

Appendix 3.

The Classic Models of Duopoly

Perhaps the first use of formal game-theoretic ideas in economics, which came well before the formal development of those ideas, came in the work of Cournot, Bertrand, and von Stackelberg, concerning equilibria in the (static) competition between two oligopolists.

Since Chapter 3 is about rivalry—and so were these models—this is the appropriate place to discuss them. That said, this appendix enlists a whole bunch of concepts that are introduced only in Parts II, III, and IV of the textbook, and if you find yourself wondering “what is that?,” you might want to set this appendix aside until you have gone deeper into the text.

The story begins with Problem 2.7, which I reproduce here for convenience: Imagine two firms, labeled A and B , producing products that are substitutes but not perfect substitutes. The inverse demand functions for their two goods are

$$p_A = a - x_A - bx_B, \quad \text{and} \quad p_B = a - x_B - bx_A$$

for a positive constant a and for $0 < b < 1$, with the same constants a and b appearing in both inverse demand functions. Each firm has a constant marginal cost of production c .

(a) Suppose the two firms must simultaneously and independently choose quantities to produce, each without knowing what the other has chosen. What are (or is) the Nash equilibria of this game?

(b) Suppose that the two firms must simultaneously and independently choose prices to charge, each without knowing what price the other has chosen. What are (or is) the Nash equilibria of this game?

(c) Suppose firm A can choose its production quantity first. Firm B sees this choice, then chooses its production quantity. What do you predict would happen?

(d) Suppose firm A can choose its price first. Firm B sees this choice, then B chooses its price. What do you predict would happen?

And here, in gory detail, is the solution to this problem:

Begin by noting that parts a and b are about simultaneous-move games and Nash equilibria. Parts c and d are about games where firm A moves first and firm B responds. So for parts c and d, we will use backward-induction techniques.

(a) In part a, the strategies for the players are their quantity choices, x_A and x_B . An equilibrium is a pair of quantities (x_A^*, x_B^*) such that each is a best response to the choice by the other firm. Suppose, for instance, that firm B chooses x_B . Firm A's profit, if it chooses x_A , is

$$(p_A - c)x_A = (a - x_A - bx_B - c)x_A = (a - bx_B - c)x_A - x_A^2,$$

which is maximized for a given x_B when the derivative in x_A is 0, or when

$$x_A = \frac{a - bx_B - c}{2}.$$

(To be a bit more precise, this is true when $a - bx_B - c \geq 0$. Firm A's optimal choice is $x_A = 0$ when $a - c \leq bx_B$.) By a symmetric argument, if firm A chooses x_A , firm B's best response is

$$x_B = \frac{a - bx_A - c}{2}.$$

So a Nash equilibrium is where these two equations hold simultaneously:

$$x_A^* = \frac{a - bx_B^* - c}{2} \quad \text{and} \quad x_B^* = \frac{a - bx_A^* - c}{2}.$$

This is two equations in two unknowns; the unique solution is

$$x_A^* = x_B^* = \frac{a - c}{2 + b}.$$

Figure A3.1 depicts the solution, with a graph of pairs (x_A, x_B) . For each value of x_A , we graph the *best-response function for firm B*, $x_B(x_A) = (a - bx_A - c)/2$ as a solid line (with the extension $x_B(x_A) = 0$ if $bx_A \geq a - c$), and we graph the *best-response function for firm A*, $x_A(x_B) = (a - bx_B - c)/2$

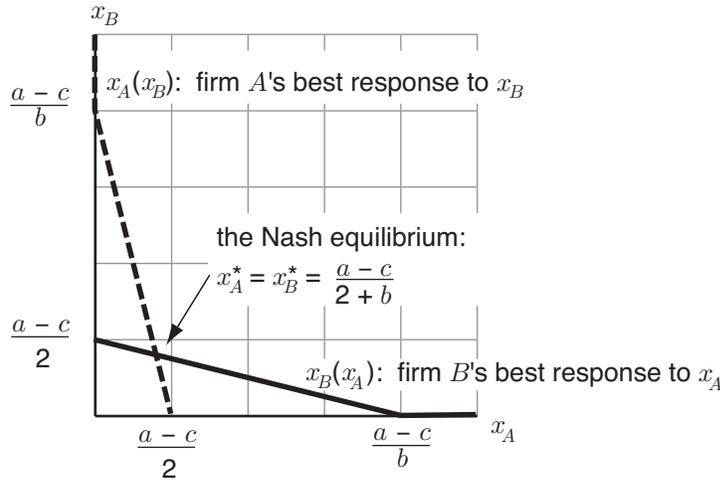


Figure A3.1. The solution to Problem 2.7(a) in a picture. For each level of x_A , the function $x_B(x_A)$, drawn with a solid line, gives B 's best response to A 's choice of quantity. For each level of x_B , the function $x_A(x_B)$, drawn with dashes, gives A 's best response to B 's choice of quantity. Where they intersect, at $x_A^* = x_B^* = (a - c)/(2 + b)$, is the unique Nash equilibrium, the solution to part a.

as a dashed line. Where the two intersect, where each firm chooses the best response to what the other firm has chosen, is the Nash equilibrium. (The graph is drawn for the values $a = 10$, $b = 0.5$, and $c = 2$.)

