Microeconomics IV (Game Theory)

Lecture 3 – Equilibria of Strategic-Form Games

Lasha Chochua

TA: Luka Makharadze

2026

Strategic-Form Games and Rationalizability

Strategic-Form Game

Definition 1: Strategic-Form Game

A strategic-form game is a triple $\Gamma = (N, (C_i)_{i \in N}, (u_i)_{i \in N})$ where:

- N: set of players
- C_i: set of strategies for player i
- $u_i : C \to \mathbb{R}$: utility function of player i
- $C = \times_{i \in N} C_i$: set of all strategy profiles

- Γ is assumed to be a **finite game** unless stated otherwise

Solution Concepts

- The simplest solution concept specifies the set of strategies each player might reasonably be expected to use
- Without assessing the probability of various strategies
- Formally: for each player i, a set $D_i \subseteq C_i$, nonempty, interpreted as the strategies player i might actually choose

Rationality and Intelligence

- Players are assumed to be:
  - **Rational:** maximize expected utility
  - **Intelligent:** understand the game and others' rationality
- If each player j chooses a strategy in D_j, then player i chooses a strategy in $G_i(D_{-i})$, where:

$$D_{-i} = \bigtimes_{j \in N \setminus \{i\}} D_j$$

Best Response Set

- Let $\Delta(Z)$ denote the set of probability distributions over Z
- **Best response set** $G_i(D_{-i})$: $d_i \in G_i(D_{-i})$ if there exists $\eta \in \Delta(D_{-i})$ such that:

$$d_i \in \arg\max_{c_i \in C_i} \sum_{d_{-i} \in D_{-i}} \eta(d_{-i})\, u_i(d_{-i}, c_i)$$

- (d_{-i}, c_i) denotes the strategy profile where player i plays c_i and all others play as in d_{-i}

Rationalizability Condition

- A solution concept should satisfy:

$$D_i \subseteq G_i(D_{-i}), \quad \forall\, i \in N \tag{1}$$

- Let $C_i^{(\infty)}$ denote strategies surviving iterated elimination of strongly dominated strategies
- Then:

$$C_i^{(\infty)} = G_i\left(\bigtimes_{j \in N \setminus \{i\}} C_j^{(\infty)}\right)$$

Iteratively Undominated Strategies

- Every strategy surviving iterated elimination is a best response to some distribution over surviving strategies
- Condition (1) implies:

$$D_i \subseteq C_i^{(\infty)}$$

- Shown by induction: $D_i \subseteq C_i^{(k)}$ for all k
- Rational and intelligent players will not choose strongly dominated strategies
- Game outcome must lie in $\times_{i\in N} C_i^{(\infty)}$

Weak vs. Strong Domination

- **Iterative strong domination** is a weak solution concept
  - Rules out clearly suboptimal strategies
  - Cannot eliminate strategies that depend on beliefs about others
- **Iterative weak domination** eliminates more outcomes
  - The result may depend on the order of elimination
  - One fix: look for the largest sets $(D_i)_{i\in N}$ where no D_i contains weakly dominated strategies against $\times_{j\neq i} D_j$
  - However, such sets may not exist in all games

Rationalizability with Independent Randomization

- Players choose strategies **independently**; strategies are independent random variables
- Let $H_i(D_{-i})$ be the set of d_i such that there exist $\sigma_j \in \Delta(D_j)$ for all $j \neq i$ with:

$$d_i \in \arg\max_{c_i \in C_i} \sum_{d_{-i} \in D_{-i}} \left(\prod_{j \neq i} \sigma_j(d_j) \right) u_i(d_{-i}, c_i)$$

- Solution concept requires:

$$D_i \subseteq H_i(D_{-i}), \quad \forall\, i \in N \tag{2}$$

Rationalizable Strategies

Definition 2: Rationalizable Strategies

There exist sets $(D_i^*)_{i \in N}$ satisfying (2) that contain all other such sets. These are the **rationalizable strategies**, and:

$$D_i^* = H_i(D_{-i}^*)$$

- Any rationalizable strategy is a best response to independent randomizations over others' rationalizable strategies
- In general: $D_i^* \subseteq C_i^{(\infty)}$; with equality for **two-player games**

Mixed Strategies and Nash Equilibrium

Randomized Strategies

Definition 3: Mixed (Randomized) Strategy

A **mixed strategy** for player i is a probability distribution $\sigma_i \in \Delta(C_i)$. Elements of C_i are called **pure strategies**.

