GAME THEORY

László Gulyás
Associate Professor, ELTE, AI Department
lgulyas@inf.elte.hu

Tamás Takács
PhD student, ELTE, AI Department
tamastheactual@inf.elte.hu
tamastheactual.github.io

Game Theory

Lecture 7

Bayesian Games and Mechanism Design

ELTE
EÖTVÖS LORÁND UNIVERSITY

Stochastic Games: Formal Definition

A finite two-player **stochastic game** is:

$$\mathcal{G} = (S, A, B, P, r, \gamma)$$

State space: $S = \{1,2, \dots, |S|\}$ (finite)
Action sets: Row $A = \{1, \dots, m\}$, Column $B = \{1, \dots, n\}$
Transition kernel: $P(s' \mid s, a, b) =$ probability of next state s' given current state s and joint action (a, b)
Reward function: $r(s, a, b) =$ immediate payoff to row player
Discount factor: $\gamma \in (0,1)$

In zero-sum games, column receives $-r(s, a, b)$.

How Play Proceeds

At each round t:

1. Observe current state s_t
2. Row chooses action $a_t \in A$, Column chooses $b_t \in B$(simultaneously)
3. Players receive rewards: $r(s_t, a_t, b_t)$ to row
4. State transitions: $s_{t+1} \sim P(\cdot | \, s_t, a_t, b_t)$
5. Repeat from new state s_{t+1}

The Shapley Equation: Intuition

At state s, row wants to:

- **Maximize** current reward + discounted future value
- Knowing column will **minimize** it

This gives us the **Shapley operator** T:

$$(TV)(s) = \max_{p \in \Delta(A)} \min_{q \in \Delta(B)} \left\{ r(s,p,q) + \gamma \sum_{s'} P(s' \mid s,p,q) V(s') \right\}$$

where:

- $r(s,p,q)$ = expected immediate reward
- $\sum_{s'} P(s' \mid s,p,q) V(s')$ = expected value of next state
- $V(s')$ = value function for state s' (to be determined)

Value Iteration Algorithm

Initialize: Start with any V_0 (e.g., all zeros)

Iterate: For $k = 0, 1, 2, \dots$

$$V_{k+1}(s) = \max_p \min_q \left\{ r(s,p,q) + \gamma \sum_{s'} P(s' \mid s,p,q) V_k(s') \right\}$$

Stop: When $|V_{k+1} - V_k|_\infty < \epsilon$ (convergence tolerance)

Extract policy: The (p, q) that achieve the max-min at each state form **stationary minimax policies**

Q-Functions: An Alternative View

Define the **state-action value** (Q-function):

$$Q_V(s,a,b) = r(s,a,b) + \gamma \sum_{s'} P(s' \mid s,a,b) V(s')$$

Then the Shapley operator becomes:

$$(TV)(s) = \max_p \min_q \sum_{a,b} p(a) q(b) Q_V(s,a,b)$$

Interpretation: Q-values capture "how good is action pair (a,b) at state s?"

Familiar from reinforcement learning? Q-learning extends this to unknown games!

POSG Definition

A two-player POSG adds:

Observation sets: Ω_1, Ω_2 for each player
Observation kernels: $O_i(\, o_i \mid s, a, b, s' \,) =$ probability player i observes o_i after transition

Play:

1. True state is s_t (hidden)
2. Players choose actions based on their **observation histories** $h_t^i = (o_{i,1}, a_{i,0}, o_{i,2}, a_{i,1}, \dots)$
3. Receive private observations $o_{i,t+1} \sim O_i(\,\cdot \mid s_t, a_t, b_t, s_{t+1}\,)$

Challenge: Players must **infer** the state and **coordinate** without shared information!