(b) To solve part b, we first must invert the two inverse demand functions to get the corresponding pair of demand functions. It takes a bit of algebra, but these are

$$x_A = \frac{a(1 - b) - p_A + bp_B}{1 - b^2} \quad \text{and} \quad x_B = \frac{a(1 - b) - p_B + bp_A}{1 - b^2}.$$

The strategies for the two firms are now their price choices, and a Nash equilibrium is a pair of prices (p_A^*, p_B^*) where each firm chooses the best response to the price choice of the other firm. As in part a, for each price choice p_B by firm B , we can find firm A 's best response, which I denote by $p_A(p_B)$: If firm A chooses p_A , and given firm B 's choice p_B , firm A 's profit is

$$(p_A - c) \times \left[\frac{a(1 - b) - p_A + bp_B}{1 - b^2} \right] = \frac{[a(b - 1) - bp_B]c + p_A[a(1 - b) + bp_B + c] - p_A^2}{1 - b^2}.$$

Set the derivative in p_A to 0, to get

$$p_A(p_B) = \frac{a(1 - b) + bp_B + c}{2}.$$

Similarly and symmetrically,

$$p_B(p_A) = \frac{a(1-b) + bp_A + c}{2}.$$

The Nash equilibrium occurs where each firm chooses the best response to each other, or where

$$p_A^* = \frac{a(1-b) + bp_B^* + c}{2} \quad \text{and} \quad p_B^* = \frac{a(1-b) + bp_A^* + c}{2},$$

the solution of which is

$$p_A^* = p_B^* = \frac{a(1-b) + c}{2-b}.$$

Once again, we can draw a picture of the equilibrium as in Figure A3.2. This time the graph is of pairs of prices. Firm A 's best response function $p_A(p_B)$ once again is the dashed line and firm B 's best response function $p_B(p_A)$ is the solid line. Where they intersect is the Nash equilibrium.

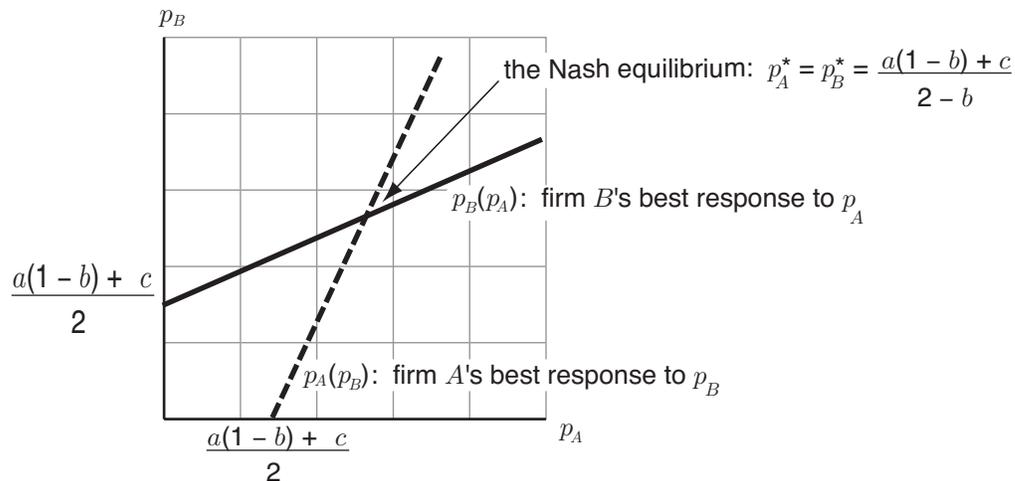


Figure A3.2. The solution to Problem 2.7(b) in a picture. For each level of p_A , the function $p_B(p_A)$, drawn with a solid line, gives B 's best response to A 's choice of price. For each level of p_B , the function $p_A(p_B)$, drawn with dashes, gives A 's best response to B 's choice of price. Where they intersect, at $p_A^* = p_B^* = [a(1-b) + c]/(2-b)$, is the unique Nash equilibrium, the solution to part b.

