نظریه بازیها

Game Theory

ارائه کننده: امیرحسین نیکوفرد

مهندسی برق و کامپیوتر دانشگاه خواجه نصیر

دانشگاه صنعتی خواجه نصیرالدین طوسی

Non-zero sum games

Material

- Dynamic Non-cooperative Game Theory: Second Edition
  - Chapter4: Sections 4:1–4:3.

Infinite Games

- Zero sum games
- Non-zero sum games
- Infinite Games
  - **Countably Infinite Actions**
  - **Continuous Action Sets**
    - **Introduction**
    - **Reaction Curves and Pure Strategy NE**
    - **Existence of NE**

Finite v/s Infinite

- What does 'infinite' refer to?
- Infinite Action Spaces:
- Example 1: Countably infinite actions

P_2

P_1						
	2	1/2	1/3	1/4	1/5	...
	0	1/2	2/3	3/4	4/5	...

  - Example 2: Continuous action sets

    Players choose actions $u \in [0,1], u \in R, \ldots$

- Infinite Stages:
  - Example 1: Repeated games (Repeated prisoner's dilemma)
  - Example 2: Infinite horizon games

Examples of Infinite Matrix Games

Pick the Largest Number: Two players simultaneously choose a natural number each. The player who has chosen the highest number wins and receives a payoff of 1 from the other player. If both players choose the same number, the outcome is a draw.

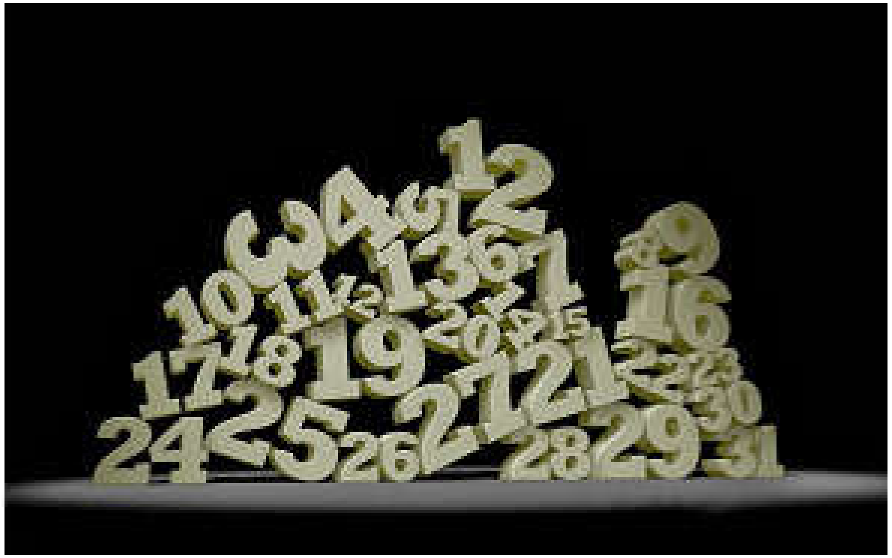

Examples of Infinite Matrix Games

Pick the Largest Number: Two players simultaneously choose a natural number each. The player who has chosen the highest number wins and receives a payoff of 1 from the other player. If both players choose the same number, the outcome is a draw.

P_2

P_1

0	1	1	1	$\dots$
-1	0	1	1	$\dots$
-1	-1	0	1	$\dots$
-1	-1	-1	0	$\dots$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	

Examples of Infinite Matrix Games

Consider the 2-player zero-sum game:

P_2

P_1	2	1/2	1/3	1/4	1/5	...
	0	1/2	2/3	3/4	4/5	...

Mixed Strategy NE:

- P_1 chooses $y_i, i = 1, 2$,such that $y_i \geq 0, \sum_{i=1}^{2} y_i = 1$
- P_2 chooses $z_j, j = 1, 2, \ldots$, such that $z_j \geq 0, \sum_{j=1}^{\infty} z_j = 1$

Mixed Security Strategy of P_1?

Examples of Infinite Matrix Games

Consider the 2-player zero-sum game:

P_2

P_1	y_1	2	1/2	1/3	1/4	1/5	...
	y_2	0	1/2	2/3	3/4	4/5	...

