



ELTE
EÖTVÖS LORÁND
UNIVERSITY



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](https://github.com/tamastheactual)

Lecture 3

Best Response, Nash Equilibrium and ε -Nash

1 Best Response

2 Nash Equilibrium

3 ε -Nash

What is a Game?

A normal-form game is a formal structure for interactive decision-making

- **Players:** $N = \{1, \dots, n\}$ (list the players)
- **Action sets:** A_i for each player i (choices available)
- **Outcomes:** Joint action $a = (a_1, \dots, a_n) \in A = \prod_i A_i$
- **Payoffs:** $R_i(a)$ is how much player i values outcome a

Formal Definition

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
- $R_i: A \rightarrow R$: payoff

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

Battle of the Sexes: The Game

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

Best Response Analysis:

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

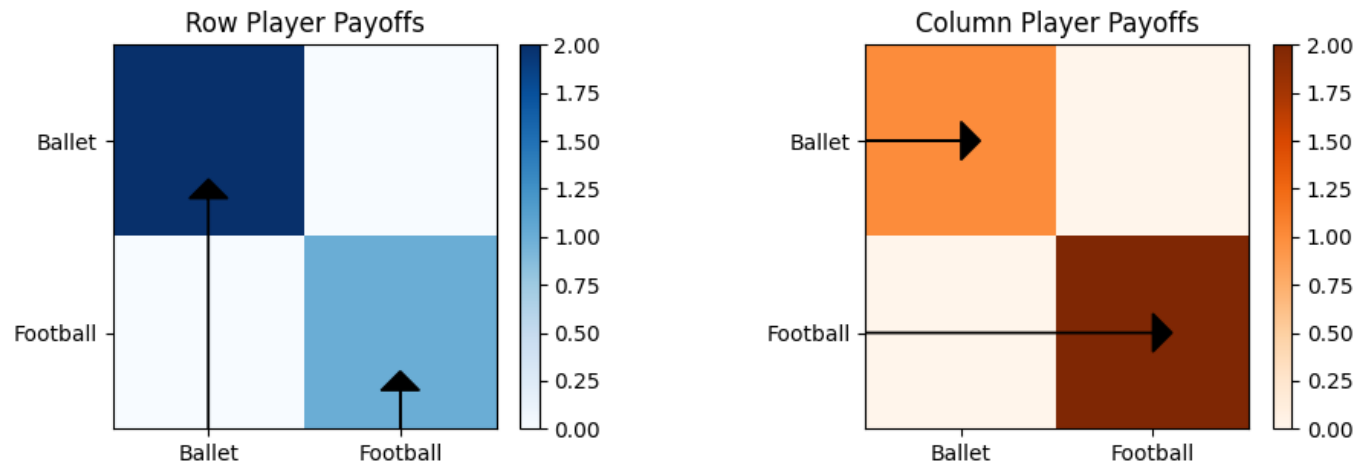
Mutual Best Responses: (Ballet, Ballet) and (Football, Football)

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: Payoffs and Best Responses



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)



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.

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

Lecture 3

Best Response, Nash Equilibrium and ε -Nash

1 Best Response

2 Nash Equilibrium

3 ε -Nash

Pure and Mixed Strategies

- **Pure strategy:** Choose a single action deterministically.
- **Mixed strategy:** Probability distribution over available actions.
- Mixed strategies expand the strategy space to convex sets (simplices).
- Many games (e.g., Matching Pennies) have **no pure NE**.
- Mixed strategies guarantee existence of equilibrium [1].

Why might randomization be rational in some games?

[1] Nash, J. F. (1950). "Equilibrium Points in n-Person Games." *Proceedings of the National Academy of Sciences*, **36**(1), 48–49.

Example: Pure vs Mixed Strategies

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

No pair of pure strategies is stable (one player always wants to deviate).

Mixed strategy solution: Each player randomizes: H with 0.5, T with 0.5.

What happens if you try to "outguess" your opponent in this game?

Best Response: Definition

For player i and opponents' mixed strategy π_{-i} ,

$$BR_i(\pi_{-i}) = \arg \max_{\pi_i \in \Delta(A_i)} R_i(\pi_i, \pi_{-i}).$$

- May be multi-valued
- Always nonempty for finite games
- Contains all optimal mixtures against π_{-i}