(c) If firm A chooses its quantity x_A , then firm B responds, a backward-induction analysis is in order. Firm B , faced with x_A , optimally responds with the choice $x_B(x_A)$ that we found in part a, namely $x_B(x_A) = (a - bx_A - c)/2$. Firm A understands that firm B will respond in this fashion, so it chooses its initial quantity to maximize its profit given this: If it chooses x_A , its profit will be

$$[a - x_A - bx_B(x_A) - c]x_A = \left[a - x_A - \frac{b(a - bx_A - c)}{2} - c \right] x_A =$$

$$\left[\left(a - c \right) \left(1 - \frac{b}{2} \right) \right] x_A - \left(1 - \frac{b^2}{2} \right) x_A^2.$$

This is maximized in x_A at

$$x_A^\# = \frac{(a - c)(1 - b/2)}{2 - b^2} = \frac{(a - c)(2 - b)}{2(2 - b^2)}.$$

You can work out the corresponding response by firm B by plugging $x_A^\#$ into the response function $x_B(x_A)$.

Although it is a bit confusing for some readers, let me draw some pictures that illustrate this backward-induction solution and compare it to the Nash equilibrium computed in part a. These pictures are all for the values $a = 10$, $b = 0.5$, and $c = 2$.

We begin with Figure A3.3(a) (overleaf). This figure shows firm A 's *iso-profit curves*, curves of the form

$$(a - x_A - bx_B - c)x_A = K$$

for various constants K . (The four iso-profit curves are for $K = 4, 6, 8$, and 10 . Firm A 's profit increases unambiguously as x_B decreases, so the curve the furthest "south" is for $K = 10$.)

Suppose firm B is committed to $x_B = 8$. What is the best response for firm A ? Looking across the horizontal line $x_B = 8$, we can see that the largest profit that A can achieve is where $x_A = 2$, where the iso-profit curve is tangent to the horizontal line, which just says that, for this set of parameter values, the general function $x_A(x_B) = (a - bx_B - c)/2$ has the specific value $x_A(8) = (10 - 0.5 \times 8 - 2)/2 = 2$. In this figure, and more generally, if you

