

Lecture Notes
on
Game Theory

(Revised: July 2018)

These lecture notes extend some of the basic ideas in game theory that were covered in 15.010. We will begin by explaining what we mean by rational — or *rationalizable* — strategies. We will apply this concept to a discussion of the War of Attrition, which can turn out to be a rather complicated game. We will see what it means to approach and play (or decide to not play) a War of Attrition type game rationally. Then we will turn to models of duopolistic competition. We will first consider the choice of strategic variable when duopolists compete in the sale of differentiated products — in particular, what are the implications of choosing quantities instead of prices. Next we will examine some of the issues that arise when there is asymmetric or imperfect information. In particular, we will see what happens when a firm has limited information about the costs of its competitor, or when both firms have limited information about market demand. Finally, we will discuss bargaining situations, and the Nash cooperative solution to a bargaining game.

1. Rationalizable Strategies

Our focus throughout these notes is with rational decision-making. Sometimes what we mean by rational decision-making is obvious — for example, setting a price or level of output that maximizes profits. But sometimes it is not so obvious, particularly in gaming situations. Thus we will find it useful to introduce the concept of a *rationalizable strategy*, a concept that is in fact rather simple.

In 15.010 you encountered the concepts of a dominant strategy and a Nash strategy. But in many games there is no dominant strategy, and an equilibrium in Nash strategies (i.e., a Nash equilibrium) might not exist. There are times when we need something more general:

- **Dominant Strategy:** I'm doing the best I can no matter what you are doing.
- **Nash Strategy:** I'm doing the best I can given what you are doing.
- **Bayesian Nash Strategy:** I'm doing the best I can in expected value terms, accounting for the uncertainties that affect what you do and that affect outcomes.
- **Rationalizable Strategy:** I'm doing the best I can given my expectations regarding your likely behavior, what I think are your beliefs about me and my likely behavior, and whether your beliefs and likely behavior are themselves rationalizable.

This might seem a bit abstract, but it will become clear as we proceed through some examples. Also, note that “doing the best I can” might mean deciding whether to play or not to play a particular game. This is very much the case when the game happens to be a war of attrition, which we turn to next.

2. The War of Attrition

Wars of attrition often arise in business (and in other settings as well). The game arises when two (or more) firms compete with each other, each one losing money but hoping that the competitor will eventually give up and exit the industry. When playing the game, each firm must decide whether to cut its losses and exit, or alternatively, tough it out in the hope that the competitor will soon exit.

An example of this game is the competition that took place in the U.K. in the late 1980s in the satellite television market. The competing firms were British Satellite Broadcasting (BSB), a consortium, and Sky Television, which was part of Rupert Murdoch's news corporation. Through October 1990, the two firms accumulated losses in excess of £1 billion as they fought for control of the satellite broadcasting business. The war ended in November

1990, when BSB and Sky announced that they would merge into a single firm, BSkyB, with control split evenly among the shareholders of the original entities.

Another example of the War of Attrition is the building cascades that sometimes occur in new shopping malls or other urban settings. Each firm buys land or other property rights and starts construction, knowing that several other firms are doing the same thing, and that all of the firms will lose money unless some of them drop out. Sometimes some of the firms do indeed drop out, but often there is over-building, and all of the firms end up with large losses. You have likely seen shopping malls where such over-building has occurred. And you can probably think of other examples of the War of Attrition.

Many people view a War of Attrition as something to stay far away from. Management consultants and other practitioners of corporate strategy often promote the view that a War of Attrition almost always ends badly, and thus should be avoided. But if this is indeed how most people feel, then a War of Attrition can be very attractive. If that seems counterintuitive, read on.

A Simple Example. To understand the War of Attrition, let's consider the following simple example. Suppose two companies, *A* and *B*, must decide each month whether to spend \$10 million. If in the first month one company spends the \$10 million and the other does not, the game is over: the first company becomes a monopolist worth \$100 million, and the second company looks for something else to do. If neither company invests \$10 million in the first month, the game is likewise over, with neither company losing or making money. However, if *both* companies spend \$10 million in the first month, neither one wins anything. We then move to the second month, where again each company must decide whether to spend \$10 million. If both companies again spend \$10 million, we move to the third month, and so on. If, at the start of some month, one of the companies spends \$10 million and the other does not, the first company wins the \$100 million prize. But of course many months (and much money) could go by before this happens.

Suppose you are Company *A*, and one of your classmates is Company *B*. Collusion is not permitted. All you can do is decide whether (and how) to play the game. What should you do in this situation? Think carefully about the following questions:

1. Is it rational to spend \$10 million in the first round of this game? Why or why not? How would you decide whether or not to spend the \$10 million?
2. Would it be rational to start playing the game with a plan to exit if, after three or four rounds, your opponent has not yet exited? Why or why not?
3. Is there anything that you could say to your opponent, or that your opponent could say to you, that would affect your decision to start playing the game? Is there anything you or your opponent could say that would affect your decision to continue playing if, after two or three rounds, neither of you has dropped out?

Again, think *carefully* about these questions. If the answers seem obvious, think harder. And then think about the following: If it seems obvious that you shouldn't play this game — that no rational person would play the game — then you should probably play the game. Do you understand why?

3. The Repeated Prisoners' Dilemma

A simple example of the Prisoners' Dilemma is given by the payoff matrix shown below. Obviously both firms would do best if they *both* set a high price. But setting a low price is a *dominant strategy* — it yields the highest payoff no matter what your opponent does. Assuming you are going to play this game once and only once, that your objective is to *maximize your payoff*, and collusion is not allowed, what would you do?