Belief States

Each player maintains a **belief** over hidden states:

$$b_i(s) = \Pr\left(s_t = s \mid h_t^i \right)$$

Belief update (Bayesian):

$$b_{i,t+1}(s') \propto \sum_s P\left(s' \mid s, a_t, b_t \right), O_i\left(o_{i,t+1} \mid s, a_t, b_t, s' \right), b_{i,t}(s)$$

Problem: Each player's belief depends on their own observations only → **asymmetric information**

Consequence: Optimal policies can be very complex (history-dependent)

Game Theory

Lecture 7

Bayesian Games and Mechanism Design

1 Bayesian Games

2 Mechanism Design

3 VCG Algorithm

Incomplete Information in the Real World

Consider these scenarios:

Auctions: Bidders don't know others' valuations for the item
Negotiations: Each side is uncertain about opponent's reservation price
Job market: Employers can't observe worker ability directly
Procurement: Government uncertain about contractor's costs
Used car sales: Buyer doesn't know quality (lemons problem)

*Players have **private information** (types) unknown to others*

Incomplete Information in the Real World

Complete Information	Incomplete Information (Bayesian)
Players know all payoffs	Private types affect payoffs
Nash equilibrium	Bayesian Nash equilibrium
Extensive form	Bayesian games
Backward induction	Beliefs and types
Perfect information	Imperfect information about types

*Players have **private information** (types) unknown to others*

Why "Bayesian"?

Harsanyi's insight (1967-68): Model incomplete information as:

1. **Nature** draws type profile $\theta = (\theta_1, \dots, \theta_n)$ from common prior $p(\theta)$
2. Each player i observes only their own type θ_i
3. Players form beliefs about others' types using **Bayes rule**
4. Play proceeds as game of imperfect information

Transformation: Incomplete information → Imperfect information + chance move

Nobel Prize: Harsanyi, Nash, Selten (1994) for equilibrium analysis with incomplete info

Type Spaces and Beliefs

A **Bayesian game** is:

$$\mathcal{G} = (N, A, \Theta, p, u)$$

- $N = \{1, \ldots, n\}$: Players
- $A = A_1 \times \cdots \times A_n$: Action profiles
- $\Theta = \Theta_1 \times \cdots \times \Theta_n$: Type profiles
- $p(\theta)$: Common prior over types (often independent: θ_i)
- $u_i: A \times \Theta \to R$: Payoffs depend on actions *and* types

Private information: Each player i observes only their own type θ_i

Bayesian Nash Equilibrium

Strategy for player i: $\sigma_i: \Theta_i \to \Delta(A_i)$ (map types to action distributions)

Definition: Strategy profile $\sigma = (\sigma_1, \ldots, \sigma_n)$ is a **Bayesian Nash equilibrium (BNE)** if for all i and all $\theta_i \in \Theta_i$:

$$\sigma_i(\theta_i) \in \arg\max_{a_i \in A_i} E_{\theta_{-i} \sim p(\cdot|\theta_i)}[u_i(a_i, \sigma_{-i}(\theta_{-i}), \theta_i, \theta_{-i})]$$

Intuition: Each type of each player best responds given beliefs about others' types

Special case: When all types are public ($|\Theta_i| = 1 \; for \; all \; i$),BNE = NE

BNE: Key Properties

Existence: Bayesian Nash equilibrium exists (by similar fixed-point argument as Nash)

Computation: Generally harder than NE

- Continuous type spaces → infinite-dimensional strategy spaces
- Need to solve for function $\sigma_i(\theta_i)$ rather than finite mixed strategy

Revelation principle: Can focus on direct mechanisms where truth-telling is BNE

Example 1: First-Price Auction

Setup:

- n bidders for single item
- Each bidder i has private valuation $v_i \sim U[0,1]$ (independent)
- Bidders submit sealed bids $b_i \in [0, \infty)$
- Highest bidder wins, pays their bid
- Ties broken uniformly at random

Payoffs:

$$u_i(b_1, \dots, b_n, v_i) = \begin{cases} v_i - b_i & \text{if } b_i > \max_{j \neq i} b_j \\ 0 & \text{otherwise} \end{cases}$$

Question: What is the Bayesian Nash equilibrium?

First-Price Auction: Solution

Claim: Symmetric BNE has bidding function:

$$b^*(v_i) = \frac{n-1}{n} v_i$$

Verification: If all others bid $b^*(v)$, player i with value v_i solves:

$$\max_{b_i} \Pr[\text{ win} \mid b_i\,] \cdot (v_i - b_i)$$

Given $b_j = \frac{n-1}{n} v_j$ for $j \neq i$, we have $\Pr[\text{ win} \mid b_i\,] = \left(\frac{n}{n-1} b_i\right)^{n-1}$

FOC gives $b_i = \frac{n-1}{n} v_i$

First-Price Auction: Insights

Shading: Bid below true value (more with more bidders)

Expected revenue:

$$R = E\left[\max_i b_i\right] = \frac{n-1}{n} \cdot \frac{n}{n+1} = \frac{n-1}{n+1}$$

Winner's curse: Winner often has highest valuation, so be cautious!