Best Response in Pure Strategies

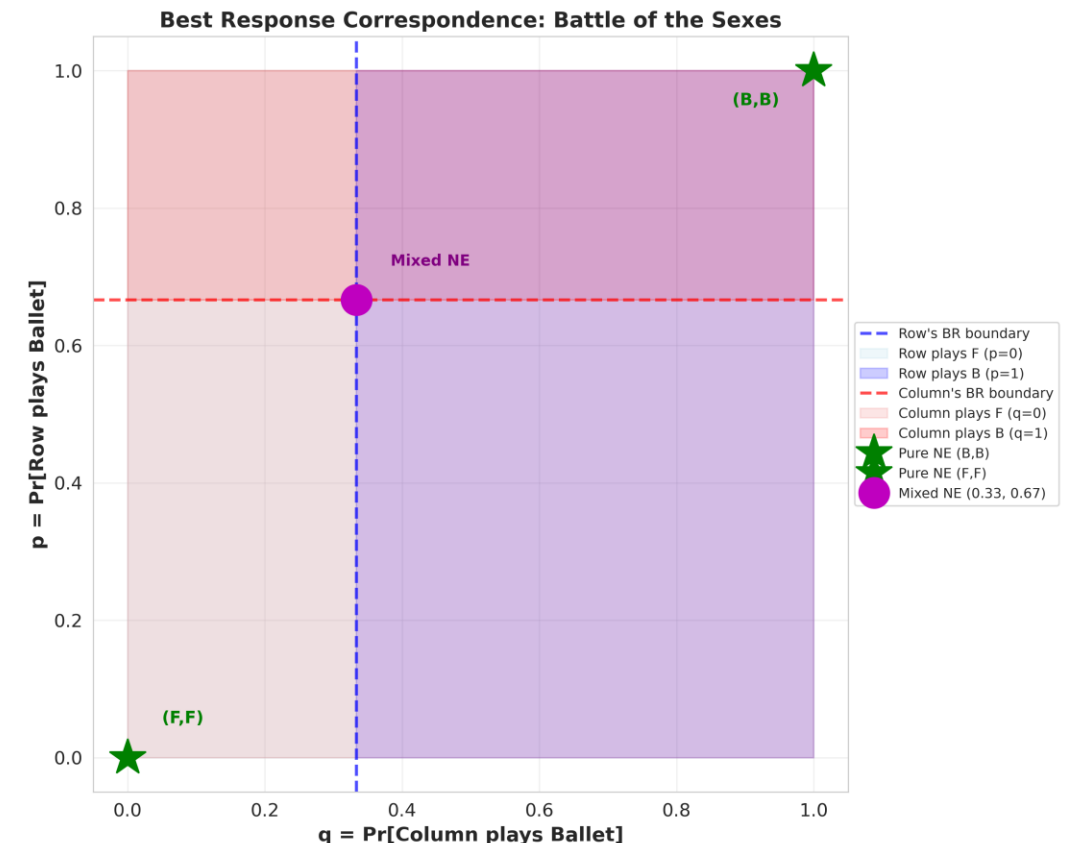
With opponents fixed at a pure action a_{-i} ,

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

- Quick to read off in a payoff matrix
- Underline the highest entry in each row or column as appropriate

Best Response as a Correspondence

- Domain: $\Delta(A_{-i})$
- Range: subsets of $\Delta(A_i)$
- For finite games: nonempty, convex-valued, upper hemicontinuous
- These properties are key to existence ideas later




Upper Hemicontinuity: Intuition

If $\pi_{-i}^k \rightarrow \pi_{-i}$ and $\pi_i^k \in BR_i(\pi_{-i}^k)$ with $\pi_i^k \rightarrow \pi_i$, then $\pi_i \in BR_i(\pi_{-i})$.

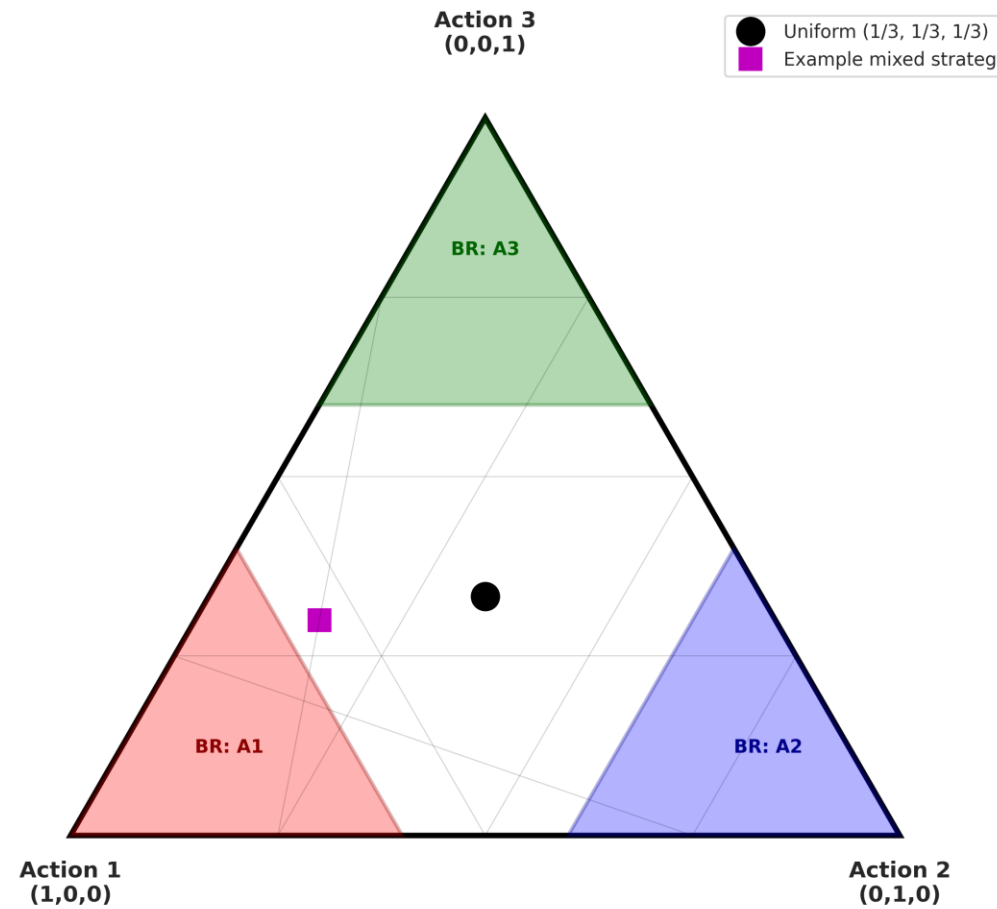
Small changes in beliefs do not create discontinuous jumps in optimal responses.