		Firm 2	
		High Price	Low Price
Firm 1	High Price	100, 100	0, 200
	Low Price	200, 0	50, 50

Most students say they will set a low price, and indeed, that is the only rationalizable strategy (that I can think of). Some students, however, respond in the following ways:

- Some version of “I want to be nice.” For example, “We will both be better off if we set high prices, so why should I take advantage of my competitor by setting a low price?” That’s a different game. In *this* game your *only* objective is to maximize *your* payoff, and the only way to do that is set a low price.
- Some version of “Yes, this is a one-shot game, but I may meet that same competitor again in another game.” Once again, you are trying to play a different game. Your objective in *this game* is to maximize your payoff *now*, not the payoff from some other game that you might or might not play in the future.

There is an important lesson here: *Always be clear about the game you are playing, including your objective.* (Think about playing tennis against an 8-year old versus playing against someone your own age.) This seems so simple, and yet people often lose sight of the game and their objective.

Play Three Times. Now suppose you are going to play this game three times, and with the same competitor. Each time you play, you and your competitor will choose a high price or a low price, and will receive a payoff given by the payoff matrix. Your objective is to *maximize the sum of your payoffs over the three rounds.* That’s it — that’s your only objective. We are starting with Round 1. What will you do?

Once again, most students say they will set a low price, and indeed, that strategy is easy to rationalize. The reason is that the game *unravels*. Think about the third (and last) round. No matter what happened in the first two rounds, setting a low price in the third round is a dominant strategy. Thus you know that both you and your competitor will set a low price in the third round. Now back up to the second round. Given that you will surely both set a low price in the third round, it is a dominant strategy to also set a low price in the second round. Thus you know that you and your competitor will set a low price in the second round. And now what about the first round? Once again, given that you will both set low prices in rounds 2 and 3, setting a low price in round 1 is also a dominant strategy.

Is there any way to rationalize setting a high price in round 1? Perhaps. You might have reason to believe that your competitor has been asleep for most of the semester, and hasn’t

gone to the trouble of thinking this thing through. Then you might set a high price in round 1 with the expectation you will shift to a low price in rounds 2 and 3. But note that this is risky. You better check that your competitor has indeed been asleep most of the time.

Play Ten Times. Now suppose you are going to play this game ten times, and with the same competitor. In each round, you and your competitor will choose a high price or a low price, and will receive a payoff given by the payoff matrix. Your objective is to *maximize the sum of your payoffs over the ten rounds*. We are starting with Round 1. What will you do?

Now things get a bit trickier. If you and your competitor think this through, you will both see that the unraveling again occurs: You will both set a low price in round 10, thus you will both set a low price in round 9, ..., thus it is a dominant strategy to set a low price in round 1. So setting a low price in round 1 is clearly rationalizable.

Is there any way to rationalize setting a high price in round 1? Perhaps. Wouldn't it be nice if you and your competitor didn't think too hard about the unraveling? Of course, you already have thought of it, and in all likelihood your competitor has too. But suppose you think your competitor is using a "don't worry, be happy" strategy of ignoring the unraveling. If so, setting a high price in the hope that both of you will maintain a high price for at least several rounds would be rationalizable. Likewise for your competitor, so that a high price in round 1 may indeed occur.

Note that it is not necessary that you and your competitor actually ignore the unraveling. All that matters is that you believe your competitor *will act as though she is ignoring it* — perhaps because she believes you are ignoring it — and that your competitor believes the same about you. Maybe both of you feel the game is best played after having a couple of glasses of wine. The point is that "don't worry, be happy" can be a self-fulfilling expectation that can lead to high prices for at least several rounds.

Retail Store Pricing. In the previous example I chose 10 rounds rather than 9 or 11 because it matches the number of rounds in the Strategic Oligopoly Game that you play every week. So here's a question: is a 10-round repeated Prisoners' Dilemma a "realistic" game, i.e., the kind of game you might encounter in the real world? Or is it just a mathematical exercise that might have some educational value, but has little to do with the real world?

It's extremely realistic. Just think about retail stores, and the pricing decisions they face each year in November and December. Whatever doesn't get sold by December 24 will get sharply marked down, often below cost. Each week (or day), competing stores must set and announce prices. All the players know that as December 24 approaches, the game will unravel and prices will be cut. So who will undercut first? "Cooperate" until Thanksgiving, or drop prices and undercut your competitors early, before they undercut you?

Gaming Against Consumers. But consumers make the problem for retail stores even worse. Customers *know* that prices will fall, if not before Thanksgiving, then surely shortly after Thanksgiving, so they have an incentive to wait, and shop after prices fall. As a result, retail stores must also play a game against their customers.

What might prevent consumers from waiting, and refusing to buy before the sales begin and prices fall? The perception that stores will run out of the "hot" items they want to buy, such as popular toys that kids want, the newest electronic devices, etc. Thus stores might try to create the perception that the "hot" items are indeed "hot," and will be in short supply, so that consumers don't wait for prices to fall. Some of the larger retail chains (such as Wal-Mart) will reduce the quantities they order so that some of the "hot" items indeed do run out. The idea is that consumers will remember this, and then next year they will buy early before prices are cut.

4. Nash Equilibrium in Prices Versus Quantities

Recall the example of price competition with differentiated products from Chapter 12 of Pindyck & Rubinfeld, *Microeconomics*. Two duopolists have fixed costs of \$20, but zero variable costs, and they face the same demand curves:

$$Q_1 = 12 - 2P_1 + P_2 \tag{1a}$$

$$Q_2 = 12 - 2P_2 + P_1 \tag{1b}$$

Note that the cross-price elasticities are positive, i.e., the two products are substitutes. Also note that once the two duopolists choose their prices P_1 and P_2 , the quantities Q_1 and Q_2 that they sell are determined from equations (1a) and (1b). If instead they were to choose

the quantities Q_1 and Q_2 that they produce, then the equilibrium market prices P_1 and P_2 would likewise be determined by these equations.

