Oligopoly and Monopolistic Competition

Readings: Ch. 13

In the last two chapters, we considered two extreme forms of industry. We will now turn to more intermediate and realistic cases.

Definition
Oligopoly: An industry in which there are only a few important sellers of an identical product.

Definition
Monopolistic Competition: An industry in which there are (1) numerous firms each providing different but very similar products (close substitutes) and (2) free entry and exit.

Objective: Understand the basic characteristics of oligopoly and monopolistic competition.

With perfect competition, firms are price takers. They do not believe individual actions can influence the equilibrium price and quantity.

With monopoly, one firm provides output. Therefore, the monopolist can choose its equilibrium price and quantity given market demand.

With oligopoly and monopolistic competition, firms recognize that their decisions influence, but do not ultimately determine the equilibrium price and quantity. The choice of one firm affects the profit potential of other firms, which results in strategic interactions among firms. It is these strategic interactions that make oligopoly and monopolistic competition so interesting to study.

Objective: Understand firm behavior in the Cournot Duopoly model.

The first example of oligopoly we will explore is a duopoly where there are only two firms producing the exact same product. As mentioned previously, the characteristic of a duopoly is that the choice of one firm affects the profit potential of another. To understand how, we must make some decisions regarding what those choices are and how they affect a competitor.

In the Cournot model, we assume the choice a firm can make is its output. We also assume that its choice of output affects the demand for its competitor's output.

For example, suppose the demand curve is given by $P = a - bQ$, where Q is industry output. With two firm's in the industry we can write $Q = Q_1 + Q_2$ where Q_1 is the output choice of firm 1 and Q_2 is the output choice of firm 2, so $P = a - b(Q_1 + Q_2)$.

Question: How does Firm 1's choice of output affect the demand for Firm 2's output?

Rewrite the demand function as $P = (a - bQ_1) - bQ_2$. In this form, we see that the intercept of Firm 2's demand curve is $(a - bQ_1)$, while the slope is $-b$. Therefore, the more Firm 1 produces the lower Firm 2's demand.

For example, suppose a = 100 and b = 5, Figure 1 shows the demand for Firm 2's output when Firm 1 produces 5, 10, and 15 units of output. Figure 1 shows that the demand for Firm 2's output shifts down as Firm 1 increases output.

Firm 2's choice of output has the exact same affect on the Firm 1's demand.

So, how is a firm suppose to make a decision when its demand depends on its competitors output, which is not known when the firm makes its choice?

As with the perfect competitor and monopolist, a duopolist maximizes profit by setting its marginal cost equal to its marginal revenue. But, what is the duopolist's marginal revenue? Total revenues for Firm 1 equals price multiplied by its output where price will depend on each firm's level of output:

$TR_1 = PQ_1 = (a - b(Q_1 + Q_2))Q_1 = aQ_1 - bQ_1^2 - bQ_2Q_1$.

Marginal Revenue will then be equal to the derivative of total revenue:

$MR_1 = TR_1' = a - 2bQ_1 - bQ_2$.

Similarly, we can show that total and marginal revenue for firm 2 is

$TR_2 = PQ_2 = (a - b(Q_1 + Q_2))Q_2 = aQ_1 - bQ_1Q_2 - bQ_2^2$

and

$MR_2 = TR_2' = a - bQ_1 - 2bQ_2$.

To keep things relatively simple, lets assume that both firms have no cost, such that $MC_1 = MC_2 = 0$. Firm 1's profit maximizing level of output is determined by $MC_1 = MR_1 \Rightarrow 0 = a - 2bQ_1 - bQ_2 \Rightarrow Q_1^* = \frac{a - bQ_2}{2b}$. Similarly, firm 2's profit maximizing level of output will be determined by $Q_2^* = \frac{a - bQ_1}{2b}$.

But we are still left with the problem of knowing how much each firm should produce when its optimal level of output depends on its competitor's output, a quantity that is not known ahead of time.

Hypothesis: Since the two firms are identical, they should choose the same level of output: $Q_1^* = Q_2^*$.

If this hypothesis is correct,

$$Q_1* = \frac{a - bQ_1*}{2b} \Rightarrow Q_1* + \frac{bQ_1*}{2b} = \frac{a}{2b} \Rightarrow Q_1*\left(1 + \frac{b}{2b}\right) = \frac{a}{2b} \Rightarrow Q_1*\left(\frac{2b+b}{2b}\right) = \frac{a}{2b} \Rightarrow$$

$$Q_1*\frac{3}{2} = \frac{a}{2b} \Rightarrow Q_1* = Q_2* = \frac{a}{3b}.$$

Total output will then be $Q* = Q_1* + Q_2* = 2a/(3b)$ and the equilibrium price will be $P* = a – b(2a/3b) = a/3$. Total revenue for each firm will be $TR_1* = TR_2* = (a/3)(a/3b) = a^2/(9b)$. Since firms have no cost, this total revenue is equal to profit. Industry profit is equal to $\Pi* = TR_1* + TR_2* = a^2/9b + a^2/9b = 2a^2/(9b)$.

Objective: Understand the affect of strategic behavior in the Cournot duopoly model.

Suppose the two firms merged into a single firm, so they have a monopoly. Total and marginal revenue are then

$TR = P(Q)Q = (a – bQ)Q = aQ – bQ^2$

and

$MR = TR' = a – 2bQ$.

The profit maximizing output satisfies $MC = MR \Rightarrow 0 = a – 2bQ* \Rightarrow Q* = a/2b$ and the equilibrium price is $P* = a – b(a/2b) = a/2$. Industry profit is $\Pi* = P*Q* = (a/2)(a/2b) = a^2/(4b)$, which is greater than industry profit for the Cournot duopoly: $a^2/(4b) > 2a^2/(9b)$.

Even if the two firms do not merge, both could choose to produce half the monopoly output, $Q_1* = Q_2* = a/(4b)$, resulting in an equilibrium price of $P* = a/2$ and individual firm profits equal to $\Pi_1* = \Pi_2* = (a/2)(a/(4b)) = a^2/(8b)$, which is greater than individual firm profit in the Cournot Duopoly model because $a^2/(8b) > a^2/(9b)$.

Question: Why won't firms produce the monopoly output, $Q_1* = Q_2* = a/(4b)$, instead of the Cournot output, $Q_1* = Q_2* = a/(3b)$, so they can earn higher profits?

To make things more concrete, lets again assume a = 100 and b = 5. If each firm produces the duopoly output, $Q_1* = Q_2* = 100/(3\times5) = 20/3$ and $\Pi_1* = \Pi_2* = 100^2/(9\times5) = 2,000/9 = 222.2$. If each firm produces the monopoly output, $Q_1* = Q_2* = 100/(4\times5) = 20/4$ and $\Pi_1* = \Pi_2* = 100^2/(8\times5) = 2,000/8 = 250$.

Suppose instead that one firm produced the Cournot output, while the other produced the monopoly output: $Q* = 20/3 + 20/4 = (80 + 60)/12 = 140/12$. The equilibrium price would then be $P* = 100 – 5\times(140/12) = (1200 – 700)/12 = 500/12$. Profit for the firm that produced the

Cournot output would be $\Pi_C^* = (500/12)(20/3) = 10{,}000/36 = 277.8$. Profit for the firm that produced the monopoly output would be $\Pi_M^* = (500/12)(20/4) = 10{,}000/48 = 208.3$.

We can now summarize the duopoly firm's choice between the Cournot and Monopoly Output as a game:

		Firm 2's Output	
		$Q_2 = 20/4$	$Q_2 = 20/3$
Firm 1's Output	$Q_1 = 20/4$	**250**, *250*	**208.3**, *277.7*
	$Q_1 = 20/3$	**277.7**, *208.3*	**222.2**, *222.2*

In this game, Firm 1 chooses the row by choosing either the Cournot ($Q_1 = 20/3$) or Monopoly output ($Q_1 = 20/4$). Firm 2 chooses the column by choosing either the Cournot ($Q_2 = 20/3$) or Monopoly output ($Q_2 = 20/4$). Firm 1's profits are in bold, while Firm 2's profits are in italics. If both choose the monopoly output both earn 250. If both choose the Cournot output, both earn 222.2. If Firm 1 chooses the Cournot output and Firm 2 chooses the Monopoly output, Firm 1 earns 277.7, while Firm 2 earns 208.3. If Firm 1 chooses the Monopoly output and Firm 2 chooses the Cournot output, Firm 1 earns 208.3, while Firm 2 earns 277.7.

