

Non-Zero Sum Games

R&N Section 17.6

Matrix Form of Zero-Sum Games

m_{11}	m_{12}
m_{21}	m_{22}

m_{ij} = Player A's payoff if Player A follows pure strategy i and Player B follows pure strategy j

Results so far

- 2 players, perfect information, zero-sum:
 - The game has always a *pure* strategy solution given by the minimax procedure

$$\underset{\text{Rows } i}{\text{Max}} \underset{\text{Columns } j}{\text{Min}} m_{ij}$$

- ~~2 players, perfect information, zero-sum:~~
 - The game has always a *mixed* strategy solution given by the minimax procedure

$$\max_p \min(p \times m_{11} + (1-p) \times m_{21}, p \times m_{12} + (1-p) \times m_{22})$$

- ~~2 players, perfect information, zero-sum:~~
 - ????????

Prisoner's Dilemma

- Two persons (A and B) are arrested with enough evidence for a minor crime, but not enough for a major crime.
- If they *both* confess to the crime, they each know that they will serve 5 years in prison.
- If only one of them testify, he will go free and the other prisoner will serve 10 years.
- If neither of them confess, they'll each spend 1 year in prison

Matrix normal form for non-zero-sum games

Player B

↔

Player A		Testify	Refuse
Testify	-5,-5	0,-10	
Refuse	-10,0	-1,-1	

Matrix normal form for non-zero-sum games

Player B

↔

Player A		Testify	Refuse
Testify	-5,-5	0,-10	
Refuse	-10,0		

Player A's payoff for the pair of strategies A:Testify, B: Testify

Player B's payoff for the pair of strategies A: Testify, B: Refuse

Why this example?

- Although simple, this example models a huge variety of situations in which participants have similar rewards as in this game.
- *Joint work*: Two persons are working on a project. Each person can choose to either work hard or rest. If A works hard then prefers to rest, but the outcome of both working is preferable to the outcome of both resting (the project does not get done).
- *Duopoly*: Two firms compete for producing the same product and they both try to maximize profit. They can set two prices, "High" and "Low". If both firms choose High, they both make a profit of \$1000. If they both choose Low, they both make a lower profit of \$600. Otherwise, the High firm makes a profit of \$1200 and the Low firm takes a loss of \$200.
- *Arms race, Robot detection, Use of common property*.....

Matrix normal form for non-zero-sum games

		Player B	
		↔	↔
Player A	↑	Testify	Refuse
	Testify	-5,-5	0,-10
Refuse	-10,0	-1,-1	

- This not a zero-sum game → The interests (payoffs) of the "players" are no longer opposite of each other
- What is the best strategy to follow for each player, assuming that they are both rational

Player B

↔

Player A		Testify	Refuse
	Testify	-5,-5	0,-10
	Refuse	-10,0	-1,-1

Note: In the original image, the 'Refuse' column and the bottom row are crossed out with blue and red lines respectively.

Dominant Strategies

Player B

↔

Player A		Testify	Refuse
	Testify	-5,-5	0,-10
	Refuse	-10,0	-1,-1

- Player A's payoff is greater if he testifies than if he refuses, no matter what strategy B chooses
- Therefore Player A does not need to consider strategy "refuse" since it cannot possibly yield a higher payoff

Dominant Strategies

		Player B	
		Testify	Refuse
Player A	Testify	-5,-5	0,-10
	Refuse	-10,0	-1,-1

- The same reasoning can be applied to Player B:
 - Player B's payoff is greater if he testifies than if he refuses, no matter what strategy A chooses
 - Therefore Player B does not need to consider strategy "refuse" since it cannot possibly yield a higher payoff

Dominant Strategies

		Player B	
		Testify	Refuse
Player A	Testify	-5,-5	0,-10
	Refuse	-10,0	-1,-1

- We say that a strategy *strictly dominates* if it yields a higher payoff than any other strategy for *every* one of the possible actions of the other player.
- Key result → If both players have strictly dominating strategies, they provide a solution for the game (i.e., predict the outcome of the game) → a *dominant strategy equilibrium*
 - Testify is a strictly dominant strategy for A
 - Testify is a strictly dominant strategy for B
 - Therefore (Testify, Testify) is the solution

I	3,0	4,1	5,9	5,6
II	5,3	5,8	9,7	9,0
III	2,3	8,4	6,2	6,3
IV	3,8	3,1	2,3	4,5

Iterated Elimination of Dominated Strategies

I	3,0	4,1	5,9	5,6
II	5,3	5,8	9,7	9,0
III	2,3	8,4	6,2	6,3
IV	3,8	3,1	2,3	4,5

- More generally: We can safely remove any strategy that is strictly dominated → It will never be selected as a solution for the game
- Iteratively removing dominated strategies is the first step in simplifying the game toward a solution
- Is it sufficient? Did we get lucky earlier?

	IB	IIB	IIIB
IA	-1,6	6,-1	5,4
IIA	6,-1	-1,6	5,4
IIIA	4,5	4,5	7,7

How would the two players play this game?

