Price Theory

Lecture 8: Oligopoly and Strategic Behavior

I. Basic Notions of Game Theory

Game theory is an approach to modeling social situations (and sometimes other situations, as in evolutionary biology) as though they are games. We ask who the players are, what strategies they can adopt, and so forth in order to make predictions about the outcomes of games that represent real-world problems in abstract form.

Game theory is useful for many reasons, but primarily because it allows us to consider the interdependent nature of many economic situations: your choice of action depends on other people's actions, and their choice of actions depends on yours. Our models so far (perfect competition and monopoly) have relied on somewhat unrealistic assumptions to justify treating people's decisions as though they are largely independent of others' decisions. A monopoly need not worry about the actions of other firms, and the actions of consumers are neatly summarized in a demand curve. A perfectly competitive firm does not have to deal with other firms indirectly, because the cumulative actions of all firms and all consumers lead to prices -- but since no firm is large enough to affect prices alone, we can act like the price is just given. Game theoretic market models do away with these assumptions.

Any “game” in a game-theoretic model consists of three parts:

- Players. We generally say there are N players in the game, numbered 1, 2, 3, ..., n. We let *i* stand for any one player, when we don't need to be specific which of the N players we're talking about. In most simple games, N = 2.
- Strategies. Each player can choose among multiple strategies for playing the game. You can think of a strategy as a plan of action: a strategy specifies what action the player will take at any given point in time. More importantly, a strategy specifies what action a player will take at any given point in time, for any given history of the game up to that point. For example, if you're playing O's in tic-tac-toe, your strategy is not just, “I'll play in the center square.'” It's more like, “I'll play in the center square when X does this, but I'll play in a corner square when X does that.” So a strategy is a set of actions you plan to take, contingent on what's happened so far. In a one-shot game (a game where you only play once, and you play simultaneously with all the other players, like in a game of rock-paper-scissors), a strategy is the same as an action because there is nothing to make your choice of action contingent on. For a while, all the games we look at will be one-shot games.
- Payoffs. The payoff that a player gets from a game is the expected utility he expects to get as result of the strategies he and all the other players have chosen. This is where the truly interdependent element comes in, because a player’s payoff depends not just on your own strategy, but everyone else’s as well.

II. Game #1: The Prisoners' Dilemma

This is probably the single most famous game in the history of game theory. The story goes like so: Two partners in crime are arrested by the authorities and placed in separate rooms. Each one is visited by the district attorney. The DA says that if both criminals remain silent ("mum"), he'll only have enough evidence to convict them of minor offenses, jailing them for one year each. If one of them rats on the other ("fink") while the other guy stays silent, the rat will get off scot-free while the other guy goes to jail for 15 years. Finally, if both of them rat on each other, they will both be convicted and go to jail for 10 years each.

Let's define the elements of the game: N = 2, strategies are Mum and Fink, and payoffs are as described in the following matrix. Player 1 chooses the row, and player 2 chooses the column. The first payoff in each cell is player 1's, and the second payoff is player 2's.

	Mum	Fink
Mum	**-1, -1**	**-15, 0**
Fink	**0, -15**	**-10, -10**

By looking at this picture, we can see that for player 1, F is the best response to M (0 > -1), and F is also the best response to F (-10 > -15). The same is true of player 2. So we predict that the outcome of this situation will be that both players fink. This is an example of a dominant strategy equilibrium.

Definition: A dominant strategy is a strategy that is a best response to any strategies chosen by the remaining players. In the Prisoners' Dilemma, note that Fink is a dominant strategy for each player.

Definition: A dominant strategy equilibrium is a set of strategies, one for each player, such that every player's strategy is a dominant strategy. Note that this is a very strong concept of equilibrium. It means that nobody would ever want to change their decision, regardless of what anyone else does.

The key feature of the Prisoners' Dilemma, and what makes it so intriguing, is that there is a dominant strategy equilibrium whose payoffs are clearly worse (for both players) than the payoffs that would have resulted from different actions by the players. Thus, it is a situation in which individual rationality leads to an outcome that is socially irrational (from any player's perspective). In other words, the equilibrium is not Pareto efficient. There are numerous examples of this type of phenomenon, which we'll talk about later.

> Application: Duopoly Cartelization. Each firm can choose either a high price or a low price. They would like to become a cartel by agreeing to set high prices. The payoffs are profits as given below (again, player 1 chooses the row, player 2 chooses the column, the first payoff in each cell is player 1's, and the second payoff is player 2's).