Question: Does it ever make sense for both firms to choose the monopoly output?

Notice that if Firm 2 chooses the Monopoly output, Firm 1 maximizes its profit by choosing the Cournot output. If Firm 2 chooses the Cournot output, Firm 1 maximizes its profit by again choosing the Cournot output. Therefore, regardless of what Firm 2 does, Firm 1 maximizes its profits by choosing the Cournot output. Similarly, regardless of what Firm 1 does, Firm 2 maximizes its profit by choosing the Cournot output. Together, these results suggest both firms will choose the Cournot output and earn 222.2.

In this context, choosing the monopoly output seems to make little sense even though both firms would ultimately be better off.

Objective: Understand reaction functions and the Nash equilibrium.

We used a little trick to solve the Cournot model in the previous example. We argued that since the two firms were identical, they would choose to produce the same output. But what if the two firms are not identical? Suppose for instance that Firm 1's marginal cost of production is $MC_1 = c_1$, while Firm 2's marginal cost of production is $MC_2 = c_2$ where $c_1 \neq c_2$. Does it still make sense to assume both firms will produce the same level of output? Probably not. So how do we find each firm's profit maximizing level of output?

We can solve this problem by again setting marginal cost equal to marginal revenue: $c_1 = a - 2bQ_1 - bQ_2$ for Firm 1 and $c_2 = a - bQ_1 - 2bQ_2$ for Firm 2. Firm 1's profit maximizing output is then determined by $Q_1* = \frac{a - bQ_2 - c_1}{2b} = R_1(Q_2)$, while Firm 2's profit maximizing output is determined by $Q_2* = \frac{a - bQ_1 - c_2}{2b} = R_2(Q_1)$. The functions $R_1(Q_2)$ and $R_2(Q_2)$ tell us Firm 1's and Firm 2's optimal level of output given their competitors level of output and are referred to as reaction functions or best response functions.

Definition
Reaction/Best Response Function: A curve that tells the profit maximizing level of output for one oligopolist for each quantity supplied by others.

To find out how much each firm will produce, we want to find a level of output for each firm so that no firm has an incentive to change its output. That is, we want to find a combination of outputs such that each firm can do no better by changing its level of output individually. This idea is attributable to a famous mathematician named John Nash, who subsequently won the 1994 Nobel Prize in economics. Yes, he also had an award-winning movie made about his life called "A Beautiful Mind." Unfortunately, the movie was more about him being crazy and less about his path breaking contributions in math and economics. But I suppose crazy sells more movie tickets than math and economics.

Definition
Nash Equilibrium: A combination of outputs such that each firm's output maximizes its profit given the output chosen by other firms.

What this means is that Firm 1's Nash equilibrium output should maximize its profits given Firm 2's Nash equilibrium output, $Q_1* = \frac{a - bQ_2 * - c_1}{2b} = R_1(Q_2 *)$, and Firm 2's Nash equilibrium output should maximize its profits given Firm 1's Nash equilibrium output, $Q_2* = \frac{a - bQ_1 * - c_2}{2b} = R_2(Q_1 *)$. What we need to do now is find Q_1* and Q_2*. The easiest way to accomplish this is to rewrite $Q_1* = \frac{a - bQ_2 * - c_1}{2b} = \frac{a - c_1}{2b} - \frac{1}{2}Q_2 *$ and $Q_2* = \frac{a - bQ_1 * - c_2}{2b}$ as $Q_1* = \frac{a - c_2}{b} - 2Q_2 *$. Setting our two new equations equal to each other, $\frac{a - c_1}{2b} - \frac{1}{2}Q_2 * = \frac{a - c_2}{b} - 2Q_2 *$, we can now solve for Q_2*: $\frac{3}{2}Q_2* = \frac{a - c_2}{b} - \frac{a - c_1}{2b} = \frac{a - 2c_2 + c_1}{2b}$ or $Q_2* = \frac{a - 2c_2 + c_1}{3b}$. Finally, we need to substitute Q_2* back into one of our reaction functions to find Q_1*: $Q_1* = \frac{a - c_1}{2b} - \frac{1}{2}\frac{a - 2c_2 + c_1}{3b} = \frac{2a - 2c_1 + 2c_2}{6b} = \frac{a - 2c_1 + c_2}{3b}$.

Graphically, what we are doing is plotting the reactions functions and then looking for where they intersect. Again assuming a = 100, b = 5, $c_1 = 50$, and $c_2 = 45$, $Q_1^* = \frac{a - c_1}{2b} - \frac{1}{2}Q_2^* = 5 - \frac{1}{2}Q_2^*$ and $Q_1^* = \frac{a - c_2}{b} - 2Q_2^* = 11 - 2Q_2^*$. These two equations are graphed in Figure 2. The intersection at Point **A** tells us the Nash equilibrium output for each firm: $Q_1^* = 3$ for Firm 1 and $Q_2^* = 4$ for Firm 2.

Objective: Understand firm behavior in the Bertrand Duopoly model.

The Cournot Duopoly model assumes firms choose how much to produce and then the market determines the price. Bertrand however challenged this assumption by arguing firms are more likely to choose prices instead of quantities, which is the assumption used in the Bertrand model. While it is not immediately clear why this would make a difference, it does. It makes a huge difference.

Consider our example of only two firms with no costs and the demand curve P = a – bQ.

Question: How does the price chosen by one firm affect the demand for the other's output?

Suppose Firm 1 chooses a price P_0. If Firm 2 chooses a price above P_0, no one will want to buy from it because they can get the product cheaper elsewhere. If Firm 2 chooses a price equal to P_0, it will split the market demand with Firm 1 since consumers will not care who they buy from. If Firm 2 chooses a price below P_0, it will have the market all to itself because all consumers can get the product cheaper from Firm 2. Figure 3 shows an example of Firm 2's demand for a = 100 and b =5 when Firm 1 chooses a price of 75, 50, and 25.

Bertrand then argued that a firm always has an incentive to just undercut its competitor's price as long as its price exceeds the marginal cost, and average variable cost in the short run or average total cost in the long run. If firms are identical and both act in this manner, both firms will lower their price until $P = MC_1 = MC_2$, which is 0 in our example. Each firm's profit and industry profit are then also equal to 0 in our example.

This result is very different from the Cournot Model where firms restricted output and were able to capture an economic profit. In a Bertrand Duopoly, more will be produced at a lower price and no economic profits will exist.

Objective: Understand firm behavior in the Stackelberg Duopoly model.

In the Cournot Duopoly model, before a firm could choose its profit maximizing output, it needed to make some assumption about what its competitor was going to produce. The Nash equilibrium then provided some rationale for the firm's choice. But suppose one of the firms knew how much its competitor was going to produce and that its competitor knew it knew. This is the question posed by Stackelberg. The basic idea was that one firm had an advantage and could produce its output first because its competitor faced some sort of production delay. In this instance, the firm with delayed production is forced to respond optimally given its competitors choice of output. Does this matter? Yes! Let see why.

Again, we will assume there are only two firms with no cost and demand equal to P = a –bQ.

Question: How much should a Duopolist produce if it knows its competitor is faced with some sort of production delay?