- A randomized-strategy profile: $\sigma \in \times_{i \in N} \Delta(C_i)$
- For each $c_i \in C_i$: $\sigma_i(c_i)$ is the probability player i chooses c_i, with:

$$\sum_{c_i \in C_i} \sigma_i(c_i) = 1, \quad \forall\, i \in N$$

- Under independent play, probability of pure profile $c = (c_i)_{i \in N}$ is $\prod_{i \in N} \sigma_i(c_i)$

Expected Utility Under σ

- Expected utility for player i under profile σ:

$$u_i(\sigma) = \sum_{c \in C} \left(\prod_{j \in N} \sigma_j(c_j) \right) u_i(c), \quad \forall\, i \in N$$

- If player i deviates to $\tau_i \in \Delta(C_i)$ while others play σ_{-i}:

$$u_i(\sigma_{-i}, \tau_i) = \sum_{c \in C} \left(\prod_{j \neq i} \sigma_j(c_j) \right) \tau_i(c_i)\, u_i(c)$$

Nash Equilibrium Condition

- Let $[d_i]$ denote the degenerate mixed strategy placing probability 1 on pure strategy d_i
- Player i's expected payoff from pure strategy d_i against σ_{-i}:

$$u_i(\sigma_{-i}, [d_i]) = \sum_{c_{-i} \in C_{-i}} \left(\prod_{j \neq i} \sigma_j(c_j) \right) u_i(c_{-i}, d_i)$$

- Intelligent players put positive probability only on payoff-maximizing strategies:

$$\text{if } \sigma_i(c_i) > 0 \;\Rightarrow\; c_i \in \arg\max_{d_i \in C_i} u_i(\sigma_{-i}, [d_i]) \tag{3}$$

Definition of Nash Equilibrium

Definition 4: Nash Equilibrium

A randomized-strategy profile σ^* is a **Nash equilibrium** if it satisfies (3) for every player i. Equivalently:

$$u_i(\sigma^*) \geq u_i(\sigma^*_{-i}, \tau_i), \quad \forall\, i \in N,\ \forall\, \tau_i \in \Delta(C_i) \tag{4}$$

No player can improve by **unilateral deviation**.

- A **pure-strategy Nash equilibrium** $c \in C$ satisfies:

$$u_i(c) \geq u_i(c_{-i}, d_i), \quad \forall\, i \in N,\ \forall\, d_i \in C_i$$

Lemma 1 – Mixed vs. Pure Best Responses

Important

Lemma 1

For any σ and player i:

$$\max_{d_i \in C_i} u_i(\sigma_{-i}, [d_i]) = \max_{\tau_i \in \Delta(C_i)} u_i(\sigma_{-i}, \tau_i)$$

Moreover, ρ_i maximizes $u_i(\sigma_{-i}, \cdot)$ **if and only if** $\rho_i(c_i) = 0$ for every $c_i \notin \arg\max_{d_i} u_i(\sigma_{-i}, [d_i])$.

Proof of Lemma 1

- For any $\tau_i \in \Delta(C_i)$:

$$u_i(\sigma_{-i}, \tau_i) = \sum_{d_i \in C_i} \tau_i(d_i)\, u_i(\sigma_{-i}, [d_i])$$

- $u_i(\sigma_{-i}, \tau_i)$ is a **weighted average** of $\{u_i(\sigma_{-i}, [d_i])\}$ with weights $\tau_i(d_i) \geq 0$ summing to 1

- A weighted average $\leq$ maximum of the values being averaged
- It is **strictly below** the maximum if any non-maximizing d_i gets positive weight

- **Interpretation:** the highest payoff player i can achieve does not depend on whether i uses pure or mixed strategies

Comparison: Rationalizability vs. Nash Equilibrium

Rationalizability	Nash Equilibrium
Identifies which strategies are *reasonable*	Specifies *numerical probabilities* for each strategy
No mutual consistency requirement	Strategies are mutually consistent
Weaker: larger set of predictions	Stronger: usually a smaller set

Example – Rationalizability vs. Predictive Power

	x_2	y_2	z_2
x_1	$(3,0)$	$(0,2)$	$(0,3)$
y_1	$(2,0)$	$(1,1)$	$(2,0)$
z_1	$(0,3)$	$(0,2)$	$(3,0)$

- No strategies are dominated – all strategies are rationalizable
- Any strategy is a best response to some strategy of the other player
- Example: player 1 plays x_1 believing 2 plays x_2; 2 plays x_2 believing 1 plays z_1 – but then 1 would not play x_1

- Only $([y_1],[y_2])$ forms a **self-consistent** reasoning loop: the unique pure-strategy Nash equilibrium

Nash Existence Theorem

Theorem 1 – Nash's Existence Theorem

Important

Theorem 1 (Nash 1950)
Given any finite game Γ in strategic form, there exists at least one equilibrium in $\times_{i \in N} \Delta(C_i)$.

