Game Theory & Oligopoly

Noah Lyman

June 16, 2023

Motivation

- When discussing markets, we think about many firms making decisions in tandem
- One thing we've implicitly been assuming: all firms ultimate payoff only depends on their own choice of action
- Often, it makes sense to think that your payoff depends on both:
    - What you choose to do
    - What everybody else chooses to do
- In this setting, firms (or agents more generally), behave *strategically*, recognizing their action influences others' actions (and vice versa)
- Game theory studies the strategic interactions of economic agents

Preliminaries

- There are N agents indexed by $i = 1, \ldots, N$
- Each agent has a set of *strategies* S_i which they can pick from
- We'll let capital S_i denote agent i's set of possible strategies, and lowercase s_i to denote the one they actually selected
- A *strategy profile* $s = \{s_1, s_2, \ldots, s_N\}$ specifies the strategy selected by each agent
- We'll often let $s_{-i} = \{s_1, \ldots, s_{i-1}, s_{i+1}, \ldots, s_N\}$ denote the strategies selected by everybody except for agent i

Preliminaries

- Let S be the set of all possible strategy profiles
- Agents' payoff (utility) function $u_i(s)$ is:

$$u_i : S \rightarrow \mathbb{R}$$

- $u_i(s)$ assigns a utility value to any possible strategy profile
- Each agent's payoff depends on their strategy, as well as everybody else's strategy
- We can now define what a "game" actually is
  - A *simultaneous-move* game at least

Normal Form Game

- A *normal form game* G consists of:
  - Set of N agents $i = \{1, 2, \ldots, N\}$
  - Set of strategies for each agent $S_1, S_2, \ldots, S_N$
  - Payoff (utility) functions: $u_i : S \to \mathbb{R}$
- The game G summarizes all we need to know about the strategic environment
- Given the information in G, we can make predictions about the outcome of the strategic interaction in question
- What we seek are *equilibria* of the game

Nash Equilibrium

- A strategy profile s is a *Nash Equilibrium* of game G if for every agent $i = 1, \ldots, N$ and every alternative strategy s_i':

$$u(s_i, s_{-i}) \geq u(s_i', s_{-1})$$

- Given everybody else's strategy s_{-i}, nobody has an incentive to deviate from their own selected strategy s_i

Example: The Prisoner's Dilemma

- The most commonly used introductory example of a simultaneous-move game is the *prisoner's dilemma*
- Consider two agents A and B:
- Both have just committed a crime and are being interrogated by the police
- A and B have two "strategies" available to them:
  - Admit to the crime
  - Don't admit in the crime
- However, A and B are being interrogated in different rooms, so when they decide whether to admit or not, they are unsure what their partner is doing

Example: The Prisoner's Dilemma

- Let's formalize this a bit
- Again, two agents A and B
- Strategy sets:
  - $S_A = \{H, D\}$
  - $S_B = \{H, D\}$
- H refers to the case when the agent "holds out" and doesn't admit to the crime
- D refers to the case when the agent "defects" on their partner and admits to the crime
- We have agents and strategies, we just need to define payoffs to have a complete characterization of the game

Example: The Prisoner's Dilemma

- Let $s = \{s_A, s_B\}$ denote a generic strategy profile
  - s_A and s_B respectively denote agent A and B's selected strategies
- Four possible strategy profiles:

$$\begin{matrix} \{H, H\} & \{H, D\} \\ \{D, H\} & \{D, D\} \end{matrix}$$

- For each agent, need to assign payoffs following any strategy profile

Example: The Prisoner's Dilemma

	H	D
H	1, 1	-1, 2
D	2, -1	0, 0

- One way to neatly write out payoffs is through a *payoff matrix*
- Agent A is the "row player" while agent B is the "column player"
- The first coordinate of each payoff denotes the row player's payoff, while the second denotes the column player's payoff

Example: The Prisoner's Dilemma

	H	D
H	1, 1	-1, 2
D	2, -1	0, 0

- For example, if both agents hold out, each agent receives utility $= 1$
- If agent A holds out but agent B defects, agent A gets utility -1 and agent B gets utility 2
- If both defect, both agents get utility $= 0$

Example: The Prisoner's Dilemma

	H	D
H	1, 1	-1, 2
D	2, -1	0, 0

- How do we find the equilibrium of this game?
- A strategy profile constitutes a Nash Equilibrium if nobody has an incentive to deviate from it
- Let's start by taking note of each agent's *best response* to the other agent's strategy