Dominant Strategies

	IB	IIB	IIIB
IA	-1,6	6,-1	5,4
IIA	6,-1	-1,6	5,4
IIIA	4,5	4,5	7,7

How would the two players play this game?

- It is *not* the case that we are guaranteed that both players, or even that one player has a dominant strategy
- We can still use the rule for simplifying the game: Get rid of the strictly dominated strategies because they will never be selected in the solution
- However, we need a more general way of finding a solution to the game (i.e., to predict how rational players would play the game)
- We need a definition that generalizes the earlier definition for zero-sum games (using minimax)

	IB	IIB	IIIB
IA	-1,6	6,-1	5,4
IIA	6,-1	-1,6	5,4
IIIA	4,5	4,5	7,7

$u_A(\text{IIIA}, \text{IIIB}) \geq u_A(X, \text{IIIB})$ For any strategy X of Player A
 $u_B(\text{IIIA}, \text{IIIB}) \geq u_B(\text{IIIA}, Y)$ For any strategy Y of Player B

(IIIA,IIIB) is an equilibrium because:

- Player A cannot find a better strategy given that Player B uses strategy IIIB
- Conversely, Player B cannot find a better strategy given that Player A uses strategy IIIA

Side Note: More than 2 Players?

- The formalism extends directly to more than 2 players.
- If we have n players, we need to define n payoff functions $u_i, i=1, \dots, n$.
- Payoff function u_i maps a tuple of n strategies to the corresponding payoff for player i
- $u_i(s_1, \dots, s_n) =$ payoff for player i if players $1, \dots, n$ use pure strategy s_1, \dots, s_n .
- Everything else (definition of dominating strategies, etc. remains the same)

More formal definition

- A tuple of pure strategies $(s_1^*, s_2^*, \dots, s_n^*)$ is a pure equilibrium if, for all i 's:

$$u_i(s_1^*, \dots, s_{i-1}^*, s_i, s_{i+1}^*, \dots, s_n^*) \leq u_i(s_1^*, \dots, s_{i-1}^*, s_i^*, s_{i+1}^*, \dots, s_n^*)$$

for any strategy s_i .

- In words: Player i cannot find a better strategy than s_i^* if the other player use the remaining strategies in the equilibrium
- Technically, called a pure Nash Equilibrium (NE)

More formal definition (equivalent)

- A tuple of pure strategies $(s_1^*, s_2^*, \dots, s_n^*)$ is a pure equilibrium if, for all i 's:

$$s_i^* = \arg \max_{s_i} u_i(s_1^*, \dots, s_{i-1}^*, s_i, s_{i+1}^*, \dots, s_n^*)$$

- In words: Player i cannot find a better strategy than s_i^* if the other player use the remaining strategies in the equilibrium
- Technically, called a pure Nash Equilibrium (NE)

Examples and Question

- So, we've generalized our concepts for solving games to non zero-sum games → NEs
- Basic questions:
 - Is there always a NE?
 - Is it unique?

Example with multiple NEs

	Left	Right
Left	+1,+1	-1,-1
Right	-1,-1	+1,+1

- Two vehicles are driving toward each other; they have 2 choices: Move right or move left.
- Why is having multiple NEs a problem?

Example with multiple NEs

	Hockey	Movie
Hockey	+2,+1	0,0
Movie	0,0	+1,+2

- Two friends have different tastes, A likes to watch hockey games but B prefers to go see a movie. Neither likes to go to his preferred choice alone; each would rather go the other's preferred choice rather than go alone to its own.

Example with no pure NE

	I	II
I	0,1	1,0
II	1,0	0,1

- Even very simple games may not have a pure strategy equilibrium → This is not surprising since we saw earlier that we had a similar problem with zero-sum games, which did not necessarily have a pure strategy solution
- Solution: Same trick as with zero-sum games → Allow the players to randomize and to use *mixed* strategies

Mixed Strategy Equilibrium

- The concept of equilibrium can be extended to mixed strategies.
- In that case, a mixed strategy for each player i is a vector of probabilities $\mathbf{p}_i = (p_{ij})$, such that player i chooses pure strategy j with probability p_{ij}
- A set of mixed strategies $(\mathbf{p}_1^*, \dots, \mathbf{p}_n^*)$ if player i (for any i) gets a lower payoff by changing \mathbf{p}_i^* to any other mixed strategy \mathbf{p}_i

Example

	Hockey	Movie
Hockey	+2,+1	0,0
Movie	0,0	+1,+2

- An example mixed strategy is:
 - A chooses Hockey with probability: $p = 2/3$
 - B chooses Hockey with probability: $q = 1/3$
- In fact, this is a mixed strategy equilibrium for this game
- The expected payoff is $2/3$ for both A and B