Upper Hemicontinuity of Best Response Correspondence



Visualizing BR: Probability Simplex

Probability Simplex: 3-Action Game with BR Regions



Dominance and BR

- **Strictly dominated actions are never best responses.**
- **Iterated elimination of strictly dominated strategies simplifies BR maps.**
- **Order independence holds for strict dominance.**

Example 1: Coordination BRs

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

- Mutual BR at (U,L) and at (D,R).
- Two pure NE arise.

Example 2: Prisoner's Dilemma BRs

	C	D
C	(3, 3)	(0, 5)
D	(5, 0)	(1, 1)

- D strictly dominates C for both.
- Unique mutual BR at (D,D).

Example 3: Chicken BRs

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

- Two off-diagonal pure NE when M is large.
- No dominant strategies.

Example 4: Matching Pennies BRs

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

- BRs cycle.
- No pure NE.
- Mixed NE required.

Lecture 3

Best Response, Nash Equilibrium and ε -Nash

1 Best Response

2 Nash Equilibrium

3 ε -Nash

From BR to NE: Core Idea

A strategy profile π^* is a Nash equilibrium if for all i ,

$$\pi_i^* \in BR_i(\pi_{-i}^*)$$

- No unilateral profitable deviation
- Mutual best response characterization

Nash Equilibrium: Formal Definition

In finite $(N, \{A_i\}, \{R_i\})$, π^* is a Nash equilibrium if for all i and all π_i ,

$$R_i(\pi_i^*, \pi_{-i}^*) \geq R_i(\pi_i, \pi_{-i}^*)$$

- Pure NE if each π_i^* is a point mass
- Mixed NE if some π_i^* is a distribution

NE as Fixed Point of BR

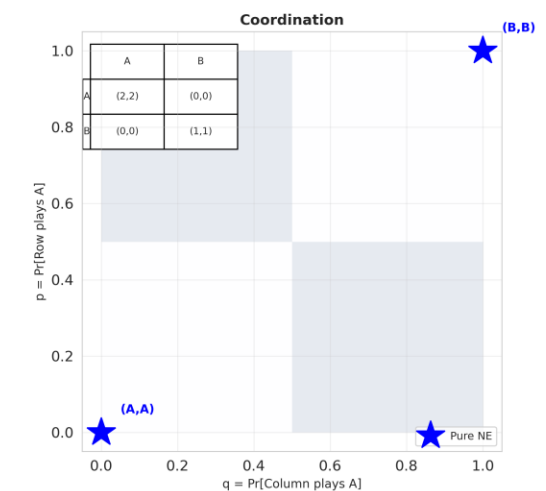
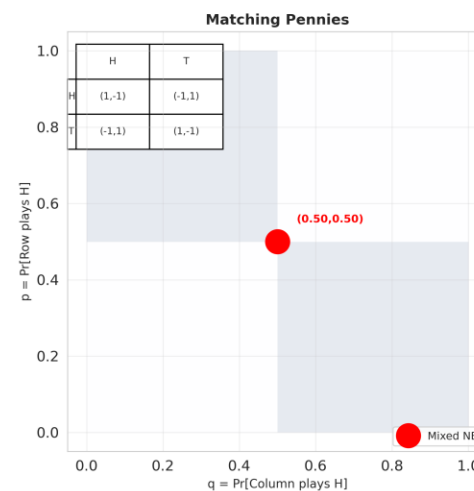
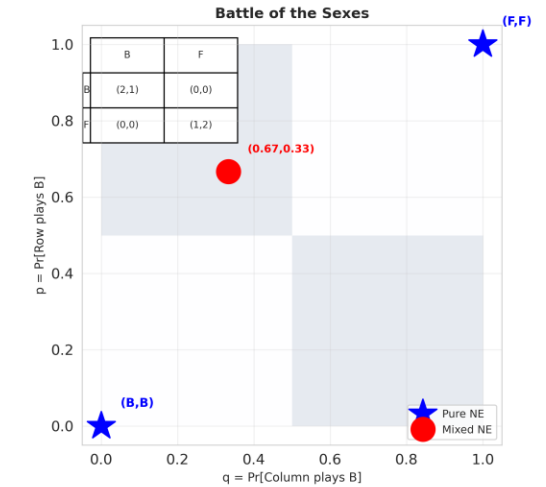
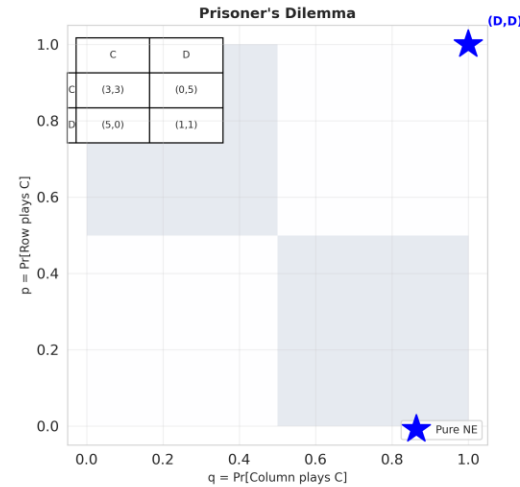
Define $BR(\pi) = \times_i BR_i(\pi_{-i})$ across players.
Nash equilibria are fixed points of BR :

$$\pi^* \in BR(\pi^*)$$

NE as Fixed Point of BR (Cont.)

