GAME THEORY

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

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

Game Theory

Lecture 2

Introduction to Normal-Form Games (NFGs)

1 NFGs

2 Dominance

3 Repeated NFGs

What is a Game?

- Competition? → How?
- Is every **competition** a game?
- When would **competing** not involve a "game"?

An interesting problem (Hotelling, 1929)

Two competing shops

- located along the length of a street
- selling the same good at the same price
- with customers spread equally along the street

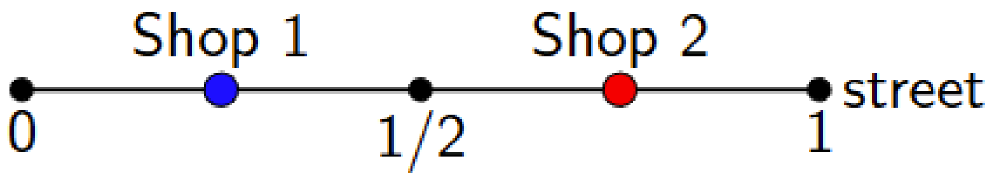

Both shop owners want

- to position their shops to be where they will get most customers

WHERE THE SHOPS WILL BE LOCATED?

Customers

- are indifferent between the shops,
- go to what is closest

Outcome depends on others

- Decisions
- Different values (valuations) of outcomes
- **Outcome / Value of outcome depends on others**

Behavioral assumptions

- Rational (maximizing / optimal) behavior

- What else?
  - **Random?**
  - **"The wilderness of irrationality"**
  - **Learning or experience**

Decision Making - Three major elements

1. **Who is in charge to make the decision? The decision maker (DM):**
   - one or
   - more

2. **What choices the DM has? Alternatives:**
   - Finitely many (discrete problem), $A_1, A_2, \dots, A_m$
   - Described by continuous variables (continuous problem), like $X = \{x \mid x \in R^m, g(x) \leq 0\}$

3. **What are the consequences of the decision?**
   - Objective functions, $\phi_1, \phi_2, \dots, \phi_n$.

Many Cases – Decision Scenarios

Number of DMs and number of objectives

	1 DM	Multiple DMs
1 objective (each)	Single objective optimization	Game
Multiple objectives	Multi-objective optimization	Pareto Game

One-off or repeated (iterated) games, etc.

Two approaches to Game Theory

Bottom-Up

- Game ⟶ (Equilibrium) Outcome

Top-Down

- Problem & (rational) DMs ⟶ Game
    - **Implementation Theory** - design the game to get desired outcomes

Which approach seems more useful to you as a student or researcher?

Game Theory

Lecture 2

Introduction to Normal-Form Games (NFGs)

1 NFGs

2 Dominance

3 Repeated NFGs

Game Theory

Game theory studies how strategic agents interact, how their choices affect each other, and how to analyze their outcomes.

- **Predict choices** in economic, social, biological systems
- **Design algorithms/protocols** in computer science, cryptography, and networks
- **Analyze stability**, efficiency, and robustness

Where do you see strategic interaction outside of games?

What is a Game?

A normal-form game (NFG) is $(N, \{A_i\}_{i \in N}, \{R_i\}_{i \in N})$ with finite N and finite action sets A_i:

- N: players
- A_i: finite actions
- $R_i: A \to \mathbb{R}$ payoff

Play: All players choose $a_i \in A_i$ at once; tuple $a = (a_1, \dots, a_n)$ yields payoffs $R_i(a)$.

Preferences and Utilities

- **Preferences** are binary relations: $x \succcurlyeq_i y$ if player i likes x at least as much as y.
- **Utility functions** U_i map outcomes to this ordering.
- **Utilities**: unique up to positive affine transformation: for any $U, U'(x) = a + bU(x)$ with $b > 0$ represents the same preferences.
- **Ordinal:** Order is all that matters (rankings).
- **Cardinal:** The magnitude of difference is meaningful (risk, expectation).

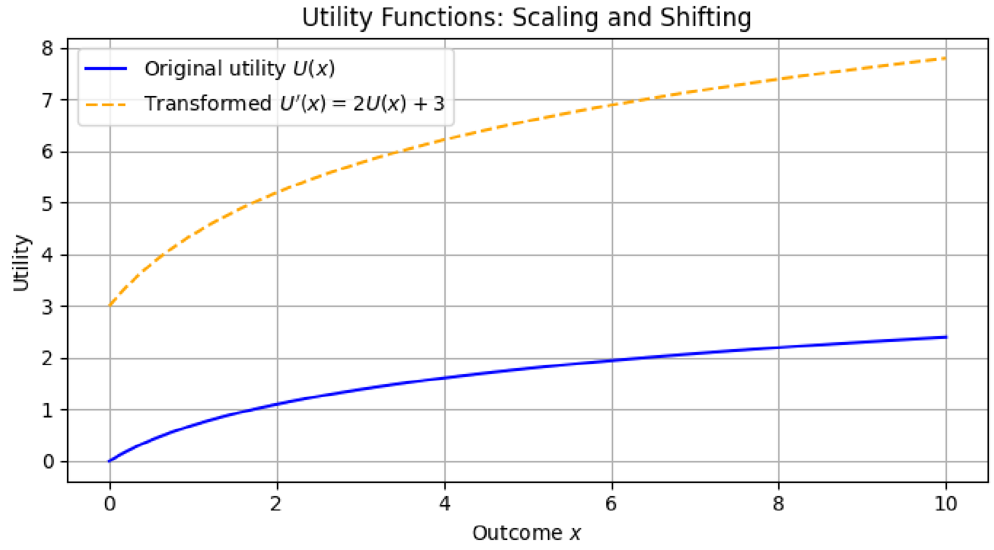

Notation Summary

- **Players:** $N,\ i\ \in N$
- **Actions:** A_i, joint space $A = \prod_i A_i$
  - In other terminology: strategies
- **Joint action:** $a = (a_1, \dots, a_n)$
- **Payoffs:** $R_i: A\ \rightarrow \mathbb{R}$
- **Strategy:** distribution π_i over A_i with $\sum_{a_i} \pi_i(a_i) = 1$
  - In other terminology: *mixed* strategy

In game theory, "strategy" can mean a choice at a single point (normal form), or more complex rules (in repeated/sequential games).