Example: The Prisoner's Dilemma

	H	D
H	1, 1	-1, 2
D	$\underline{2}$, -1	0, 0

- If agent B holds out, agent A is better off defecting
- Why? Conditional on agent B choosing H, agent A gets 1 from H and 2 from D

Example: The Prisoner's Dilemma

	H	D
H	1, 1	-1, 2
D	$\underline{2}$, -1	$\underline{0}$, 0

- If agent B defects, agent A is again better off defecting
- Conditional on agent B choosing H, agent A gets -1 from H and 0 from D

Example: The Prisoner's Dilemma

	H	D
H	1, 1	<u>-1</u>, 2
D	<u>2</u>, -1	<u>0</u>, 0

- Next for agent B's best responses
- If agent A holds out, agent B is better off defecting
- Conditional on agent A choosing H, agent B gets 1 from H and 2 from D

Example: The Prisoner's Dilemma

	H	D
H	1, 1	-1, 2
D	2, -1	0, 0

- Lastly, if agent A defects, agent B is again better off defecting
- Conditional on agent A choosing D, agent B gets -1 from H and 0 from D

Example: The Prisoner's Dilemma

	H	D
H	1, 1	$\underline{-1}$, 2
D	$\underline{2}$, -1	$\underline{0}$, $\underline{0}$

- The only time in which no agent has an incentive to deviate is when both agents select D
- $s^* = \{D, D\}$ is the Nash Equilibrium of this game
- Despite $\{H, H\}$ yielding the highest "welfare," this strategy profile cannot be supported in equilibrium
- Why? Because conditional on their partner selecting H, both players have an incentive to deviate

Example: The Prisoner's Dilemma

	H	D
H	1, 1	$\underline{-1}$, 2
D	$\underline{2}$, -1	$\underline{0}$, $\underline{0}$

- In this example, the Nash Equilibrium was unique
  - This doesn't always happen
- Why did it happen here? Notice for each agent, no matter what their partner does, they are always better off choosing D
- For both agents, D is their *dominant strategy*

Dominant Strategies

- Given a set of strategies S_i, s_i is a *strictly dominant strategy* if for all $s_{-i} \in S_{-i}$:
$$u(s_i, s_{-i}) > u(s_i', s_{-i}) \quad \text{for all} \quad s_i' \in S_i$$
- In words: no matter what the other agents do (no matter the s_{-i}), s_i is strictly better than any alternative strategy s_i' for agent i
- Of course, if all agents play a dominant strategy, this will constitute a Nash Equilibrium
  - Referred to as a *dominant strategy equilibrium*
- However, this is not always possible

Example: Mutually Assured Destruction

- Consider two countries: A and B
- Each country has two possible strategies:
  - N - nuke the other country
  - D - don't do that
- Let's define payoffs for this game

Example: Mutually Assured Destruction

	N	D
N	-10, -10	10, -20
D	-20, 10	20, 20

- Again, country A is the "row player" and country B is the "column player"
- What are the equilibria of this game?

Example: Mutually Assured Destruction

	N	D
N	$\underline{-10}$, -10	10, -20
D	-20, 10	20, 20

- If country B chooses N, it is in country A's interest to also choose N

Example: Mutually Assured Destruction

	N	D
N	$\underline{-10}$, -10	10, -20
D	-20, 10	$\underline{20}$, 20

- If country B chooses D, it is in country A's interest to also choose D

Example: Mutually Assured Destruction

	N	D
N	$\underline{-10}$, $\underline{-10}$	10, -20
D	-20, 10	$\underline{20}$, 20

- Next, let's check country B's best responses
- If country A chooses N, it is in country B's interest to also choose N

Example: Mutually Assured Destruction

	N	D
N	$\underline{-10}$, $\underline{-10}$	10, -20
D	-20, 10	$\underline{20}$, $\underline{20}$

- If country A chooses D, it is in country B's interest to also choose D

Example: Mutually Assured Destruction

	N	D
N	$\underline{-10}$, $\underline{-10}$	10, -20
D	-20, 10	$\underline{20}$, $\underline{20}$

- This game has 2 Nash Equilibria
- Either both countries strike or both countries don't
- For both countries, neither N nor D are strictly dominant strategies
- As long as neither country nukes, no country has any incentive to deviate
  - If they launch a nuke, they'll also get nuked
  - "Mutually assured destruction"

Best Response Function

- Let's step in agent i's shoes
- Again, let $s_{-i} = \{s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_N\}$ denote the strategies selected by everybody except for agent i
- Agent i wants to ensure they select the strategy which makes them best off as possible
- Agent i's *best response function* gives the optimal strategy in response to any strategies s_{-i} selected by their opponents