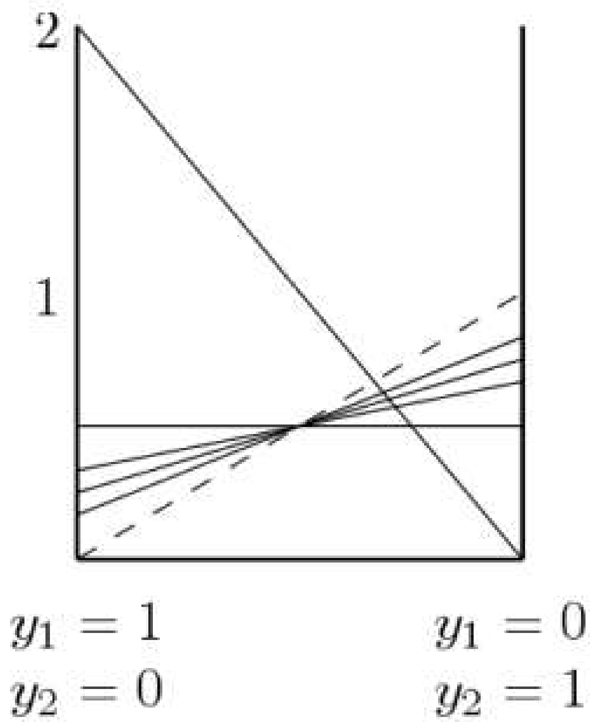

The upper envelope is not well defined!

Mixed Security Strategy of P1

Truncate the matrix to one comprised of k columns:

P_2

P_1							
	2	1/2	1/3	1/4	1/5	...	$1/k$
	0	1/2	2/3	3/4	4/5	...	$(k-1)/k$

- Mixed Security Strategy for P_1 : $y_1 = \frac{k-1}{3k-2}, y_2 = \frac{2k-1}{3k-2}$
- Average Security Level for P_1 : $\bar{V}_m = \frac{2k-2}{3k-2}$
- Mixed Security Strategy for P_2 : $z_1 = \frac{k-2}{3k-2}, z_k = \frac{2k}{3k-2}$
- Average Security Level for P_2 : $\underline{V}_m = \frac{2k-2}{3k-2}$

Mixed Security Strategy of P1

Truncate the matrix to one comprised of k columns:

P_2

P_1							
	2	1/2	1/3	1/4	1/5	...	$1/k$
	0	1/2	2/3	3/4	4/5	...	$(k-1)/k$

- In the Limit as $k \to \infty$:

$$y^*_1 = \frac{1}{3}, y^*_2 = \frac{2}{3}, z^*_1 = \frac{1}{3}, z^*_\infty = \frac{2}{3}, \quad \bar{V}_m = \underline{V}_m = \frac{2}{3}$$

- The security level of P_2 cannot be achieved!

Existence of NE in Infinite Games

- P_2 cannot achieve its security level of $2/3$ as it cannot choose column ∞
- By choosing a sufficiently large column k, P_2 can secure an average lower value arbitrarily close to $2/3$!
- ε-mixed security strategy: The mixed security strategy employed when P_2 chooses k such that $\underline{V}_m > \frac{2}{3} - \varepsilon$. This is also an ε-mixed saddle-point strategy
- Non-existence of a saddle point within the class of mixed strategies!
- Same effect occurs for a pure strategy saddle point in zero-sum games and for Nash equilibria in nonzero-sum games.

ε-Saddle Point for Zero-Sum Games

Definition:

For a given $\varepsilon \geq 0$, the pair $\{u_\varepsilon^{1*}, u_\varepsilon^{2*}\} \in U^1 \times U^2$ is called an ε-saddle point if

$$J(u_\varepsilon^{1*}, u^2) - \varepsilon \leq J(u_\varepsilon^{1*}, u_\varepsilon^{2*}) \leq J(u^1, u_\varepsilon^{2*}) + \varepsilon$$

for all $\{u^1, u^2\} \in U^1 \times U^2$.

- Lower Value of a two person zero-sum infinite game:

$$\underline{V} = \sup_{u^2 \in U^2} \inf_{u^1 \in U^1} J(u^1, u^2)$$

- Upper Value of a two person zero-sum infinite game:

$$\bar{V} = \inf_{u^1 \in U^1} \sup_{u^2 \in U^2} J(u^1, u^2)$$

ε-Saddle Point for Zero-Sum Games

If $\bar{V} = \underline{V}$, then $V = \bar{V} = \underline{V}$ is called the value of the game.