Information and Timing

- **Simultaneous move (normal-form):** All act at once, unaware of others' choices.
- **Sequential move (extensive-form):** Players can observe prior moves.
- **Complete info:** All payoffs and choices are known.
- **Incomplete info:** Hidden actions, private payoffs, uncertainty.
- **Perfect monitoring vs noisy monitoring:** Is everyone watching?

- This lecture assumes simultaneous moves and complete information in the one-shot model.

Is chess a simultaneous or sequential move game? What about email negotiations or public auctions?

Notation Summary

Feature	Example
Zero-Sum	Matching Pennies, Rock-Paper-Scissors (RPS)
General-Sum	Prisoner's Dilemma, Coordination
Common-Payoff	Pure Coordination, Group Planning
Symmetric	RPS, Battle of the Sexes (with swap)
Assymetric	"Game of Pigs", Stackelberg Duopoly

Matrix Representation and Exercise

Structure of 2-player games as payoff matrices:

Table:

	C1	C2
R1	(a,b)	(c,d)
R2	(e,f)	(g,h)

- Row player chooses row; column player chooses column. Each cell gives both payoffs.

Draw a small matrix for "Choosing a movie with a friend" scenario (2 choices each).

Battle of the Sexes

Real-Life Story:

A couple wants to spend the evening together. She prefers ballet, he prefers football. Both prefer being together over being apart, but each has their own preference.

VS.

Battle of the Sexes: The Game

	Ballet	Football
Ballet	$(2, 1)$	$(0, 0)$
Football	$(0, 0)$	$(1, 2)$

- **Row player (She):**
  - If Column plays Ballet → Row's best: Ballet $(2 > 0)$
  - If Column plays Football → Row's best: Football $(1 > 0)$
- **Column player (He)**
  - If Row plays Ballet → Column's best: Ballet $(1 > 0)$
  - If Row plays Football → Column's best: Football $(2 > 0)$

Mutual Best Responses: (Ballet, Ballet) and (Football, Football)

Battle of the Sexes: The Game

	Ballet	Football
Ballet	$(2, 1)$	$(0, 0)$
Football	$(0, 0)$	$(1, 2)$

- Coordination with conflicting preferences
- Two pure profiles are efficient, each favors a different player

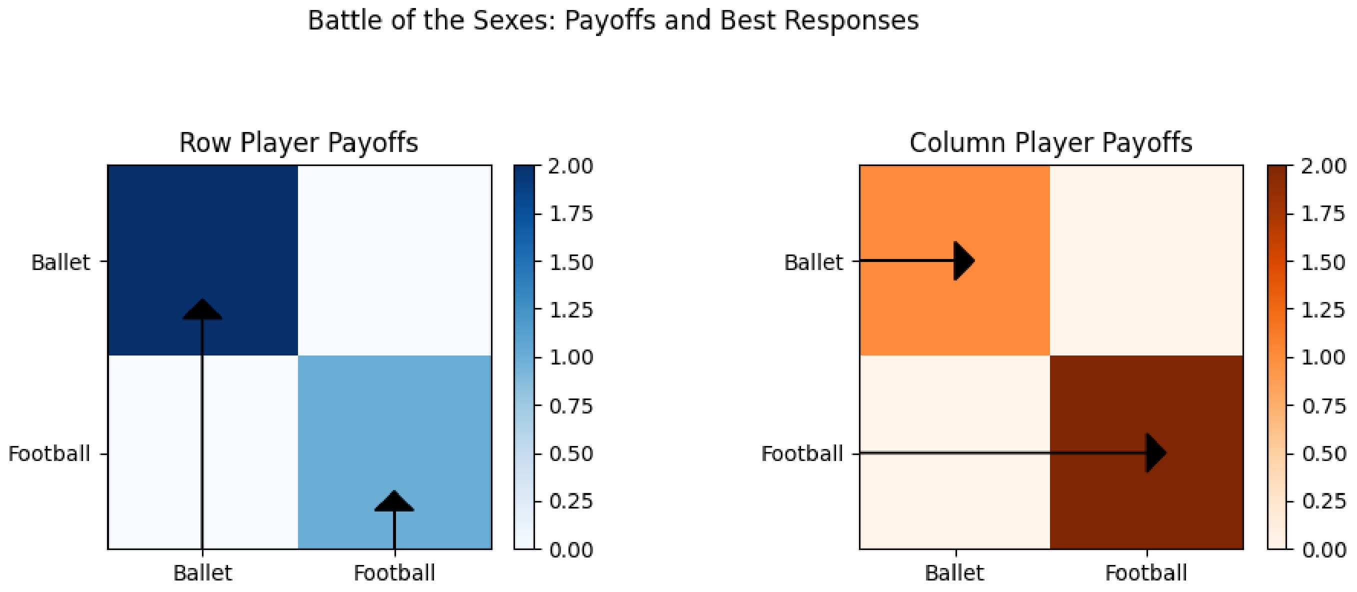

Battle of the Sexes: The Game

	Ballet	Football
Ballet	$(2, 1)$	$(0, 0)$
Football	$(0, 0)$	$(1, 2)$

1. If you're Row, how do you decide what to choose without communication?
2. Is there a "fair" outcome? Who gets what they want in each mutual best response?
3. What real-life mechanisms help people coordinate here?

Coordination Game

Real-Life Story:

Two friends want to watch something together on streaming. Alice slightly prefers sports; Bob slightly prefers comedy. Miscoordination (different choices) means they can't enjoy it together.

VS.

Coordination Game: The Game

	Bob: Sport	Bob: Comedy
Alice: Sport	$(3, 2)$	$(0, 0)$
Alice: Comedy	$(1, 1)$	$(2, 3)$

Best Response Analysis:

- **Alice (Row):**
  - If Bob plays Sport → Alice's best: Sport $(3 > 0)$
  - If Bob plays Comedy → Alice's best: Comedy $(2 > 1)$

- **Bob (Column):**
  - If Alice plays Sport → Bob's best: Sport $(2 > 0)$
  - If Alice plays Comedy → Bob's best: Comedy $(3 > 1)$

Mutual Best Responses: (Sport, Sport) and (Comedy, Comedy)

Coordination Game: The Game

	Bob: Sport	Bob: Comedy
Alice: Sport	$(3, 2)$	$(0, 0)$
Alice: Comedy	$(1, 1)$	$(2, 3)$

- Is (Sport, Sport) better than (Comedy, Comedy)? For whom?
- What coordination mechanisms exist in real life?
- If you played this repeatedly, could you establish a pattern?