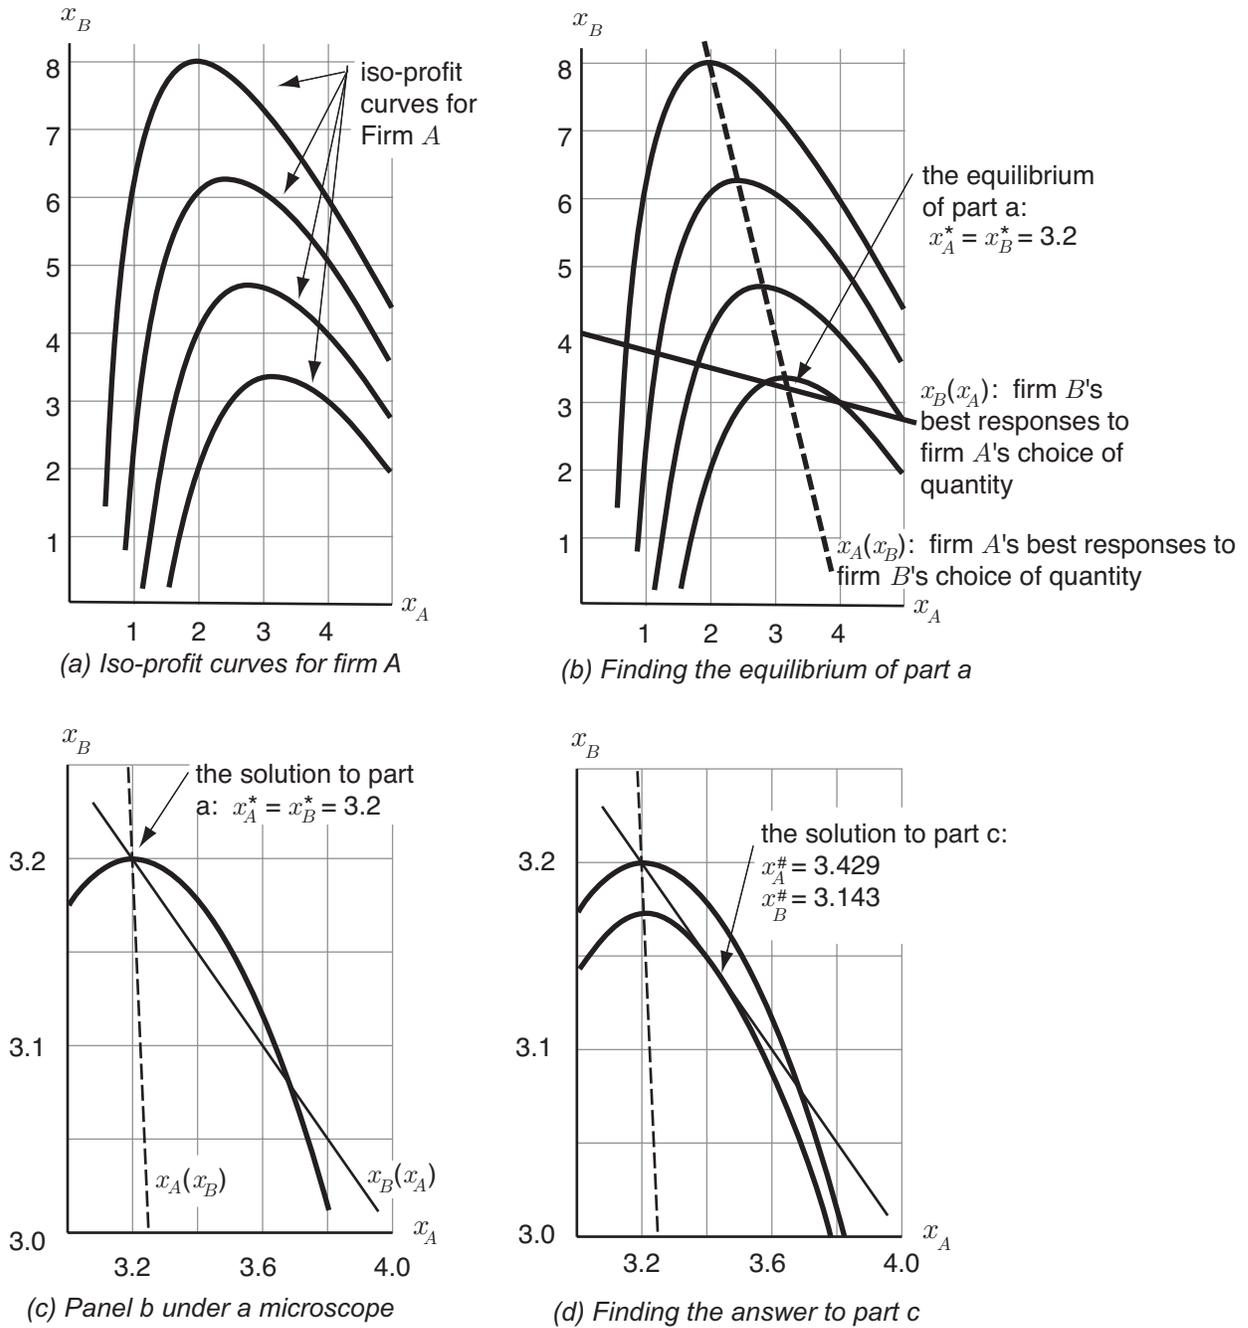


Figure A3.3. Comparing parts a and c. Panel a shows the iso-profit curves for firm A. In panel b, the (dashed-line) function $x_A(x_B)$ is constructed from the iso-profit curves and the (solid-line) function $x_B(x_A)$ is added. Where they intersect is the Nash equilibrium of part a. Panel c shows a close-up of the equilibrium, with the iso-profit curve of firm A at the equilibrium drawn in. If firm A moves first and firm B responds, firm B responds to x_A with $x_B(x_A)$, so firm A can choose any pair $(x_A, x_B(x_A))$ that it wants, and the choice that maximizes firm A's profit is shown in panel d.

connect the “tops” of the isoprofit curves, you get the full function $x_A(x_B)$, shown in panel b with the usual dashed line.

The answer to part a, the Nash equilibrium (x_A^*, x_B^*) , occurs where $x_A(x_B)$ intersects $x_B(x_A)$. For these parameter values, this occurs at $x_A^* = x_B^* = 3.2$. I supply $x_B(x_A)$ as a solid line as well in panel b.

In panel c, you get a close-up view of the Nash equilibrium, together with firm A’s iso-profit curve passing through the equilibrium. (For these parameter values, this profit level is 10.24.) Now we turn from part a of the problem, where the two firms choose quantities simultaneously, to part c, where firm A goes first. If firm A goes first, then firm B responds to x_A with $x_B(x_A)$. Firm A, recognizing this, knows that it can select whichever pair $(x_A, x_B(x_A))$ gives it the highest profit. It is apparent from panel c that firm A can increase its profit to some extent by increasing x_A from x_A^* , since the response of firm B to decrease x_B benefits firm A. How much should firm A increase x_A ? It should do so until it attains the highest possible profit level along the line $(x_A, x_B(x_A))$. In panel d of Figure A3.3, you see this level of x_A , which is $x_A^\#$, and the corresponding iso-profit curve for firm A, which for these parameter values is a profit level of 10.28571429.

Note that $x_B(x_A)$ is a decreasing function of x_A (until x_B hits and stays at 0). This means that *the more aggressive is firm A in its choice of quantity, the less aggressive is the optimal response of firm B*. From firm A’s perspective, less aggressive responses by firm B are always in its (firm A’s) favor. So, relative to the equilibrium we found in part a, in this part of the problem, firm A is going to be more aggressive, or $x_A^\# > x_A^*$. The picture makes this clear; to verify it algebraically, note first that (since $b > 0$) $4 - b^2 > 4 - 2b^2$, therefore $(2 - b)(2 + b)(a - c) > 2(2 - b^2)(a - c)$ and hence

$$x_A^\# = \frac{(a - c)(2 - b)}{2(2 - b^2)} > \frac{(a - c)}{(b + 2)} = x_A^*.$$

A final thing to note about this part of the problem is that, while we have produced a Nash equilibrium for the game where firm A chooses its quantity first and firm B responds, there are other Nash equilibria to this game. We know that this is a Nash equilibrium because we produced it via backward induction and a general result tells us that the result of backward induction always is a Nash equilibrium. But imagine that firm B issued the following threat.