Compare to second-price: We'll see later that second-price has $b^*(v_i) = v_i$ (truthful!)

Example 2: Entry Game with Incomplete Information

Setup:

- **n** potential entrants to a market
- Each has private entry cost $c_i \sim U[0,1]$
- Profit if enter: $\pi(k) = \max\{0, 1 - 0.2k\}$ where k = number of entrants
- Payoff: $\pi(k) - c_i$ if enter, 0 if stay out

Symmetric BNE: Enter if $c_i \leq c^*$ where c^* solves:

$$\pi(\text{expected \# entrants}) = c^*$$

Threshold: $c^* \approx 0.5$ for moderate n

Insight: With uncertainty, get probabilistic entry (smoother than all-or-nothing)

Sequential Rationality with Beliefs

Combine **Bayesian games** (types) with **extensive form** (timing):

Perfect Bayesian Equilibrium (PBE) is (σ, μ) where:

- **Sequential rationality**: At each information set I, strategy $\sigma_i(\cdot \mid I, \theta_i)$maximizes expected continuation payoff given beliefs μ and opponents' strategies
- **Belief consistency**: Beliefs μ derived by Bayes rule on path (given σ and prior p)
- **Off-path beliefs**: Specified to satisfy sequential rationality

Extension of:

- SPE + incomplete information
- BNE + sequential structure

Example 3: Entry Deterrence with Incomplete Information

Setup:

- Incumbent is **Tough** (likes fighting, prob p) or **Soft** (dislikes fighting, prob $1 - p$)
- Entrant chooses **In** or **Out** (doesn't know incumbent type)
- Incumbent observes entry and chooses **Fight** or **Accommodate**

Payoffs:

	Out	(In, Accommodate)	(In, Fight)
Entrant	0	2	-1
Tough	10	3	$\mathbf{5}$
Soft	10	3	$\mathbf{1}$

Tough type prefers fighting $(5 > 3)$, *Soft prefers accommodation* $(3 > 1)$

Entry Deterrence: PBE Analysis

Candidate equilibria depend on p:

Case 1: p high → Entrant stays **Out** (expects fight)

Case 2: p low → Entrant goes **In**, **Tough Fights**, **Soft Accommodates** (separating)

Case 3: p intermediate → Pooling equilibrium where **Soft** type mimics **Tough**!

- **Soft** fights with some probability to maintain tough reputation
- Entrant mixes between **In** and **Out**

Reputation effect: Soft incumbent has incentive to build tough reputation

Example 4: Job Market Signaling

Sequence:

1. **Nature** draws worker ability $\theta \in \{H, L\}$ with $\pi(H) = 0.3$
2. **Worker** observes θ, chooses education $e \in \{0, 2\}$
3. **Firm** observes e (not θ), offers wage $w \in \{w_L, w_H\}$

Payoffs: Worker gets $w - c(e, \theta)$ where $c(e, H) = e/2, c(e, L) = e$ (education costly!)

Question: When can education credibly signal high ability?

Job Market Signaling: Separating PBE

Separating candidate: H chooses $e = 2$, L chooses $e = 0$

Beliefs: $\mu(H \mid e = 2) = 1, \mu(H \mid e = 0) = 0$

Firm responses: $w(e = 2) = w, w(e = 0) = w_L$

Incentive constraints:

- H prefers education: $w_H - 1 \geq w_L$ (cheap for high type)
- L prefers no education: $w_L \geq w_H - 2$ (expensive for low type)

Separating exists if: $w_H - w_L \in [1,2]$ (education costly enough to separate, not too costly)

Even unproductive education can signal if costs differ by type!