- Randomized strategies are **essential** for this existence result
- Many games have no equilibrium in pure strategies
- The proof uses the **Kakutani Fixed-Point Theorem** applied to the best-response correspondence

Best-Response Correspondence

- Define the best-response set for player i:

$$B_i(\sigma_{-i}) = \{\sigma_i \in \Delta(C_i) \mid u_i(\sigma_i, \sigma_{-i}) \geq u_i(\sigma_i', \sigma_{-i}) \ \forall \, \sigma_i' \in \Delta(C_i)\}$$

- Define the joint best-response correspondence:

$$B(\sigma) = (B_1(\sigma_{-1}), \dots, B_n(\sigma_{-n}))$$

- σ^* is a Nash equilibrium $\iff$ $\sigma^* \in B(\sigma^*)$ – a **fixed point** of $B(\cdot)$

Kakutani Fixed-Point Theorem

Important

Kakutani's Theorem

A correspondence $B : X \to X$ has a fixed point if:

1. X is compact, convex, and nonempty (subset of Euclidean space)
2. $B(\sigma)$ is nonempty for all σ
3. $B(\sigma)$ is convex for all σ
4. $B(\cdot)$ is upper hemi-continuous

- To prove Nash's theorem, verify all four conditions hold for $B(\cdot)$ on $\times_{i \in N} \Delta(C_i)$

Proof of Theorem 1 – Conditions 1–3

- **Condition 1:** $\times_{i \in N} \Delta(C_i)$ is closed and bounded, hence compact and convex
- **Condition 2:** $u_i(\cdot, \sigma_{-i})$ is continuous (linear) in σ_i; by the Weierstrass theorem a continuous function on a compact set attains its maximum
- **Condition 3 (convexity):** Suppose $\sigma_i', \sigma_i'' \in B_i(\sigma_{-i})$. Then for any $\lambda \in [0, 1]$:

$$u_i(\lambda\sigma_i' + (1-\lambda)\sigma_i'', \sigma_{-i}) = \lambda\, u_i(\sigma_i', \sigma_{-i}) + (1-\lambda)\, u_i(\sigma_i'', \sigma_{-i}) \geq u_i(\sigma_i, \sigma_{-i})$$

(using linearity of u_i in σ_i), so $B_i(\sigma_{-i})$ is convex

Proof of Theorem 1 – Condition 4 (UHC)

- **By contradiction:** suppose $B(\cdot)$ is not upper hemi-continuous
- Then $\exists$ a sequence $(\hat{\sigma}^n, \sigma^n) \to (\hat{\sigma}, \sigma)$ with $\hat{\sigma}^n \in B(\sigma^n)$ but $\hat{\sigma} \notin B(\sigma)$
- So for some i: $\exists\, \varepsilon > 0$ and $\sigma_i' \in \Delta(C_i)$ with:

$$u_i(\sigma_i', \sigma_{-i}) \geq u_i(\hat{\sigma}_i, \sigma_{-i}) + 3\varepsilon$$

- By continuity and $\sigma^n \to \sigma$, for large n:

$$u_i(\sigma_i', \sigma_{-i}^n) > u_i(\hat{\sigma}_i^n, \sigma_{-i}^n) + \varepsilon$$

- This contradicts $\hat{\sigma}_i^n \in B_i(\sigma_{-i}^n)$ $\blacksquare$

Corollary 1 – Finite Games

Important

Corollary 1
All finite games have a Nash equilibrium.

Corollary 2 – Infinite Games

Important

Corollary 2

All infinite games have a Nash equilibrium provided:

(a) S_i are nonempty, compact, convex subsets of Euclidean space

(b) $u_i(\cdot)$ is continuous in s and quasi-concave in s_i

- (a) gives conditions 1 and 2; quasi-concavity gives condition 3; continuity gives condition 4

Finding Nash Equilibria – The 5-Step Method

General Procedure

- To find Nash equilibria in games with more than two actions, use the **5-step support method**
- A player's **support** $S_i \subseteq C_i$ is the set of pure strategies played with positive probability
- The method searches over all possible support combinations and checks consistency
- At each step, verify: **consistency**, **nonnegativity**, and **no profitable deviation**
- The process is **systematic** and **finite** – only finitely many support profiles exist

Five-Step Method

1. **Guess a support:** for each player i, guess $S_i \subseteq C_i$ as the set of strategies used with positive probability
2. **Solve for probabilities:** find a mixed strategy over S_i making the opponent indifferent
3. **Check existence:** if the indifference equations have no solution, return to step 1
4. **Check nonnegativity:** if any probability is negative, the support is invalid
5. **Check no deviation:** verify no player wants to deviate to a strategy outside the guessed support