Should it make any difference whether these duopolists choose prices (so that the quantities they sell are determined by the market) or whether they choose quantities (so that their prices for their products are determined by the market)? It might seem as though it makes no difference, but as we will see, it makes a big difference.

Choosing Prices. In 15.010, you examined the Nash equilibrium that results when the two firms set their *prices* at the same time. It will help to begin by briefly summarizing the derivation of that equilibrium. For Firm 1, profit is:

$$\pi_1 = P_1Q_1 - 20 = 12P_1 - 2P_1^2 + P_1P_2 - 20$$

Taking P_2 as fixed, Firm 1's profit-maximizing price is given by:

$$\frac{\Delta\pi_1}{\Delta P_1} = 12 - 4P_1 + P_2 = 0,$$

so Firm 1's reaction curve is given by:

$$P_1^* = 3 + \frac{1}{4}P_2 \tag{2a}$$

Likewise for Firm 2:

$$P_2^* = 3 + \frac{1}{4}P_1 \tag{2b}$$

Solving for the Nash equilibrium, $P_1 = P_2 = \$4$, so $Q_1 = Q_2 = 8$, and $\pi_1 = \pi_2 = \$12$. Note that the *collusive outcome* is $P_1 = P_2 = \$6$, so that $Q_1 = Q_2 = 6$, and $\pi_1 = \pi_2 = \$16$. The reaction curves, Nash equilibrium, and collusive outcome are shown in Figure 1, which in Pindyck & Rubinfeld, 9th Edition, is Figure 12.6.

Choosing Quantities. Now suppose the firms choose *quantities* instead of prices. Everything is the same as before, except that now each firm chooses its quantity, taking its competitor's quantity as fixed. To find the Nash (Cournot) equilibrium, we must first rewrite the demand curves (1a) and (1b) so that of each price is a function of the two quantities. Eqns. (1a) and (1b) can be rearranged as:

$$P_1 = 6 - \frac{1}{2}Q_1 + \frac{1}{2}P_2$$




Figure 1: Nash Equilibrium in Prices




Figure 2: Nash Equilibrium in Quantities

$$P_2 = 6 - \frac{1}{2}Q_2 + \frac{1}{2}P_1$$

Combining these two equations and rearranging the terms,

$$P_1 = 12 - \frac{2}{3}Q_1 - \frac{1}{3}Q_2 \quad (3a)$$

$$P_2 = 12 - \frac{2}{3}Q_2 - \frac{1}{3}Q_1 \quad (3b)$$

Note that eqns. (3a) and (3b) represent the same demand curves as eqns. (1a) and (1b). We have simply rearranged the equations so that price is on the left side and the quantities are on the right side. Using eqn. (3a), the profit for Firm 1 can be written as:

$$\begin{aligned} \pi_1 &= P_1Q_1 - 20 \\ &= 12Q_1 - \frac{2}{3}Q_1^2 - \frac{1}{3}Q_1Q_2 - 20 \end{aligned}$$

Maximize this profit with respect to Q_1 , taking Q_2 as fixed:

$$\frac{\Delta\pi_1}{\Delta Q_1} = 12 - \frac{4}{3}Q_1 - \frac{1}{3}Q_2 = 0$$

$$Q_1^* = 9 - \frac{1}{4}Q_2 \quad (4a)$$

Likewise,

$$Q_2^* = 9 - \frac{1}{4}Q_1 \quad (4b)$$

The reaction curves (4a) and (4b) can be combined to find the Nash equilibrium: $Q_1 = Q_2 = 7\frac{1}{5}$, so that $P_1 = P_2 = 4\frac{4}{5}$, and $\pi_1 = \pi_2 = 14.56$. The reaction curves and Nash equilibrium are shown in Figure 2.

Observe that compared to the Nash equilibrium with price as the strategic variable, *both firms now make higher profits*. All we have done is change the strategic variable from price to quantity, and yet the outcome is quite different.

Now try to answer the following questions: (1) *Why* do the firms make higher profits when they choose quantities instead of prices? (2) Should the two firms “agree” to choose

quantities rather than prices? (3) How does this relate to the problem you face each week in the Strategy Game?

Asymmetric Choice of Strategic Variable. We have considered a situation in which both firms choose prices and compared it to the situation in which both firms choose quantities. Suppose, instead, that one firm chooses price and the other chooses quantity as the strategic variable. In particular, suppose that *Firm 1 chooses price* but *Firm 2 chooses quantity*. What will happen in this case?

Firm 1 takes P_2 as fixed, and thus has the reaction function:

$$P_1^* = 3 + \frac{1}{4}P_2.$$

Firm 2 takes Q_1 as fixed, and thus has the reaction function:

$$Q_2^* = 9 - \frac{1}{4}Q_1.$$

From Eqn. (1a):

$$\begin{aligned} Q_1 &= 12 - 2\left(3 + \frac{1}{4}P_2\right) + P_2 \\ &= 6 + \frac{1}{2}P_2 \end{aligned}$$

Likewise, from Eqn. (3b):

$$\begin{aligned} P_2 &= 12 - \frac{2}{3}\left(9 - \frac{1}{4}Q_1\right) - \frac{1}{3}Q_1 \\ &= 6 - \frac{1}{6}Q_1 \end{aligned}$$

We can combine these two equations to solve for Q_1 and P_2 . Doing so, we find that $Q_1 = 8.31$ and $P_2 = \$4.62$. Now, use the reaction functions, to find P and Q_2 :

$$P_1 = 3 + \frac{1}{4}(4.62) = \$4.16$$

$$Q_2 = 9 - \frac{1}{4}(8.31) = 6.92$$

We now know each firm's price and quantity, and thus can calculate that the profits for the two firms are given by $\pi_1 = \$14.57$ and $\pi_2 = \$11.97$. We see that Firm 1 does better than Firm 2, and it makes approximately the same profit that it did when both firms used

quantities as their strategic variables. Firm 2, however, does worse — slightly worse than it did when both firms chose prices as their strategic variables.