Stag Hunt

Real-Life Story:

Two hunters can cooperate to hunt a stag (high reward, requires both) or individually hunt hare (lower reward, guaranteed). Based on Jean-Jacques Rousseau's philosophy.

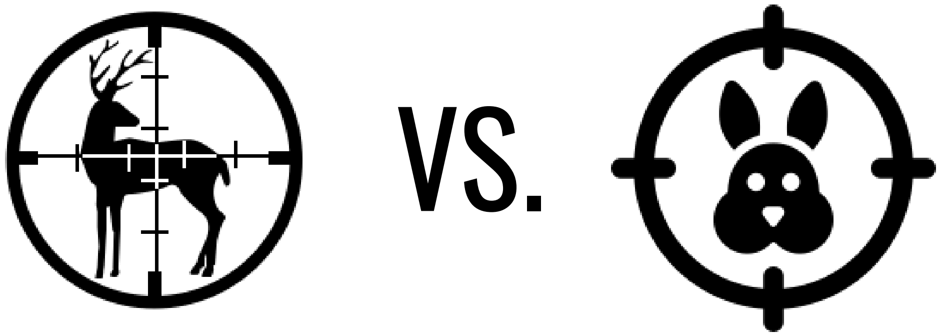

Stag Hunt: The Game

	Stag	Hunt
Stag	$(3,3)$	$(0,2)$
Hunt	$(2,0)$	$(2,2)$

Best Response Analysis:

- **Row player:**
  - If Column plays Stag → Row's best: Stag $(3 > 2)$
  - If Column plays Hare → Row's best: Hare $(2 > 0)$

- **Column player:**
  - If Row plays Stag → Column's best: Stag $(3 > 0)$
  - If Row plays Hare → Column's best: Hare $(2 > 2)$

Mutual Best Responses: (Stag, Stag) and (Hare, Hare)

Stag Hunt: The Game

	Stag	Hunt
Stag	$(3, 3)$	$(0, 2)$
Hunt	$(2, 0)$	$(2, 2)$

- Payoff-dominant vs risk-dominant equilibria
- Captures trust and assurance problems

Stag-Stag is payoff dominant (highest payoffs), but Hare-Hare is risk dominant (safer).

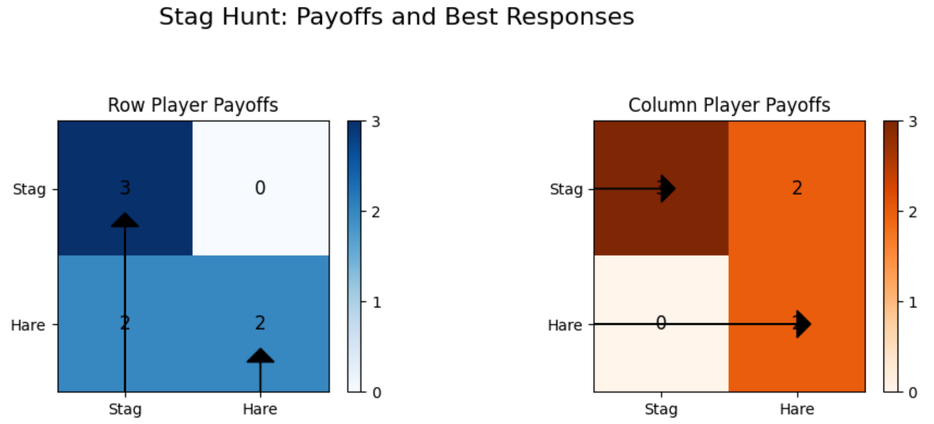

Stag Hunt: The Game

	Stag	Hunt
Stag	$(3,3)$	$(0,2)$
Hunt	$(2,0)$	$(2,2)$

1. If you don't trust your partner, what do you choose? Why?
2. What would change if you could communicate beforehand?
3. Can you think of international relations examples (climate change, arms control)?
4. How might repeated interaction affect your choice?

Chicken (Hawk-Dove)

Real-Life Story:

Two drivers speed toward each other. The first to swerve "loses face" but avoids catastrophe. If neither swerves, disaster.

Chicken (Hawk-Dove): The Game

	Swerve	Straight
Swerve	$(0, 0)$	$(-1, 1)$
Straight	$(1, -1)$	$(-10, -10)$

Best Response Analysis:

- **Row player:**
  - If Column Swerves → Row's best: Straight $(1 > 0)$
  - If Column goes Straight → Row's best: Swerve $(-1 > -10)$

- **Column player:**
  - If Row Swerves → Column's best: Straight $(1 > 0)$
  - If Row goes Straight → Column's best: Swerve $(-1 > -10)$

Mutual best responses: (Swerve, Straight) and (Straight, Swerve)

BRs are *anti-diagonal*: you want to do the opposite of opponent!