Example – Game with No Pure-Strategy Equilibrium

	L	M	R
T	$(7,2)$	$(2,7)$	$(3,6)$
B	$(2,7)$	$(7,2)$	$(4,5)$

- No pure-strategy equilibrium exists (verify by inspection)
- Each player's best response to any pure strategy of the other eliminates pure-strategy equilibria
- Both players must randomize – we apply the 5-step method

Case 1 – Support $\{T, B\}$ and $\{L, M, R\}$

- Let $\sigma_1 = p[T] + (1-p)[B]$ and $\sigma_2 = q[L] + (1-q-r)[M] + r[R]$
- Indifference for player 1: $7q + 2(1-q-r) + 3r = 2q + 7(1-q-r) + 4r$
- Indifference for player 2: $2p + 7(1-p) = 7p + 2(1-p) = 6p + 5(1-p)$
- Solving yields $p = 1/2$ **and** $p = 3/4$ – contradiction
- $\Rightarrow$ **no equilibrium** with this support

Case 2 – Support $\{T, B\}$ and $\{M, R\}$

- Let $\sigma_2 = (1-r)[M] + r[R]$ (so $q = 0$)
- Indifference for player 1: $2(1-r) + 3r = 7(1-r) + 4r \;\Rightarrow\; r = 5/4$
- Invalid (probability exceeds 1)
- $\Rightarrow$ **no equilibrium** with this support

Case 3 – Support $\{T, B\}$ and $\{L, M\}$

- Let $\sigma_2 = q[L] + (1-q)[M]$ (so $r = 0$)
- Indifference for player 1: $7q + 2(1-q) = 2q + 7(1-q) \;\Rightarrow\; q = 1/2$
- Indifference for player 2: $2p + 7(1-p) = 7p + 2(1-p) \;\Rightarrow\; p = 1/2$
- **Check deviation to R:**

$$E_{u_2}(R \mid \sigma_1) = 6p + 5(1-p) = 5.5 > E_{u_2}(L \mid \sigma_1) = 3.5$$

- Player 2 would deviate to R – $\Rightarrow$ **not an equilibrium**

Case 4 – Support $\{T, B\}$ and $\{L, R\}$

- Let $\sigma_2 = q[L] + (1-q)[R]$
- Indifference for player 1: $7q + 3(1-q) = 2q + 4(1-q) \;\Rightarrow\; q = 1/6$
- Indifference for player 2: $2p + 7(1-p) = 6p + 5(1-p) \;\Rightarrow\; p = 1/3$
- **Check deviation to M:**

$$E_{u_2}(M \mid \sigma_1) = 7p + 2(1-p) = 11/3 < E_{u_2}(L \mid \sigma_1) = 3.5 \quad \checkmark$$

- **Nash equilibrium:**

$$\sigma_1^* = \tfrac{1}{3}[T] + \tfrac{2}{3}[B], \qquad \sigma_2^* = \tfrac{1}{6}[L] + \tfrac{5}{6}[R]$$

$$E_{u_1} = 11/3 \approx 3.67 \qquad E_{u_2} = 16/3 \approx 5.33$$

Example – The Simple Card Game

- No pure-strategy equilibrium exists
- Strategies Ff and Fr are dominated (eliminate first)
- Let $\sigma_1 = q[\text{Rr}] + (1-q)[\text{Rf}]$ and $\sigma_2 = s[\text{M}] + (1-s)[\text{P}]$

	C_2	
C_1	M	P
Rr	0,0	1,−1
Rf	0.5,−0.5	0,0
Fr	−0.5,0.5	1,−1
Ff	0,0	0,0

Mixed-Strategy Equilibrium in the Card Game

- Player 1 indifferent between Rr and Rf:

$$0 \cdot s + 1 \cdot (1 - s) = 0.5s + 0 \cdot (1 - s) \;\Rightarrow\; s = \tfrac{2}{3}$$

- Player 2 indifferent between M and P:

$$0 \cdot q + (-0.5)(1 - q) = -1 \cdot q + 0 \cdot (1 - q) \;\Rightarrow\; q = \tfrac{1}{3}$$

- **Unique equilibrium:**

$$\sigma_1^* = \tfrac{1}{3}[\mathsf{Rr}] + \tfrac{2}{3}[\mathsf{Rf}], \qquad \sigma_2^* = \tfrac{2}{3}[\mathsf{M}] + \tfrac{1}{3}[\mathsf{P}]$$

- Player 1 raises for sure with red; raises with probability $1/3$ with black
- Player 2 calls with probability $2/3$ upon seeing a raise
- Payoffs: player 1 gets $1/3$, player 2 gets $-1/3$

Efficiency and Multiplicity

- A game may have Nash equilibria that are **inefficient**
- A game may also have **multiple equilibria**