Suppose both firms are free to choose between price and quantity as the strategic variables. What outcome would you expect? What does this tell you about pricing and output decisions in the airline industry? The automobile industry? The Strategy Game you play every week?

5. Incomplete Information — Bayesian Nash Equilibrium

In the real world, firms rarely have complete information about demand, their competitors' costs, or even their own costs. We now turn to the problems that arise when a firm has limited information about its competitors. To do this, we will extend the simple example of Cournot equilibrium that you studied earlier. The basic idea is that each player is doing the best he/she can in expected value terms, given the expected behavior of the competitors. Let's see how this works.

5.1 Cost Uncertainty.

Two firms produce a homogenous good, and face the following market demand curve:

$$P = 30 - Q$$

The firms' marginal costs are c_1 and c_2 . Each firm chooses its quantity, taking the quantity of its competitor as given.

For Firm 1, revenue is

$$\begin{aligned} R_1 = PQ_1 &= (30 - Q_1 - Q_2)Q_1 \\ &= 30Q_1 - Q_1^2 - Q_1Q_2 \end{aligned}$$

So its marginal revenue is

$$RM_1 = 30 - 2Q_1 - Q_2$$

Setting $RM_1 = c_1$ gives the reaction curve for Firm 1:

$$Q_1^* = 15 - \frac{1}{2}Q_2 - \frac{1}{2}c_1 \tag{5a}$$

Likewise for Firm 2:

$$Q_2^* = 15 - \frac{1}{2}Q_1 - \frac{1}{2}c_2 \quad (5b)$$

1. Note that if $c_1 = c_2 = 0$, we get $Q_1 = Q_2 = 10$, $P = \$10$ and $\pi_1 = \pi_2 = \$100$, a result you might recall from 15.010.
2. Suppose $c_1 = 0$, but $c_2 = 6$. Then:

$$Q_1^* = 15 - \frac{1}{2}Q_2$$

$$Q_2^* = 12 - \frac{1}{2}Q_1$$

You can check that in this case, $Q_1 = 12$, $Q_2 = 6$, $P = \$12$, $\pi_1 = \$144$, and $\pi_2 = \$36$. Firm 2 has a higher marginal cost than Firm 1, and thus produces less and makes a smaller profit.

3. Now suppose that $c_1 = 0$ and both firms know this. However, c_2 is either 0 or 6. Firm 2 can observe its own cost and thus *knows* what c_2 is, but Firm 1 doesn't. Firm 1 therefore assigns a probability of $\frac{1}{2}$ to each possibility. What is the equilibrium in this case? We will assume that each firm maximizes its *expected profit* — the result is a *Bayesian Nash equilibrium* (BNE).

Start with Firm 2. If $c_2 = 0$, Firm 2 will have the reaction curve

$$Q_2^*(0) = 15 - \frac{1}{2}Q_1 \quad (6a)$$

If instead $c_2 = 6$, Firm 2 will have the reaction curve

$$Q_2^*(6) = 12 - \frac{1}{2}Q_1 \quad (6b)$$

What is Firm 1's reaction curve? The answer depends on Firm 1's objective. We will assume that Firm 1 maximizes its *expected profit*. Firm 1 does not know Firm 2's reaction curve because it does not know c_2 . There is a probability of $\frac{1}{2}$ that Firm 2's cost is zero so that its reaction curve is $Q_2^*(0)$, and there is a probability of $\frac{1}{2}$ that it is 6 so that Firm 2's reaction curve is $Q_2^*(6)$. Thus, Firm 1's expected profit is:

$$\begin{aligned}
E(\pi_1) &= \frac{1}{2}[30 - Q_1 - Q_2^*(0)]Q_1 + \frac{1}{2}[30 - Q_1 - Q_2^*(6)]Q_1 \\
&= 30Q_1 - Q_1^2 - \frac{1}{2}Q_2^*(0)Q_1 - \frac{1}{2}Q_2^*(6)Q_1
\end{aligned}$$

To maximize this expected profit, differentiate with respect to Q_1 , *holding each possible Q_2^* fixed*, and set the derivative equal to zero:

$$30 - 2Q_1 - \frac{1}{2}Q_2^*(0) - \frac{1}{2}Q_2^*(6) = 0$$

or,

$$Q_1^* = 15 - \frac{1}{4}Q_2^*(0) - \frac{1}{4}Q_2^*(6) \quad (6c)$$

To find the equilibrium, solve (6a), (6b), and (6c) for Q_1 , $Q_2(0)$, and $Q_2(6)$:

$$Q_1 = 11, \quad Q_2(0) = 9\frac{1}{2}, \quad Q_2(6) = 6\frac{1}{2}$$

Compare this result to the case (1) where $c_2 = 0$ and *both firms know it*, and case (2) where $c_2 = 6$ and *both firms know it*. Note that Firm 2 does *better* (by having superior information) if $c_2 = 6$, but it does *worse* if $c_2 = 0$. Does this seem surprising? Think about the following:

- When $c_2 = 0$, Firm 2 produces *less* when only it knows its cost than it does when Firm 1 *also* knows that $c_2 = 0$. Why is this? And why does Firm 2 produce *more* when $c_2 = 6$ and only it knows this than it does when its cost is common knowledge?
- Suppose $c_2 = 0$ and Firm 1 does not know this. Can Firm 2 do better by *announcing its cost to Firm 1*? Should Firm 1 believe Firm 2? What would you do if you were Firm 2? If you were Firm 1?