Chicken (Hawk-Dove): The Game

	Swerve	Straight
Swerve	$(0, 0)$	$(-1, 1)$
Straight	$(1, -1)$	$(-10, -10)$

- Strategic aggression and brinkmanship
- Mutual stubbornness is catastrophic

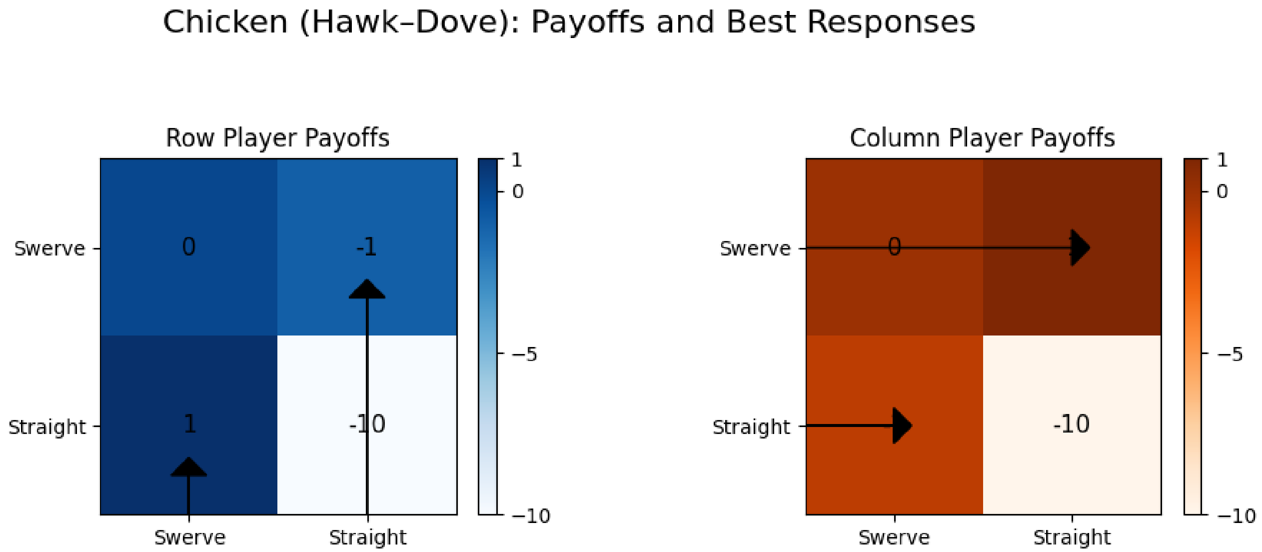

Chicken (Hawk-Dove): The Game

	Swerve	Straight
Swerve	$(0, 0)$	$(-1, 1)$
Straight	$(1, -1)$	$(-10, -10)$

1. Why is (Straight, Straight) so bad? Is it ever rational?
2. What gives a player credibility in "not swerving"?
3. Can you commit to a strategy before your opponent? How does that help?
4. Compare this to Prisoner's Dilemma. What's different about the incentives?

Matching Pennies (Zero-Sum)

Real-Life Story:

Two players simultaneously show either heads or tails of a coin. If they match, Row wins. If they differ, Column wins.

Examples from real life:

- Penalty kicks in soccer (keeper vs striker)
- Rock-paper-scissors
- Hide-and-seek, poker bluffing

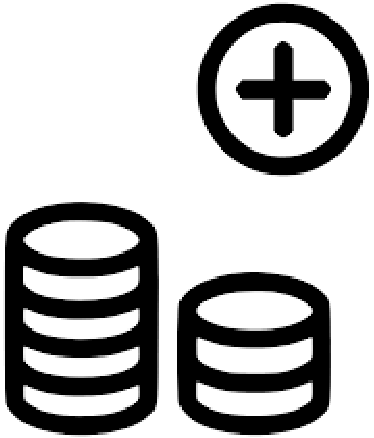

Matching Pennies: The Game

	H	T
H	$(1, -1)$	$(-1, 1)$
T	$(-1, 1)$	$(1, -1)$

Best Response Analysis:

- **Row player:**
  - If Column plays H → Row's best: H $(1 > -1)$
  - If Column plays T → Row's best: T $(1 > -1)$
- **Column player:**
  - If Row plays H → Column's best: T $(1 > -1)$
  - If Row plays T → Column's best: H $(1 > -1)$

Mutual Best Responses: NONE! BRs cycle forever.

This is a zero-sum game (payoffs sum to zero in every cell).

Matching Pennies: The Game

	H	T
H	$(1, -1)$	$(-1, 1)$
T	$(-1, 1)$	$(1, -1)$

- Zero sum
- Pure best responses cycle
- Requires mixing for stability (details later in the course)

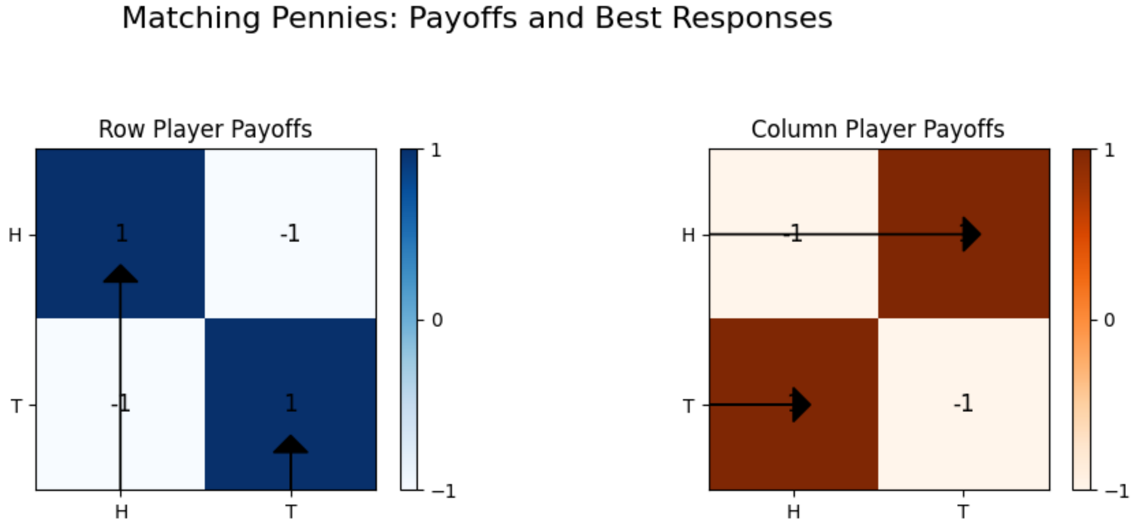

Matching Pennies: The Game

	H	T
H	$(1, -1)$	$(-1, 1)$
T	$(-1, 1)$	$(1, -1)$

- Can you ever find a "stable" pure strategy pair? Why not?
- What does this tell us about predictability?
- How would you play this game?
- Is this game fair?