If you, firm A, choose the quantity $x_A = (a - c)/10$, I will respond with

$9(a - c)/20$. If you choose any other quantity x_A , I plan to respond with $(a - c - x_A)/b$.

More to the point, suppose firm B adopts the strategy implicit in this threat. How should firm A respond? If firm A believes that B would carry out the threatened strategy, A can respond with $x_A = (a - c)/20$, which will result in a positive profit when B chooses $x_B = 9(a - c)/20$. On the other hand, if A chooses any other quantity, B 's response of $(a - c - x_A)/b$ means firm A 's profit is 0. So firm A 's best response is $(a - c)/20$. And, if A chooses $(a - c)/20$, firm B 's best response is $x_B((a - c)/20)$ which is, in fact, $9(a - c)/20$. Firm B 's threatened responses to all other choices by A are costless, because firm A chooses $(a - c)/20$. This is indeed a Nash equilibrium.

But, suppose firm A chooses some quantity other than $(a - c)/20$. In particular, suppose firm A chooses $x_A^\#$. Would firm B carry out its threat to produce $(a - c - x_A^\#)/b$? This is not firm B 's best response, given the *fait accompli* choice of $x_A^\#$ by firm A . This is indeed a Nash equilibrium but one based on incredible threats by firm B . As such, it is not a very credible prediction for this game.

(d) The answer to part d parallels the answer to part c, so I do this more quickly. Following a choice of p_A by firm A , firm B chooses $p_B(p_A) = [a(1 - b) + bp_A + c]/2$. Hence, if firm A chooses p_A , its profit is

$$(p_A - c) \left[\frac{a(1 - b) - p_A + bp_B(p_A)}{1 - b^2} \right] =$$

$$(p_A - c) \left[\frac{a(1 - b) - p_A + b[a(1 - b) + bp_A + c]/2}{1 - b^2} \right].$$

Since the denominator is a constant, we can drop it: Firm A chooses p_A to maximize

$$(p_A - c) \left(a(1 - b) - p_A + \frac{ab(1 - b)}{2} + \frac{b^2 p_A}{2} + \frac{bc}{2} \right).$$

Taking the derivative, setting it equal to 0, and denoting the solution by p_A^b gives

$$0 = a(1 - b) - p_A^b + \frac{ab(1 - b)}{2} + \frac{b^2 p_A^b}{2} + \frac{bc}{2} + c - p_A^b + \frac{b^2(p_A^b - c)}{2} \quad \text{or}$$

$$p_A^b(2 - b^2) = a(1 - b) + \frac{ab(1 - b)}{2} + \frac{bc}{2} + c - \frac{b^2c}{2},$$

which is

$$p_A^b = \frac{2a - ab - ab^2 + bc + 2c - b^2c}{2(2 - b^2)}.$$

The corresponding level of p_B is obtained by computing $p_B(p_A^b)$.

————— o —————

Problem 2.7 and its solution provide the mechanics of the classic models of oligopoly: the Cournot, Bertrand, and von Stackelberg models. For many years, those models constituted the economic theory of oligopoly (generalizing to more than two firms); for many economists today, those models continue to constitute the theory. Chapter 3 is supposed to be about the theory of oligopoly, yet the classic models are not mentioned in the text. So in the discussion to follow, I want to provide connections between Chapter 3 and these models.

This discussion has three parts. First, I provide “stories” that support the game structures of the classic models. Then, while exploring the appropriateness of Nash equilibrium analysis of these games (what happens in the solution to Problem 2.7), I explain why I have little faith a priori in these models. Finally, I indicate why, despite my lack of faith, if you continue to study economics relevant to management, you are quite likely to encounter these models.

The Cournot Model and Equilibrium

Imagine two firms, A and B , that sell their product in a marketplace some distance from where they produce. Think in terms of a marketplace in the 19th century. The marketplace is the central square in the local market town. One day each week, people with stuff to sell load their carts, travel to the marketplace, and set up to sell what they brought. Customers also come into town this one day each week to buy.