Examples of Nash equilibria in classic 2×2 games:

Nash Equilibria in Classic 2×2 Games



Invariance and Normalization

- Positive affine transformations of a player's payoff preserve BR and NE
- Normalize scales for convenience
- Cross-player mixing of scales does not matter for equilibrium structure

Mixed Strategies: Expected Utility

- Extend payoffs linearly in mixed strategies
- For 2×2 , let row play U with probability p and column play L with q
- Compute expected payoffs for each pure action, then use indifference

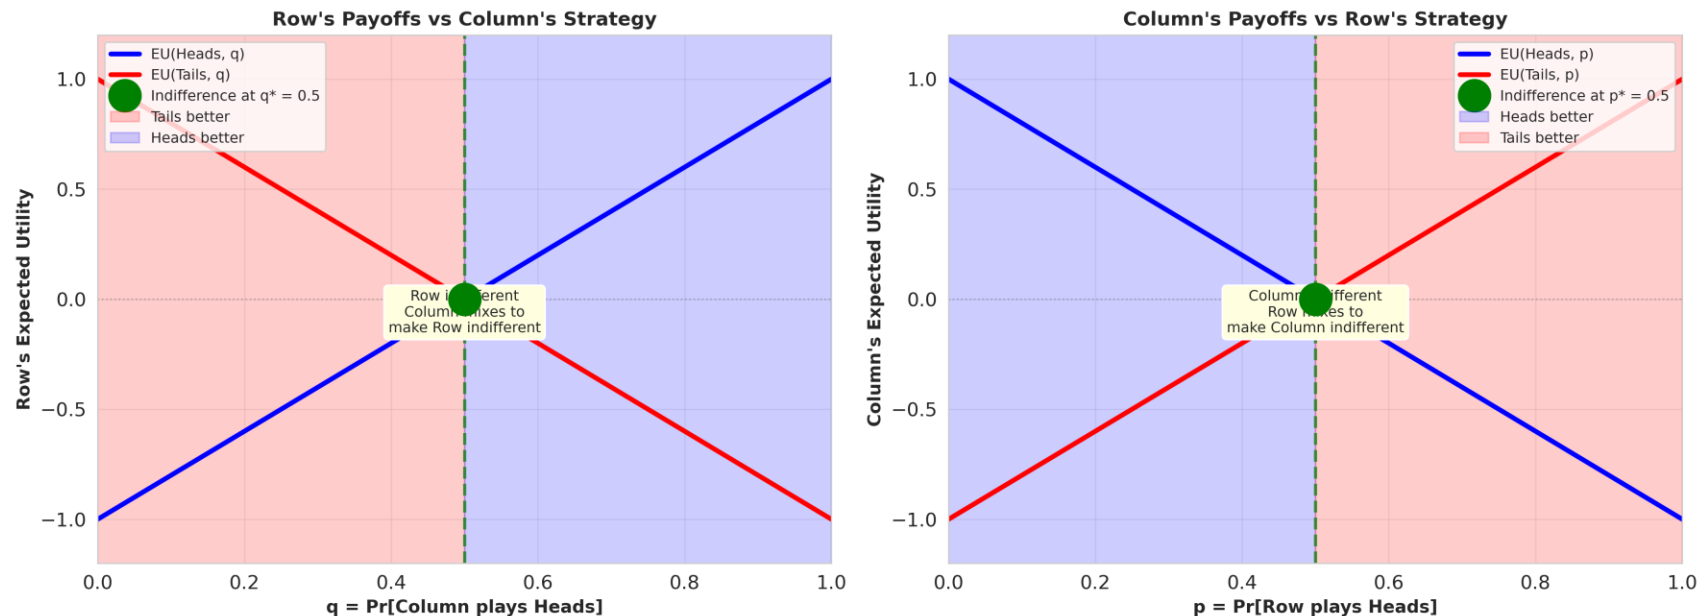
	L	R
U	(3, 2)	(0, 0)
D	(0, 0)	(2, 3)

Indifference Principle

At a mixed NE, each action in the support yields the same expected payoff.

- Equalize payoffs of supported actions.
- Out-of-support actions do not exceed that payoff.

Indifference Principle: Matching Pennies Mixed NE



Mixed NE: Matching Pennies

Row plays H with p , column plays H with q :

$$u_R(H) = 2q - 1, \quad u_R(T) = 1 - 2q \Rightarrow q = \frac{1}{2}$$

By symmetry, $p = \frac{1}{2}$

- Value: 0 for both players

Mixed NE: Battle of the Sexes

Payoffs:

- Ballet together: (2,1)
- Football together: (1,2)

Let row pick Ballet with p , column pick Ballet with q :

$$u_R(B) = 2q, \quad u_R(F) = 1(1 - q) \Rightarrow q = \frac{1}{3},$$

$$u_C(B) = 1 \cdot p, \quad u_C(F) = 2(1 - p) \Rightarrow p = \frac{2}{3}$$

Mixed NE is $(p^*, q^*) = \left(\frac{2}{3}, \frac{1}{3}\right)$.