Both firms will produce where marginal revenue equals marginal cost. What differs between the two firms is marginal revenue. If Firm 1 gets to produce first, total revenue is $TR_1 = PQ_1 = aQ_1 - bQ_1^2 - bQ_2Q_1$, but we can go further. Since Firm 1 produces first, it gets to show Firm 2 its level of output before Firm 2 chooses its output. From earlier, we know that Firm 2's profit maximizing level of output is $Q_2{}^* = \frac{a - bQ_1}{2b}$, which Firm 1 can use to determine its total revenue:

$$TR_1 = aQ_1 - bQ_1^2 - bQ_1\left(\frac{a - bQ_1}{2b}\right) = \frac{a}{2}Q_1 - \frac{b}{2}Q_1^2 .$$

Taking the derivative of total revenue, marginal revenue is

$$MR_1 = TR_1' = \frac{a}{2} - bQ_1$$

If we set marginal revenue equal to marginal cost, $a/2 - bQ_1^* = 0 \Rightarrow Q_1^* = a/(2b)$. We can now use $Q_2{}^* = \frac{a - bQ_1}{2b}$ to find Firm 2's profit maximizing level of output: $Q_2{}^* = \frac{a - ba/2b}{2b} = \frac{a}{4b}$. Industry output is $Q^* = Q_1^* + Q_2^* = a/2b + a/4b = 3a/(4b)$. Substituting back into the demand function tells us the equilibrium price: $P^* = a - b(3a/4b) = a/4$. Profit for Firm 1 is $\Pi_1^* = TR_1 = P^*Q_1^* = (a/4)(a/2b) = a^2/(8b)$. Profit for Firm 2 is $\Pi_2^* = TR_2 = P^*Q_2^* = (a/4)(a/4b) = a^2/(16b)$. Total industry profit is $\Pi^* = \Pi_1^* + \Pi_2^* = a^2/8b + a^2/16b = 3a^2/(16b)$.

If a = 100 and b = 5, $Q_1^* = 100/(2\times5) = 10$, $Q_2^* = 100/(4\times5) = 5$, $Q^* = 10 + 5 = 15$, $P^* = 100/4 = 25$, $\Pi_1^* = 100^2/(8\times5) = 250$, $\Pi_2^* = 100^2/(16\times5) = 125$, and $\Pi^* = 250 + 125 = 375$.

Objective: Understand how price, industry output, and industry profit compare for monopoly, Cournot duopoly, Stackelberg duopoly, Bertrand duopoly, and perfect competition.

We can now summarize our results for each of the industry structures we have considered to date when P = a – bQ and marginal costs are 0 (see Table 1).

Table 1

Model	Industry Output (Q*)	Market Price (P*)	Industry Profit (Π*)
Monopoly	$Q_M^* = a/(2b)$	$P_M^* = a/(2b)$	$\Pi_M^* = a^2/(4b)$
Cournot	$(4/3)Q_M^*$	$(2/3)P_M^*$	$(8/9)\Pi_M^*$
Stackelberg	$(3/2)Q_M^*$	$(1/2)P_M^*$	$(3/4)\Pi_M^*$
Bertand	$2Q_M^*$	0	0
Perfect Competition	$2Q_M^*$	0	0

What Table 1 shows is that the monopoly quantity is the lowest followed by the Cournot, Stackelberg, and Bertrand and Perfect Competition.

Alternatively, the monopoly price and industry profit is the highest followed by the Cournot, Stackelberg, and Bertrand and Perfect Competition.

What is most important to remember about Table 1 is that it shows how the assumptions we use to characterize firm interactions in an oligopoly market are very important. Whether firms choose quantity or price is important. What firms know and when they know it is also important.

Objective: Understand the classical model of monopolistic competition.

Recall that two important conditions for monopolistic competition are (1) numerous firms each providing different but very similar products (close substitutes) and (2) free entry and exit.

The book provides several interesting examples of how monopolistic competition may arise in the restaurant industry. In these examples, what differentiates each restaurant's product is its location. Customers have to choose between restaurants offering the same menu, but some were further away than others. Another way restaurants may choose to differentiate themselves from competitors is by their menu or the quality of food they serve.

The basic idea of monopolistic competition is that firms can convince you that what they have to offer is some how different and better than their competitors. If they can accomplish this objective, they will be able to charge a higher price and increase profits. However, since there are close substitutes, if the price charged is too high they will begin to lose customers and profit will fall.

In the short run, we can think of our monopolistic competitor almost like any other monopolist. There are however two important differences. First, the firm's demand and marginal revenue will depend on what other monopolistic competitors choose to do. For example, recall that demand shifts down when the price of a substitute decreases. Therefore, the demand for the monopolistic competitor will be lower when other monopolistic competitors choose to charge a lower price. A monopolistic competitor must choose its output and price while keeping this fact in mind, which is very similar to problem faced by oligopolists. Second, contrary to monopolists, monopolistic competitors cannot sustain economic profits in the long run.

Suppose we have two monopoly firms whose products are close substitutes, so each firm's demand depends on its competitor's price. Given its competitors price, each firm maximizes profit by setting marginal cost equal to marginal revenue: $MC_1 = MR_1(P_1, P_2)$ and $MC_2 = MR_2(P_2, P_1)$. Figure 4 illustrates for Firm 1. When Firm 2's price is lower, Firm 1's demand is lower, which leads it to offer less for sale at a lower price. We can write Firm 1's profit maximizing price as a function of Firm 2's price, which gives us Firm 1's reaction function: $P_1 = R_1(P_2)$. Similarly, we can write Firm 2's profit maximizing price as a function of Firm 1's price, which gives us Firm 2's reaction function: $P_2 = R_2(P_1)$. We can then graph these reaction functions and look for the combination of prices that satisfies both simultaneously (see Figure 5 for an example): $P_1^* = R_1(P_2^*)$ and $P_2^* = R_2(P_1^*)$.

In the short run, a monopolistic competitor will shut down if average variable cost is always above demand (See Figure 6 for example). In the long run, a monopolistic competitor will shut down if long run average cost is always above demand (See Figure 7 for example).

The final thing to note about monopolistic competition is that there can be no economic profit in the long run because there is free entry and exit. Therefore, long run average cost will equal the price (see Figure 8 for example). What this means is that if our demand curve is downward sloping, long run average cost will not be at a minimum as it was in a perfectly competitive market in the long run.

We like the fact that the long run perfectly competitive equilibrium resulted in production at the minimum long run average cost. So should we be unhappy about the fact that monopolistic competitors typically will not produce at the minimum long run average cost?

Some say yes, while others say no. Monopolistic competition leads to product diversity. If we didn't have it all cars, houses, restaurants, etc. would be identical. Therefore, those in favor of monopolistic competition argue that costs above the minimum long run average are simply the price we pay for product diversity.

Two More Important Notes on Monopolistic Competitors:

As with any monopoly, monopolistic competitors will never produce where demand is inelastic because when demand is inelastic an increase in price increases total revenue and decreases the quantity demanded, which also decreases total cost. If increasing price increases revenue and decreases cost, profit must rise. Therefore, as long as demand is inelastic a monopolistic competitor can increase profit by increasing the price it charges.

A MONOPOLISTIC COMPETITOR ALSO DOES NOT HAVE A SUPPLY CURVE!

Objective: Understand the basic concepts of economic games and how they are solved.

Throughout our discussions of oligopoly and monopolistic competition, I have emphasized the strategic interactions competitors must deal with in order to determine how to behave. These strategic interactions and their importance have lead to an increasing emphasis in economics on games and how they are played.

So, what is a game?

Games have players that interact with one another based on some set of rules. The rules of the game tell us what people are allowed to do and when. They also tell us what people know and when. Finally, they tell us how people are rewarded based on what each player chooses to do.

To play a game, a player must formulate a strategy.

Definition
Strategy: A complete description of what a player does given what it knows.

A strategy must be detailed enough for you to give to someone else so they can play the game for you without having to ask you any questions. This means it must specify what you want to do given all possible contingencies.

Consider the familiar game of rock/paper/scissors. In this game, there are two players. Each player must choose simultaneously whether to play rock, paper, or scissors. Rock beats scissors, scissors beats paper, and paper beats rock. The winner receives some prize from the loser. If there is a tie, both players get nothing.

Each player for this game has three possible strategies: (i) choose rock, (ii) choose paper, and (iii) choose scissors.

Now lets change the rules of the game a little. Suppose the tallest player has to choose his strategy first and reveal it to the shorter player before he chooses. For these rules, the tallest player still has only three strategies: (i) choose rock, (ii) choose paper, and (iii) choose scissors. The shorter player however has lots more than just three strategies because it gets to choose after knowing what the taller player did. An example of one such strategy is: if the tall player chooses rock, choose paper; if the tall player chooses paper, choose scissors; and if the tall player chooses scissors, choose rock. All that is just one strategy. It provides a complete description of what the short person will do for each possible choice the tall player can make. The strategy has to be so detailed because you do not know what the tall player will actually choose before the game starts, but once it does you will find out before you make your choice.