The critical assumption in this story is that the two producers of this particular good (call it cheese) must decide each week how much of their cheese to bring with them, to sell that day. Their quantity decisions are made simultaneously and are irrevocable (each week)—unsold cheese spoils—so if the first duopolist arrives with x_A units of cheese and the second with x_B units, the prices p_A and p_B that the two set are the prices that clear the market. Assume that the market for these two varieties of cheese is symmetric and that

market-clearing prices are given by the pair of inverse demand equations

$$p_A = a - x_A - bx_B \quad \text{and} \quad p_B = a - x_B - bx_A,$$

where a and b are constants, $a > 0$, $0 \leq b \leq 1$. If $b > 0$, the two types of cheese are substitutes: The more firm B brings to the market, the lower is the market-clearing price for firm A . If $b = 1$, the two are perfect substitutes, while if $b < 1$, the two types of cheese are only imperfect substitutes; still, if firm B brings enough of its cheese to the market, it can drive the price of firm A 's cheese to 0. (Its own price would have hit 0 at a lower quantity, so it is unlikely to do this.) The marginal cost of production for each firm is a constant c between 0 and a , and both firms have fixed costs of 0.

Assume this market reconvenes week after week and the two firms do not discount their weekly profits too heavily. Then, we have an ideal setting for folk-theorem-style collusion. Each firm brings to the market the same amount of cheese x^M that maximizes the joint profits of the two firms—I leave it to you to verify that this is $x^M = (a - c)/(2 + 2b)$, so the market-clearing prices are $p_A = p_B = (a + c)/2$ —where each firm sticks to this quantity out of fear of a price war if it tries to take short-run advantage by bringing more cheese to the market in any given week. Please note that this is the symmetric collusive “solution”; the folk theorem, in the usual way, provides us with many equilibria, some collusive and asymmetric and some not very collusive at all.

But, nothing guarantees that the two firms arrive at a collusive solution at all. The cheese makers might not trust each other. It might never occur to them that they can trade off short-run profits for larger long-run profits, if they act with mutual restraint. So, the French economist Augustin Cournot asked (in 1838), “What would happen in this market, if each firm attempts, each week, to maximize its own profit that week, given the actions of the other firm?”

Although Cournot wrote long before game theory was developed, in modern terminology, he found the Nash equilibrium of the one-time, one-week game, in which firms A and B simultaneously choose x_A and x_B , each seeking the maximal one-week's profit, given the choice of the other firm. From Problem 2.7(a), we know that the Nash equilibrium of this game, known in the literature as the Cournot equilibrium, is

$$x_A^* = x_B^* = \frac{a - c}{2 + b},$$

which means higher quantities, lower prices, and lower profits than if the two firm managed the symmetric collusive scheme. (In Problem 2.7, $0 < b < 1$ was assumed, but the analysis presented works perfectly well for $b = 0$ or 1.)

Two special cases are worthy of note. If $b = 0$, the two types of cheese are not substitutes at all; the market-clearing price of one depends only on how much of it was brought to market that week and not at all on how much of the other type is being sold. And, in this case, the Cournot-equilibrium quantities are the monopoly quantities. At the other end of the spectrum, if $b = 1$, the two types of cheese are perfect substitutes, and the Cournot-equilibrium prices are significantly lower than the collusive prices, although not so low that the two firms make profits of 0.

Cournot–von Stackelberg Equilibrium

Now change the story. Imagine that firm A can somehow commit to its quantity of cheese before firm B produces, and firm B sees how much cheese firm A has produced before choosing its own level of production. In game-theoretic terms, this is the game described and analyzed in Problem 2.7(c), and we know that, in this case, backward induction leads to the prediction that firm A chooses the quantity

$$x_A^\# = \frac{(a - c)(2 - b)}{2(2 - b^2)} > x_A^*$$

and firm B responds with a quantity smaller than x_B^* . In the borderline case $b = 0$, both firms choose the monopoly quantities. In the other borderline case $b = 1$ of perfect substitutes, firm A chooses its monopoly quantity, but firm B produces a positive amount, so that firm A does not make its monopoly profit. This is known in the literature as the *Von Stackelberg* or *Cournot–Von Stackelberg* equilibrium, after the German economist Heinrich Von Stackelberg, who suggested in 1938 that one firm might take a first-mover or leadership position of this sort.

Bertrand Equilibrium

Another change to the story is to assume, as in the original, that the two firms move simultaneously but choose the prices they charge rather than the quantities they sell. To make a parable out of this, imagine the firms advertise the prices they will charge on market day (Saturday) in the local newspaper on Friday. The two firms phone in the prices they will charge on Wednesday afternoon, see what prices their rival picked when the paper arrives on Friday, and bring to the market as much as they are able to sell,

given the pair of prices in the newspaper and the demand functions derived from the inverse demand functions given earlier, which are

$$x_A = \frac{a(1-b) - p_A + bp_B}{1-b^2} \quad \text{and} \quad x_B = \frac{a(1-b) - p_B + bp_A}{1-b^2}.$$

That price announcements are phoned in on Wednesdays is not quite consistent with this being an 19th century town-square market, but this is (just) a parable.