5.2 Demand Uncertainty.

We have already seen that having better information can sometimes make a firm better off, and sometimes make it worse off. The example above focused on uncertainty over one of

the firm's cost, but there could just as well be uncertainty over demand. Once again, more information may or may not make firms better off. By now, you should be able to understand this intuitively. To make sure you do, think through the following problem, which appeared a recent 15.013 Final Exam:

Artic Cat and Yamaha Motors compete in the market for snowmobiles. Each company is concerned about the extent of cross-brand substitutability (i.e., the extent to which consumers would choose one brand over the other in response to a small price difference). Neither firm knows the extent of substitutability, and each firm therefore operates under the assumption that the brands are moderately substitutable. In fact, the brands are *highly* substitutable, but *only we know this* — not the firms. The firms compete by setting prices at the same time.

(a) Suppose that both firms conduct statistical demand studies and learn the truth, i.e., that the brands are highly substitutable. Would this knowledge make the firms better off, i.e., lead to higher profits? Explain briefly.

(b) Suppose that the only Artic Cat conducts a study and learns that the brands are highly substitutable. Should it announce this finding to Yamaha? Explain briefly.

Are the answers obvious to you? In (a), if both firms learn that the brands are highly substitutable, they will both set lower prices. (Their reaction curves will shift because each firm gains more by undercutting its competitor.) Thus, both firms will be worse off. In (b), Artic Cat should *not* announce the findings of its study to Yamaha. Artic Cat will lower its price, to the surprise of Yamaha, and earn greater profits. Of course these greater profits may not last long, as Yamaha eventually figures out what is going on.

This example is fairly simple, and quite limited in its scope. But it makes the point once again that having more information can make firms worse off. In the next section, we will examine the implications of imperfect or asymmetric information about demand in more detail.

6. The Strategic Oligopoly Game

In the Strategic Oligopoly Game, you and your competitors all have the same information. You know how much you have produced and how much they have produced, so you can

deduce your own and their marginal costs. You also know how much inventory they hold. As for demand, the “shocks” to the market demand curve are announced in advance each week, and thus are known to you and your competitors. What you *don't know* is what your competitors are going to do, and how they will react to what you do. You and your competitors are in a repeated Prisoners' Dilemma, but that Dilemma is complicated by the weekly demand shocks and by the ability to hold inventory.

6.1 Responding to Demand Shocks.

The weekly demand shocks are announced and known to all of you, but how should you respond to them? Suppose you and your competitors could collude. Then, given the demand shock, you could figure out the price and individual output levels that maximizes total profit (taking into account the learning curve), and agree to set that price and output. But you can't collude. Of course you and your competitors could independently figure out the collusive (i.e., joint profit-maximizing) price and output, and each of you set that price and output with the hope that the others will do the same. But how do you know that what you have figured out is the same as what they have figured out?

To make this clear, let's take an example. Suppose the in the first two weeks the demand shocks were zero, and you and your competitors all set the collusive (i.e., “don't worry, be happy” joint profit-maximizing) price and output levels. Very nice. In Week 3, however, there is a large negative demand shock, say -400 . What price and output level should you set now?

Suppose you work through the math and conclude that the joint profit-maximizing price should be \$7.50 less than it was in Weeks 1 and 2, and the individual joint-profit maximizing output levels for each firm should be 40 less than before. Given this calculation, should you go ahead and reduce your price and output by \$7.50 and 40, respectively? The problem is that maybe your competitors didn't do the same calculation you did. Maybe they are not as smart as you. Maybe they did the calculation and got a different answer, perhaps a smaller price reduction. Now what will they think when they see you have lowered your price by \$7.50? Might they think that you are trying to undercut them, and respond with an even

larger price reduction in Week 4?

You might decide the way to play it safe is to reduce your output level by even more than 40, but leave your price unchanged. This way your competitors can't possibly think you are trying to undercut them. But suppose your two competitors do figure out that the joint profit-maximizing price is \$7.50 less than it was before, and go ahead and reduce their prices by that amount. In that case *you* will be sharply undercut, and sell much less than you produce.

So what should you do? Reduce your price by only half the amount you calculated is optimal? Unfortunately there is no clear answer. But there *is* a lesson to be learned from this. Hopefully you now understand why in real-world oligopolies, firms are reluctant to change their prices in response to demand fluctuations. Prices are *sticky* — they don't move the way you might expect them to during booms and busts.

6.2 The Use of Inventory.

As you play the game, it is likely that at least in some weeks, you will accumulate inventory. This might be the result of bad planning; you simply produced more than you thought you could sell. But it also might be intentional; perhaps you *purposely* produced more than you thought you could sell in order to build up inventory. Why do this? What is the value of inventory, and what can you do with it?

Cost Reduction. One reason to accumulate inventory is to reduce production cost over time. This is relevant when marginal cost is increasing with the rate of production (as it is in the Strategic Oligopoly Game). To see this, suppose there is a positive demand shock, which increases the amount you can sell, and thus the amount you would like to produce. But if you produce more, your average and marginal costs will go up. One way around this problem is to produce more than you think you can sell when there is a negative demand shock, and then sell the resulting inventory later when there is a positive demand shock. By “smoothing” production this way, firms can use inventory to reduce production costs over time.