2x2 Mixed NE: Template

For

$$A = \begin{pmatrix} a & b \\ c & d \end{pmatrix}, B = \begin{pmatrix} e & f \\ g & h \end{pmatrix},$$

let the row player play U with probability p , and the column player play L with probability q .

- **Row indifference:** $aq + b(1 - q) = cq + d(1 - q) \Rightarrow$ solve for q
- **Column indifference:** $ep + g(1 - p) = fp + h(1 - p) \Rightarrow$ solve for p

Mutual BR Characterization

- In **pure strategies**: a cell is an NE if it is a best response for both players
- In **mixed strategies**: supports produce equal expected payoffs

Degeneracy and Tie Breaking

- Degenerate equilibria have multiple BR at the boundary
- Small perturbations can select a unique equilibrium
- Good practice for robustness checks

Equilibrium Selection

	Description	Example
Risk Dominance	Largest "basin of attraction". Risk of others' mistakes → risk-dominant strategy	Hare in Stag Hunt: lower, but guaranteed payoff
Payoff Dominance	Payoffs are as good as in other NEs, but someone is strictly better off	Stag in Stag Hunt: significantly higher reward, if successful

Zero-Sum Preview: Minimax

- In 2-player zero-sum, NE equals minimax solution.
- Value of the game is the row player's equilibrium payoff.
- Solvable by linear programming.
- Details in a later week.

Linear Programming Formulation (Row Player)

Let A be the row player's payoff matrix. Solve

$$\max_{x,v} v \quad \text{s.t.} \quad A^T x \geq v \mathbf{1}, x \geq 0, \mathbf{1}^T x = 1$$

- The dual problem gives the column player's formulation
- The solution returns both the equilibrium strategy and the game value

Battle of the Sexes (2×2 , mixed + 2 pure)

```
import numpy as np, nashpy as nash

A = np.array([[2,0],[0,1]])
B = np.array([[1,0],[0,2]])
G = nash.Game(A,B)

def show(eqs):
    for sa, sb in eqs:
        payA, payB = G[sa, sb]
        print(f"Row: {np.round(sa,3)}, Col: {np.round(sb,3)} ->
              f"payoffs ({payA:.3f}, {payB:.3f})")
eqs = list(G.support_enumeration())
show(eqs)
# Mixed NE should be (p,q) = (2/3, 1/3)
```

- Enumerates supports and checks indifference
- Returns pure and mixed equilibria where they exist

Chicken / Hawk–Dove (2×2 , mixed + 2 pure)

```
import numpy as np, nashpy as nash

A = np.array([[ 0,-1],[ 1,-10]])
B = np.array([[ 0, 1],[-1,-10]])

G = nash.Game(A,B)
for sa, sb in G.support_enumeration():
    payA, payB = G[sa, sb]
    print(f"Row: {np.round(sa,3)}, Col: {np.round(sb,3)} ->
    \\    ({payA:.3f}, {payB:.3f})")
# Expect two pure off-diagonal + one interior mixed.
```

Matching Pennies (2×2 zero-sum, unique mixed only)

```
import numpy as np, nashpy as nash

A = np.array([[ 1,-1],[-1, 1]])
G = nash.Game(A, -A)

for sa, sb in G.support Enumeration():
    payA, payB = G[sa, sb]
    print(f"(Row,Col) = ({np.round(sa,3)}, {np.round(sb,3)}) -> \\
          {payA:.3f}, {payB:.3f}")
# Unique mixed: ([0.5, 0.5], [0.5, 0.5])
```

Rock–Paper–Scissors (3×3 zero-sum, unique mixed)

```
import numpy as np, nashpy as nash

A = np.array([[ 0,-1, 1],
              [ 1, 0,-1],
              [-1, 1, 0]])
G = nash.Game(A, -A)

for sa, sb in G.vertex_enumeration():
    payA, payB = G[sa, sb]
    print(f"Row: {np.round(sa,3)}, Col: {np.round(sb,3)} ->
    \\    ({payA:.3f}, {payB:.3f})")
# Unique mixed: each player (1/3, 1/3, 1/3)
```

Non-uniform mixed in 3×3 zero-sum

```
import numpy as np, nashpy as nash

# Slightly biased RPS; equilibrium stays interior but not uniform.
A = np.array([[ 0, -1,  1.2],
              [ 1.0, 0, -1],
              [-1.2, 1,  0]])
G = nash.Game(A, -A)

for sa, sb in G.vertex_enumeration():
    payA, payB = G[sa, sb]
    print(f"Row: {np.round(sa,3)}, Col: {np.round(sb,3)} -> \\  
      ({payA:.3f}, {payB:.3f})")
# Expect mixed with probabilities != 1/3
```