Lecture 7

Bayesian Games and Mechanism Design

The Reverse Game Theory Problem

Classical game theory:

- **Given**: Game (players, actions, payoffs)
- **Find**: Equilibria and outcomes

Mechanism design:

- **Given:** Desired outcomes (efficiency, fairness, revenue)
- **Design:** Game rules (who moves when, what they know, how payoffs determined)
- **Find:** Mechanisms where equilibrium achieves desired outcome

"Reverse engineering": Choose rules so that rational play implements social goals

Examples of Mechanism Design

Auctions:
- **Designer**: Seller/Auctioneer
- **Goal**: Revenue efficiency
- **Constraints**: Bidders have private values

Voting:
- **Designer**: Society
- **Goal**: Aggregate preferences
- **Constraints**: Voters have private types

Public Goods:
- **Designer**: Government
- **Goal**: Efficient provision
- **Constraints**: Citizens have private values

Direct Mechanisms

A **direct mechanism** asks each player to report their type $\widehat{\theta_i}$, then:

1. **Outcome function** $f\colon \Theta \rightarrow X$ determines outcome based on reports
2. **Payment function** $t_i\colon \Theta \rightarrow R$ determines transfers

Example: In auction, ask for valuations, give item to highest reporter, charge based on reports

Why direct? Revelation principle says we can focus on these WLOG!

Revelation Principle

Theorem (Revelation Principle): Any outcome achievable by any mechanism can be achieved by a **direct mechanism** where:

1. Players truthfully report their types
2. Mechanism computes outcome based on reports

Proof sketch:

- Original mechanism: Players use strategies $s_i(\theta_i)$
- New mechanism: Ask for types, internally compute $s_i(\theta_i)$, run original mechanism
- Truthful reporting is best response (by construction)!

Focus on incentive-compatible direct mechanisms WLOG

Incentive Compatibility

Definition: A direct mechanism(f, t) is **incentive-compatible** (IC) if truth-telling is a (Bayesian Nash) equilibrium:

For all i, all θ_i, all $\widehat{\theta}_i$:

$$E_{\theta_{-i}}[u_i(f(\theta_i, \theta_{-i})) - t_i(\theta_i, \theta_{-i})] \geq E_{\theta_{-i}}\left[u_i\left(f(\widehat{\theta}_i, \theta_{-i})\right) - t_i(\widehat{\theta}_i, \theta_{-i})\right]$$

Dominant-strategy IC (DSIC): Truth-telling optimal **regardless** of others' reports (stronger!)

Bayesian IC (BIC): Truth-telling optimal in expectation over others' types

Individual Rationality

Definition: A mechanism is **individually rational (IR)** if each player's expected payoff from participating is non-negative:

$$E_{\theta_{-i}}\left[u_i\big(f(\theta_i, \theta_{-i})\big) - t_i(\theta_i, \theta_{-i})\right] \geq 0 \quad \forall i, \theta_i$$

Voluntary participation: Players must prefer participating to opting out

Ex-post IR: Non-negative for every type realization (stronger)

Interim IR: Non-negative in expectation over others' types (as above)

Efficiency

Pareto efficiency: Outcome x is PE if no reallocation makes someone better off without hurting others

Ex-post efficiency: For reported types θ, outcome $f(\theta)$ maximizes:

$$\sum_{i=1}^{n} u_i(f(\theta), \theta_i)$$

(utilitarian social welfare)

Goal: Design mechanisms that are IC, IR, **and** efficient

Challenge: Not always possible! (See impossibility theorems)

Game Theory

Lecture 7

Bayesian Games and Mechanism Design

1 Bayesian Games

2 Mechanism Design

3 VCG Algorithm

Vickrey-Clarke-Groves Mechanism

Setup: Social choice among alternatives X based on private valuations

Mechanism:

1. Each player i reports valuation function $\hat{v}_i: X \to R$
2. Choose outcome $x^* = \arg\max_{x \in X} \sum_i \hat{v}_i(x)$ (maximize reported total value)
3. Player i pays **Clarke pivot tax**:

$$t_i = \sum_{j \neq i} \hat{v}_j(x^*) - \max_{x \in X} \sum_{j \neq i} \hat{v}_j(x)$$

Intuition: Player pays the **externality** they impose on others

VCG: Dominant Strategy Incentive Compatible

Theorem: In VCG, truth-telling is a **dominant strategy**

Proof: Player i's utility when reporting $\hat{v}_i$:

$$u_i = v_i(x^*) - t_i = v_i(x^*) - \sum_{j \neq i} \hat{v}_j(x^*) + \max_x \sum_{j \neq i} \hat{v}_j(x)$$

Since $x^* = \arg\max_x \sum_j \hat{v}_j(x)$, reporting $\hat{v}_i = v_i$ maximizes:

$$v_i(x^*) + \sum_{j \neq i} \hat{v}_j(x^*)$$

This is **independent** of others' reports (only changes x^* optimally for i)

Payment is independent of own report in the right way!