Another potential strategy is if the tall player chooses rock, choose rock; if the tall player chooses paper, choose rock; and if the tall player chooses scissors, choose rock.

The traditional rock/paper/scissors is an example of a simultaneous move game, as is the Cournot and Bertrand Duopoly models. In a simultaneous move game, players receive no new information during the course of the game. Everybody makes their decision without knowing precisely what the other players will do.

Our modified version of rock/paper/scissors is an example of a sequential game, as is the Stackelberg model. In a sequential move game, some players receive new information during the course of the game and have the opportunity to respond to this new information.

Simultaneous move games are easily described using a game matrix. There is a different row in the matrix for each strategy the first player has. There is a different column in the matrix for each strategy the second player has. For example, here is another example of the prisoner's dilemma game.

		Your Partner's Choice	
		A	*B*
Your Choice	**A**	*100* **100**	*200* **0**
	B	*0* **200**	*50* ***50***

Your strategies are listed in the rows, while your opponent's strategies are listed in the columns. In each cell of the matrix are each players' rewards for the corresponding strategy choices. Yours are denoted in bold in the lower left-hand corner. Your opponents are denoted in italics in the upper right-hand corner. We will continue to use this convention throughout the rest of the class.

Sequential games can also be described using a game matrix, but as we saw with our modified version of rock/paper/scissors, sequential games lead to a proliferation of strategies that can be difficult to fit into a reasonable game matrix. Also, using a game matrix makes it harder to see how players should play the game.

Another way to describe a sequential game is a game tree. Figure 9 offers an example for our modified rock/paper/scissors game when the reward is $10 and rock is denoted by R, paper by P, and scissors by S. Rewards are denoted in parentheses. The first reward corresponds to the tall player, while the second corresponds to the short player. The tall player moves first. After making his selection the short player gets to move knowing what the tall player has already done.

Now that we have a way to describe the strategic interactions among firms or consumers, or consumers and firms, or students and teachers or girl friends and boy friends or husbands and wives… we are ready to talk in more detail about how we can solve these games to say what players should do and predict what players will do. We will call these solutions equilibrium.

To start, we will focus on simultaneous move games. The most appealing way to solve a simultaneous move games is to look for dominant strategies.

Definition
Dominant Strategy: The strategy in a game that produces the best results irrespective of the strategy chosen by an opponent.

If a player has a dominant strategy, they should always play it and usually do. Our prisoner's dilemma offers a good example of a game where both players have a dominant strategy. You were always better off choosing B regardless of what your opponent did and your opponent was always better off choosing B regardless of what you did. Therefore, the dominant strategy equilibrium is for both players to choose B.

Unfortunately, there are lots of strategic interactions where one or more players do not have a dominant strategy. For example, lets modify our prisoner's dilemma just a bit, so your opponent only earns 75 instead of 200 when you choose A and he chooses B.

		Your Partner's Choice	
		A	*B*
Your Choice	**A**	*100* **100**	*75* **0**
	B	*0* **200**	*50* ***50***

Now your partners profit maximizing choice depends on whether you choose A or B. He is better off choosing A when you choose A and B when you choose B.

Since your partner no longer has a dominant strategy, there is no longer a dominant strategy equilibrium and we must look further. This is where the Nash equilibrium becomes useful again. To find the Nash equilibrium, we need to look at each strategy combination and ask ourselves if either player would want to change their strategy. For example, if you and your opponent both chose A, would either of you want to change your strategy. You would because you could earn 200 instead of 100, but your opponent wouldn't because he would only earn 75 instead of 100. However, since someone would want to change, both of you choosing A is not a Nash equilibrium. For **B** and *A*, your partner would want to change, so this is not a Nash equilibrium. For **A** and *B*, you would want to change, so this is not a Nash equilibrium. For **B** and *B*, neither of you would want to change, so this is a Nash equilibrium.

For the most part, the Nash equilibrium takes us pretty far. What is nice about it is that we are assured there will always be at least one. What is not so nice about it is that there can be many. For example, lets modify our prisoner's dilemma once again so that you also only earn 75 instead of 200 when you choose B and your opponent chooses A.

		Your Partner's Choice	
		A	*B*
Your Choice	**A**	*100* **100**	*75* **0**
	B	*0* **75**	*50* ***50***

For **A** and *A*, neither of you would want to change, so this is a Nash equilibrium. For **A** and *B*, your opponent wants to change. For **B** and *A*, you want to change. For **B** and *B*, neither of you would want to change, so this is a Nash equilibrium. Here there are two Nash equilibrium: (1) **A** and *A* and (2) **B** and *B*.

For situations like this, we must dig even deeper. How? Well, there are lots of different ways that have been proposed. We will not go into them in this class, but one example, is to ask if one of the equilibrium makes both players better off? Indeed, this is the case for **A** and *A*, so we might expect that to be the answer.

For simultaneous move games, you need to know how to find the dominant strategy equilibrium when there is one or the Nash equilibrium when there is no dominant strategy equilibrium. Actually, any dominant strategy equilibrium is a Nash equilibrium, so all you really need to know is how to find a Nash equilibrium.

To solve sequential games we can take advantage of the game tree and what it tells us about what people can do when. The important thing to remember about solving sequential games is that we want to work backwards. But what do I mean by work backward?

Consider the game in Figure 10. Firm 1 and 2 get to choose whether to produce a High output or Low output. However, Firm 1 gets to choose first and firm 2 gets to choose after seeing what Firm 1 does. The first number in parentheses is Firm 1's profit, while the second is Firm 2's profit.

To solve this game working backward, we start by asking the question, what should Firm 2 do if Firm 1 chooses a Low output. If Firm 1 chooses a Low output, Firm 2 earns 420 from choosing a High output and 400 from choosing a low output, so it maximizes its profit by choosing the High output. This allows us to eliminate Firm 2's choice of a Low output when Firm 1 chooses a Low output (See Figure 11).

Now we ask the question, what should Firm 2 do if Firm 1 chooses a High Output? If Firm 1 chooses a High output, Firm 2 earns 260 from also choosing a High output and 280 from choosing a low output, so it maximizes its profit by choosing the Low output. This allows us to eliminate Firm 2's choice of a High output when Firm 1 chooses a High output (See Figure 12).

Now that we know what Firm 2 will do in response to Firm 1's choice we are ready to ask what Firm 1 should do. We know from our answer to the first two questions that Firm 2 will choose a High output if Firm 1 chooses a Low output and a Low output if Firm 1 chooses a High output. Therefore, if Firm 1 chooses a Low output it will earn 420. If it chooses a High output it will earn 460. Since 460 is more than 420, Firm 1 should choose a High output.

The equilibrium for this game is for Firm 1 to choose a High output and Firm 2 to respond by choosing a Low output. Actually, we should be more specific. The equilibrium for this game is for Firm 1 to choose a High output and for Firm 2 to choose Low output if Firm 1 chooses a High output and a High output if Firm 1 chooses a Low output.

That's all there is to it.

This sequential game is the same as the Stackelberg model we looked at earlier except we only allow firms to choose between two different levels of output. Remember that in the Stackelberg model we substituted Firm 2's reaction function (or profit maximizing level of output given Firm 1's output) into Firm 1's total revenue function before calculating its marginal revenue. This is the same as crossing out the strategies that do not maximize Firm 2's profit given Firm 1's choice of output in Figure 12.

If you check you will see that this equilibrium is a Nash equilibrium, but it is also more than that. It is more than a Nash equilibrium because it eliminates what are called incredible threats. To see what I mean, lets construct a game matrix based on Figure 10. To do this we must first describe each player's strategies. For Firm 1 this is simple, it can choose a Low output or High output. For Firm 2 things are more complicated because it must plan for contingencies. What I mean is that Firm 2's strategies must say what it will do if Firm 1 chooses a Low output and what it will do if Firm 1 chooses a High output. Therefore, Firm 2 will have four possible strategies: (i) Choose High if Firm 1 chooses Low & High if Firm 1 Chooses High, (ii) Choose High if Firm 1 chooses Low & Low if Firm 1 Chooses High, (iii) Choose Low if Firm 1 chooses Low & High if Firm 1 Chooses High, and (iv) Choose Low if Firm 1 chooses Low & Low if Firm 1 Chooses High. With these strategies we can know construct our game matrix.