Example

	Hockey	Movie
Hockey	+2,+1	0,0
Movie	0,0	+1,+2

- Let A choose Hockey with probability p and B choose Hockey with probability q
- The expected payoff for Player A is:

$$u_A = (+2) \times pq + (+1) \times (1-p) \times (1-q) = 1 - p - q + 3pq$$
- The expected payoff for Player B is:

$$u_B = (+1) \times pq + (+2) \times (1-p) \times (1-q) = 2 - 2p - 2q + 3pq$$
- At the equilibrium, the derivative of u_A with respect to p is zero (because $u_A(p^*, q^*)$ is greater than u_A for any other value (p, q^*)), therefore: $3q^* - 1 = 0 \rightarrow q^* = 1/3$
- Similarly, the derivative of u_B with respect to q must be 0 at the equilibrium, therefore: $3p^* - 2 = 0 \rightarrow p^* = 2/3$
- Of course this example is constructed specifically so that these equations can be solved very easily.

Key results

- *Theorem:* For any game with a finite number of players, there exists at least one equilibrium
- There might not exist an equilibrium with only *pure* strategies, but at least one mixed strategy equilibrium exists
- Any equilibrium survives iterated elimination of dominated strategies

(we'll use mostly pure strategies here for convenience)

How to compute the equilibrium: Example

- The same product is produced by two firms A and B
- The unit production cost is c , so the cost to produce q_A unit for firm A is $C = cq_A$
- The market price depends on the total number of units produced: $P = \alpha - (q_A + q_B)$
- Therefore firm A's revenue is $q_A(\alpha - c - (q_A + q_B))$
- Problem: How to figure out the "optimal" output for firm A and B?
- If they produce too much, the price will go down and so would the revenue for each firm
- If they produce too little, the revenue will be small

Example

- Each possible value of q_A is a pure strategy for firm A (and similarly for B).
- At equilibrium, A's revenue is maximum as we vary $q_A \rightarrow$ The derivative of $q_A(\alpha - c - (q_A + q_B))$ with respect to q_A is zero at the NE
- Similarly, B's revenue is maximum as we vary $q_B \rightarrow$ The derivative of $q_B(\alpha - c - (q_A + q_B))$ with respect to q_B is zero at the NE
- Therefore (q_A^*, q_B^*) is solution of the system:
$$\alpha - c - 2q_A - q_B = 0 \quad \alpha - c - 2q_B - q_A = 0$$
- With the solution: $q_A^* = q_B^* = (\alpha - c)/3$
- And revenue for each firm: $(\alpha - c)^2/9$

Note: We ignored the fact that the price must be set to 0 for $q_A + q_B > \alpha$

Is NE really the best that the 2 players can do?


- Suppose that instead of trying to find an equilibrium for A and B independently, we try to maximize the total revenue:
- The maximum is reached for $q_A = q_B = (\alpha - c)/4$ (just take the derivative of the total revenue with respect to the total output $q_A + q_B$)
- This corresponds to a revenue per firm of $(\alpha - c)^2/8$, which is *greater* than the revenue we get from the NEs. So the firms are doing better than our wonderful theory predicts. What is going on??? Is there something wrong with the theory??

Coordination vs. No coordination

- There is nothing wrong with the theory. The reason is that, in the second calculation, the two firms are cooperating instead of deciding their strategy independently
- In general, in any game, the players would get a greater payoff if they agree to cooperate (coordinate, communicate).
- For example, in the prisoner's dilemma, the obvious solution is for the prisoners to both refuse to testify, *if* they agree in advance to coordinate their actions.
- We have considered only games without coordination, thus the seemingly paradoxical result.

Tragedy of the Commons

- The previous example is one example of a more general situation, illustrated by the canonical example:
 - n farmers use a common field for grazing goats
 - Because the common field is a *finite resource shared* among all the farmers, the larger the total number of goats, the less food there is, and their unit value goes down
 - Each individual farmer gets a higher profit if they all cooperate (maximize total profit) than if they use the NE equilibrium, acting “rationally” → In the latter case, they tend to each try to “exhaust” the common resource.
- Note: Replace the silly example by changing common field → energy resources, communication bandwidth, oil,.. and farmers → customers, robots, vehicles, firms,...



For those interested in historical references:
 "The Tragedy of the Commons," Garrett
 Hardin, Science, 162(1968):1243-1248

NE Recipes

- The NEs are within the strategies that survive iterated removal of dominated strategies → Iterated removal is one way to get at the NEs
- For strategies described by continuous variables (see previous example), the NEs are found by solving the system of equations obtained by writing that the derivative of u_i with respect to s_i is zero (i.e., $u_i(s_1^*, \dots, s_n^*)$ is an extremum with respect to s_i):

$$\frac{\partial u_i}{\partial s_i}(s_1^*, \dots, s_n^*) = 0$$

- and retain solutions that are maxima

Summary

- Matrix form of non-zero-sum games and basic concepts for those games
- Strict dominance and its use
- Definition of game equilibrium
- Key result: Existence of (possibly mixed) equilibrium for any finite game
- Understand the difference between cooperating and non-cooperating situations
- Continuous games and corresponding recipes