General-sum 3×3 with a support-size-2 mixed NE

```
import numpy as np, nashpy as nash

A = np.array([[3, 0, 2],[0, 2, 3],[2, 3, 0]], dtype=float)
B = np.array([[2, 3, 0],[3, 0, 2],[0, 2, 3]], dtype=float)
G = nash.Game(A,B)

for sa, sb in G.support_enumeration():
    payA, payB = G[sa, sb]
    support_row = np.flatnonzero(sa > 1e-9)
    support_col = np.flatnonzero(sb > 1e-9)
    print(f"Row: {np.round(sa,3)} (supp {support_row}), "
          f"Col: {np.round(sb,3)} (supp {support_col}) "
          f"-> ({payA:.3f}, {payB:.3f})")
# Typically finds a mixed with support size 2 for each.
```

Validating Returned Equilibria

Given (π_1, π_2) :

1. Compute $u_i = R_i(\pi_1, \pi_2)$.
2. Compute $u_i^{BR} = \max_{\pi'_i} R_i(\pi'_i, \pi_{-i})$.
3. Check that $u_i^{BR} - u_i = 0$ up to tolerance.


1. Banach Fixed Point Theorem (Contraction Mapping Theorem)

Statement:

Let (X, d) be a nonempty complete metric space and $f: X \rightarrow X$ a contraction (i.e., there exists $0 < c < 1$ such that $d(f(x), f(y)) \leq c \cdot d(x, y)$ for all x, y). Then there exists a unique fixed point x^* in X such that $f(x^*) = x^*$.

- Fundamental for proving convergence, but *not* generally used for Nash equilibrium existence, because game-theoretic best response correspondences are not contractions or even single-valued.

1. Banach Fixed Point Theorem



Attracting fixed point: $|f'(x^*)| < 1$.
 Repulsive fixed point: $|f'(x^*)| > 1$.

Theorem: If (\mathcal{X}, d) is a complete metric space,
 $f : \mathcal{X} \mapsto \mathcal{X}$ with $d(f(x), f(y)) \leq \kappa d(x, y)$ and $\kappa < 1$
 then $\exists! x^*$ such that $f(x^*) = x^*$.

$$x_{k+1} \stackrel{\text{def.}}{=} f(x_k) \xrightarrow{k \rightarrow +\infty} x^*$$

$$d(x_k, x^*) \leq \kappa^k d(x_0, x^*)$$

2. Brouwer Fixed Point Theorem


Statement:

Let $S \subseteq \mathbb{R}^n$ be nonempty, compact, and convex. If $f: S \rightarrow S$ is continuous, then there exists $x^* \in S$ such that $f(x^*) = x^*$.

- Applies to single-valued, continuous functions.
- Nash's early proofs for continuous strategies (but not for general, set-valued best responses) used Brouwer.

Most best response functions in games are not continuous or are set-valued, so Brouwer does not directly apply in general.

2. Brouwer Fixed Point Theorem



3. Kakutani Fixed Point Theorem

Statement:


Let $S \subseteq R^n$ be nonempty, compact, and convex. Let $\Phi: S \rightarrow 2^S$ be an upper hemicontinuous set-valued function (i.e., a correspondence) such that

1. For all $x \in S$: $\Phi(x)$ is nonempty, convex, and closed.
2. The graph of Φ is closed: If $x^k \rightarrow x$, $y^k \rightarrow y$, and $y^k \in \Phi(x^k)$, then $y \in \Phi(x)$.

Then: There exists $x^* \in S$ such that $x^* \in \Phi(x^*)$.

- This is the key fixed-point theorem for proving Nash equilibrium in games with mixed strategies, because the best response correspondence may be set-valued and upper hemicontinuous.

3. Kakutani Fixed Point Theorem



Relationship & Hierarchy

- **Brouwer** is a special case of **Kakutani** (when the correspondence is single-valued, i.e., a function).
- **Banach** is fundamentally different, deals with iterative contractions. Not generally present in game theoretic contexts.
- In Nash's context, **Kakutani** is needed due to set-valuedness of the best response correspondence.

Nash Equilibrium Existence: Proof Sketch (Cont.)

Verification of Kakutani's Conditions:

- For each player, $BR_i(\sigma_{-i})$ is nonempty (maximum of a continuous function over a compact simplex exists).
- **Convexity:** The set of maximizers of a linear function over a simplex is convex.
- **Upper hemicontinuity (Closed Graph):** Follows from Berge's Maximum Theorem because the payoff functions are continuous.
- **Strategy space** is a nonempty, compact, convex subset of R^k .

Applying Kakutani:

- By Kakutani's Theorem, there exists $\sigma^* \in S$ with $\sigma^* \in BR(\sigma^*)$.
- Therefore, every finite game has at least one mixed Nash equilibrium.

Why Is Proving Existence Hard?

- **Best response** maps are *set-valued* (correspondences), not single-valued functions.
- **Continuity issues:** Best response is not continuous as a function.
- **Compactness** and **convexity** of the strategy space is essential; non-convexity can lead to non-existence.
- **Upper hemicontinuity** and nonempty convex values are subtle properties, mathematically nontrivial to prove for arbitrary correspondences.
- Early attempts (pure strategies only) often fail, since no NE need exist in pure strategies.

Upper Hemicontinuity and Graph

- The graph of a set-valued function $\Phi: S \rightarrow 2^S$ is $\{(x, y) \mid x \in S, y \in \Phi(x)\}$.
- Φ is *upper hemicontinuous* at x if, whenever $x^k \rightarrow x$ and $y^k \in \Phi(x^k)$ with $y^k \rightarrow y$, it follows that $y \in \Phi(x)$.
- This ensures "no sudden jumps," which is crucial for fixed point existence.