Definition 5: Pareto Efficiency

An outcome is **(weakly) Pareto efficient** if no other outcome makes all players strictly better off.

	g_2	f_2
g_1	$(5,5)$	$(0,6)$
f_1	$(6,0)$	$(1,1)$

- Unique Nash equilibrium: (f_1, f_2) – also the only undominated outcome
- But (g_1, g_2) Pareto dominates it: both players prefer the cooperative outcome

Applications

Application 1 – Political Competition (Hotelling Model)

- Two candidates, Democrat (D) and Republican (R), choose a policy position on $[0, 1]$
- Voters' ideal points are uniformly distributed on $[0, 1]$; each votes for the closest candidate
- Tie broken by fair coin toss; winner gets payoff $+1$, loser gets -1

Results to show:

1. There exists a **unique pure-strategy Nash equilibrium**: both candidates choose $x = 0.5$
2. With **three candidates**, no pure-strategy Nash equilibrium exists

Eliminating Asymmetric Profiles

- Let $x_i \in [0, 1]$ denote candidate i's position; suppose $x_D < x_R$
- **Case 1:** $x_D < x_R < 0.5$ – D can gain by moving just right of x_R; not an equilibrium
- **Case 2:** $0.5 < x_D < x_R$ – R can gain by moving just left of x_D; not an equilibrium
- **Case 3:** $x_D < 0.5 < x_R$ – both want to move toward the center; not an equilibrium

- All asymmetric profiles are eliminated

Eliminating Non-Centered Symmetric Profiles

- Suppose $x_D = x_R = x \neq 0.5$; each wins with probability 0.5
- **Case 1:** $x < 0.5$ – either candidate can gain by moving slightly right
- **Case 2:** $x > 0.5$ – either candidate can gain by moving slightly left

- Symmetric non-median profiles are not stable

- **Unique equilibrium:** $x_D = x_R = 0.5$
  - Each wins exactly half the votes; unilateral deviation yields strictly fewer votes ✓

Three Candidates – No Pure-Strategy Equilibrium

- Suppose $\exists$ equilibrium (x_D^*, x_R^*, x_I^*); consider cases:
- **All equal** $x_D^* = x_R^* = x_I^* = x^*$: each wins $1/3$; any candidate gains by deviating toward center
- **Two equal, one different** $x_i^* = x_j^* \neq x_k^*$: candidate k can gain by moving to $x^* \pm \varepsilon$; always a profitable deviation
- **All different** $x_i^* < x_j^* < x_k^*$: outermost candidate x_k^* wins only voters in $\left[\frac{x_j^*+x_k^*}{2}, 1\right]$ – less than half; can gain by moving to $x_j^* + \varepsilon$
- All cases admit a profitable deviation – **no pure-strategy Nash equilibrium exists** ■

Why Corollary 2 Is Not Enough

Recall – Corollary 2

Corollary 2

All infinite games have a Nash equilibrium provided:

(a) S_i are nonempty, compact, convex subsets of Euclidean space
(b) $u_i(\cdot)$ is **continuous** in s and **quasi-concave** in s_i

- Both conditions are substantive – they are not just technical fuss
- Many natural economic models violate continuity

Note

Question: does Bertrand competition have a Nash equilibrium? Corollary 2 cannot answer this – we need a sharper tool.

Bertrand Competition – Setup

Definition: Bertrand Duopoly

Two firms $i \in \{1, 2\}$ simultaneously choose prices $p_i \in [\underline{p}, \bar{p}]$ where $0 < \underline{p} < \bar{p} < \infty$. Market demand is $D(p) > 0$, strictly decreasing and continuous. Marginal cost is $c \in (\underline{p}, \bar{p})$, identical for both firms.

- Consumers buy from the cheapest firm; ties split demand equally
- Firm i's **profit** given prices (p_1, p_2):

$$\pi_i(p_i, p_j) = \begin{cases} (p_i - c)\, D(p_i) & \text{if } p_i < p_j \\ (p_i - c)\, \frac{1}{2} D(p_i) & \text{if } p_i = p_j \\ 0 & \text{if } p_i > p_j \end{cases} \tag{5}$$

The Discontinuity – Diagnosis

- Fix any $p_j > c$ and consider firm i's profit as a function of p_i:

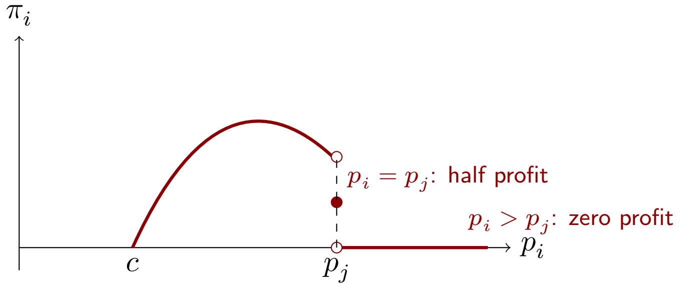

- Left limit at p_j: $(p_j - c)D(p_j)$ – full profit (open circle)
- Value at p_j: $\frac{1}{2}(p_j - c)D(p_j)$ – half profit (filled circle)
- Right limit at p_j: 0 – zero profit (open circle)