VCG: Efficiency and Limitations

Efficiency: VCG is **ex-post efficient** (by construction, maximizes social welfare)

Individual rationality: May **violate** IR! (Player may pay more than they value outcome)

Budget balance: Sum of payments may not equal zero (often negative = deficit!)

Vulnerability: Can be manipulated by coalitions (not group-strategy-proof)

Bottom line: VCG achieves DSIC + efficiency, but sacrifices budget balance and sometimes IR

Example 5: Second-Price Auction (Vickrey Auction)

Special case of VCG:

- Alternatives: $X = \{1, \dots, n\}$ (who gets item)
- Valuations: $v_i(x) = v_i$ if $x = i$, else 0

Mechanism:

1. Each bidder reports $\hat{v}_i$
2. Highest bidder wins (allocate efficiently)
3. Winner pays **second-highest bid**

VCG payment:

$$t_i = \max_{j \neq i} \hat{v}_j - \max_j \hat{v}_j = (\text{2nd highest}) - (\text{1st highest}) = -(\text{2nd price})$$

Result: Bidding true value $b_i = v_i$ is dominant strategy!

Second-Price vs First-Price

Second Price:

- Equilibrium Strategy: $b^*(v_i) = v_i(truthful)$
- Revenue: $E[\text{2nd highest } v]$
- Properties: DSIC, simple

First Price:

- Equilibrium Strategy: $b^*(v_i) = \frac{n-1}{n} v_i(shading)$
- Revenue: $\frac{n-1}{n+1}$
- Properties: BIC, strategic

Revenue equivalence theorem: Both auctions yield same expected revenue (in symmetric IPV model)!

Shading in first-price exactly compensates for difference in payment rule

Example 6: Public Project

Setup:
- Decide whether to build public project (bridge, park, etc.)
- Cost c to build
- Each player i has private value v_i for project

Efficient outcome: Build if $\sum_i v_i \geq c$

VCG mechanism:
1. Ask for values $\hat{v}_i$
2. Build if $\sum_i \hat{v}_i \geq c$
3. If built, player i pays:

$$t_i = \max\left\{0, c - \sum_{j \neq i} \hat{v}_j\right\}$$

Result: Truth-telling is dominant strategy, project built if and only if efficient!

Spectrum Auctions

Problem: Allocate radio spectrum licenses to telecom companies

Challenge:

- Companies have private valuations
- Want efficient allocation (licenses to highest-value users)
- Government wants revenue

FCC spectrum auctions:

- Simultaneous ascending auction (related to VCG)
- Raised billions in revenue (e.g., $\$19B$ in 2015 AWS-3 auction)
- Achieved relatively efficient allocation

Mechanism design in practice!

Keyword Auctions (Google AdWords)

Setup:

- Advertisers bid for keywords (e.g., "car insurance")
- Multiple ad slots with different click-through rates (CTRs)
- Advertisers have private values-per-click

Generalized Second-Price (GSP) auction:

1. Rank advertisers by bid $\times$ quality score
2. i-th highest bidder gets i-th slot
3. Pays $(i+1)$-th highest bid per click

Not exactly VCG, but approximately incentive-compatible in practice

Billions per year for Google!

Kidney Exchange

Problem: Patients need kidney transplants, have willing donors, but blood types incompatible

Solution: Pairwise or chain exchanges

Mechanism design:

- Patients/donors report preferences over potential matches
- Algorithm finds efficient matching (cycles or chains)
- Top Trading Cycles algorithm is strategy-proof!

Real-world impact: Thousands of lives saved annually

Nobel connection: Al Roth (2012 Nobel) pioneered kidney exchange design

Matching Markets

Examples:

- Medical residents to hospitals (NRMP)
- Students to schools
- Workers to firms

Deferred Acceptance (Gale-Shapley):

- Participants report preferences
- Algorithm finds stable matching
- One side's optimal stable matching

Strategy-proofness: Truthful preference revelation is optimal for one side

NRMP, school choice in NYC/Boston, etc.