	High Price	Low Price
High Price	**100, 100**	**0, 150**
Low Price	**150, 0**	**10, 10**

Each firm has a dominant strategy, Low, and so there is a dominant strategy equilibrium (Low, Low). So we expect the firms' attempt at cartelization to fail, because each firm has an incentive to cheat on the agreement. This outcome is borne out by many historical observations, such as OPEC (in which member nations started producing more oil for the world market than their quotas allowed) and railroad companies in the 19th century (where railroads repeatedly made explicit agreements to hold up price, and then turned around and cut under-the-table deals with their customers).

We'll talk about several other examples in later lectures.

III. Game #2: Drug Lord and Mule

Consider a modified version of the prisoners' dilemma, in which the authorities are much more interested in capturing one criminal (the drug lord) than another (the mule who carries the drugs for him). The drug lord has two strategies, Plea and Don't Plea. The mule has two strategies, Testify or Don't Testify against his boss. The authorities offer the mule a deal in which he goes free for agreeing to testify (say, by putting his story in writing), regardless of what the drug lord does. Otherwise, the mule will be jailed for 2 years. The drug lord will only be convicted in trial if the mule testifies. If he takes a plea, he'll be jailed for 5 years; if he's convicted at trial, he'll be jailed for 15. The payoff matrix is like so:

	Plea	Don't Plea
Testify	**0, -5**	**0, -15**
Don't Testify	**-2, -5**	**-2, 0**

Notice that the mule has a dominant strategy, Testify. But the drug lord does not have a dominant strategy, so we can't have a dominant strategy equilibrium. We need a different equilibrium concept. We will use the following:

Definition: A Nash Equilibrium (NE) is a set of strategies, one for each player, such that each player's strategy is a best response to the equilibrium strategies of the remaining players. Note that this is a weaker concept than dominant strategy equilibrium, which required that each player's strategy had to be a best response to <u>any</u> strategies of the other players. NE only requires that a strategy be a best response to the equilibrium strategies of the other players.

In the Drug Lord & Mule game, the drug lord can predict that the mule will Testify. So the drug lord compares -5 and -15, and hence decides to Plea. So we have a NE of

(Testify, Plea). No other pair of strategies is a NE, because at least one player would wish to change his strategy, given what the other player is doing.

Application: Offensive and Defensive Advertising. Suppose there is dominant company in the soft drink industry, Coke, that will get the whole market (worth $100) if no one advertises. He has a competitor, Joe Bob Cola, who will only get market share if he advertises. By advertising, Coke can get some (but not all) of the market back. Advertising costs $10. The payoffs (in profits) are given below, with Coke picking rows and Joe Bob picking columns:

	Ad	No Ad
Ad	**70, 10**	**90, 0**
No Ad	**60, 30**	**100, 0**

Joe Bob has a dominant strategy (Ad), but Coke does not, so we don't have a DSE. But we do have a NE: (Ad, Ad). Given that Joe Bob advertises, it makes sense for Coke to advertise.

IV. Game #3: Battle of the Sexes

Two lovers intend to meet after work and attend an event. They can either see the Opera or the Fight. Pat prefers the Fight to the Opera, whereas Terry prefers the Opera to the Fight. But both prefer to be together at either event than alone at either one. In the matrix below, Terry chooses rows and Pat chooses columns:

	Opera	Fight
Opera	**3, 1**	**0, 0**
Fight	**0, 0**	**1, 3**

No one has a dominant strategy here. But there are two NE: (Opera, Opera) and (Fight, Fight). This is a case of multiple equilibria. Which one will occur? At least with the minimal structure we've given here, we can't make a prediction. Which equilibrium occurs will presumably depend on the beliefs of the players, their relative stubbornness, what they've done in the past, and so forth.

Note: This example also shows that you can have a NE without even one dominant strategy.

Application: Market Segmentation. There are two peanut butter makers, Jif and Skippy, who can choose to produce Crunchy peanut butter or Creamy peanut butter. Some consumers like Crunchy, some like Creamy, but more like Crunchy than Creamy.

	Crunchy	Creamy
Crunchy	**3, 3**	**8, 5**
Creamy	**5, 8**	**2, 2**

There are two NE in this game: (Crunchy, Creamy) and (Creamy, Crunchy).

The example above shows one type of a <u>first-mover advantage</u>. A first-mover advantage exists when one player can improve his position by being able to make the first move. In this case, if firm 1 could commit itself in advance to Crunchy, say by buying crunchy-specific capital or having a crunchy advertising campaign, then firm 2 will have to pick creamy as its best response. In a sense, moving first allows a firm to choose which equilibrium will occur.