Note, however, that this has nothing to do with the learning curve. One does *not* ac-




Figure 3: Over-Producing

accumulate inventory as a way of producing more in order to move down the learning curve faster. This is illustrated in Figure 3, where MC_0 is your initial marginal cost curve, and MC_{10} is your “final” marginal cost curve, i.e., where you think marginal cost will be in Week 9 or 10. The discount rate is zero, so MC_{10} is the relevant marginal cost curve, and suppose that using that marginal cost curve, Q^* is the optimal output level, i.e., the output level that will maximize cumulative profits. If you produce Q^* , you will move down the learning curve, but marginal cost will fall slowly. What if you try to move down the learning curve faster by producing Q_1 (with the intention of accumulating an amount of inventory equal to $Q_1 - Q^*$). The problem is that to do this, you will incur a very high marginal cost, as shown in the figure, and your cumulative profit will be lower than if you produced Q^* .

Inventory as a Signal. Did you ever hear the expression “Speak softly but carry a big stick?” Inventory can be that big stick. It can act as a threat against possible undercutting

by competitors.

To see how this could work, suppose that for the past three or four weeks, you and your competitors have been setting prices at close to the joint profit-maximizing level. You are worried, however, that the game could unravel sooner than you expect, as one or both of your competitors start to undercut you. You would like to threaten your competitors by signaling to them that if they do undercut you, your response will hurt them badly.

How might you respond to a price cut by a competitor? You could sharply reduce your own price, and produce and sell a large quantity. The problem is that producing a large quantity will push you up your marginal cost curve, and thus is unlikely to be profitable for you. But suppose you have accumulated a substantial amount of inventory. Then you could sharply reduce your price and sell most or all of your inventory. This is likely to be profitable for you and thus is a credible threat.

Of course you can't talk to your competitors, but there is no need to say anything. They know how much inventory you have accumulated, so the threat is implicit.

To make sure you understand the use of inventory, look at the following questions, which appeared on a recent 15.013 Final Exam:

When you played the Strategic Oligopoly Game every week you might have accumulated inventories, perhaps intentionally and perhaps not.

(a) How might inventories be useful as a way of reducing costs over the 10 weeks of the game? Explain clearly.

(b) How might inventories be useful as a way of sending a strategic signal to your competitors? Suppose you accumulated a large amount of inventory during the first few weeks of the game. What signal might that convey?

A number of students were confused by the first question, and gave answers along the lines of "By producing more and accumulating inventories, we can move down the learning curve faster." No. Over-producing in that way will increase costs (and reduce profits) over the long run, as explained above.

7. Price Competition With Asymmetric Information

Note: This section is optional. Consider a situation in which two firms compete by setting prices. For simplicity, we will take the demand curves to be linear:

$$Q_1 = a_{10} - a_{11}P_1 + a_{12}P_2 \quad (7a)$$

$$Q_2 = a_{20} + a_{21}P_1 - a_{22}P_2 \quad (7b)$$

(Later we will see how with information about elasticities and equilibrium prices and quantities, the six parameters in the above equations can be estimated.) For the time being, suppose that each firm knows its demand curve and its competitor's demand curve, and that prices are chosen simultaneously. It is then easy to compute the Nash equilibrium for this pricing problem. If we are dealing with the short run and variable costs are relatively small, we can focus on revenues. The revenue for Firm 1 is given by:

$$R_1 = P_1Q_1 = a_{10}P_1 - a_{11}P_1^2 + a_{12}P_1P_2$$

Maximizing this with respect to P_1 gives:

$$dR_1/dP_1 = a_{10} - 2a_{11}P_1 + a_{12}P_2 = 0$$

Hence the reaction function for Firm 1 (i.e., its price as a function of Firm 2's price) is:

$$P_1^*(P_2) = \frac{a_{10}}{2a_{11}} + \frac{a_{12}}{2a_{11}}P_2 \quad (8a)$$

Likewise for Firm 2:

$$P_2^*(P_1) = \frac{a_{20}}{2a_{22}} + \frac{a_{21}}{2a_{22}}P_1 \quad (8b)$$

Since the firms are setting prices simultaneously, we can solve these two equations for P_1 and P_2 . Defining $\Delta \equiv 4a_{11}a_{22} - a_{12}a_{21}$, the solution for prices will be:

$$P_1 = (2a_{10}a_{22} + a_{12}a_{20})/\Delta \quad (9a)$$

$$P_2 = (2a_{20}a_{11} + a_{21}a_{10})/\Delta \quad (9b)$$

It will be helpful at this point to introduce a numerical example. Suppose that $a_{10} = a_{20} = 12$, $a_{11} = a_{22} = 2$, and $a_{12} = a_{21} = 1$. Then $\Delta = 15$, and $P_1 = P_2 = \$4$, $Q_1 = Q_2 = 8$,

and $R_1 = R_2 = \$32$. Also note that because the demands are symmetric and information is symmetric, each firm will have a 50 percent market share.

Incomplete Information. Now suppose that each firm knows its own demand curve, but does not know exactly how price-sensitive its competitor's demand is. In particular, suppose that Firm 1 does not know the value of a_{22} , and Firm 2 does not know the value of a_{11} . Each firm instead relies on an estimate of this parameter of its competitor's demand. Suppose that the *true* parameters are:

$$a_{11}^* = a_{11} + \epsilon_1$$

$$a_{22}^* = a_{22} + \epsilon_2$$

Firm 1 knows ϵ_1 , but not ϵ_2 ; Firm 2 knows ϵ_2 , but not ϵ_1 . The expected value of ϵ_2 (for Firm 1) is 0, and likewise for ϵ_1 .