Why Corollary 2 Fails

- **Continuity fails:** $\pi_i(\cdot, p_j)$ has a discontinuity at $p_i = p_j$ with two one-sided jumps:
  - From the left: limit is full profit, value is half profit
  - From the right: value is half profit, right limit is zero
- This alone is enough to make Corollary 2 inapplicable – **continuity is the binding condition**

Note

Whether π_i is quasi-concave depends on the shape of $(p - c)D(p)$. If monopoly profit is single-peaked, upper contour sets may still be convex away from p_j. The failure of **continuity** is what matters here, not quasi-concavity.

Dasgupta-Maskin Existence Theorem

Upper Semicontinuity (USC)

Definition : Upper Semicontinuity (USC)

$\pi_i(p_i, p_j)$ is **USC** in p_i at p^* if the limiting profit from any direction does not exceed the profit at p^*:

$$\limsup_{p_i \to p^*} \pi_i(p_i, p_j) \leq \pi_i(p^*, p_j)$$

The value at the point is **at least as large** as nearby limiting values.

- In general: a function is USC if it cannot jump **up** at any point – downward jumps are permitted
- In Bertrand: USC **fails** at $p^* = p_j$ because $\limsup_{p_i \to p_j} \pi_i = (p_j - c)D(p_j) > \frac{1}{2}(p_j - c)D(p_j) = \pi_i(p_j, p_j)$ – the left limit exceeds the value at the tie.

Weak Lower Semicontinuity

Definition: Weak Lower Semicontinuity (Weak LSC)

$\pi_i(p_i, p_j)$ is **weakly LSC** in p_i at a tie point $p_i = p_j = p$ if $\exists\, \lambda \in [0,1]$ such that:

$$\lambda \liminf_{p_i \nearrow p} \pi_i(p_i, p_j) + (1-\lambda) \liminf_{p_i \searrow p} \pi_i(p_i, p_j) \geq \pi_i(p, p_j)$$

- A **convex combination** of the left and right limiting profits must dominate the profit at the tie – weaker than requiring both limits to dominate separately

Note

In general: the condition is required at every point of discontinuity, for every configuration of opponents' strategies that triggers the discontinuity. In Bertrand the only such points are ties $p_i = p_j$.

Weak LSC – Verification for Bertrand

- At any tie $p_i = p_j = p \in [\underline{p}, \bar{p}]$, compute the one-sided limits:

$$\liminf_{p_i \nearrow p} \pi_i(p_i, p_j) = (p - c)D(p), \qquad \liminf_{p_i \searrow p} \pi_i(p_i, p_j) = 0$$

- Take $\lambda = \frac{1}{2}$:

$$\frac{1}{2}(p - c)D(p) + \frac{1}{2} \cdot 0 = \frac{1}{2}(p - c)D(p) = \pi_i(p, p_j) \quad \checkmark$$

- This holds for **all** $p \in [\underline{p}, \bar{p}]$, including $p \leq c$ where $(p - c)D(p) \leq 0$

Note

Weak LSC holds because the payoff at the tie is exactly a convex combination of the left and right limiting payoffs – not because the two limits are equal in magnitude (the left limit is $(p - c)D(p)$ and the right limit is 0).

The Discontinuity Sets (1/2)

Definition: Permitted and Actual Discontinuity Sets

- $A^*(i) = \{(p_1, p_2) : p_1 = p_2\}$ – the **diagonal**: the set of price profiles where discontinuities of π_i are **permitted** to occur
- $A^{**}(i) \subseteq A^*(i)$ – the set of profiles where π_i is **actually discontinuous**

In Bertrand: $A^{**}(i) \subseteq A^*(i) = \{(p_1, p_2) : p_1 = p_2\}$ – actual discontinuities are contained in the diagonal; at $p_1 = p_2 = c$ all limits equal zero so the payoff is continuous there

- $A^*(i)$ is defined by a one-to-one continuous function: $f_{ij}(p_i) = p_i$ (the identity) – firm j's price equals firm i's price
- The diagonal has **Lebesgue measure zero** in $[\underline{p}, \bar{p}]^2$ – discontinuities are rare

The Discontinuity Sets (2/2)

Note

In general: DM allow $A^*(i)$ to be defined by arbitrary one-to-one continuous functions $f_{ij}^d : \mathbb{R} \to \mathbb{R}$, capturing richer relationships between players' strategies. In Bertrand only the identity matters.