Application: <u>Sides of the Road</u>. There are two drivers who will approach each other from opposite directions. Each can decide on whether to drive Left or Right. Both sides are equally good, but there's a wreck if the players don't drive on the same side. There are two NE: (Left, Left) and (Right, Right). Turns out this is true even if one side of the road really is better than the other. This is what's known as a <u>coordination game</u>: there are two equilibria, but one is clearly better for both players than the other. The challenge for the players is to coordinate their behavior so they end up in the superior equilibrium.

V. The Cournot Model of Duopoly

This is our first serious attempt to model oligopoly using game theory. In this structure (outlined by the French economist Auguste Cournot), each firm chooses a quantity, and the total quantity between the two firms results in a market price according to the following demand curve: $P = 130 - Q$, where Q is the sum of both firms' quantities, q_1 and q_2. Each firm faces a constant marginal cost of 10. What will be the firms' NE quantities?

Suppose that firm 1 knows that the other player will pick quantity q_2. Then the demand curve that remains is $P = (130 - q_2) - q_1$. Now we can have the firm maximize profits in the usual fashion. $MR = (130 - q_2) - 2q_1$. Setting this equal to MC = 10,

$$(130 - q_2) - 2q_1 = 10$$
$$120 - q_2 = 2q_1$$
$$q^R_1 = 60 - q_2/2$$

We put the "R" on this to show that it's a best-response to the quantity chosen by firm 2, whatever that quantity might be. But we haven't determined it yet. Now we go to firm 2, and using the same process we find that 2's best-response function is

$$q^R_2 = 60 - q_1/2$$

Now we can just solve the system of equations. Subbing in with firm 2's quantity, we get

$$q_1 = 60 - (60 - q_1/2)/2 = 30 - q_1/4$$
$$(3/4)q_1 = 30$$

$$q^*_1 = 40$$

and by symmetry of the problem, $q^*_2 = 40$.

What we've found here is a Nash Equilibrium. When both of the best-response functions are satisfied, it means each firm is setting a quantity that is a best-response to the other's quantity; that's the definition of Nash Equilibrium. The implied market price is P = 130 - 40 - 40 = 50.

What's interesting about this result is that it's intermediate between the perfectly competitive outcome and the monopoly outcome. Under PC, the market price would be equal to the marginal cost, 10. Under monopoly, the one firm would set Q = 60 and the price would be 70. The Cournot price is in between.

The two firms would like to be able to act like a cartel and set a higher price (the monopoly price). The problem is that if the other guy is setting the monopoly price, your best response is to set a <u>low</u> price. You'll get lots of customers while the other guy loses them. Of course, he has the same incentive if you set the monopoly price. So the result is that cartelization is not an equilibrium of this game.

It turns out that if we do the same thing with 3 firms, the price will be closer to the marginal cost. If we keep on increasing the number of firms, the price gets closer and closer to the marginal cost. So we can see perfect competition as a special case of Cournot competition, with a very large number of firms.

We've assumed here that both firms face the same marginal cost. Obviously, if the marginal costs differ, each firm sets its MR equal to its <u>own</u> MC. The resulting Nash equilibrium quantities will not be identical.

VI. The Bertrand Model of Duopoly

In the Cournot model, we had firms choose their quantities and then let demand conditions set the price. But what if it works the other way around? What if firms choose their prices, and then let the demand conditions determine how much they sell? If we assume that firms compete by choosing price, not quantity, then we have the Bertrand model.

Suppose once again that the demand curve is P = 130 - Q. But this time, if the firms set different prices, consumers only buy from the firm that sets the lower price. If they set the same price, then Q is split between the two firms. We continue to assume that the marginal cost is constant at 10.

What happens in this model? Suppose that firm 1 sets a price of, say, 20. Then what price should firm 2 set? If it sets a price of $19.99, it can get the whole market and make a profit. But then firm 1 could set its price at $19.98 and get the whole market back. Continue this process of price-cutting until it stops. But where does it stop? As long as firm 1 has a price above 10, firm 2 can set a price between firm 1's price and 10. So the process won't stop until P = 10. Or, more generally, the process continues until price is equal to marginal cost.

Note that this is a somewhat counterintuitive result. It predicts that two competing firms will produce the perfectly competitive outcome, in that P = MC. This contrasts with the more intuitive Cournot result, which predicts a price between the competitive price and the monopoly price. This is an example of the indeterminacy of the game theoretic approach to modeling oligopoly: the predictions are very sensitive to the initial assumptions.