Theorem:

A two-person zero-sum infinite game has a finite value if, and only if, for every $\varepsilon > 0$, an ε -saddle point exist

- Proof : First, suppose that the game has a finite value ($V = \bar{V} = \underline{V}$) . Then, given an $\varepsilon > 0$, one can find $u_\varepsilon^{1*} \in U^1$ and $u_\varepsilon^{2*} \in U^2$ such that

$$J(u^1, u_\varepsilon^{2*}) > \underline{V} - \tfrac{1}{2}\varepsilon \quad , \forall u^1 \in U^1$$

$$J(u_\varepsilon^{1*}, u^2) < \bar{V} + \tfrac{1}{2}\varepsilon \quad , \forall u^2 \in U^2$$

which follow directly from the definitions of $\underline{V}$ and $\bar{V}$, respectively.

ε-Saddle Point for Zero-Sum Games

Now, since $V = \bar{V} = \underline{V}$, by adding ε to both sides of the first inequality we obtain

$$J(u^1, u_\varepsilon^{2*}) + \varepsilon > \mathrm{V} + \tfrac{1}{2}\varepsilon > J(u_\varepsilon^{1*}, u_\varepsilon^{2*}) \quad , \forall u^1 \in U^1 \qquad \text{(i)}$$

where the latter inequality follows from the second inequality by letting $u^2 = u_\varepsilon^{2*}$.

Similarly, if we now add $-\varepsilon$ to both sides of the second inequality, and also make use of the first inequality with $u^1 = u_\varepsilon^{1*}$, we obtain

$$J(u_\varepsilon^{1*}, u^2) - \varepsilon < V - \tfrac{1}{2}\varepsilon < J(u_\varepsilon^{1*}, u_\varepsilon^{2*}) \quad , \forall u^2 \in U^2 \qquad \text{(ii)}$$

ε-Saddle Point for Zero-Sum Games

If (i) and (ii) are collected together, the result is the set of inequalities

$$J(u_\varepsilon^{1*},u^2)-\varepsilon \le J(u_\varepsilon^{1*},u_\varepsilon^{2*}) \le J(u^1,u_\varepsilon^{2*})+\varepsilon \qquad \text{(iii)}$$

for all $\{u^1,u^2\}\in U^1\times U^2$, which verifies the sufficiency part of the theorem.

Second, suppose that for every $\varepsilon>0$, an ε saddle point exists, that is, a pair $\{u_\varepsilon^{1*},u_\varepsilon^{2*}\}\in U^1\times U^2$ can be found satisfying (iii) for all $u^1\in U^1$ and $u^2\in U^2$. Let the middle term be denoted as J_ε. We now show that the sequence $\{J_{\varepsilon 1},J_{\varepsilon 2},\ldots\}$, with $\varepsilon_1>\varepsilon_2>\ldots>0$ and $\lim_{i\to\infty}\varepsilon_i=0$ is Cauchy. Toward this end, let us first take $\varepsilon=\varepsilon_k$ and $\varepsilon=\varepsilon_j$, $j>k$, in subsequent order in (iii) and add the resulting two inequalities to obtain

$$J(u_{\varepsilon_k}^{1*},u^2)-J(u^1,u_{\varepsilon_j}^{2*})+J(u_{\varepsilon_j}^{1*},u^2)-J(u^1,u_{\varepsilon_k}^{2*})\le 2(\varepsilon_k+\varepsilon_j)<4\varepsilon_k$$

ε-Saddle Point for Zero-Sum Games

Now, substituting first $\{u^1 = u^{1*}_{\varepsilon_j}, u^2 = u^{2*}_{\varepsilon_k}\}$ and then $\{u^1 = u^{1*}_{\varepsilon_k}, u^2 = u^{2*}_{\varepsilon_j}\}$ the result is the set of inequalities

$$-4\varepsilon_k < J_{\varepsilon_k} - J_{\varepsilon_j} < 4\varepsilon_k$$

for any finite *k* and *j*, with $j > k$, which proves that $\{J_{\varepsilon_k}\}$ is indeed a Cauchy sequence. Hence, it has a limit in R, which is the value of the game.

It should be noted that, although the sequence $\{J_{\varepsilon_k}\}$ converges, the sequences $\{u^{1*}_{\varepsilon_k}\}$ and $\{u^{2*}_{\varepsilon_k}\}$ need not have limits.

ε-NE solution for Nonzero-Sum Games

Definition:

For a given $\varepsilon > 0$, an N-tuple $\{u_\varepsilon^{1*},\ldots,u_\varepsilon^{N*}\} \in U^1 \times \ldots \times U^N$ is called a pure ε-Nash Equilibrium solution if

$$J^i(u_\varepsilon^{1*},\ldots,u_\varepsilon^{N*}) \le J(u_\varepsilon^{1*},\ldots,u_\varepsilon^{i-1*},u^i,u_\varepsilon^{i+1*},\ldots,u_\varepsilon^{N*}) + \varepsilon, i \in N$$

Definition:

We define a mixed strategy for P_i as a probability distribution (probability measure) μ^i on U^i. Furthermore, $\mu^i \in M^i$, where M^i denotes the class of all such probability distributions (or measures).