The reaction functions are of the same form as before, except that now Firm i cannot predict its competitor's reaction function $P_j^*(P_i)$. In other words, it does not know exactly what price its competitor will charge, even as a function of its own price. The reaction functions are now:

$$P_1^* = \frac{a_{10} + a_{12}P_2}{2(a_{11} + \epsilon_1)} \quad (10a)$$

$$P_2^* = \frac{a_{20} + a_{21}P_1}{2(a_{22} + \epsilon_2)} \quad (10b)$$

But, once again, Firm 1 is uncertain as to what P_2 will be, and Firm 2 is uncertain as to what P_1 will be.

We now have a strategic pricing problem in which there is *incomplete information*, but no asymmetry of information, because each firm is equally in the dark about its competitor. A natural solution is for each firm i to assume that $\epsilon_j = 0$, and to assume that Firm j thinks that $\epsilon_i = 0$. In effect, Firm 1 assumes that Firm 2's price will be given by Eqn. (9b), and Firm 2 assumes that Firm 1's price will be given by Eqn. (9a). Substituting Eqn. (9b) into Eqn. (10a), and (9a) into (10b), we get the following solution for the firms' prices:

$$P_1 = \frac{a_{10}\Delta + a_{12}(2a_{20}a_{11} + a_{21}a_{10})}{2(a_{11} + \epsilon_1)\Delta} \quad (11a)$$

$$P_2 = \frac{a_{20}\Delta + a_{21}(2a_{10}a_{22} + a_{12}a_{20})}{2(a_{22} + \epsilon_2)\Delta} \quad (11b)$$

Note that this is *not* a Bayesian Nash equilibrium. In a BNE, each firm sets its price to maximize its expected profit, using a probability distribution for its competitor's ϵ_j . We have simplified matters by having each firm assume that $\epsilon_j = 0$.

As an example, suppose that a_{10} , a_{20} , etc., have the same values as before, and that $\epsilon_1 = \epsilon_2 = 2$. (Thus each firm will underestimate its competitor's demand elasticity and overestimate its competitor's price.) Plugging in the numbers, we find that in this case $P_1 = P_2 = \$2$, $Q_1 = Q_2 = 6$, and $R_1 = R_2 = \$12$. For comparison, if each firm *knew* that $\epsilon = 2$ for its competitor, the prices would be given by Eqns. (9a) and (9b), but with $a_{11} = a_{22} = 4$. In this case, $P_1 = P_2 = \$1.71$. Likewise, $Q_1 = Q_2 = 6.87$, and $R_1 = R_2 = \$11.75$. Thus the firms do better as a result of this (symmetric) lack of information. The reason is that it leads each firm to *overestimate* its competitor's price, and thereby induces the firm to set a higher price than it would otherwise. Thus each firm is "misled," but in a direction that helps both firms.

Asymmetric Information. Now consider what happens when the information is *asymmetric*. Suppose, once again, that $\epsilon_1 = \epsilon_2 = 2$, but that this time Firm 1 *knows* that $\epsilon_2 = 2$, and hence knows that Firm 2 is going to charge \$2, and *not* \$1.71. What should Firm 1 do in this case? It should set price according to Eqn. (10a), with $\epsilon_1 = 2$ and $P_2 = \$2$. Plugging in the numbers, we can see that in this case Firm 1 will charge a price $P_1 = \$1.75$. Thus, it will undercut its competitor. Then the quantities sold will be:

$$Q_1 = 12 - 4(1.75) + 2 = 7$$

$$Q_2 = 12 - 4(2) + 1.75 = 5.75,$$

and the revenues will be $R_1 = \$12.25$ and $R_2 = \$11.50$. Clearly this informational asymmetry gives Firm 1 an advantage. It obtains a larger market share than its competitor (even though the demand curves are completely symmetric), and it earns more revenue than its competitor.

Now suppose that $\epsilon_1 = \epsilon_2 = -1$. In this case, each firm *overestimates* its competitor's elasticity of demand, and hence *underestimates* the price that its competitor will charge.

Using eqns. (11a) and (11b) as before, we find that $P_1 = P_2 = \$8$, $Q_1 = Q_2 = 12$, and $R_1 = R_2 = \$96$. (The negative value for ϵ_1 and ϵ_2 means that demand is much less elastic, so both firms can end up charging much higher prices and earning higher revenues.)

For comparison, if each firm *knew* that $\epsilon = -1$ for its competitor, the prices would be given by eqns. (9a) and (9b), but now with $a_{11} = a_{22} = 1$. In this case, $P_1 = P_2 = \$12$, $Q_1 = Q_2 = 12$, and $R_1 = R_2 = \$144$. Thus in this case the firms do worse when they have a (symmetric) lack of information. Again, the reason is that it leads each firm to underestimate its competitor's price, and therefore set a lower price than it would otherwise.

As before, let us again consider what happens when the information is asymmetric. Suppose that $\epsilon_1 = \epsilon_2 = -1$, but that this time Firm 1 knows that $\epsilon_2 = -1$ and hence knows that Firm 2 is going to charge \$8. Then, Firm 1 will price according to Eqn. (4a), with $\epsilon_1 = -1$ and $P_2 = \$8$. In this case we can see that Firm 1 will charge a price $P_1 = \$10$, i.e., it will price *above* its competitor's price. Then the quantity sold will be:

$$Q_1 = 12 - 1(10) + 1(8) = 10$$

$$Q_2 = 12 - 1(8) + 1(10) = 14,$$

and the revenues will be $R_1 = \$100$ and $R_2 = \$112$. Now both firms do better than when they both lacked information (recall that then they both made revenues of \$96), but Firm 2 does better than Firm 1, even though Firm 1 has more information. Why? Because the lack of information leads Firm 2 to underestimate its competitor's price, and thus set its own price at a level below that which it would otherwise. Firm 1 knows that Firm 2 will set this low price, and the best it can do in this situation is to set a somewhat higher price.