There are some interesting variants of the Bertrand model that produce different results. First, what happens if the two firms have different marginal costs? In that case, the firm with the lower marginal cost will set its price just below the other firm's marginal cost, so that the other firm will choose not to produce. The higher-cost firm's marginal cost is an upper bound on the price that will occur on the market. Second, what happens if firms have to choose production capacities (plant sizes) before they set their prices? I won't prove this result, but it turns out that this model will produce results identical to the Cournot model. The production capacities they choose will be the Cournot quantities, and they will set prices that cause them to produce at capacity. Third, what if the firms do not produce perfectly homogeneous goods? If their products are differentiated, then not all consumers will choose the firm with the lower price regardless of how small the difference is. Instead, the size of the price difference will determine how many consumers are lost to the other firm. The result of allowing differentiated products in the Bertrand model is, again, to produce results more similar to those of the Cournot model: prices are not as high as they would be for a monopoly, nor as low as they would be under perfect competition.

VII. Game #4: The Mad Bomber

All of the previous games have been one-stage games – games in which each player makes a one-time choice. But now we want to consider multi-stage games, meaning games in which actions take place in a sequence of turns.

Suppose a man comes up to you on the street with a bomb in his hand. He demands that you give him your wallet, or else he'll drop the bomb and kill both of you. What do we predict as the outcome of this game?

[Draw tree-game picture. The first choice is made by you: either pay or not pay. If you don't pay, there is a choice to be made by the bomber: either bomb or don't bomb. If you pay, the outcome is -100 for you and +100 for the bomber. If you don't pay and he drops the bomb, both of you get negative infinity. If you don't pay and he doesn't drop the bomb, the outcome is 0, 0.]

To predict the outcome of the game, we're going to use a process called backward induction. This means we look at the last decision or decisions to be made in the game, and figure out what the decision-maker would do. Then we'll go backwards to the previous choice in the game. The decision-maker there will predict the outcome of her actions based on what will happen next.

In this example, the bomber has the final choice in the game. If you have refused to pay, he gets to choose between 0 and negative infinity. Naturally, he chooses not to drop the bomb. In other words, his promise to drop the bomb if you didn't pay was a non-credible threat. Now we go back to your choice of whether to pay or not. You can predict that efven if you don't pay, the bomber will not bomb. So your choice is between -100 and 0, and you choose 0 (not paying).

What we have found here is called a subgame perfect Nash Equilibrium. This is the equilibrium concept we use in multi-stage games. It is the same as Nash Equilibrium, with the added requirement that we must also have a Nash Equilibrium in every "subgame." A subgame is the remainder of a game that is left over after some decisions have already been made, and everyone knows what's already happened. In the Mad Bomber game, there was one subgame: the game in which you've already refused to pay.

> Application: The Market Entry Game. There is a market in which there is a single firm, which we call the "incumbent" or the monopolist. There is also another firm, a potential entrant into the market. If the incumbent is alone in the market, he will make monopoly profits. If both are in the market, they will have Cournot competition (as described earlier). To enter the market, the entrant must make a one-time investment of $1400. The incumbent makes a threat: if the entrant decides to enter, the incumbent will viciously increase his quantity and lower the price, causing the entrant to make profits that are too small to cover the $1400 entry cost. Does this threat succeed?
>
> Probably not. Say that monopoly profits are $3600. (This is the profit a monopolist would make with the demand curve and marginal cost used in the Cournot game above, assuming fixed cost is zero.) If they are both in the market, and they play the Cournot quantities, they'll make $1600 each. (This is the profit duopolists would make with the demand curve and marginal cost used in the Cournot game above, again assuming fixed cost is zero.) From this, the entrant must subtract his entry cost, so his anticipated profit is actually $200. If the incumbent carries out his threat to set a higher quantity, the incumbent will make profits of $700, and the entrant will make profits of -$1000 once the entry cost is taken into account. [Draw a tree-game picture, with payoffs of (0, 3600) after no entrance, (200, 1600) after entrance with no fight, and (-1000, 700) after entrance with a fight. In each case, the first payoff is the entrant's.] It turns out that the incumbent's promise to fight is not credible, because he's better of with 1600 than 700. Therefore, the entrant will predict that the incumbent won't fight, and he will therefore choose to enter. To put this another way, subgame perfection requires a NE in the after-entrance subgame. We already know that the only NE of the Cournot game is the Cournot quantities found earlier, and the entrant can predict that as well as we can. So the outcome is that the entrant enters, and then they play like Cournot competitors.

As in our previous examples, this outcome is dependent on the set-up of the game. It's possible to come up with alternate versions of the game which can allow the incumbent to deter entrance. But these games are substantially more complex. One example is allowing the incumbent to make a commitment prior to the entrant's choice. He could invest in additional production capacity, for instance, that would lower his marginal cost of production. This would reduce the firm's monopoly profits (because of the costliness of installing the additional production capacity), but it may also keep out the entrant.