			Firm 2		
		(i)	*(ii)*	*(iii)*	*(iv)*
Firm 1	**Low**	*420* **420**	*420* **420**	*400* **500**	*400* **500**
	High	*260* **340**	*280* **460**	*260* **340**	*280* **460**

To find the Nash equilibrium for this game we can ask ourselves what would Firm 1 do if Firm 2 chose *(i)*, *(ii)*, *(iii)*, or *(iv)*. If Firm 2 chooses *(i)*, Firm 1 should choose **Low**. If Firm 2 chooses *(ii)*, Firm 1 should choose **High**. If Firm 2 chooses *(iii)*, Firm 1 should choose **Low**. If Firm 2 chooses *(iv)*, Firm 1 should choose **Low**. Now what should Firm 2 do if Firm 1 chooses **Low** or **High**. If Firm 1 chooses **Low**, Firm 2 should choose *(i)* or *(ii)*. If Firm 1 chooses **High**, Firm 2 should choose *(ii)* or *(iv)*. Therefore, there are two Nash equilibrium: (1) **Low** & *(i)* and (2) **High** & *(ii)*. Yet, before we argued there was only one equilibrium: (2) **High** & *(ii)*. When we worked backward, we were able to rule out **Low** & *(i)* as an equilibrium not because it wasn't a Nash equilibrium, but because it was a Nash equilibrium that relied on what we call an incredible threat. If Firm 2 threatens to choose a High output if Firm 1 choose a High output and Firm 1 believes it, then Firm 1 will always be better off choosing a Low output, so **Low** & *(i)* is a Nash equilibrium. But, should Firm 1 believe Firm 2's threat? Not if Firm 2 is really interested in maximizing its profit because if Firm 1 chooses a Low output, Firm 2 maximizes profit by choosing a Low, not High, output. By working backward through the game, we assume these types of threats are not credible and should be eliminated.