ε-NE in Mixed Strategies

Average Cost Function:

$$\bar{J}^i(\mu^1,\ldots,\mu^N)=\int_{U^1}\cdots\int_{U^N} J^i(u^1,\ldots,u^N)\,d\mu^1(u^1)\ldots d\mu^N(u^N)$$

Specialization to Semi-Infinite Bimatrix Case:

$y_i, i=1,\ldots,m$, such that $y_i \ge 0, \sum_{i=1}^{m} y_i = 1$,and $z_j, j=1,2,\ldots$ such that $z_j \ge 0, \sum_{j=1}^{\infty} z_j = 1$

$$\bar{J}^1(y,z)=\sum_i \sum_j y_i a_{ij} z_j \qquad \bar{J}^2(y,z)=\sum_i \sum_j y_i b_{ij} z_j$$

assuming that the above sums are absolutely convergent.

ε-NE in Mixed Strategies

Theorem :

For each $\varepsilon > 0$, the semi-infinite bimatrix game *(A;B)* , with $A = \{a_{ij}\}_{i=1}^{m}{}_{j=1}^{\infty}, B = \{b_{ij}\}_{i=1}^{m}{}_{j=1}^{\infty}, m < \infty$ and the entries a_{ij}, b_{ij} being bounded, admits an ε equilibrium solution in mixed strategies.

Sketch of Proof:

- Let $(\hat{y}, \hat{z})$ be a mixed strategy equilibrium point of the finite bimatrix game with the truncated matrices

  $$A_n = \{a_{ij}\}_{i=1}^{m}{}_{j=1}^{n}, B_n = \{b_{ij}\}_{i=1}^{m}{}_{j=1}^{n}$$

- Let $\bar{z} = (\hat{z}, 0, 0, \ldots)^T$
- It is easy to show that $(\hat{y}, \hat{z})$ is a mixed εequilibrium solution for (A, B)

ε-NE in Mixed Strategies

Corollary :

Every bounded semi-infinite zero-sum matrix game has a value in the class of mixed strategies.

These results do not apply to infinite bimatrix games. (A counterexample will be shown in the previous section.)

Infinite Games

- Zero sum games
- Non-zero sum games
- Infinite Games
  - Countably Infinite Actions
  - **Continuous Action Sets**
    - **Introduction**
    - **Reaction Curves and Pure Strategy NE**
    - **Existence of NE**

Examples of Games with Continuous Action Sets

- Mixed Strategies: Players choose to play u_i with probability $y_i \in [0,1]$
- Stopping Time Games: An investor must decide when to sell a particular asset.

Examples of Games with Continuous Action Sets

- Mixed Strategies: Players choose to play u_i with probability $y_i \in [0,1]$
- Stopping Time Games: An investor must decide when to sell a particular asset. The decision to be selected is a time $t \in [0, \infty)$

Examples of Games with Continuous Action Sets

- The Cournot Game: Firms 1and 2must choose to produce quantities u^1 and u^2 of some commodity. The market price is $p = c_1 - u^1 - u^2$ and the production costs are $c_2 u^i, i = 1,2$

Reaction Curves

Definition:

In an N-person nonzero-sum game, let the minimum of the cost function of P_i, $J^i(u^i, u^{-i})$, with respect to $u^i \in \mathrm{U}^i$ be attained for each $u^{-i} \in \mathrm{U}^{-i}$. Then the optimal response or rational reaction set of P_i, $R^i(u^{-i}) \in \mathrm{U}^i$, is defined by

$$R^i(u^{-i}) = \{\xi \in \mathrm{U}^i : J^i(\xi, u^{-i}) \leq J^i(u^i, u^{-i}) \quad \forall u^i \in \mathrm{U}^i\}$$

If $R^i(u^{-i})$ is a singleton for every $u^{-i} \in \mathrm{U}^{-i}$, then it is called the reaction curve or reaction function of P_i.