Rock-Paper-Scissors (Zero-Sum)

	R	P	S
R	(0,0)	$(-1,1)$	(1,-1)
P	$(1,-1)$	(0,0)	(-1,1)
S	(-1,1)	(1,-1)	(0,0)

Best Response Analysis:

- **Row player:**
  - If Column plays Rock → Row's best: Paper $(1 > 0, -1)$
  - If Column plays Paper → Row's best: Scissors $(1 > 0, -1)$
  - If Column plays Scissors → Row's best: Rock $(1 > 0, -1)$

Column player: (By symmetry, analogous)

Mutual Best Responses: NONE! Cyclic dominance: Rock $<$ Paper $<$ Scissors $<$ Rock.

Rock-Paper-Scissors (Zero-Sum)

- Cyclic dominance
- No pure equilibrium
- No pure strategy is safe
- Each action beats one and loses to one

	R	P	S
R	(0,0)	(−1,1)	(1,-1)
P	(1, −1)	(0,0)	(-1,1)
S	(-1,1)	(1,-1)	(0,0)

Modern examples:

- Evolutionary biology (species competition cycles)
- Market competition with cyclic advantages
- Combat strategies in games (tank-infantry-artillery cycles)

Rock-Paper-Scissors (Zero-Sum)

	R	P	S
R	(0,0)	$(-1,1)$	(1,-1)
P	$(1,-1)$	(0,0)	(-1,1)
S	(-1,1)	(1,-1)	(0,0)

- Why can't you find a safe pure strategy?
- What would happen if you always played Rock?
- How is this different from Matching Pennies?
- Can you predict what a human opponent will do?

Symmetry and Relabeling

- **Symmetric game:** Swapping player identities leaves payoffs unchanged
  - Examples: Matching Pennies, RPS, symmetric Prisoner's Dilemma
- **Asymmetric game:** Players have fundamentally different roles
  - Examples: Game of Pigs, Battle of the Sexes
- Many population models use symmetric games
- Relabeling actions does not change strategic structure

For each game above, determine if it's symmetric or asymmetric. Justify your answer.

Asymmetric Game: Game of Pigs

- Two pigs: big and small
- A lever to press for food
- Food appears on the other end of pen
- There is a cost for pressing (energy)

	Big: Press	Big: Wait
Small: Press	$(4, 2)$	$(2, 3)$
Small: Wait	$(6, -1)$	$(0, 0)$

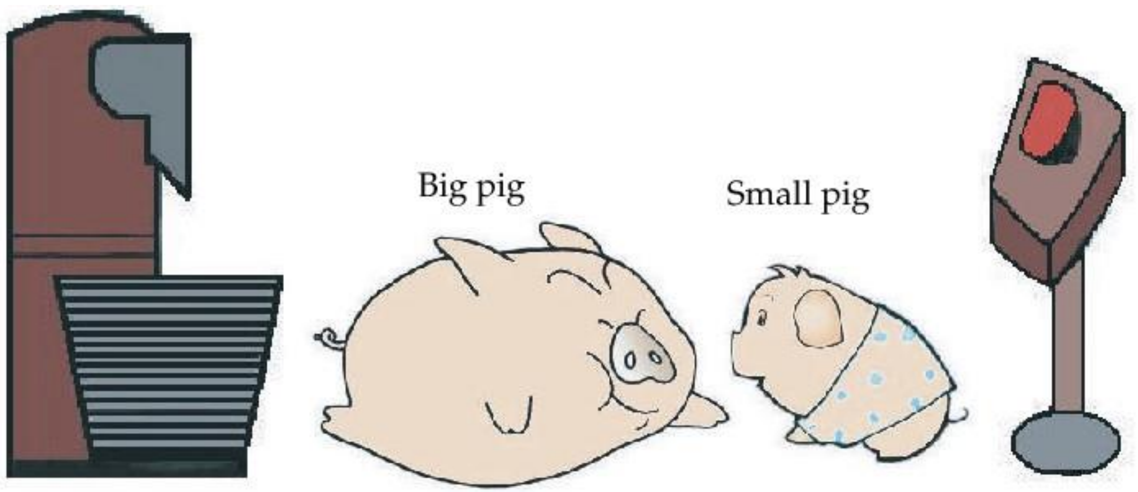

Game of Pigs: BR Analysis

	Big: Press	Big: Wait
Small: Press	$(4, 2)$	$(2, 3)$
Small: Wait	$(6, -1)$	$(0, 0)$

- **Small Pig:**
  - If Big presses: **Press** (4) vs **Wait** (6) $\rightarrow$ **Wait** is better
  - If Big waits: **Press** (2) vs **Wait** (0) $\rightarrow$ **Press** is better
  - No dominant strategy

- **Big Pig:**
  - If Small presses: **Press** (2) vs **Wait** (3) $\rightarrow$ **Wait** is better
  - If Small waits: **Press** (-1) vs **Wait** (0) $\rightarrow$ **Wait** is better
  - Dominant strategy for Big Pig: **Wait**

Mutual Best Response: (Small: **Press**, Big: **Wait**) $\rightarrow (2, 3)$

Game of Pigs: Questions

	Big: Press	Big: Wait
Small: Press	$(4, 2)$	$(2, 3)$
Small: Wait	$(6, -1)$	$(0, 0)$

1. Why does the small pig press when the big pig **Waits**, but **Wait** when the big pig presses?
2. Is the equilibrium (**Press**, **Wait**) fair? How might you adjust the numbers to share benefits more evenly?
3. Can you think of a workplace or social situation where one person does all the work and the other benefits?
4. What happens if the big pig's pressing cost increases further? When would big no longer have **Wait** as a dominant strategy?