Figure 1: Demand For Firm 2’s Output Given Firm 1’s Output

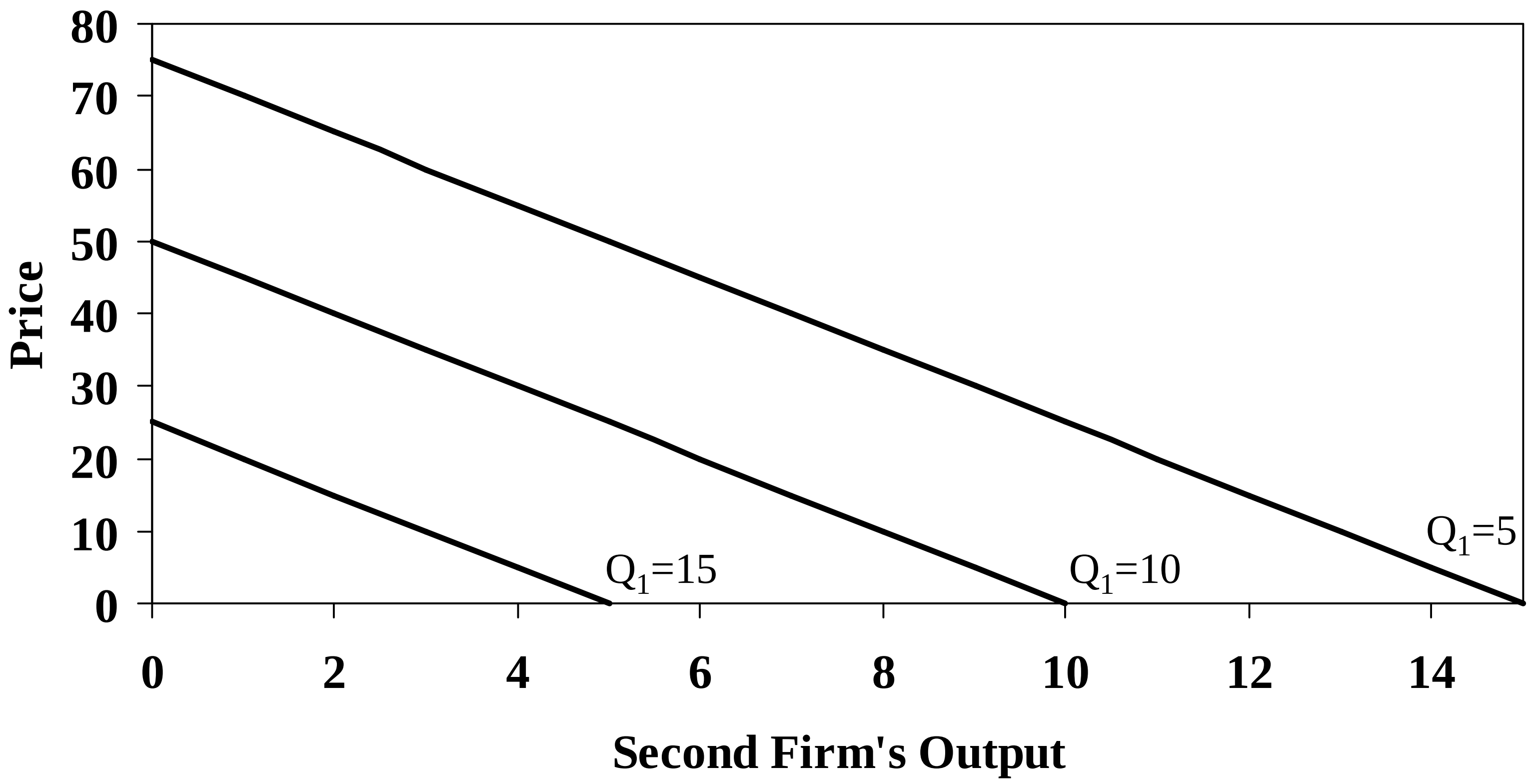

Figure 2: Reaction Functions For Duopoly Example With Different Marginal Costs

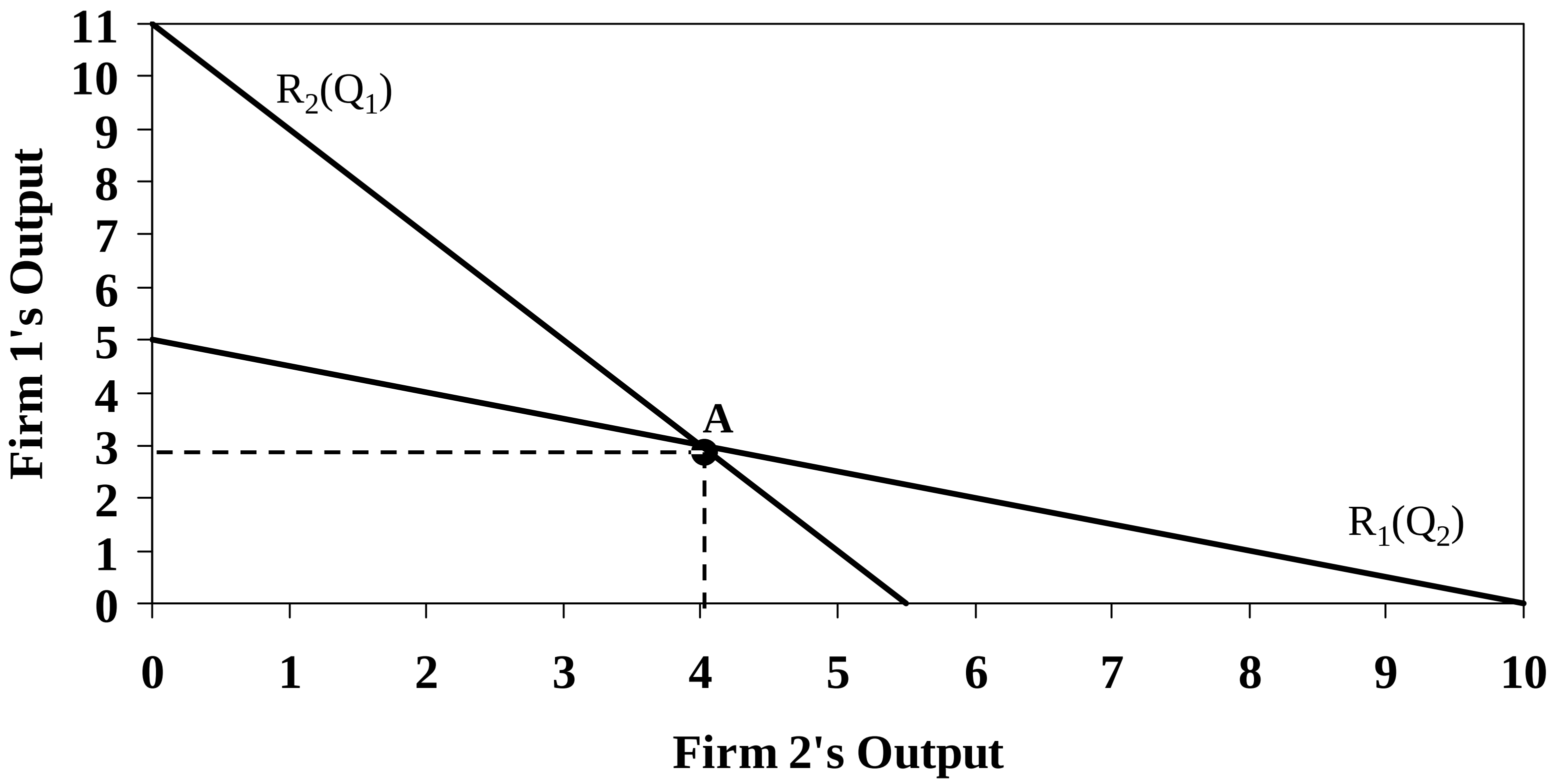

Figure 3: Firm 2's Demand Given Firm 1's Price

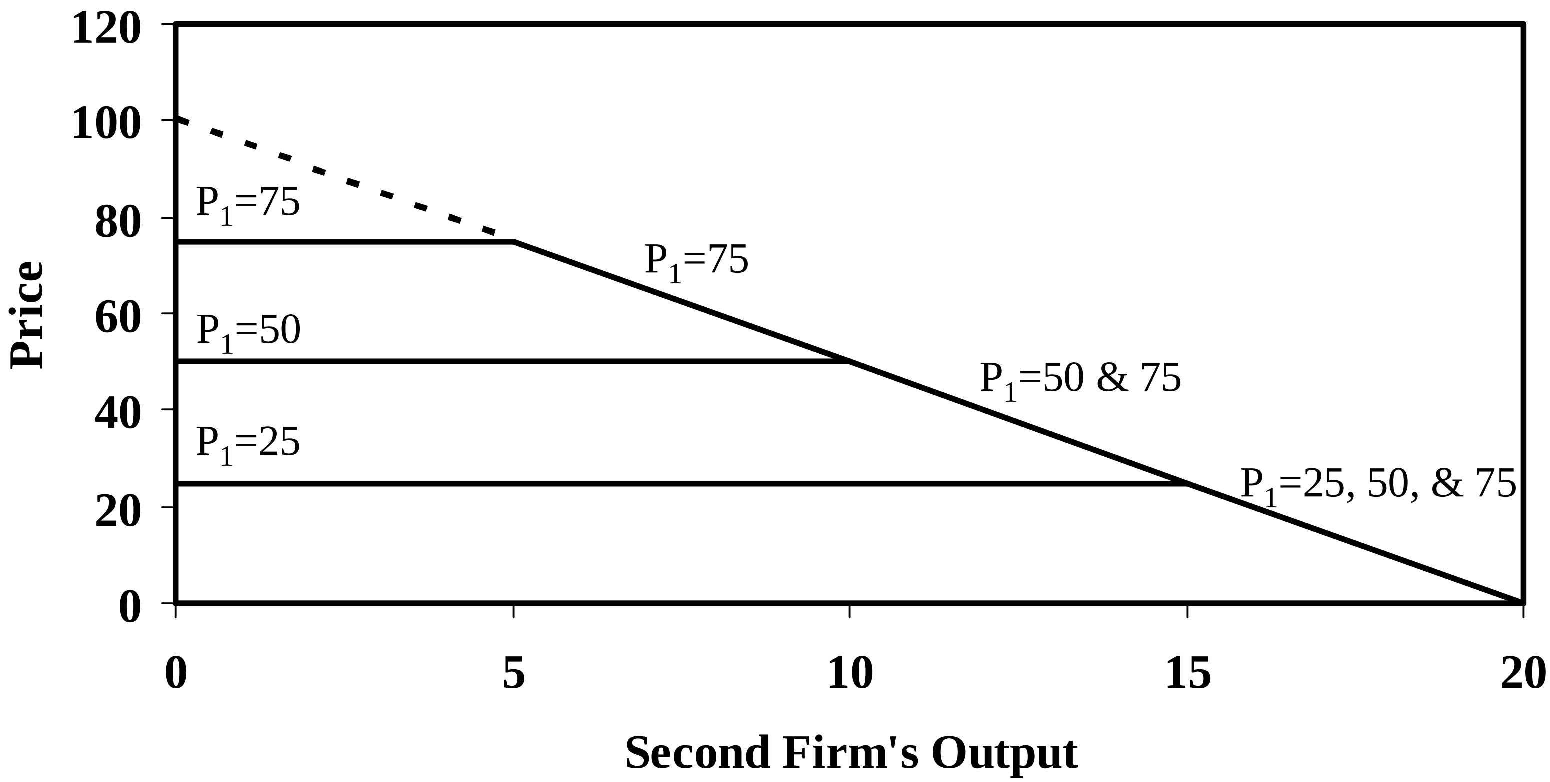

Figure 4: Monopolistic Competitor In the Short Run.

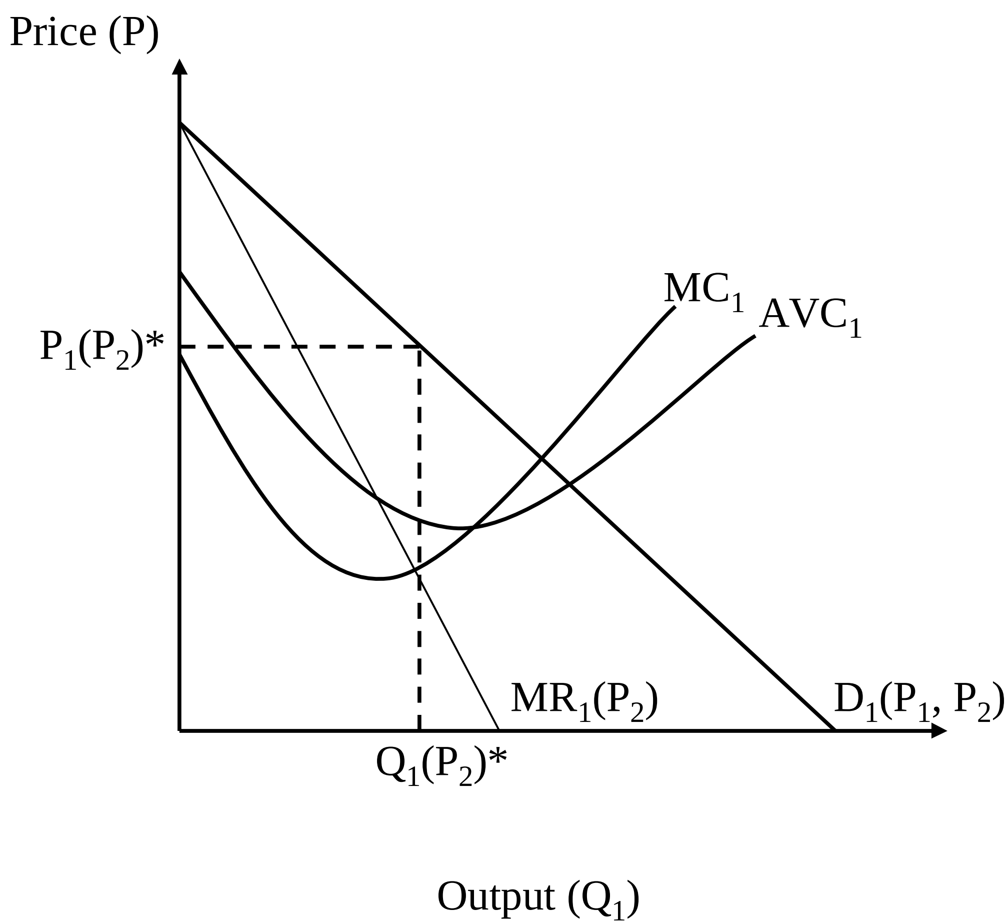

Figure 5: Reaction Functions For Monopolistic Competitors.

