic Game Theory, 2016. (Chapters 3-6)