Reaction Curves (Example)

- The Cournot Game: Firms1and 2 must choose to produce quantities u^1 and u^2 of some commodity. The market price is $p = c_1 - u^1 - u^2$ and the production costs are $c_2 u^i, i = 1,2$
- The payoff functions are

$$J^i(u^1,u^2) = (c_1 - u^1 - u^2)u^i - c_2 u^i, i = 1,2$$

- The reaction curves are obtained as follows:

$$\frac{\partial J^i(u^1,u^2)}{\partial u^i} = 0 \Rightarrow \quad R^i(u^{-i}) = \frac{c_1 - c_2 - u^{-i}}{2}$$

Reaction Curves (Example)

- The payoff functions are

$$J^i(u^1,u^2) = (c_1 - u^1 - u^2)u^i - c_2 u^i, i = 1,2$$

- The reaction curves are obtained as follows:

$$\frac{\partial J^i(u^1,u^2)}{\partial u^i} = 0 \Rightarrow \quad R^i(u^{-i}) = \frac{c_1 - c_2 - u^{-i}}{2}$$

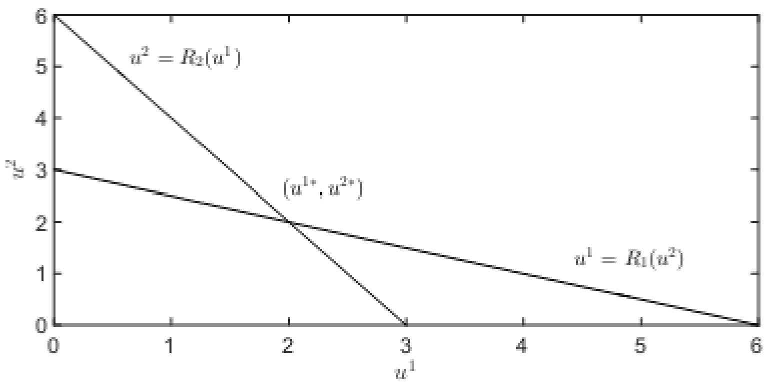

- Pure Strategy NE solution: $u^{i*} = \frac{c_1 - c_2}{3}$

Pure Strategy NE

- Pure strategy NE solution = intersection points of the reaction curves. (The Nash equilibrium solution must lie on both the reaction curves.)
- Condition for uniqueness of the pure strategy NE solution: The reaction curves must have only one point of intersection
- Reading Reaction Curves:
  - Reaction curves intersect at a single point: unique NE
  - Reaction curves do not intersect: No NE or ε -NE
  - Reaction curves intersect at multiple points: multiple NE
  - Reaction curves intersect along an interval: a continuum of NE

Stability of nonunique NE

- Adjustment Scheme:
  - P_1 and P_2 begin at a nonunique equilibrium (u_0^1, u_0^2)
  - One player, say P_1 deviates to $(u_0^1 + \delta)$
  - P_2 observes this and reacts.
  - P_1 reacts to P_2's reaction and so on …..
- Possible Consequences:
- (Globally) Stable Equilibrium: Infinite sequence of moves converges to (u_0^1, u_0^2) fall $\delta \geq 0$. Implies uniqueness of NE.
- Locally Stable Equilibria: Infinite sequence of moves converges to (u_0^1, u_0^2) for $\delta \in \Delta$. Multiple such NE can exist.
- Unstable Equilibria: Infinite sequence of moves never converges to (u_0^1, u_0^2) for any $\delta \geq 0$.

Infinite Games

- Zero sum games
- Non-zero sum games
- Infinite Games
  - Countably Infinite Actions
  - **Continuous Action Sets**
    - Introduction
    - Reaction Curves and Pure Strategy NE
    - **Existence of NE**

Existence of Pure Strategy NE

- Existence of pure strategy NE = existence of well defined reaction functions with a common point of intersection.

Theorem:

An N -person nonzero-sum game admits a Nash Equilibrium in pure strategies if, for each $i \in N$

- U^i is a closed, bounded and convex subset of a finite-dimensional Euclidean space
- $J^i : U^i \times U^{-i} \mapsto R$ is jointly continuous in all its arguments,
- J^i is strictly convex in u^i for every $u^{-i} \in U^{-i}$.

NE in Zero-sum Infinite Games

Theorem:

A two-person zero-sum infinite game admits a **unique** pure strategy NE if:

- U^1 and U^2 are closed, bounded (or effectively closed, bounded) and convex subsets of a finite dimensional space
- $J(u_1, \mathrm{u}_2)$ is jointly continuous in its arguments,
- $J(u_1, \mathrm{u}_2)$ is strictly convex in u^1 for every $u^2 \in U^2$.
- $J(u_1, \mathrm{u}_2)$ is strictly concave in u^2 for every $u^1 \in U^1$.