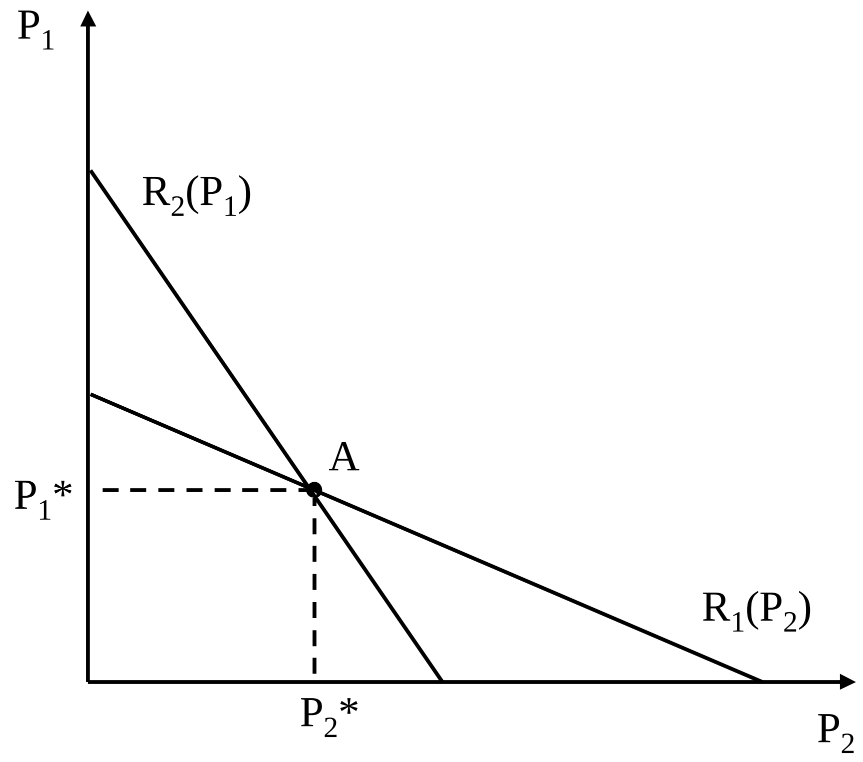

Figure 6: Monopolistic Competitor Who Will Not Operate In the Short Run.

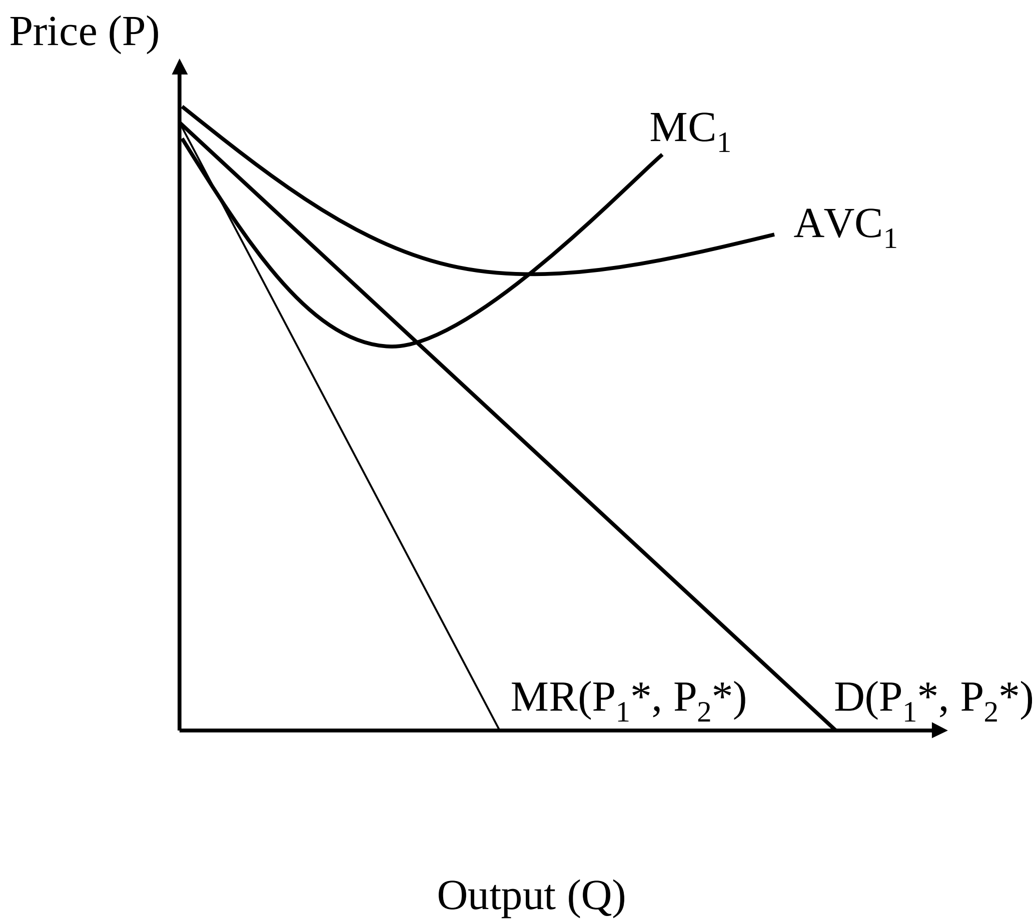

Figure 7: Monopolistic Competitor Who Will Not Operate In the Long Run.