With this latest story, we look for a Nash equilibrium in the name-the-prices-simultaneously game, which is the Nash equilibrium computed in Problem 2.7(b),

$$p_A^* = p_B^* = \frac{a(1-b) + c}{2-b}.$$

Although it is not obvious until you do the algebra, these prices are lower than the prices one gets in the Cournot equilibrium, and profits are correspondingly diminished; price competition of this sort is “more competitive” than quantity competition. As for the two extreme cases, if $b = 0$, this is once again the monopoly outcome. But for $b = 1$, well, that story takes another paragraph at least.

For the case $b = 1$, where the goods are perfect substitutes, the demand functions are not well defined; there is a 0 in the denominator. But the equilibrium prices are quite well defined; substitute $b = 1$ into $[a(1-b)+c]/(2-b)$ and out pops c . The story that goes with $b = 1$, which gives equilibrium prices of c , runs as follows. If the two goods are perfect substitutes, then the firm that names a lower price than the other gets all the demand. If the two name the same price, we assume, they split the entire demand 50–50. In such a situation, the only possible simultaneous-price equilibrium is $p = c$, for the following reasons:

- It is not an equilibrium for the two firms to name different prices if the lower of the two prices is less than c , because the firm naming the lower price takes a loss, which it can turn into 0 profit by naming price c .
- It is not an equilibrium for the two firms to name different prices if the lower of the two prices is c , because the firm naming the lower price of c makes profit 0 but would make a strictly positive profit by naming a price halfway between c and the price its rival names.

- It is not an equilibrium for the two firms to name different prices if the lower of the two prices strictly exceeds c , because the firm naming the higher price makes profit 0 but would make a strictly positive profit by naming a price halfway between c and the price named by its rival.
- It is not an equilibrium for the two firms to name the same price less than c , because then each firm takes a loss but would make 0 profit by naming price c instead.
- It is not an equilibrium for the two firms to name the same price more than c . In this case, let p be the price each firm names, and X the total demand at that price. Each firm gets half the demand, for a profit of $(p - c)X/2$. If one firm instead names the price $(3p + c)/4$, it has demand at least X (it gets all the old demand and more) and its profit per unit is $(3p + c)/4 - c = 3(p - c)/4$, for a profit of more than $3(p - c)X/4 > (p - c)X/2$.

This rules out everything except each firm naming the price c . And this is an equilibrium: Each firm makes profit 0, but if a firm raises the price it names, it gets no demand so continues to make profit 0; if it lowers the price it names, it loses money.

Whether in the extreme case of $b = 1$ or not, the simultaneous-price equilibrium is called the *Bertrand equilibrium*, named for Joseph Louis Francois Bertrand, who wrote (in 1883) that simultaneous-price competition was a more reasonable parable for real-life competition than Cournot's story of simultaneous choices of quantity.

Bertrand–Von Stackelberg and Comparisons

A final variation is if firm A names its price p_A and firm B responds. This is known as *Bertrand–von Stackelberg*, because it mixes price or Bertrand competition with one firm moving first, Von Stackelberg style. This is part d of Problem 2.7, and the backward induction analysis gives firm A naming the price

$$p_A^b = \frac{2a - ab - ab^2 + bc + 2c - b^2c}{2(2 - b^2)} > p_B^*$$

and firm B responding with its best response price.

This gives us four different “predictions”—five if you count the prediction that the firms settle on the symmetric collusive outcome. To compare them, Table A3.1 provides quantities, prices, and profit levels for both firms, for each of the five predictions. In panel a, these values are given for $a = 10$, $b = 0.5$, and $c = 2$; panel b gives the values for $a = 10$, $b = 1$, and $c = 2$.

(For Bertrand–Von Stackelberg and the case $b = 1$, I provide no data. If you evaluate p_A^b for $b = 1$, you get c , and it is a Nash equilibrium for firm A to name the price c and for firm B to respond with that price, but this is a very degenerate situation, and one should not put too much credence in this as the solution.)

	firm A quantity	firm B quantity	firm A price	firm B price	firm A profit	firm B profit
collusive	2.667	2.667	\$6.000	\$6.000	\$10.667	\$10.667
Cournot	3.200	3.200	\$5.200	\$5.200	\$10.240	\$10.240
Bertrand	3.556	3.556	\$4.667	\$4.667	\$9.481	\$9.481
Cournot–von Stackelberg	3.429	3.143	\$5.000	\$5.143	\$10.286	\$9.878
Bertrand–von Stackelberg	3.333	3.619	\$4.857	\$4.714	\$9.524	\$9.651

(a) Parameter values $a = 10$, $b = 0.5$, and $c = 2$

	firm A quantity	firm B quantity	firm A price	firm B price	firm A profit	firm B profit
collusive	2.000	2.000	\$6.000	\$6.000	\$8.000	\$8.000
Cournot	2.667	2.667	\$4.667	\$4.667	\$7.111	\$7.111
Bertrand	4.000	4.000	\$2.000	\$2.000	\$0.000	\$0.000
Cournot–von Stackelberg	4.000	2.000	\$4.000	\$4.000	\$8.000	\$4.000

(b) Parameter values $a = 10$, $b = 1$, and $c = 2$

Table A3.1. Comparing the different predictions. For two sets of parameter values, this table gives the quantities, prices, and profit levels of the two firms, assuming (symmetric) collusion, Cournot equilibrium, Bertrand equilibrium, Cournot–Von Stackelberg, and (for the first set of parameter values) Bertrand–Von Stackelberg equilibrium.