Dasgupta-Maskin Theorem (Theorem 5)

Important

Theorem 5 (Dasgupta-Maskin 1986)
Let $A_i \subseteq \mathbb{R}$ be a compact interval for all i. Suppose:

- (a) U_i is continuous except on $A^{**}(i) \subseteq A^*(i)$, where $A^*(i)$ is defined via one-to-one functions f_{ij}^d
- (b) U_i is **bounded** and **weakly lower semicontinuous** in a_i
- (c) $\sum_{i=1}^{N} U_i(\mathbf{a})$ is **upper semicontinuous**

Then the game possesses a **mixed-strategy Nash equilibrium**.

- Condition (c) is USC of the **sum** – weaker than continuity of the sum
- The individual condition (b) is **weak LSC**, not USC
- Together they ensure the limit of finite-game equilibria remains an equilibrium

Proof Sketch – Finite Approximations

- **Step 1:** Approximate each compact A_i by a sequence of finite grids A_i^n, with $A_i^n \to A_i$ as $n \to \infty$
- **Step 2:** Each finite game $[(A_i^n, U_i); i = 1, \dots, N]$ has a mixed-strategy equilibrium μ^n by Nash's theorem (finite games always have equilibria)
- **Step 3:** The space of (Borel) probability measures on A_i is sequentially compact under weak convergence; extract a convergent subsequence $\mu^n \to \mu^*$
- **Step 4:** Use USC of $\sum_i U_i$ and weak LSC of individual U_i to show μ^* is an equilibrium of the original game

Note

Step 4 is the hard part. USC of the sum controls aggregate payoffs in the limit; weak LSC ensures no player can find a profitable deviation at μ^*.

Bertrand Equilibrium

Verifying Dasgupta-Maskin Conditions for Bertrand (1/2)

- **Condition (a) – Discontinuity structure:** discontinuities occur only at $p_1 = p_2$ – the diagonal, defined by $f_{ij}(p_i) = p_i$ (identity, one-to-one) ✓

- **Condition (b) – Weak LSC:** at any tie $p_i = p_j = p \in [\underline{p}, \bar{p}]$, the left limit equals $(p - c)D(p)$ and the right limit equals 0; the $\lambda = \frac{1}{2}$ convex combination gives exactly $\frac{1}{2}(p - c)D(p) = \pi_i(p, p)$ – this holds for all p, including $p \leq c$ where $(p - c)D(p) \leq 0$ ✓

- **Condition (c) – USC of payoff sum:**

$$\pi_1(p_1, p_2) + \pi_2(p_1, p_2) = (p_{\min} - c)\, D(p_{\min})$$

  This is **continuous** in (p_1, p_2) (hence USC) – the minimum function and $(p - c)D(p)$ are both continuous ✓

- All conditions of Theorem 5 are satisfied $\Rightarrow$ a mixed-strategy Nash equilibrium **exists**

Verifying Dasgupta-Maskin Conditions for Bertrand (2/2)

Note

For this simple model, existence can also be verified directly: (c, c) is clearly a Nash equilibrium. DM is useful because it explains why discontinuity of this type is not fatal for equilibrium existence **in general**.

The Bertrand Equilibrium – Characterization (1/2)

Important

Proposition (Bertrand Paradox)
The unique **pure-strategy** Nash equilibrium of the Bertrand duopoly is:

$$p_1^* = p_2^* = c$$

Both firms price at marginal cost; each earns zero profit.

- **Existence:** verified above (directly or via Theorem 5)
- **Uniqueness:** we eliminate all other pure-strategy candidates by case analysis
  - No equilibrium with $p_i > c$ for some i (undercutting argument)
  - No equilibrium with $p_i < c$ for some i (negative profit argument)

The Bertrand Equilibrium – Characterization (2/2)

Note

The proof below establishes pure-strategy uniqueness. In the standard homogeneous-good Bertrand model, the unique mixed-strategy equilibrium is also the degenerate one placing probability one on c; the full argument is omitted here.

Proof of Uniqueness – No Equilibrium with $p_i > c$

- **Case 1:** $p_1^* = p_2^* = p^* > c$ – each earns $\frac{1}{2}(p^* - c)D(p^*)$; either firm deviates to $p^* - \varepsilon$ and nearly doubles profit $\Rightarrow$ not an equilibrium
- **Case 2:** $p_1^* < p_2^*$ and $p_1^* > c$ – firm 2 earns zero; deviates to $p_1^* - \varepsilon > c$ and earns positive profit $\Rightarrow$ not an equilibrium
- **Case 3:** $p_1^* = c < p_2^*$ – firm 1 earns zero at marginal cost; deviates to $p_2^* - \varepsilon > c$ and earns positive profit $\Rightarrow$ not an equilibrium

Note

Case 3 is the often-missed case: even pricing at cost is not an equilibrium when the rival prices strictly above cost, because the firm at cost can profitably raise its price.