Game of Pigs: BRs

	Big: Press	Big: Wait
Small: Press	$(4, 2)$	$(2, 3)$
Small: Wait	$(6, -1)$	$(0, 0)$

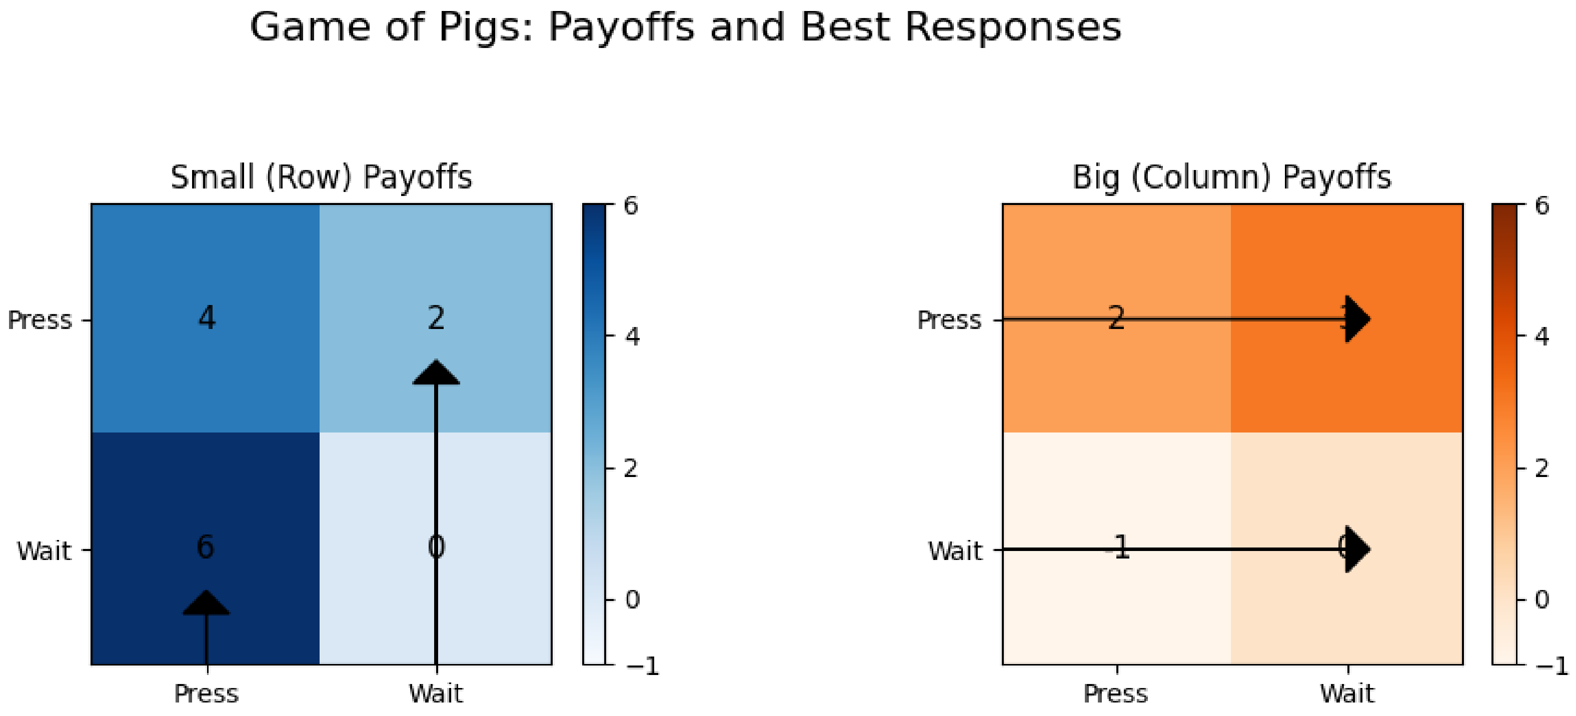

Coordination-Conflict Spectrum

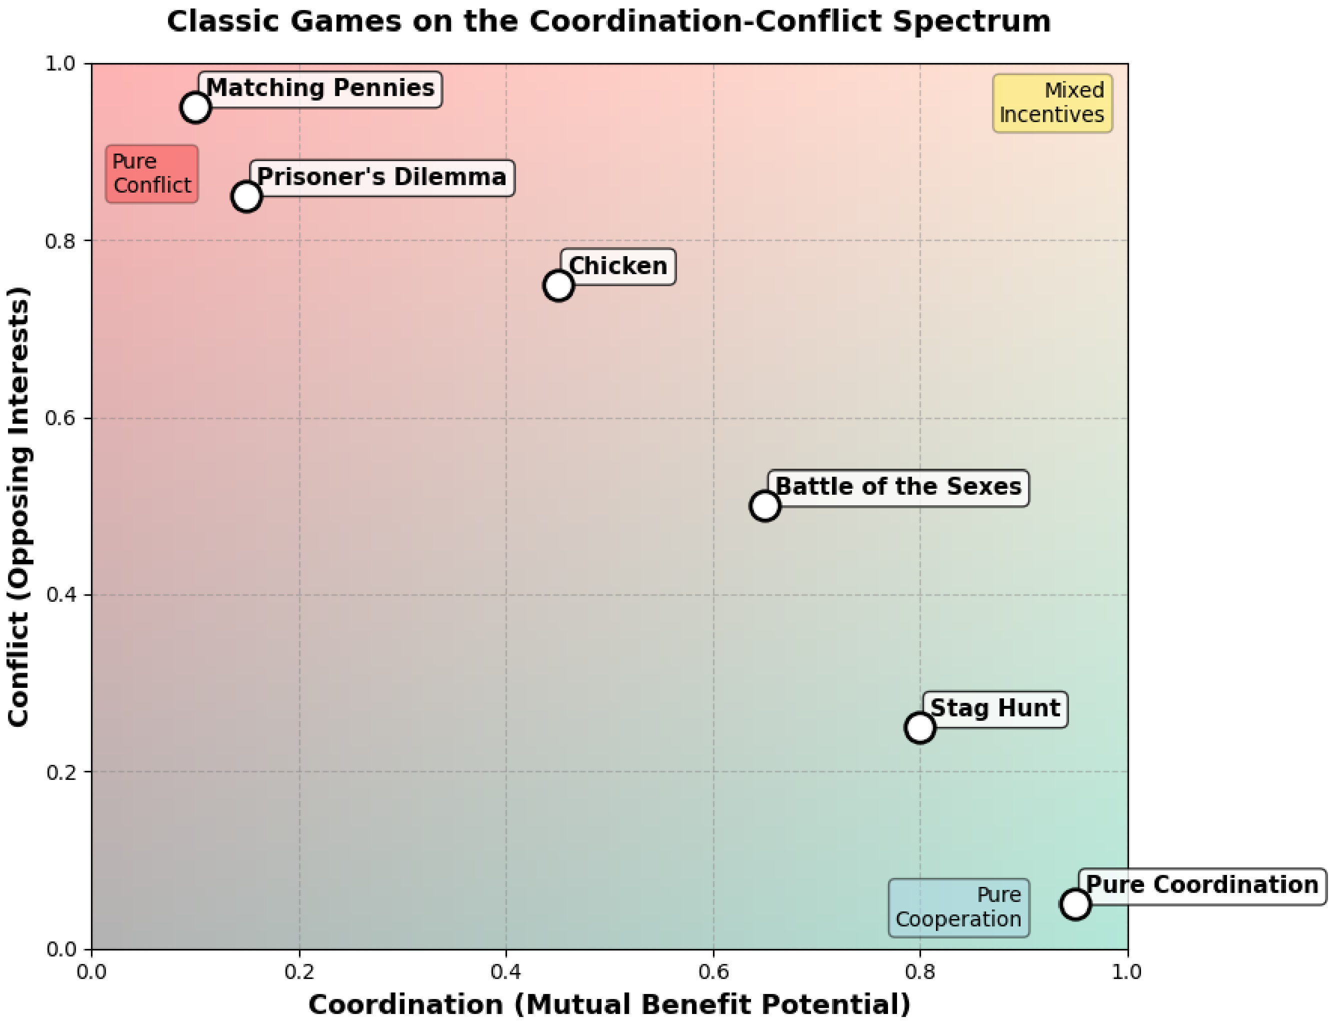

Transformations and Equivalence

- Adding a constant to R_i does not affect best responses
- Positive affine transformations preserve argmax structure
- Constant-sum vs zero-sum conversions

Interpersonal utility comparisons are not meaningful without common scale assumptions.

Game Theory

Lecture 2

Introduction to Normal-Form Games (NFGs)

1 NFGs

2 Dominance

3 Repeated NFGs

Dominance: Definitions

- Action a_i **strictly dominates** b_i if $R_i(a_i, a_{-i}) > R_i(b_i, a_{-i})$for all a_{-i}
- **Weak dominance**: $R_i(a_i, a_{-i}) \geq R_i(b_i, a_{-i})$for all a_{-i},and strictly better for some a_{-i}
- Dominated actions are never rational to play

For the Game of Pigs, identify any strictly or weakly dominated strategies for each player.

	Big: Press	**Big: Wait**
Small: Press	$(4, 2)$	$(2, 3)$
Small: Wait	$(6, -1)$	$(0, 0)$

Iterated Elimination of Strictly Dominated Strategies (IESDS)

- Remove strictly dominated actions for any player
- Repeat on the reduced game
- Order of elimination does not affect the final reduced game under strict dominance