NE in Mixed Strategies

Theorem:

An N -person nonzero-sum game admits a Nash Equilibrium with mixed strategies if:

- its finite-dimensional action spaces U^i are compact
- cost functional J_i are continuous on $U^1 \times U^2 \times \ldots \times U^N$

❑ Convexity not required.

❑ Extends directly to two-person zero-sum games and mixed strategy saddle-point equilibrium.

❑ Weakest conditions for a mixed strategy saddle point equilibrium: semi-continuity conditions (Glicksberg, 1950).

Continuous Action Sets

Existence of NE for games with continuous action sets: Does this picture still apply?

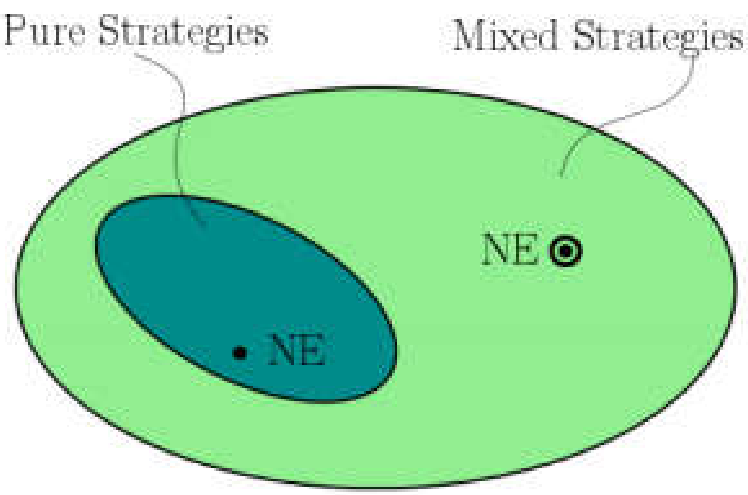

Continuous Action Sets

Existence of NE for games with continuous action sets: Does this picture still apply?

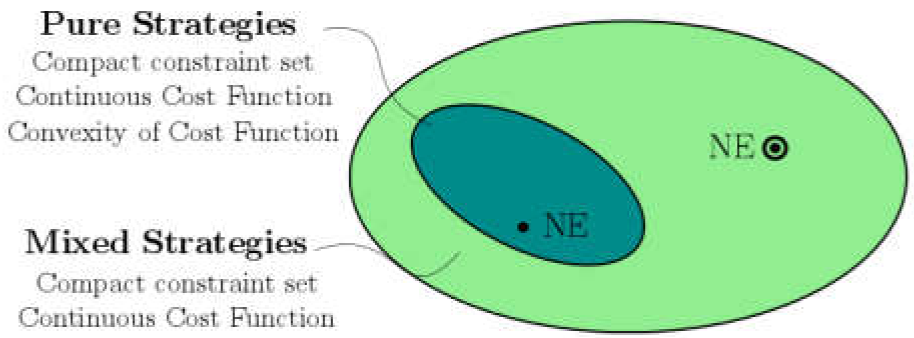

InfiniteDynamic Games

- Zero sum games
- Non-zero sum games
- Infinite Games
  - Countably Infinite Actions
  - Continuous Action Sets
    - Introduction
    - Reaction Curves and Pure Strategy NE
    - Existence of NE
- **Infinite Dynamic Games**

Static v/s Dynamic

What does 'dynamic' refer to?

- Multi-stage games: Players gain dynamic information throughout the decision process
- Stages/Levels/Time: These games require a notion of time
  - Discrete-time Dynamic Games: Finite or countably infinite number of stages/levels of play
  - Continuous-time Dynamic Games: continuum of stages/levels of play

Examples of Infinite Dynamic Games
Example 1

- Repeated Cournot Game: Firms1and 2 must choose to produce quantities u^1 and u^2 of some commodity. The market price is $p = c_1 - u^1 - u^2$ and the production costs are $c_2 u^i, i = 1, 2$. How should they refine their strategies in a repeated format game?

Examples of Infinite Dynamic Games
Example 2

- A negotiation game: Alice negotiates a job with Google, after receiving an offer *r* from Yahoo. They use alternate bargaining.
  - When Alice offers a, Google can accept and hire her, or reject and continue to bargain.
  - When Google offers g, Alice can accept and work for Google, reject and work for Yahoo, or reject and continue bargaining.
  - What should Alice's offer be?

Learning Each Other's
Bargaining Power (Signalling)