We have worked out this example for linear demand curves, but we could just as well have worked it out with, say, isoelastic demand curves (although the algebra would be a bit messier). Sticking with these linear curves, there are six parameters that must be determined if we wanted to fit this model to data. Those parameters are a_{10} , a_{20} , a_{11} , a_{22} , a_{12} , and a_{21} . We can obtain all six parameters if we have estimates of the market elasticity of demand, the elasticity of demand for each firm, and the equilibrium prices and quantities. Suppose that the elasticity of *market* demand is -1 . (This means that if P_1 and P_2 rise by 1 percent, Q_1

and Q_2 will fall by 1 percent.) This gives *two conditions*. Next, suppose that the own-price elasticity of demand for each firm is -3 . (This means that if P_1 rises by 1 percent and P_2 remains fixed, Q_1 will fall by 3 percent, and likewise for a 1-percent rise in P_2 .) This also provides *two conditions*. Finally, the equilibrium values of P_1 and Q_1 provide a condition, and the equilibrium values of P_2 and Q_2 provide a condition. Thus we can imagine building a simple spreadsheet model in which one inputs the elasticities and the equilibrium values of the prices and quantities, and the various parameter values a_{ij} are automatically calculated.

We could likewise calculate a range for ϵ . For example, it might be reasonable to think that the actual market demand elasticity lies somewhere between -0.6 and -1.4 , with an expected value of -1.0 . Assuming symmetry, this implies a range for ϵ_1 and ϵ_2 .

It is important to point out once again that the equilibria that we have calculated here are *not* Bayesian Nash equilibria. To obtain a Bayesian Nash equilibrium, we would want to find the reaction function for Firm 1 corresponding to every possible value of Firm 2's ϵ_2 , and likewise find a reaction function for Firm 2 corresponding to every possible value of Firm 1's ϵ_1 . We would then calculate the expected revenue for each firm as a function of the expected value of its competitor's reaction functions. We would then pick a price to maximize this expected revenue. If ϵ_1 and ϵ_2 have simple distributions (e.g., uniform), this would not be very difficult to do. Nonetheless, the equilibria that we have calculated above are much simpler, and are based on a simpler assumption — each firm takes the expected value of its competitor's ϵ_j , and finds an optimal price accordingly.

8. Nash Cooperative (Bargaining) Solution

The Nash bargaining solution -completely different from the Nash non-cooperative equilibrium you studied in 15.010- is an important concept that can help us understand the kinds of outcomes that can result from bargaining by rational players. To see how it works, consider a situation in which two individuals are trying to reach an agreement. For example, they might be bargaining over the division of a sum of money. We will assume that Player 1 gets utility u from the agreement, and Player 2 gets utility v . If there is *no agreement*, they get utilities u_0 and v_0 , respectively. This is called the *threat point*. It might be a Nash




Figure 4: Nash Cooperative Solution

noncooperative equilibrium, or a maximin equilibrium, or it might be that both players get nothing if there is no agreement, in which case $u_0 = v_0 = 0$.

John Nash demonstrated that there is a unique solution to this bargaining problem that satisfies certain axioms that one would reasonably think should hold when rational people are engaged in a bargaining situation. (The axioms are individual rationality, feasibility of the outcome, Pareto optimality, independence of irrelative alternatives, symmetry, and independence with respect to linear transformations of the set of payoffs.) Furthermore, the solution that Nash arrived at is quite simple; it maximizes the following function of the two players' utilities:

$$g(u, v) = (u - u_0)(v - v_0)$$

This is illustrated in Figure 4. Note that the Nash solution maximizes the area of the shaded rectangle.

An example will help to illustrate this. Suppose two individuals are trying to divide \$100. If they fail to reach an agreement, neither individual will receive any money. Player 1 is very poor, and starts with a utility of 0. Player 2, however, is rich; he has a fortune worth $F \gg \$100$. How will they divide the \$100?

Let x be the amount that Player 1 gets, so $100 - x$ is the amount that Player 2 gets. We will assume that both players have the same utility function: the logarithm of their total wealth. Hence the utilities for the two players are as follows:

$$\text{Player 1: } u = \log x, \quad u_0 = 0$$

$$\text{Player 2: } v = \log(F + 100 - x), \quad v_0 = \log F$$

Given these utilities, the function $g(u, v)$ is given by:

$$\begin{aligned} g(u, v) &= (\log x)[\log(F + 100 - x) - \log F] \\ &= (\log x) \log \left(\frac{F + 100 - x}{F} \right) \end{aligned}$$

Since F is large, $(100 - x)/F$ is small, and therefore,

$$\log \left(1 + \frac{100 - x}{F} \right) \approx \frac{100 - x}{F}$$

Hence the function $g(u, v)$ can be approximated as:

$$g(u, v) \approx (\log x) \left(\frac{100 - x}{F} \right)$$

The Nash bargaining solution is the value of x that maximizes this function. Differentiating with respect to x and setting the derivative equal to 0 gives:

$$\frac{dg}{dx} = \frac{1}{x} \cdot \frac{100 - x}{F} - \frac{1}{F} \log x = 0$$

or

$$x \log x + x - 100 = 0$$

Solving for x gives $x^* = \$24$. Hence Player 1 would get \$24, and Player 2 would get \$76.

Note that the wealthier individual walks away with the larger share of the pie. Do you see why this is? Do you expect this kind of outcome to occur in practice? Why or why not?