- **Caveat:** Weak dominance can depend on elimination order

Apply IESDS to the Prisoner's Dilemma and Stag Hunt matrices below.

Dominance Practice: Prisoner's Dilemma

	Bob: C	Bob: D
Alice: C	$(3,3)$	$(0,5)$
Alice: D	$(5,0)$	$(1,1)$

- For each player, D strictly dominates C
- IESDS yields unique outcome (D,D)

Verify inequalities cell by cell.

Dominance Practice: Stag Hunt

	Stag	Hunt
Stag	$(3, 3)$	$(0, 2)$
Hunt	$(2, 0)$	$(2, 2)$

- No strict dominance
- Context for risk vs payoff dominance

Identify safe vs efficient outcomes.

Best Response: Definition

Given opponents' actions a_{-i}, the best response set is:

$$BR_i(a_{-i}) = \arg\max_{a_i \in A_i} R_i\,(a_i, a_{-i})$$

A **Best Response Correspondence maps** opponents' actions to a set of optimal actions.

Computing Best Responses by Inspection

- Fix a column (opponent action)
- Choose row that maximizes your payoff in that column
- Repeat for each column

For the Game of Pigs, underline the best responses for each player in the matrix.

	Big: Press	Big: Wait
Small: Press	$(4, 2)$	$(2, 3)$
Small: Wait	$(6, -1)$	$(0, 0)$

Best Response Maps: Coordination Game

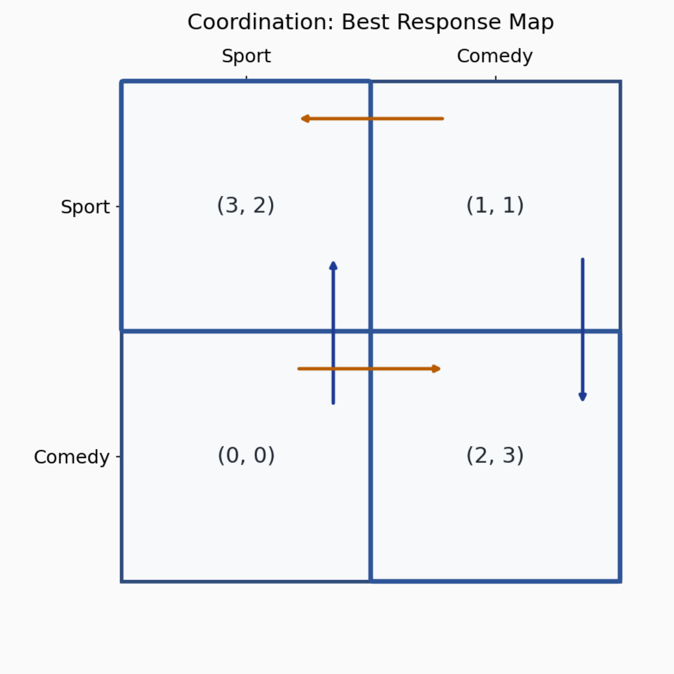

Interpretation: both players best respond by matching

- Two mutual best responses exist.