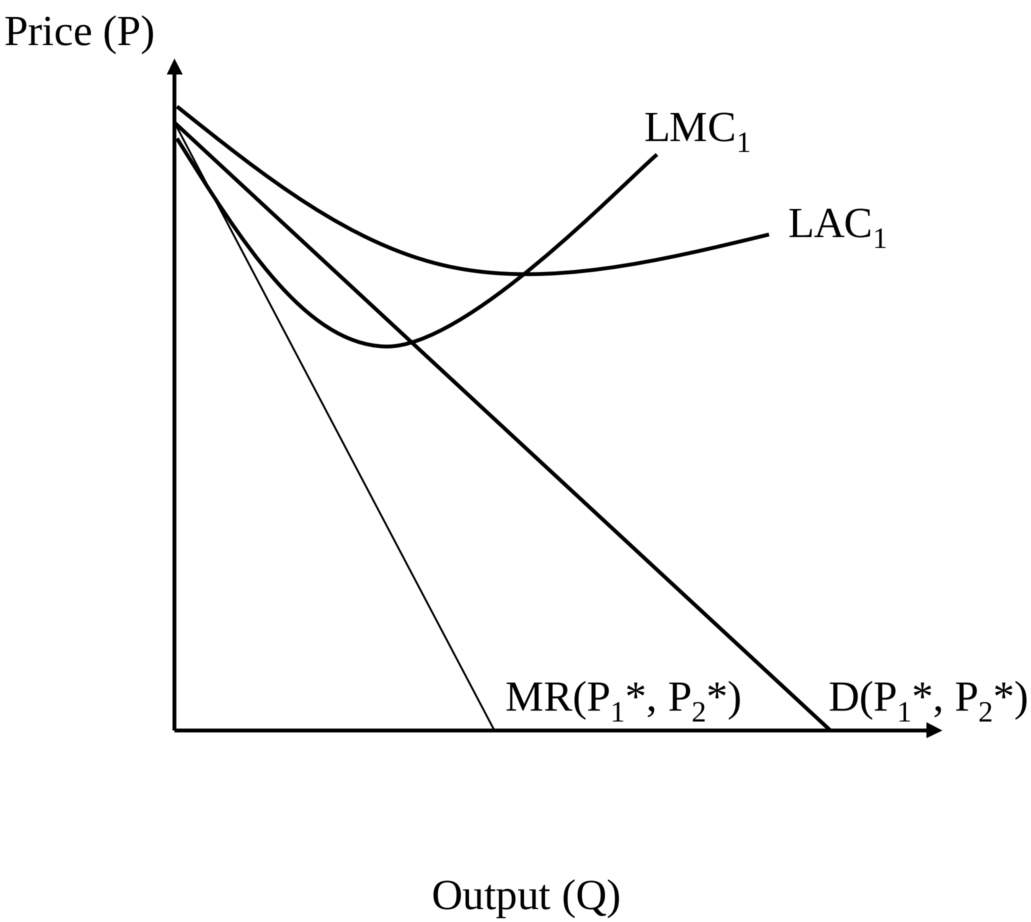

Figure 8: Monopolistic Competitor Long Run Equilibrium.

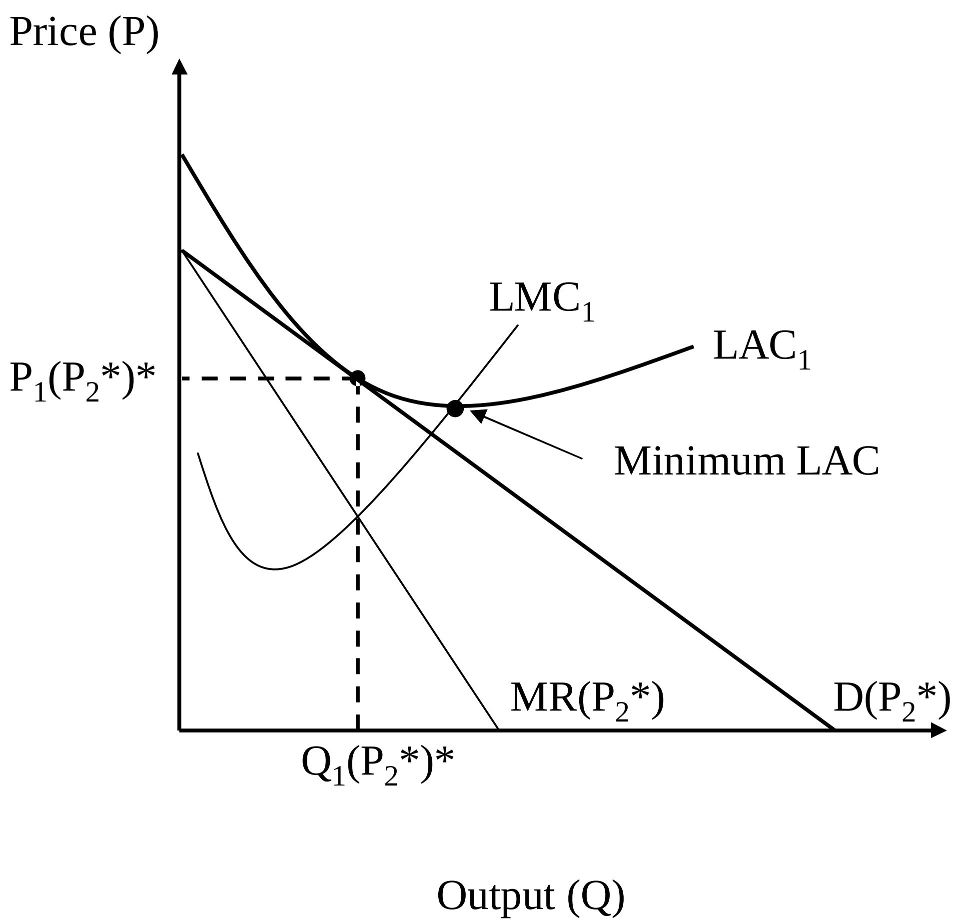

Figure 9: Game Tree For Modified Rock Paper Scissors Game

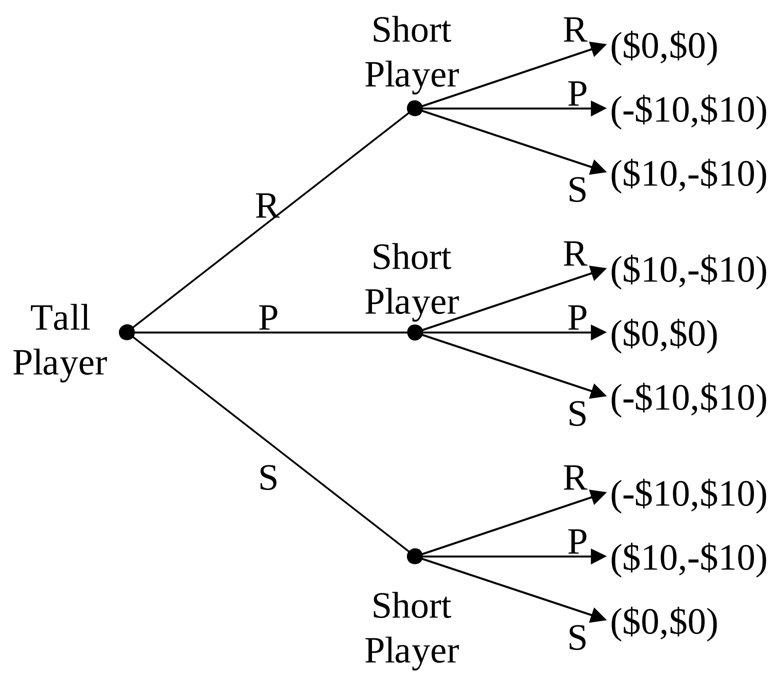

Figure 10: Sequential Game Example

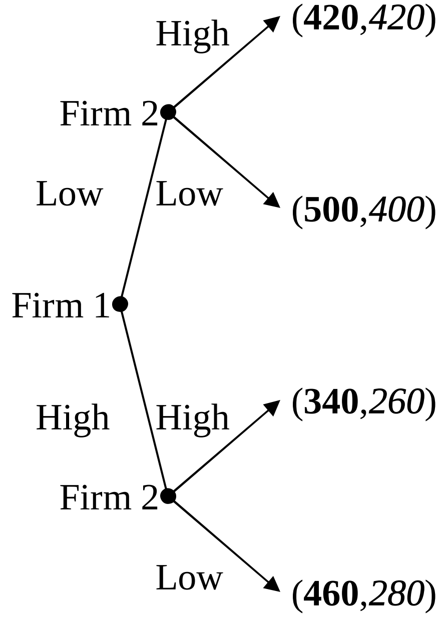

Figure 11: Firm 2 Should Never Choose a Low Output When Firm 1 Chooses a Low Output, So We Can Cross It Out

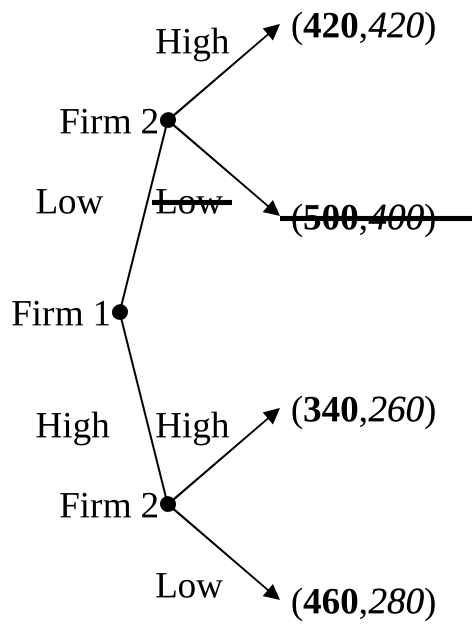

Figure 12: Firm 2 Should Never Choose a High Output When Firm 1 Chooses a High Output, So We Can Cross It Out