Intuition: Why Mixed Strategies?

- The set of pure strategies is not convex or compact, so fixed-point theorems do not apply.
- Mixed strategies complete the set, ensuring compactness and convexity of the feasible region.

Table: Summary of Theorems

Theorem	Domain Type	Map Type	Existence	Uniqueness
Banach	Complete metric space	Contraction	Yes	Yes
Brouwer	Compact convex subset	Continuous function	Yes	No
Kakutani	Compact convex subset	Upper hemicontinuous correspondence	Yes	No

In what games is the best response function not single-valued? Give an example.

Why does convexity of the mixed strategy space matter for the application of Kakutani's theorem?

Can you give a simple set-valued function that violates upper hemicontinuity, and show what goes wrong?

Lecture 3

Best Response, Nash Equilibrium and ε -Nash

1 Best Response

2 Nash Equilibrium


3 ε -Nash

Epsilon-Nash: Definition

A profile π is an ϵ -Nash equilibrium if for all i ,

$$R_i(\pi) \geq \max_{\pi'_i} R_i(\pi'_i, \pi_{-i}) - \epsilon$$

- $\epsilon = 0$ gives an exact Nash equilibrium
- Useful when using numerical solvers or rounding



Why Epsilon-Nash Matters

- Rounding effects in computation
- Approximate rationality in practice
- Many algorithms converge to epsilon-NE rather than exact NE

Measuring Epsilon in Finite Games

For each player i :

1. Compute u_i at (π_1, π_2)
2. Compute the best pure-response payoff u_i^{BR}
3. Set $\varepsilon_i = u_i^{BR} - u_i$

Report $\max_i \varepsilon_i$

Example: Epsilon for Rounded Matching Pennies

True NE: $(p, q) = (0.5, 0.5)$

Use $(p, q) = (0.55, 0.45)$ for both players.

For the row player:

If Row plays H :

$$u_R(H) = 0.45 \cdot 1 + 0.55 \cdot (-1) = -0.1$$

If Row plays T :

$$u_R(T) = 0.55 \cdot 1 + 0.45 \cdot (-1) = 0.1$$

If Row mixes with $p = 0.55$:

$$u_R = 0.55 \cdot (-0.1) + 0.45 \cdot (0.1) = -0.01$$

Example: Epsilon for Rounded Matching Pennies (Cont.)

Best-response payoff is 0.1 (by playing T).

Gap is $0.1 - (-0.01) = 0.11$

By symmetry, the same holds for the column player.

NashPy: Compute Epsilon

```
import numpy as np

A = np.array([[1,-1],[-1,1]])
pi_row = np.array([0.55, 0.45])
pi_col = np.array([0.45, 0.55])

row_pure_payoffs = A @ pi_col
col_pure_payoffs = (-A).T @ pi_row

row_gap = row_pure_payoffs.max() - row_pure_payoffs @ pi_row
col_gap = col_pure_payoffs.max() - col_pure_payoffs @ pi_col
print("eps_row=", row_gap, "eps_col=", col_gap)
```

Parametric Games

- Vary a parameter in payoffs and track how the NE moves
- Useful for comparative statics
- **Example:** scale a penalty in Chicken and observe the mixed NE threshold

Parametric Example: Battle of the Sexes

Let Ballet payoffs be $(2,1)$ and Football payoffs be $(\alpha, 2)$.

- Row indifference fixes $q(\alpha)$
- Column indifference fixes $p(\alpha)$
- Plot $p(\alpha)$ and $q(\alpha)$ to visualize shifts in mixing

Parametric Example: Solution

- Row Indifference ($q(\alpha)$):

Set Row's expected payoff from Ballet and Football equal

$$q(\alpha) = \frac{\alpha}{2 + \alpha} q(2 + \alpha) = \alpha 2q + \alpha q = \alpha 2q = \alpha(1 - q)$$

- Column Indifference ($p(\alpha)$):

Set Column's expected payoff from Ballet and Football equal

$$p(\alpha) = \frac{2}{3} 3p = 2p + 2p = 2p = 2(1 - p)$$

Parametric Example: Solution (Cont.)

- Row's mixing probability: $q(\alpha) = \frac{\alpha}{2+\alpha}$
- Column's mixing probability: $p(\alpha) = \frac{2}{3}$

*As α increases, $q(\alpha)$ increases from 0 toward 1.
 $p(\alpha)$ remains fixed at $2/3$.*

Practice 1

Identify all mutual BR cells in:

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

1. Does a pure NE exist?
2. If not, solve for mixed NE.
3. Normalize payoffs and confirm NE unchanged.

Practice 2

1. Construct a 2×2 general-sum game with exactly one mixed Nash equilibrium.
2. For the zero-sum game with payoff matrix $\begin{pmatrix} 2 & -1 \\ -3 & 4 \end{pmatrix}$, compute the value and the equilibrium mixing.
3. In Chicken, replace $(-10, -10)$ with $(-M, -M)$. Find the threshold M that yields a mixed Nash equilibrium.

Practice 3

- Use NashPy to compute all equilibria for three random 2×2 general-sum games.
- For each equilibrium, round the probabilities to two decimals and compute ε .
- Create a PNG of best-response (BR) lines and mark the intersection for one game.