Best Response Maps: Chicken

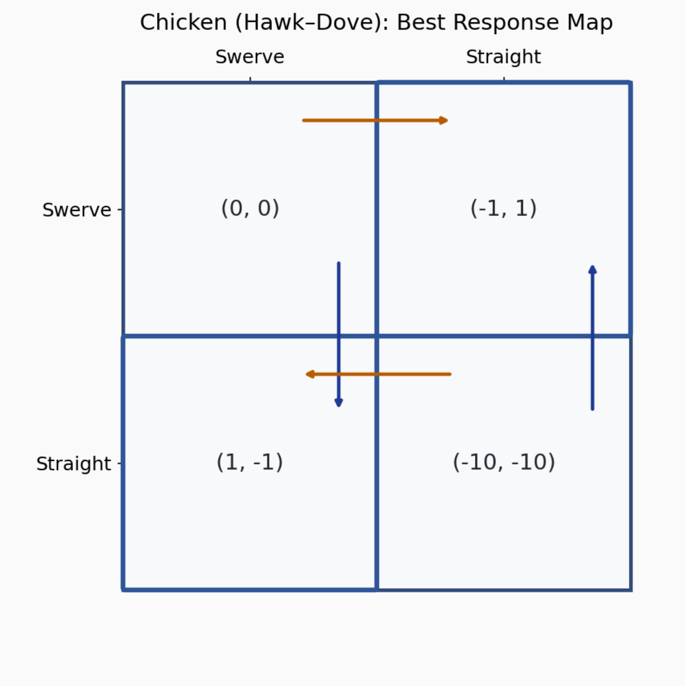

Interpretation: anti-coordination incentives

- Two mutual best responses exist; diagonal entries are not both BRs.

Best response Maps: Matching Pennies

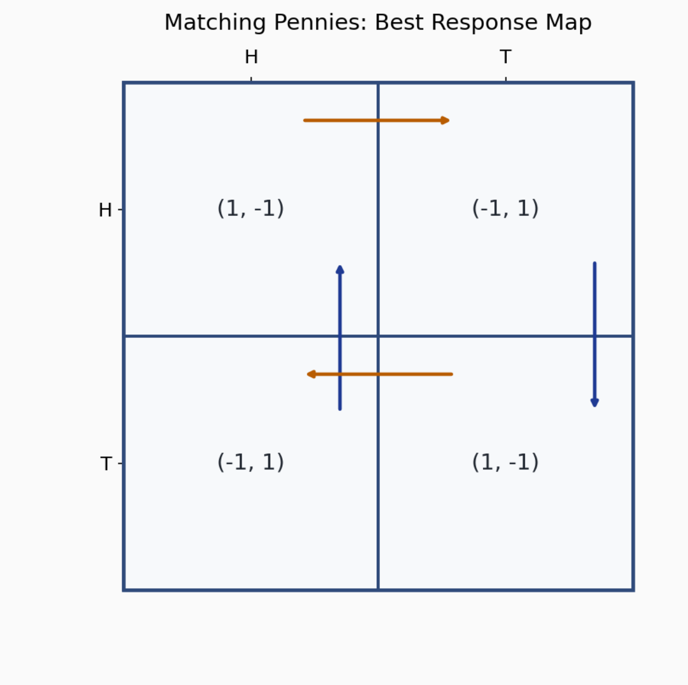

Interpretation: cycling best responses

- No mutual best response in pure strategies.

Payoff Normalization and Scaling

- Scaling by positive factor preserves BR and dominance
- Shifting by constant preserves comparisons within a player's payoffs
- Only preference ordering matters for pure-strategy reasoning

Rescale Battle of the Sexes and check BR structure.

	Ballet	Football
Ballet	$(2, 1)$	$(0, 0)$
Football	$(0, 0)$	$(1, 2)$

Battle of Sexes Normalized

Rescaled (affine) payoffs using:

- **Row:** $u' = 3u + 4$
- **Column:** $v' = 2v + 1$
  - (positive scaling + shift preserves BRs)

	Ballet	Football
Ballet	$(10, 3)$	$(4, 1)$
Football	$(4, 1)$	$(7, 5)$

Still a coordination game!

Game Theory

Lecture 2

Introduction to Normal-Form Games (NFGs)

1 NFGs

2 Dominance

3 Repeated NFGs

Repeated Normal-Form Games

Given base game $\Gamma = (N, \{A_i\}, \{R_i\})$, repeat for $t = 0,1, \dots, T-1$.

- History $h_t = (a^0, \dots, a^{t-1})$
- Strategy is a mapping from histories to actions or distributions
- Payoffs aggregated via average or discounting

Discounting and Aggregation

- Discounted return: $\sum_{t=0}^{\infty} \gamma^t r_i^t$, with $\gamma \in [0,1)$
- Average reward: $\limsup_{T \to \infty} \frac{1}{T} \sum_{t=0}^{T-1} r_i^t$

Choice affects evaluation but not one-shot best responses

- Finite vs infinite games
- Uncertainty about game ending

Axelrod's Tournament (1980)

Tournament of Iterated Prisoners' Dilemma (IPD)

- Repeated PD
  - What is the dominant strategy in IPD?
- The success of Tit-for-Tat (TFT)

Reputation and Memory in Repeated Games

- Memory-1 strategies condition only on last round
  - **Examples:** Tit-for-Tat, Grim Trigger, Pavlov (win-stay, lose-shift)

- Longer memory allows richer behavior

We do not cover learning rules today.

Examples of Repeated Strategies

- **Tit-for-Tat:**
  - Cooperate initially, then copy opponent
- **Grim Trigger:**
  - Cooperate until opponent defects once, then defect forever
- **Pavlov:**
  - Repeat previous action if you received high payoff, otherwise switch

Repeated PD Intuition

- With sufficient patience (γ high), cooperative paths can yield higher long-run payoffs
- Cooperation can be sustained by credible threat of future punishment

Formal results are covered later; focus now on interpreting incentives.

Practice Example 1

Given the matrix below, underline the best responses for each player and identify any mutual best responses.

	L	R
U	$(4, 1)$	$(0, 0)$
D	$(1, 0)$	$(2, 2)$

Practice Example 2

Classify the following game as zero-sum or general-sum and justify.

	L	R
U	$(1, -1)$	$(-1, 1)$
D	$(-1, 1)$	$(1, -1)$

Practice Example 3

Is there any strictly dominated action for either player in this game?

	L	R
U	$(2, 2)$	$(0, 3)$
D	$(3, 0)$	$(1, 1)$

Summary

- Defined games and utilities, clarified normal-form representation.
- Surveyed classic 2x2 games and their incentives, with best response analysis.
- Introduced dominance, iterated elimination, and best response correspondences.
- Outlined repeated games and history-based strategies.
- Discussed edge cases, modeling cautions, and the importance of computational tools.

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!