Best Response Function

- Agent i's problem is:

$$\max_{s_i \in S_i} \; u(s_i, s_{-i})$$

- In words: they want to maximize their utility given what everybody else has chosen to do

- The solution to the problem above is called agent i's *best response function*:

$$s_i^*(s_{-i}) = \underset{s_i \in S_i}{\text{argmax}} \; u(s_i, s_{-i})$$

- $s_i^*(s_{-i})$ specifies the best action agent i can take given everybody else's strategies s_{-i}

Best Response Function

- We can redefine Nash Equilibria in terms of best response functions
- A strategy profile $s^* = \{s_1^*, \ldots, s_N^*\}$ is a Nash Equilibrium if for all agents i:

$$s_i^* = \underset{s_i \in S_i}{\operatorname{argmax}} \; u(s_i, s_{-i}^*)$$

- In words: if everybody is best responding to everybody else's strategy, we have a Nash Equilibrium

Application: Oligopoly

Oligopoly

- Game theory is used extensively within the context of antitrust & competition
- In particular, we'll use it to think about *oligopolies*
- Oligopolies are somewhere in between the two cases of monopolies and perfectly-competitive markets
- There are a few firms, each has some impact on market prices through their supply decisions,
- But, each firm's optimal supply decision depends on what the competing firms do

Cournot Competition

- We'll focus on the *Cournot model*, which is the workhorse model of quantity competition
- Two briefly motivate, consider two firms A and B, who are competing for market share
- Each has influence over market prices
- A would like to produce a lot to gain a higher market share than their competitor (and vice versa)
- However, if they produce too much, they'll drive prices too low
- How should the two firms optimally compete against one another?

Cournot Competition

- Let's formalize this
- Assumptions:
  - All firms produce identical goods
  - Firms simultaneously decide how much to produce
  - Market price depends on total market supply
- For now, we'll restrict attention to duopolies (i.e. two firms)
- Very easy to extend it to arbitrarily many firms, but we we'll stick with two for now

Cournot Competition

- For simplicity, let's start by imagining two firms: Firm 1 and Firm 2
- Each firm has the linear cost function:

$$c(q_i) = aq_i$$

- Each firm's profits are given by:

$$\pi(q_i, q_{-i}) = p(Q)q_i - c(q_i)$$

- $P(Q)$ is the inverse market demand curve
- Q is the market supply: $Q = q_1 + q_2$

Cournot Competition

- Let's assume market demand is given by:

$$p(Q) = 100 - Q = 100 - q_1 - q_2$$

- Firm 1's profits are given by:

$$\begin{aligned}\pi(q_1, q_2) &= p(Q)q_1 - c(q_1) \\ &= (100 - q_1 - q_2)q_1 - aq_1 \\ &= 100q_1 - q_1^2 - q_1 q_2 - aq_1\end{aligned}$$

- We can derive Firm 1's best response function by maximizing with respect to q_1 and solving for q_1^* as a function of q_2

Cournot Competition

$$\pi(q_1, q_2) = 100q_1 - q_1^2 - q_1 q_2 - aq_1$$

- Maximizing with respect to q_1 yields:

$$q_1^* = \frac{100 - q_2 - a}{2}$$

- q_1^* is Firm 1's best response function, which depends on the quantity q_2 which the competing firm selects
- The game is in equilibrium when both players best respond
- Let's now derive Firm 2's best response function

Cournot Competition

- Firm 2's profit is given by:

$$\begin{aligned}\pi(q_2, q_1) &= p(Q)q_2 - c(q_2) \\ &= (100 - q_1 - q_2)q_2 - aq_2 \\ &= 100q_2 - q_2^2 - q_1 q_2 - aq_2\end{aligned}$$

- Maximizing with respect to q_2 yields:

$$q_2^* = \frac{100 - q_1 - a}{2}$$

$$q_1^* = \frac{100 - q_2 - a}{2}$$
$$q_2^* = \frac{100 - q_1 - a}{2}$$

- The two best response functions are symmetric
- We can obtain the Nash equilibrium by using the two best response functions to solve for q_1^* and q_2^*
- First, we'll plug q_2^* into q_1^* to obtain Firm 1's optimal quantity

Cournot Competition

- Plugging q_2^* into q_1^* yields:

$$
\begin{aligned}
q_1^* &= 50 - \frac{a}{2} - \frac{q_2}{2} \\
q_1^* &= 50 - \frac{a}{2} - \frac{1}{2}\left(50 - \frac{a}{2} - \frac{q_1}{2}\right) \\
q_1^* &= 50 - \frac{a}{2} - 25 + \frac{a}{4} + \frac{q_1}{4} \\
\frac{3}{4}q_1^* &= 25 - \frac{a}{4} \\
q_1^* &= \frac{100 - a}{3}
\end{aligned}
$$