True or False?

1. Every 2×2 game has a pure NE.
2. A strictly dominated action can be part of a mixed NE support.
3. In zero-sum games, NE equals minimax.
4. Mixed NE always make players indifferent across all actions.

Answers

1. **False.** Matching Pennies is a counterexample.
2. **False.** Dominated actions never belong to the support.
3. **True.** By minimax duality.
4. **False.** Only across supported actions.

Proof: Best Response Upper Hemicontinuity

- The strategy space (product of mixed strategy simplices) is compact and convex.
- Each player's payoff is continuous and linear in their own mixed strategy.
- For any fixed profile of opponents' strategies, the set of best responses is the set of maximizers of a continuous linear function over a simplex, which is a nonempty convex set.
- **Berge's Maximum Theorem:** If $f(x, y)$ is continuous and the constraint correspondence $C(x)$ is continuous (upper hemicontinuous, compact-valued), then the value function $M(x) = \max_{y \in C(x)} f(x, y)$ is continuous, and the maximizer correspondence is upper hemicontinuous and compact-valued.
- Therefore, the best response correspondence in finite games is upper hemicontinuous, convex-valued, and nonempty-valued, satisfying **Kakutani's conditions for a fixed point**.

Existence in Concave Games

If u_i is continuous in all arguments and concave in i 's own strategy on a compact convex set, then a Nash equilibrium exists.

- The best response correspondence is upper hemicontinuous and convex-valued.
- Kakutani's theorem applies, so an equilibrium exists.

Potential Games: Existence and Structure

A game is a **potential game** if there exists a function Φ such that for any unilateral deviation by player i :

$$R_i(a'_i, a_{-i}) - R_i(a_i, a_{-i}) = \Phi(a'_i, a_{-i}) - \Phi(a_i, a_{-i})$$

- Every finite potential game has at least one pure strategy Nash equilibrium (since Φ attains a maximum).
- Best response dynamics converge to a pure NE.

Supermodular Games: Monotone Best Response

A game is **supermodular** if each player's payoff has increasing differences in their own strategy and others' strategies.

- The best response correspondence is monotone (increasing in others' strategies).
- **Tarski's Fixed-Point Theorem:** Any monotone function on a complete lattice has a smallest and largest fixed point.
- Supermodular games have smallest and largest pure strategy Nash equilibria.

Support Enumeration

1. For 2×2 games, enumerate all possible supports (sets of actions played with positive probability) for each player (size 1 or 2).
2. For each support pair, solve the indifference equations (equalize expected payoffs for actions in support).
3. Check feasibility (probabilities in $[0,1]$) and that out-of-support actions do not yield higher payoffs.
4. This is implemented in NashPy and other solvers.

Vertex Enumeration

- Nash equilibria correspond to vertices of best response polytopes (sets defined by best response inequalities).
- Enumerate candidate vertices and test equilibrium conditions.
- Efficient for small games and certain classes (e.g., zero-sum, symmetric).

Example: General-Sum 2×2

	L	R
U	(4, 1)	(0, 2)
D	(1, 0)	(2, 3)

Let $p = \Pr[U]$, $q = \Pr[L]$.

Row indifference:

$$4q + 0(1 - q) = 1q + 2(1 - q) \Rightarrow 4q = 2 - q \Rightarrow q = \frac{2}{5}$$

Column indifference:

$$1p + 0(1 - p) = 2p + 3(1 - p) \Rightarrow p = 3(1 - p) \Rightarrow p = \frac{3}{4}$$

Compute expected payoffs for each action and verify that out-of-support actions do not yield higher payoffs.

Example: Mixed and Pure Check

- Check if any pure cell is a mutual best response (i.e., both payoffs are maximal in their row/column).
- If none, use $(p, q) = \left(\frac{3}{4}, \frac{2}{5}\right)$ as above.
- Verify that all equilibrium conditions are satisfied.

Example: Zero-Sum 3×3 Rock Paper Scissors

$$A = \begin{pmatrix} 0 & -1 & 1 \\ 1 & 0 & -1 \\ -1 & 1 & 0 \end{pmatrix}$$

- By symmetry, the unique NE is $\left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right)$ for both players.
- Value of the game is 0.
- Vertex enumeration recovers this efficiently.

Board Problem: Compute Mixed NE

	L	R
U	(3, 4)	(0, 1)
D	(1, 0)	(2, 3)

1. Write indifference equations for p, q .
2. Solve for p^*, q^* .
3. Verify out-of-support inequalities.
4. Compute expected payoffs for each action.

Extra Questions

1. In a 2×2 game, if both players put positive probability on both actions at an equilibrium, what must be true about the expected payoffs of their actions in support
2. Define $\varepsilon_i = u_i^{BR} - u_i$. If $\varepsilon_i = 0.08$ for each player, is the profile an ε -NE with $\varepsilon = 0.08$?

Extra Answers

1. They must be equal (indifference principle).
2. Yes

Summary

- Best response correspondences define Nash equilibrium.
- Every finite game has a mixed Nash equilibrium.
- In 2×2 games, solve for mixed NE using indifference and feasibility.
- ε -Nash equilibrium measures how close a profile is to equilibrium.
- NashPy helps compute and check equilibria in practice.

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!