Impossibility Results

Gibbard-Satterthwaite Theorem: Any voting system with ≥ 3 alternatives that is:

- **Onto** (every outcome possible)
- **Deterministic**
- **Strategy-proof** (truth-telling dominant)

must be **dictatorial** (one person always decides)

Implication: Strategic voting unavoidable in most settings!

Budget Balance vs Efficiency

Impossibility (Green-Laffont): Cannot achieve all three of:

1. Ex-post **efficiency**
2. **Budget balance** ($\sum_i t_i = 0$)
3. Dominant-strategy **incentive compatibility**

Trade-offs:

- VCG: Achieves 1 & 3, sacrifices 2
- Mechanism with budget balance: Must sacrifice efficiency or DSIC

Optimal Auctions

Myerson's optimal auction: Maximize seller revenue subject to IC and IR

Result:

- Set reserve price $r > 0$ (exclude low-value bidders)
- Use virtual valuations: $\varphi_i(v_i) = v_i - \frac{1-F_i(v_i)}{f_i(v_i)}$
- Allocate to bidder with highest positive virtual valuation

Key insight: Revenue-maximizing mechanism may sacrifice efficiency!

When to Use Bayesian Game Theory

Use Bayesian games when:

- **Private information** critical to strategic interaction
- **Beliefs** about others' types matter
- **Screening** or **signaling** important

Examples:

- Auctions (private valuations)
- Negotiations (reservation prices unknown)
- Job market (worker ability private)
- Insurance (risk types private)

When to Use Mechanism Design

Use mechanism design when:

- You **control the rules** of interaction
- Want to **implement** specific outcomes (efficiency, fairness, revenue)
- Need to **elicit private information**

Examples:

- Running an auction (seller)
- Allocating resources (government, firm)
- Matching markets (clearinghouse)
- Voting systems (social planner)

Questions

1. How do beliefs about others' types affect equilibrium strategies?
2. When is separating equilibrium possible in signaling games?
3. Why is second-price auction truthful but first-price is not?
4. What trade-offs exist between efficiency and budget balance?
5. How can mechanism design improve real-world allocation problems?

Exercise 1

In first-price auction with 3 bidders and $v_i \sim U[0,1]$, verify that $b^*(v_i) = \frac{2}{3} v_i$ is BNE.

Exercise 2

In entry game with incomplete information, find the symmetric BNE threshold c^* when $n = 5$ and $\pi(k) = 1 - 0.2k$.

Exercise 3

Show that in second-price auction, bidding true value $b\_i = v\_i$ is a weakly dominant strategy.

Exercise 4

Design a VCG mechanism for allocating a public good. Show truth-telling is DSIC.

Exercise 5

Prove that if a mechanism is DSIC and efficient, it cannot be budget-balanced (for generic valuations).

Summary

- Bayesian games: type spaces, beliefs, incomplete information
- Bayesian Nash equilibrium and Perfect Bayesian equilibrium
- Mechanism design: reverse game theory
- Revelation principle and incentive compatibility
- VCG mechanism: efficiency and DSIC
- Applications: auctions, matching, public goods
- Trade-offs and impossibility results

Course Textbooks

- Bonanno, G. (2024). ***Game Theory (3rd ed.)***. University of California, Davis. Received from: GT Book
- Axelrod, R. (1984). ***The Evolution of Cooperation***. Basic Books. Received from: Axelrod Article
- Nisan, N., Roughgarden, T., Tardos, É., & Vazirani, V. V. (2007). ***Algorithmic Game Theory***. Cambridge University Press. Received from: AGT Book
- Myerson, R. B. (1991). ***Game Theory: Analysis of Conflict.*** Harvard University Press. Received from: GT Book 2F.
- Christianos et al., ***Multi-Agent Reinforcement Learning: Foundations and Modern Approaches***, 2023. Received from: MARL Book.pdf
- Shoham, Y., & Leyton-Brown, K. (2008). ***Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations***. Cambridge University Press. Received from: MARL Book.pdf
- `***nashpy***` documentation (readthedocs). Link: NashPy Docs

That's All for Today!