- This is Firm 1's equilibrium level of quantity, we can plug this into Firm 2's best response function to get their equilibrium quantity

Cournot Competition

$$q_1^* = \frac{100 - a}{3}$$
$$q_2^* = \frac{100 - a}{3}$$

- In the end, the Nash equilibrium is characterized by the two levels of quantity above
- Each firm is best responding to the other firm's strategy
- If we wanted to obtain market price, we'd simply plug this into the inverse demand curve

Cournot Competition

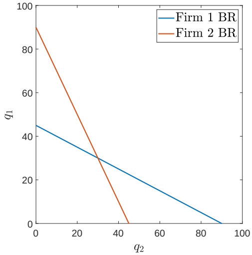

- We can plot the best-response functions graphically
- Their intersection represents the point in which both firms are best responding to each other
    - i.e. the Nash equilibrium

General Oligopoly

- Consider a more general market with N identical firms
- Inverse demand is given by:

$$p = A - Q$$

- $Q = q_1 + \ldots + q_N$ is the market quantity
- Assume that firms have identical & constant marginal costs MC, so just write cost functions as:

$$c(q_i) = MC\, q_i$$

- We seek the equilibrium strategy profile $\{q_1^*, \ldots, q_N^*\}$

General Oligopoly

- Consider firm 1, their profits are given by:

$$\pi(q) = (A - (q_1 + \ldots + q_N))q_1 - MC\, q_1$$

- To derive firm 1's best response function, maximize their profit with respect to q_1:

$$A - (q_1 + \ldots + q_N) - q_1 - MC = 0$$
$$q_1^* = \frac{A - (q_2 + \ldots + q_N) - MC}{2}$$

- Because all firms are identical, their best response functions will be symmetric

General Oligopoly

- Because all firms are symmetric:

$$q_1 = q_2 = \ldots = q_N$$

- Just call q the level of individual supply
- It can be shown then that:

$$q^* = \frac{A - MC}{N + 1}$$

- Given this individual supply, then market quantity is given by:

$$Q^* = Nq = \frac{N}{N + 1}(A - MC)$$

Quantity in an Oligopoly

$$Q^* = \frac{N}{N+1}(A - MC)$$

- As N (i.e. the number of firms) increases, the fraction $\frac{N}{N+1}$ converges to 1
- Thus, as N increases, Q^* converges to $A - MC$
- Pause for a second, and think about what the perfectly competitive level of quantity would be
- In a perfectly competitive market, firms produce until $p = MC$
- Given our assumed demand:

$$Q = A - p = A - MC$$

Quantity in an Oligopoly

- If the market were perfectly competitive, we'd have $Q = A - MC$
- In an oligopoly, the equilibrium quantity converges to the perfectly-competitive (i.e. welfare maximizing) level of quantity as the number of firms increases
- As the number of competitors increases, quantity increases until it reaches its welfare-maximizing level
- More competition $\rightarrow$ more surplus

Quantity in an Oligopoly

- Additionally, we can plug in $N = 1$ to obtain the monopoly level of quantity (verify this):
$$Q = \frac{A - MC}{2}$$
- Letting Q_m denote the monopoly quantity, Q_o denote the oligopoly quantity, and Q_c denote the perfectly-competitive quantity, we have in general that:
$$Q_m < Q_o < Q_c$$
- What happens to prices as the number of firms changes?

Price in an Oligopoly

- If we plug the expression for Q^* into the inverse demand, what we'd get is:

$$p^* = \frac{1}{N+1}A + \frac{N}{N+1}MC$$

- The equilibrium price in an oligopoly has two terms
- As N gets larger, the first term converges to zero
- As N gets larger, the second term converges to MC

Price in an Oligopoly

- In particular:

$$\lim_{N \to \infty} p^*(N) = MC$$

- In words: as the number of competitors increases, equilibrium price converges to marginal cost (i.e. the perfectly competitive price)
- Perfect competition is the limiting case of an oligopoly as the number of firms goes to infinity
- More competition $\to$ lower prices

Price in an Oligopoly

- Finally, we can again plug in $N = 1$ to obtain the monopoly price (verify this):

$$p^* = \frac{A + MC}{2}$$

- Again letting p_m, p_o, and p_c denote the monopoly, oligopoly, and PC prices, We have that in general:

$$p_m > p_o > p_c$$

- Prices decrease as more firms enter the market