A similar multi-stage game is the Predatory Pricing Game. The story goes like this: one firm threatens to drop its price so low that the other firm will make losses. It will keep doing this until the other firm goes bankrupt and quits the market. Then the remaining firm will raise its price to the monopoly level. This is what many critics said was done by Standard Oil, but that conclusion is questionable on both theoretical and evidentiary grounds. In theory, it turns out that this is typically a non-credible threat, because the losses the threatening firm must absorb are very large, and the firm would be better off pursuing a different strategy like buying the other firm. In practice, it turns out there's no good evidence that Standard Oil used predatory pricing. In most cases, Standard Oil just had lower production costs than other oil companies because of its superior distribution network and other innovations. In some cases where that was not true, Standard Oil bought up the opposition. In any case, there is no indication that Standard Oil raised its price after a competitor went out of business. In many ways, it was like the Bertrand game when one firm has a marginal cost of production smaller than the other's.

VIII. Repeated Games

In a one-period analysis of duopoly, we've concluded that the firms cannot successfully collude. In the Prisoners' Dilemma version of the problem, for example, picking the low price is a dominant strategy, even though both firms would be better off if they picked the high price.

Do these results change if the firms play the game repeatedly? Could the firms manage to cooperate by using the threat of punishment in future plays of the game?

Suppose the Prisoners' Dilemma version of the duopoly problem is repeated a finite number of times -- say, 10 times. It turns out that the firms still cannot manage to collude successfully, because doing so involves making non-credible promises or threats that are ruled out by subgame perfection. Consider the very last (10^{th}) play of the game. In that game, the outcome will be exactly the same as in the one-shot game, because there is no longer any future to use for enforcing threats or promises. Both players can predict this outcome of the 10^{th} play. Now consider the 9^{th} play of the game. Here, as well, the outcome will be just as in the one-shot game. The players could make threats or promises about behavior in the 10^{th} play -- but the threats and promises are pointless, because the result of the 10^{th} play is a forgone conclusion. Now consider the 8^{th} play of the game… All the remaining games fall down like a line of dominoes. Subgame perfection tells us that in a finitely repeated Prisoners' Dilemma, the outcome will be the same as in the one-shot game for every play of the game.

What happens if the game is repeatedly an infinite number of times? It turns out that in this case, collusion can sometimes be sustained as a Nash Equilibrium. Because there is always a future, it's always possible -- for any play of the game -- to make threats and promises based on future plays.

One method of sustaining collusion in an infinitely repeated Prisoners' Dilemma is to use a trigger strategy. A trigger strategy says to cooperate until the other player cheats, and then cheat for a long period of time if the other player cheats. An extreme trigger strategy would say to cooperate until the other player cheats, and then cheat forever after if the other player cheats.

If both players adopt a trigger strategy, and if the players care enough about the future relative to the present, then an equilibrium with collusion can be sustained. Suppose a player is thinking of cheating. If he cheats, he'll get a one-time large profit (from having the low price while the other player has the high price for one play of the game), followed by a stream of small profits (from both players setting the low price for the rest of time). If he continues to cooperate, he'll get a stream of medium profits (from both players setting the high price forever). Which outcome is better? The answer depends on how much the firms discount the streams of future profits. If the players aren't too impatient, continuing to cooperate is preferable.

Is it reasonable to think that a game will go on forever? No. But we can still use this approach if there's a certain probability that the game will end after any given play of the game. This will cause the firms to discount future streams of income more, because there is a lower chance of getting them. But for sufficiently far-sighted firms, collusion is still a possible equilibrium.

> Example: OPEC. In the 1970s, OPEC succeeded for a while in maintaining a cartel. But eventually the cartel collapsed, in part because of cheating by member nations, in part because of competition from oil producers in the North Sea and Southwestern U.S.
>
> Example: Railroads in the 19th century. They repeatedly attempted to set high shipping rates in order to escape the cut-throat competition that was severely limiting their profits. But these agreements rarely lasted more than a few weeks, because the railroads couldn't resist the incentive to cut special deals with large shippers.

In general, the historical record shows that collusive agreements are possible but (unfortunately for the firms, fortunately for the consumers) extremely fragile and difficult to maintain. One of the main reasons is that it’s difficult for a cartel member to tell whether a drop in its sales is attributable to temporary demand fluctuations or other members cheating. Thus, a temporary drop in sales because of consumer choices can be mistaken for cheating, leading to a round of punishment under the trigger strategy agreement.