Proof of Uniqueness – No Equilibrium with $p_i < c$

- **Case 1:** $p_i^* < p_j^*$ – firm i serves the whole market and earns $(p_i^* - c)D(p_i^*) < 0$; deviates to $\bar{p}$ and earns zero $\Rightarrow$ not an equilibrium
- **Case 2:** $p_i^* = p_j^* < c$ – each earns $\frac{1}{2}(p^* - c)D(p^*) < 0$; either deviates to $\bar{p}$ and earns zero $\Rightarrow$ not an equilibrium
- Therefore $p_i^* \geq c$ for all i in any equilibrium
- Combined with the previous slide: the **only surviving candidate** is $p_1^* = p_2^* = c$
- Verify (c, c) is indeed an equilibrium: each earns zero; a deviation to $p_i > c$ loses the whole market to the rival (who stays at c) and earns zero; a deviation to $p_i < c$ captures demand but earns negative profit $\Rightarrow$ no profitable deviation ■

Economic Interpretation

The Bertrand Paradox

- **Two firms are enough** to replicate the perfectly competitive outcome
- Each firm earns **zero profit** in equilibrium despite positive market demand

Cournot (quantities)	Bertrand (prices)
Equilibrium: $p^* > c$	Equilibrium: $p^* = c$
Positive profits	Zero profits
Deadweight loss	No deadweight loss
n firms needed for competition	2 firms sufficient
Payoffs continuous in strategies	Payoffs discontinuous at $p_1 = p_2$
Corollary 2 applies (under standard concavity assumptions)	Dasgupta-Maskin (Theorem 5) applies

Why Does the Paradox Arise?

- The discontinuity in π_i at $p_i = p_j$ is **not a bug – it is the mechanism**
- At any $p^* > c$, each firm faces a **discrete gain** from undercutting:
    - Charging p^*: earn $\frac{1}{2}(p^* - c)D(p^*)$
    - Charging $p^* - \varepsilon$: earn $(p^* - \varepsilon - c)D(p^* - \varepsilon) \approx 2 \cdot \frac{1}{2}(p^* - c)D(p^*)$
- The gain from undercutting is **always discrete**, no matter how small ε is
- This eliminates all $p^* > c$ as equilibrium candidates – only $p^* = c$ survives

Resolving the Paradox – Classic Extensions (1/2)

- The Bertrand paradox dissolves when the model is enriched:
- **Capacity constraints (Edgeworth 1897):** if firms cannot serve the whole market, undercutting is unprofitable beyond capacity – pure-strategy equilibrium may fail to exist; mixed-strategy equilibria arise
- **Product differentiation (Hotelling):** demand does not jump discretely at a tie – continuity is restored, $p^* > c$ is sustainable
- **Repeated interaction:** collusive prices above c can be sustained as subgame-perfect equilibria – requires repeated games

Resolving the Paradox – Classic Extensions (2/2)

Note

The Bertrand-Dasgupta-Maskin example shows that **the structure of discontinuities matters**, not just their presence. Understanding *why* a theorem's hypotheses fail is as important as knowing the theorem itself.

Summary

- **Corollary 2** requires continuity – this fails in Bertrand at the diagonal $p_1 = p_2$
- **Dasgupta-Maskin Theorem 5** applies when: discontinuities are confined to a measure-zero set; the payoff **sum** is USC; individual payoffs satisfy **weak LSC** in own strategy
- Bertrand satisfies all three conditions $\Rightarrow$ mixed-strategy equilibrium **exists**
- The **unique equilibrium** $p_1^* = p_2^* = c$ is found by eliminating all other candidates – three cases needed, including the often-missed case $p_i = c < p_j$
- The **Bertrand paradox** is a direct consequence of the discrete demand jump at a price tie
- Theorem assumptions are not decoration. When they fail, new tools – and sometimes new equilibrium predictions – are needed.

Required Reading

- **Dasgupta, P. and Maskin, E. (1986).** "The Existence of Equilibrium in Discontinuous Economic Games, I: Theory." *Review of Economic Studies*, 53(1), 1–26.
    - Theorem 5 (main existence result: USC payoff sum + weak LSC individual payoffs)
    - Section 4.2 (outline of the finite approximation argument)
- **Myerson (1997).** *Game Theory: Analysis of Conflict.*
    - Chapter 3, Sections 3.1–3.4 (background on Nash equilibrium and existence)