It is worth observing, for $b = 0.5$, that in Cournot–von Stackelberg, the leader Firm A does better than in Cournot, while Firm B does worse. But in Bertrand–von Stackelberg, while both firms do better than in Bertrand, the *follower* Firm B has a higher profit than does the leader Firm A. If you can figure out why this is, you will have learned a lot about the differences between price and quantity competition.

Is Nash Equilibrium Analysis Appropriate? Repeated Play

It is hard to think of any real-life industries that mimic the conditions of these four parable settings. But suspend your disbelief a bit longer and imagine an industry that conforms to one of these stories. The question remains, Is Nash equilibrium analysis appropriate? In this imagined situation, would we expect the two firms to behave in accordance with the Nash equilibrium highlighted?

Recall from Chapter 2 in the text that, to apply Nash equilibrium analysis, there should be some reason to suspect that the firms would find their way to the Nash equilibrium. It could be a matter of logic, general experience with this sort of situation, or experience with the specific opponent. In the context of these models, experience with the specific rival is generally the story told. Indeed, the literature concerning the Cournot model and equilibrium contains a sizeable subliteration on what is called the *stability of Cournot tâtonnement*, which is fancy talk for suppose that, at date t (think week t in an infinite sequence of weeks in which the two suppliers bring their product to the market square), each firm chooses its best response to what its rival did at date $t - 1$. Does this process converge to the Cournot equilibrium? (The answer is Yes, in the context of the specific linear example of Problem 2.7.)

But, if the two duopolists compete with one another week after week, if they do not discount the future too heavily, and if they can observe what the other party brings to the market each week, then the folk theorem kicks in and they can do a lot better (jointly) than by playing the static Cournot–Nash equilibrium quantities prescribed by our analysis.

If the two compete repeatedly, each forms beliefs about how the other side responds to moves the first makes. They can sustain the cartel outcome as a Nash equilibrium, they can sustain repeated play of the Cournot quantities, and they can sustain a lot besides. (The literature on oligopoly contains concepts such as “conjectural response” and “kinked demand” equilibria, which explore how the conjectures each side has about how the other responds affects the long-run equilibrium they reach.)

My point is that the natural rationale for looking at Nash equilibria in this context—that the two firms interact repeatedly—contains a potential internal contradiction. If the firms interact repeatedly, the sorts of ideas encountered in Chapter 3 apply and much more than the specific equilibria identified in Problem 2.7 pass the test of being Nash equilibria.

The Use of These Models by Economists

Notwithstanding what I just said, economists continue to use the classic models in at least two ways.

1. They are used as building blocks in models that tackle more complex problems, such as the proliferation of product variety, the value of excess capacity, and the adoption of new technology. Ultimately, if one believes (as economists do) that product variety, capacity, and adoption decisions

are taken by firms on a “rational” basis, then one needs a model of what would happen (what would be the profit levels of firms) if a new product were introduced, capacity were added, or a new technology were adopted. The classic models satisfy the need for such models. Moreover, they give some flexibility in these models: The Cournot model provides an example of so-called strategic substitutes, where a more aggressive stance by one firm engenders a less aggressive stance by its rivals; while the Bertrand model has so-called strategic complements, where more aggressive behavior by one firm provokes more aggressive behavior by its competitors. This fundamental difference in the nature of actions and reactions permits economists to see how that difference affects things like product variety, capacity, and technology adoption.

2. Empirical economists use the models as concrete parametric specifications to estimate what goes on in specific industries. The analogy here is to linear and constant-elasticity demand functions. No one seriously contends that the demand in a given industry is precisely linear or has constant elasticity, but it may be approximately so; these parametric models are used for statistical estimation, where the test of the model is how well the model fits the data. The data from a real-life industry can be fitted to a model that assumes, say, a Cournot equilibrium; whether this is a helpful exercise or not is an empirical question.

Speaking personally, I do not regard these models as likely to be descriptive of real-life industries. I believe that the discussion of collusion in Chapter 3 tells you more about the real world of oligopoly than the classic models. That is why, in the text, I have no chapter on oligopoly per se but instead one on collusion and cooperation. My own bottom line is that economics is unlikely to reach a point where we have the sort of precision in our predictions about oligopolies that we have about competitive markets; this is simply the nature of this particular beast. But my opinion is nothing like the “median” opinion of economists, and in your continued study of economics and related subjects, especially business strategy and corporate finance, you are likely to encounter these classic models.